

Optimizing networks

for business application performance with next-generation application switches

With the increasing use of network-based applications to drive business efficiencies, enterprise IT departments must support increased network traffic and server load while facing numerous security challenges. In the recent past, IT departments could utilize brute force approaches to solve common problems—adding more bandwidth to relieve congestion, adding more servers to improve application performance, and buying more equipment than needed to meet projected traffic growth. With stagnant or shrinking IT budgets, however, the brute force days are over and enterprises must maximize the use of existing networking infrastructure to get the most out of every IT dollar spent.

Application switches are an important tool that IT departments can use to maximize the return on investment in IT infrastructure, ensuring that networks are optimized for application performance without breaking the budget. In addition to Layer 2-3 services, application switches use sophisticated traffic management techniques to direct traffic to the optimal server or network device based on characteristics of the application and user, health of the data path, and business/network policies. In addition, these switches can provide enhanced security, quality of service, and bandwidth management.

This paper introduces application switching, highlights important aspects of application switching technology, and describes the innovative design and unique capabilities of Nortel Networks purpose-built application switches—the Nortel Networks Alteon Application Switch family. The Alteon Application Switch family is a key enabler of Nortel Networks enterprise vision: One Network, A World of Choice. With the One Network vision, Nortel Networks seeks to accelerate customers' business success by improving productivity and lowering costs through application delivery and convergence, user mobility, infrastructure security, and network resilience.

1.0 Application switching evolution

Web switches (or Layer 4-7 switches) were originally conceived to meet the load balancing and traffic management needs of rapidly growing Internet Web sites. These switches replaced separate devices such as load balancers, global load balancers, bandwidth managers, and multiple layers of local area network (LAN) switches, simplifying data center management and reducing capital and operating expenses.

Application switches are the next generation of Web switches, designed to optimize enterprise application performance. The Alteon Application Switch family was built from the ground up to improve the reliability, security, and efficiency of enterprise business applications through application-intelligent traffic management features. In addition, the Alteon Application Switch includes integrated security applications for enabling secure remote connectivity to applications.

Today, a number of factors are driving the need for application switches, including:

Proliferation of network-based enterprise applications—Enterprise software vendors are increasingly Web-enabling enterprise resource planning (ERP), customer relationship management (CRM), inventory, procurement, human resources, and many other applications for ease of use across the enterprise. These applications, because of their critical nature, require networks and server farms that are highly reliable, secure, and scalable.

Ubiquitous access to applications—Enterprise IT departments must provide secure access to many critical business applications for mobile and remote employees, customers, and business partners. Some application switches can enable high-performance secure access to applications through support for secure sockets layer (SSL) acceleration. High-end application switches take SSL support a step further with SSL virtual private networks (VPNs), enabling secure remote access for a greater number of applications and network resources.

Continued growth of Internet traffic—Despite the global economic slowdown, Internet usage continues to grow. Worldwide Internet traffic is expected to more than double annually during the next five years and the number of Internet hosts and active connections are also expected to increase significantly.¹ Increased use of the Internet will result in increased strain on servers, firewalls, and other network devices. Application switches can reduce the strain and increase performance of these devices, even as demand grows.

Security threats—Unwanted intruders, viruses, denial of service attacks, and application abuse are constant threats to critical enterprise applications. Application switches provide multi-layer security while ensuring that application performance does not suffer. By using Layer 4-7 intelligence, application switches can complement general-purpose firewalls through the use of more granular, application-specific security features and Secure Fault tolerant Web Services applications.

Increased use of multimedia applications—The use of multimedia on enterprise networks is growing rapidly. Real Time Streaming Protocol (RTSP) and Voice over IP (VoIP) protocols use separate channels for transmitting control and data traffic between a client and a server. The specific sockets used for the data transmission channels are generated dynamically and communicated between clients and server over pre-established control channels. To properly route these applications to the right servers, application switches must parse the control channels to extract the dynamic socket numbers of the data channels, so related control and data channels can be processed as a single, logical session.

Budget pressures—With stagnant or declining budgets, enterprise IT departments must squeeze as much out of existing networking infrastructure as possible while laying the foundation for exciting new applications in the future. Through intelligent traffic management, application switches allow enterprises to maximize utilization of existing infrastructure and implement new applications as business requirements dictate.

In short, enterprise IT departments need to ensure that business-critical applications perform well, remain always-on, and scale to meet increasing numbers of employees, partners, or customers who interface with the applications. The applications must also be kept secure from unwanted intruders or application abuse. And, this has to be done in a time of tight budgets, where every IT dollar spent satisfies stringent ROI analyses. Thus, any solution must simplify rather than further complicate operations. As this paper explains, application switches address all of these issues in a number of innovative ways.

¹ Worldwide Internet Traffic, Volume 2, Number 2; Probe Research; April 2002.

Alteon Application Switch

Server load balancing <ul style="list-style-type: none"> • Server load balancing • Global server load balancing • Application load balancing and health checks <ul style="list-style-type: none"> - Domain Name Server (DNS) - Real Time Streaming Protocol (RTSP) - Lightweight Directory Access Protocol (LDAP) • IP, FTP, and more 	Network device load balancing <ul style="list-style-type: none"> • Firewall • Virtual Private Network (VPN) • Intrusion Detection System (IDS) • WAN Links • Wireless Application Protocol (WAP) gateways 	Application redirection <ul style="list-style-type: none"> • Web site • Cache • SSL acceleration appliance • Streaming media • Web Services
Advanced filtering <ul style="list-style-type: none"> • Layer 2-7 attributes • VLAN filtering • Accept, Deny, Network Address Translation (NAT), Redirect 	Content intelligence <ul style="list-style-type: none"> • Layer 7 inspection • URL, HTTP header, cookie • User agent (PDA, browser) 	Embedded security services <ul style="list-style-type: none"> • Access control • Advanced Denial of Service protection (TCP, IP, UDP, ICMP) • Protection from application abuse • Integrated SSL acceleration and SSL VPN
Traffic management <ul style="list-style-type: none"> • Bandwidth management and rate limiting • Type of Service (ToS) marking • Peer to Peer application management 	Persistence support <ul style="list-style-type: none"> • Source IP • Source port • Cookies • SSL identifier • HTTP header • HTTP/HTTPS 	Network services <ul style="list-style-type: none"> • Layer 2-3 • NAT • VLAN tagging • Trunk groups

Table 1. Major functions of the Alteon Application Switch

2.0 The benefits of application switching

As highlighted in *Table 1*, the leading application switches can perform a number of important functions that provide benefits throughout an enterprise network. These benefits include:

Enhancing business productivity and simplifying operations via application-optimized networks—Application switches enhance enterprise productivity by performing health checking, intelligent load balancing, and bandwidth management to increase server/device utilization, improve application performance, increase availability, and enable differentiated services. Leading switches support multiple concurrent applications, significantly simplifying network administration. In addition, many application switches support virtual IP addresses, allowing IT administrators to seamlessly place servers and other network devices in and out of service.

Ensuring fail-safe business continuity—Through local and global load balancing with sophisticated health checking, application switches eliminate single points of failure in a network, provide device and application failover, and enhance overall application availability. Application switches allow the addition of servers or network devices to a network with no network or application downtime.

Protecting via high performance multi-layer security—Leading application switches can accelerate security applications such as firewalls, IDS, and VPN as well as utilize sophisticated filtering at Layers 4-7 to control network traffic and protect applications from abuse or denial of service attacks.

Providing scalability by design—Application switches allow IT administrators to transparently add new application servers or networking devices to load balanced clusters as capacity and

performance requirements dictate. Instead of huge upfront infrastructure investments, enterprises can scale gracefully based on business requirements.

Maximizing return on IT investment—Application switches maximize ROI by allowing enterprises to use existing IT infrastructure more effectively, optimize network design, and simplify operations. The switches enable a reduction in capital and operating expenses even as network performance is increased.

3.0 Application switching technology considerations

To obtain the greatest benefit from application switching, there are a number of key points that enterprises should consider when evaluating application switching platforms:

- Does the switch provide concurrent support for a broad range of features and applications? Enterprises should use application switches that can support multiple applications concurrently. This reduces network complexity by eliminating the need for individual devices to support individual applications. It also gives enterprises flexibility to meet constantly changing business and technical requirements.
- Is the architecture "future proof"? Inevitably, enterprises will require support for new protocols and new applications to optimize their networking and server infrastructure. Application switches should have the processing horsepower and memory to support application development far into the future.
- What is the switching performance? To fine-tune the performance and efficiency of business applications, granular information (e.g., Layer 7 information) about those applications is often required. Processing this information requires processor-intensive deep packet inspection and the flexibility to deal with multiple protocols. Application switches should have the computational horsepower required to process a large number of Layer 4-7 sessions per second for optimal network and application performance.

- Are the integrated applications business class? Some application switches support integrated applications such as SSL acceleration. If an enterprise requires the integrated application, the IT department should determine if these applications are truly business class with acceptable performance, robust feature sets, and the ability to seamlessly plug and play external appliances as required for future scalability.
- How does the switch address enterprise security requirements? In today's environment of hackers, viruses, and other security threats, an application switch should be an integrated part of an enterprise's security strategy. Application switches should support multi-layer security including security acceleration, denial of service protection, filtering, and protection from application abuse. The switches should leverage Layer 7 application intelligence to complement general purpose firewalls by providing application level security.

4.0 Alteon Application Switch

The Alteon Application Switch is based on an innovative distributed processing architecture and supports the broadest range of high-performance traffic management and control services via Alteon OS Traffic Management Software. The Alteon Application Switch extends Nortel Networks award-winning Alteon switching portfolio, which includes the Alteon Web Switch ACE director and 180 Series as well as the Passport 8600 Routing Switch with Alteon Web Switching Module, a Layer 2-7 modular switch.

Section 4.0 describes the architecture and capabilities of the Alteon Application Switch.

4.1 Architecture introduction

The Alteon Application Switch builds on the success of the market-leading Alteon Web Switches and drives the market forward in a number of key areas:

- Introduces high port density in a small form factor, with models supporting 10 to 28 ports in a single rack unit
- Provides integrated secure sockets layer (SSL) acceleration
- Supports SSL VPN for clientless remote access to applications
- Protects network investment by extending the life of existing server and network infrastructures while also providing continued performance headroom for innovative software application and feature development.
- Utilizes Alteon OS, which includes the features of Alteon Web OS 10.0 plus enhancements that add multi-layer security to networks through a host of features such as comprehensive Denial of Service (TCP, IP, UDP, ICMP) protection, intrusion detection system (IDS) load balancing, port mirroring, bandwidth management, and Peer to Peer application management.
- Provides the market's first Web Services-aware specialized traffic management features that enable secure, fault-tolerant Web Services.
- Provides the market's most powerful Layer 4-7 switch with four times the capacity, 2.5 times the performance of the Alteon Web Switch, and three to four times the performance of competitor switches, enabling deep packet inspection without adding latency to the network (Tolly, Jan '03).

Alteon Application Switch family highlights				
Alteon Switches	2424	2424-SSL	2216	2208
Total Ports	28	28	18	10
10/100 Ethernet Ports	24	24	16	8
Gigabit Ethernet Ports	4	4	2	2
IP Routing Interfaces	256	256	256	256
Virtual Server Support	1,024	1,024	1024	1024
Real Server Support	1,024	1,024	1024	1024
Policy Filters	2,048	2,048	2,048	2,048
Concurrent Sessions	2M	2M	1M	600K
Layer 7 Performance (sessions/second)	up to 51K *	30K *	15K*	
Layer 4 Performance (sessions/second)	>64K.*	40K *	20K *	
Integrated SSL Acceleration (tps.)**	No	Base: 300 Max: 1000	No	No
Integrated SSL VPN	No	Yes	No	No

* Real-world testing with zero session loss ** Real-world testing

Table 2. *Alteon Application Switch family highlights*

Table 2 highlights additional details on the Alteon Application Switch.

By taking advantage of technological advances in processor performance, the Alteon Application Switch uses switch processors (SPs) for Layer 2-7 processing. Associated with each SP are a large bank of SDRAM and a cache that holds the Forwarding Database as well as the Address Resolution Protocol (ARP) and Layer 3 routing tables.

Alteon Application Switches also include a Management Processor (MP) that supports the Alteon OS management plane functionality and computing that is not related to the data path. Like the SPs, the MP has access to a large bank of SDRAM and a cache. The entire switch has access to the MP cache and is used, where applicable, for aggregating the contents of the SP caches for availability across the entire switch.

In addition to the SPs and MP, the Alteon Application Switch architecture includes provisions for application processing within the chassis to accommodate compute-intensive applications such as SSL. Application processing within the switch simplifies network implemen-

tation and management for IT departments. The Alteon 2424-SSL is the first switch in the Alteon Application Switch family to take advantage of this integrated application processing, with support for both SSL acceleration and SSL VPN. Section 4.4 details the integrated SSL applications on the Alteon 2424-SSL.

To ensure the best of both centralized and distributed processing models, the Alteon Application Switch uses a next generation version of the proven Alteon Virtual Matrix Architecture—VMA. VMA ensures that all switch processors and the application processors share the processing load for optimal use of processing resources. Section 4.5 discusses VMA in additional detail.

VMA improves on the concurrent session capacity and sessions per second performance of the Alteon Web Switches. It also allows the Alteon Application Switch to outperform the competition. Third-party testing performed by Tolly Labs shows that the Alteon Application Switch has greater than three times the Layer 7 sessions per second performance of its closest competition.

4.2 Enabling a Dynamic Data Path

Dynamic Data Path technology—inherent in all Alteon switches—allows flexible adaptation of the network to achieve business requirements, reducing the need to adjust business applications to compensate for network limitations.

In a pure routing environment, a packet's destination is a foregone conclusion. Alteon Application Switches, on the other hand, use Dynamic Data Path technology to monitor the health and performance of the entire data path while simultaneously inspecting and classifying packets using information such as TCP port, URL, HTTP headers, and cookies. With information on the health of the data path and on the specific session being processed, Alteon Application Switches can dynamically alter the data path (locally or globally) to:

- Direct traffic to the “healthiest” or highest performing server or network device
- Redirect traffic transparently to out-of-path devices such as caches
- Apply different service levels and bandwidth policies based on attributes such as application, user, and access device.

Benefits of Dynamic Data Path technology include high availability for business continuity, network optimization, the ability to differentiate service levels to meet business objectives, and persistent application support.

4.2.1 Dynamic Data Path enables fail-safe business continuity

Alteon Application Switches provide non-stop access to business applications, ensuring easy-to-implement redundancy across the network from WAN interfaces to security infrastructure and host servers. Because the switches can use Dynamic Data Path technology to

Alteon Application Switch family health checking support

Server	UDP-based DNS	RADIUS
Direct Server Return (DSR)	FTP	HTTPS/SSL
	TFTP	
Link	POP3	WAP Gateway (WSP, WTLS)
TCP	SMTP	LDAP
ICMP	IMAP	Any application via script-based health checking
HTTP	NNTP	

Table 3. Alteon Application Switch family health checking support

intelligently direct traffic across servers or network infrastructure, a link, port, hardware device, or service failure will result in sub-second detection and failover to an alternate healthy data path.

For business applications and content that require site (physical location) redundancy, Alteon Application Switches support global server load balancing (GSLB). With GSLB enabled, the switch routes requests to the best site among multiple globally distributed sites based on server health, proximity to the client, and response time. Switches supporting the same server farm exchange performance information, ensuring that requests are delivered to the site with the best performance.

To maximize the effectiveness of Dynamic Data Path, Alteon Application Switches allow customization of health checking to ensure application availability. As highlighted in *Table 3*, the switch supports sophisticated server, link, and application health checking and allows user-scriptable health checks to enable a sequence of tests to verify application and content availability. If a switch detects failure, no new connection requests are directed to the failed server or device. However, if an individual service fails on a server, Alteon Application Switches remove the individual service from the load balancing algorithm without affecting other services provided by the server.

As an example, LDAP is a software protocol for enabling anyone to locate organizations, individuals, and other resources such as files and devices in a network. A standard way of determining the status of LDAP servers is to perform TCP health checks. However, this type of health checking does not go far enough for LDAP because it can only determine that the server is operational and that a process is listening on the port. If the LDAP process is hung, the server could respond to a TCP health check as if everything was operational. To ensure that LDAP servers are operating, an Alteon Application Switch can use special LDAP health checking, a more robust method of checking LDAP server health. It then load balances across the available LDAP servers using the specified load balancing metric. Benefits of using the Alteon Application Switch for LDAP load balancing include high availability and optimization of LDAP server resources.

True high availability requires two separate switches working in unison to ensure full failover without a single point of failure. Alteon Application Switches enable high-availability architectures via support for an advanced implementation of the Virtual Router Redundancy Protocol (VRRP). Alteon Application Switches support active-active, active-standby, and hot-standby modes. Many competitive solutions deliver half the performance in high availability designs because one switch

sits idle, waiting for a failure. In contrast, the Alteon Application Switch active-active mode allows enterprises to maximize investment by leveraging the power of both switches during normal operations.

4.2.2 Dynamic Data Path provides network optimization

Using Dynamic Data Path technology, Alteon Application Switches can optimize networks for application performance, making networks and data centers more productive for the business and eliminating wasteful use of resources. Alteon Application Switches can use Dynamic Data Path technology in many ways to achieve network optimization. Several example uses follow.

Example: Flexible content location

IT administrators can increase server utilization, extend server life, and reduce costs by eliminating the need for all servers to support all content. As highlighted in *Figure 1*, by making switching decisions based on URL, an Alteon Application Switch makes it possible to host content on servers best tuned for

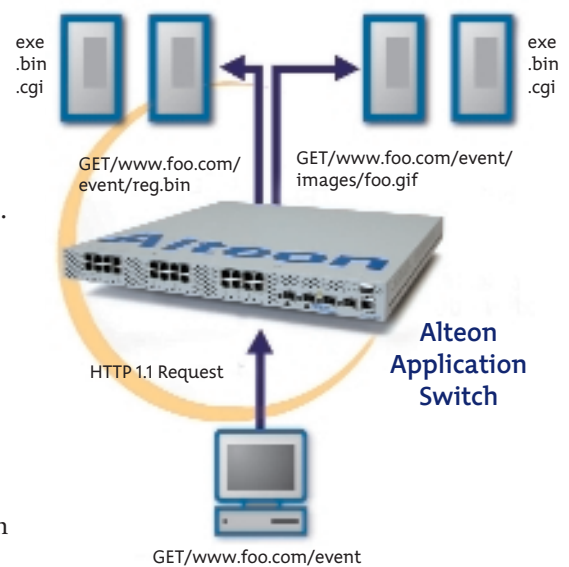


Figure 1. Dynamic Data Path technology enables application awareness for flexible content location

the content—heavy-duty applications on high-end servers and static content on low-end servers. Similarly, administrators could host content on different servers depending on the access device used—PC, personal digital assistant, or WAP-enabled phone.

Example: Differentiated services

Dynamic Data Path technology also enables differentiated services. An enterprise that wanted to provide a higher level of service and satisfaction to its best customers could use an Alteon Application Switch to differentiate customers using “cookie” information. For example, the most profitable customers could be switched to high-performance application servers to ensure a high-quality experience while standard customers could be switched to standard servers. This is illustrated in *Figure 2*.

Example: DNS optimization

The domain name system (DNS) is the way that domain names are located and resolved into IP addresses. Using Dynamic Data Path technology, Alteon Application Switches can dynamically distribute load among multiple DNS servers using the two forms of DNS queries: UDP or TCP. The switch can send TCP queries to one group of real servers and UDP to another group of real servers, providing load balancing within each group. If required in large DNS server farms, Alteon Application Switches also allow load balancing of DNS traffic based on DNS names. The switch extracts the host name from the query, processes the request, and switches the request to the appropriate real server. *Figure 3* illustrates this scenario. Utilizing DNS server load balancing allows optimization of DNS server farms, reducing the need to duplicate content on every server.

Example: Bandwidth management

Bandwidth management enables IT administrators to allocate a portion of the available bandwidth for specific users or applications so that critical business traffic receives higher priority than non-critical traffic. Administrators can allocate and meter bandwidth based on filtering rules, URLs, HTTP headers, and cookies. With URL-based bandwidth management, for example, IT administrators can allocate bandwidth based on traffic type (e.g., static HTML, graphics, and dynamic transactions) or give priority to specific applications. With cookie-based bandwidth management, IT departments can allocate more bandwidth to traffic from the most important users or prevent network abuse by allocating less bandwidth to “bandwidth-hogging” users. Alteon Application Switches provide IT administrators with granular control of bandwidth via two types of bandwidth management:

- Bandwidth management with traffic shaping
- Bandwidth management without traffic shaping (rate limiting)

Example: Streaming media

As a final example of Dynamic Data Path network optimization, Real Time Streaming Protocol is an application-level protocol for control over the delivery and presentation of data (audio, video, and text) with real-time properties. It was designed to efficiently deliver streaming media over IP networks and acts as the “network remote control” for multimedia servers, synchronizing and controlling the flow of multimedia streams. As shown in *Figure 4*, utilizing Dynamic Data Path technology, an Alteon Application Switch can load balance




Figure 2. Dynamic Data Path technology enables user awareness for differentiated services

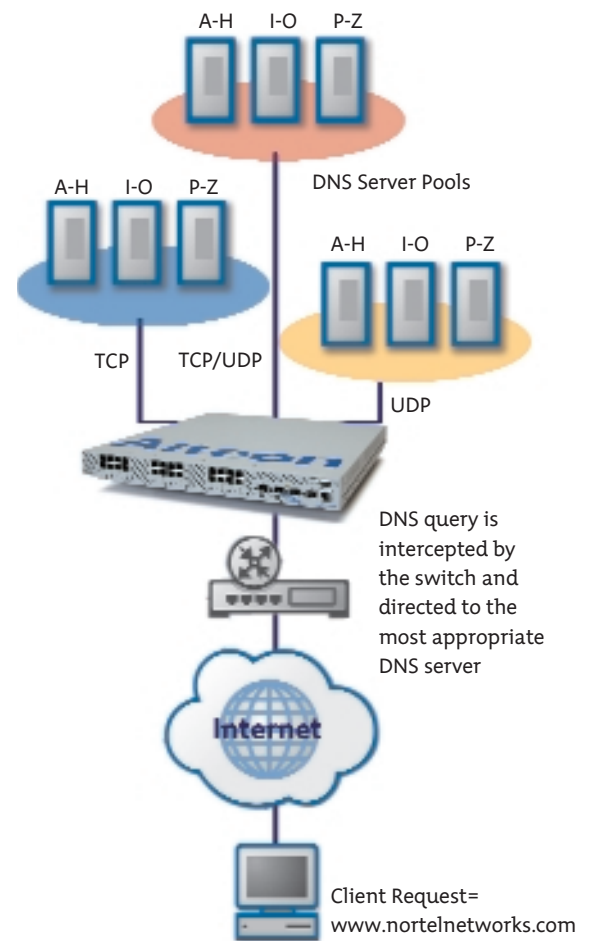


Figure 3. Dynamic Data Path technology for DNS server load balancing

RTSP requests across suitable servers or caches to improve the robustness and security of RTSP streaming implementations. RTSP load balancing with Alteon Application Switches allows private addressing on the server side and eliminates the need to open all UDP ports. It supports all Layer 4 load balancing metrics and Layer 7 URL-based load balancing (hashing or pattern-matching).

4.2.3 Persistent application support

Many applications—such as payment transactions, search displays, shopping carts, and multi-page forms—require persistent connections. This means a client must “talk” to the same real server for the duration of the transaction. If client-server association is not persistent, applications may break, resulting in lost employee productivity or disgruntled customers and partners. Dynamic Data Path technology on the Alteon Application Switch enables persistent connections and thus, transaction integrity, by using cookies embedded in HTTP sessions (or SSL session identifiers in secure HTTPS sessions) to accurately associate consecutive requests from a client with the same server.

4.3 Multi-layer security

Nortel Networks Unified Security Architecture is a holistic approach to securing the entire information technology infrastructure, including telephony, voice over IP, data, and converged networks. It provides a comprehensive set of technologies and planning tools to help chief information officers, network planners, and network operators make informed decisions regarding security applications. The Unified Security Architecture addresses each of the many layers of communications infrastructure security and promotes secure convergence and infrastructure simplicity.

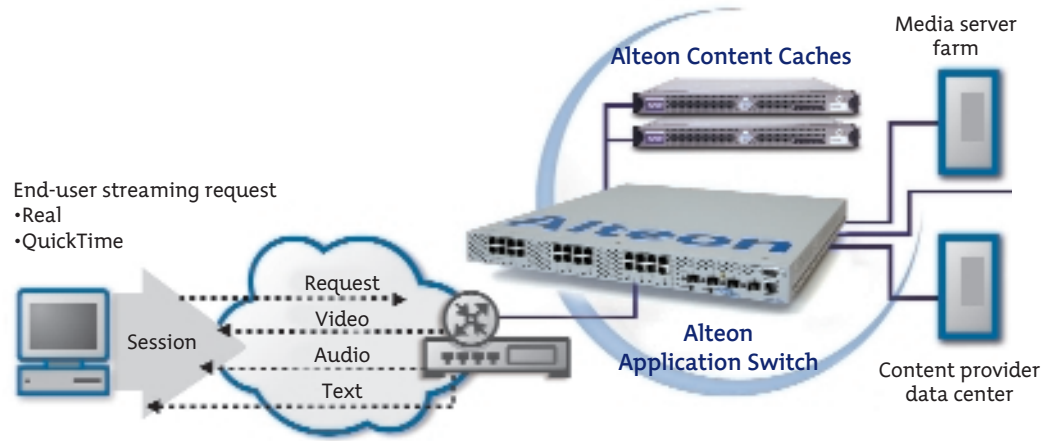


Figure 4. Dynamic Data Path technology enables RTSP load balancing

Alteon Application Switches play a key role in the Unified Security Architecture due to inherent multi-layer security features designed to protect against external and internal threats without sacrificing performance. General purpose firewalls provide security by allowing or denying traffic based on attributes including port and IP address. Alteon Application Switches complement these solutions by utilizing Layer 4-7 intelligence to add more granular, application-specific security.

Multi-layer security features include security acceleration, extensive network traffic control, and comprehensive

protection from denial of service (TCP, IP, UDP, ICMP) attacks and application abuse. The switches allow enterprises to simplify security implementations and scale efficiently to meet traffic and application requirements. Section 4.4 addresses integrated SSL acceleration and SSL VPN features—two additional security solutions available on the Alteon Application Switch.

4.3.1 Security acceleration

Alteon Application Switches allow enterprises to maintain strict network security without performance degradation through services such as firewall load balancing, VPN load balancing,

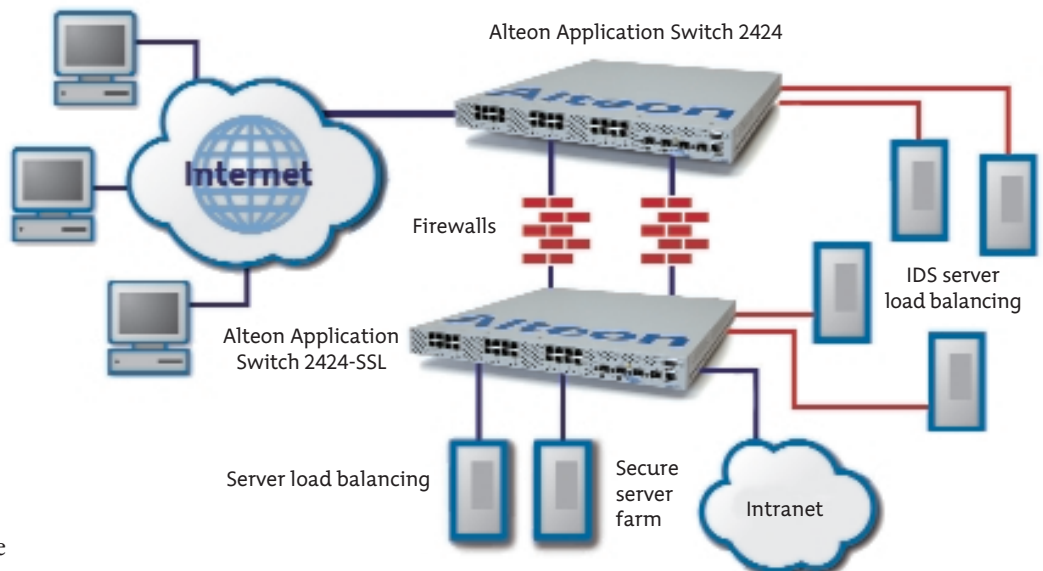


Figure 5. Security acceleration increases the performance and reliability of security applications

and IDS load balancing. These services enable enterprises to maximize utilization of security devices, accelerate performance without forklift upgrades, and eliminate security devices as single points of failure. To maintain strict security, Dynamic Data Path technology ensures that all traffic between source and destination IP address pairs flows through the same security device for the duration of the session. Health checking is performed on the entire data path to ensure highly reliable security solutions. *Figure 5* illustrates firewall and IDS load balancing.

Intrusion detection systems are critical for identifying and preventing network intrusions and attacks. IDS load balancing helps scale these systems because it is difficult for individual IDS servers to scale to the gigabit speeds of the network. During IDS load balancing, an Alteon Application Switch forwards incoming packets to an IDS server at the end of the filtering process or at the end of client processing (when filtering is not enabled). For each session entry created on the switch, an IDS server is

selected based on the server load balancing metric or the IDS hash metric in the filter menu.

Alteon Application Switches can support multiple IDS systems simultaneously. This capability is required in enterprises that use multiple vendors to perform tasks based on the strengths of each. For example, an IDS from Vendor A may be used to alert for network intrusions while an IDS from Vendor B monitors for application attacks.

Finally, Alteon Application Switches provide IDS support without impacting the ability to perform port mirroring. If an IT administrator needs to troubleshoot a network problem, he/she can do so without having to turn off IDS. This is an important security consideration not supported by all application switching vendors.

4.3.2 Network traffic control

Alteon Application Switches provide extensive network traffic control through NAT and powerful filtering capabilities. These capabilities allow Alteon Application Switches to offload firewalls from

some tasks, enabling a more efficient “DMZ” for business applications and allowing IT departments to maximize the use of existing firewalls. Alteon Application Switches support up to 2,048 filtering rules per switch. This provides administrators with extensive control of the traffic on their networks, allowing customization of the Dynamic Data Path technology to meet business requirements. Filters can be configured to allow, deny, or redirect traffic based on application type, protocol, IP source/destination addresses, Layer 7 attributes (e.g., URL, cookie, HTTP header), and VLAN ID.

4.3.3 Denial of Service protection

Alteon Application Switches can thwart advanced Denial of Service (DoS) attacks and TCP SYN attacks without blocking valid session requests. Through “delayed binding,” highlighted in *Figure 6*, client SYN requests are intercepted before they reach the server. The switch then responds to the client with a SYN ACK and does not allocate a session until a valid SYN ACK is received from the client or a valid three-way handshake is

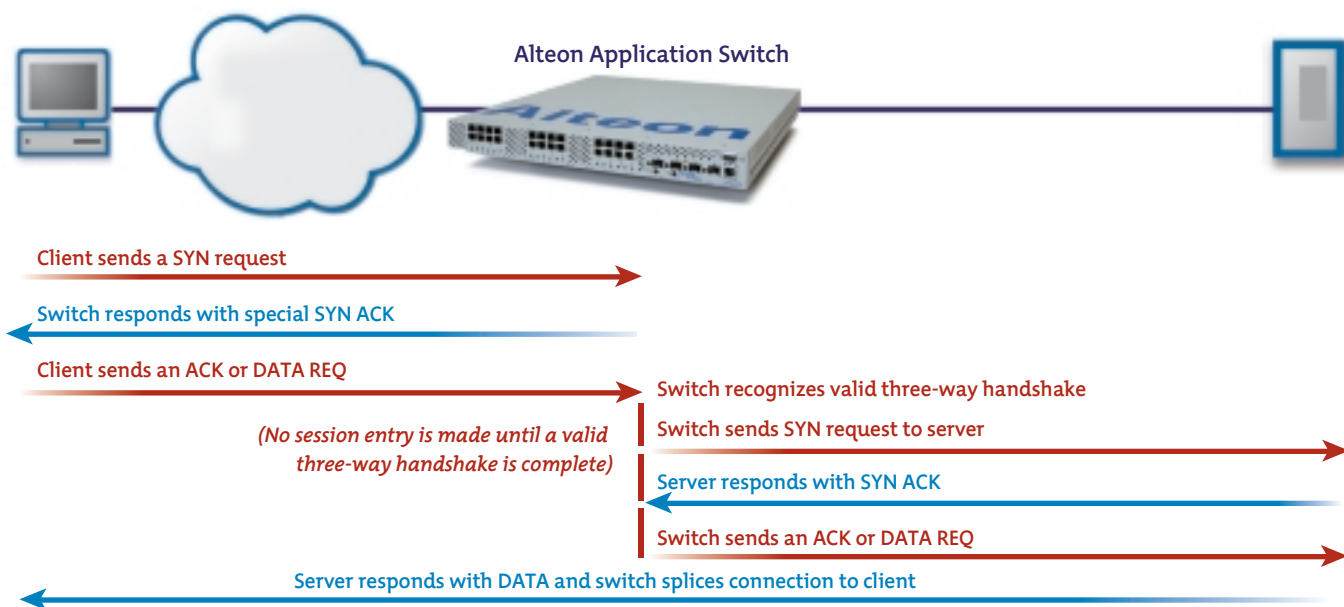


Figure 6. Delayed binding thwarts TCP Denial of Service attacks

complete. By temporarily terminating each TCP connection until content has been received, the switch prevents the server from being inundated with SYN requests. Half open sessions are the result of incomplete three-way handshakes between the client and the server. To detect SYN attacks, Alteon Application Switches track the number of new half open sessions over a period of time. If the value exceeds a specified threshold, the switch triggers a trap to notify the administrator. Using sophisticated pattern matching, Alteon Application Switches protect networks and servers against a whole host of availability-based attacks, such as ping of death, Fraggle, Pepsi, etc.

4.3.4 Protection from application abuse

Alteon Application Switches protect applications against abuse by allowing administrators to limit the rate of new TCP connections on a per-client basis, as illustrated in *Figure 7*. This can be used to limit users to a particular connection rate and to limit the number of sessions for users accessing a specific domain or application within the domain. Benefits of this capability include increased control of user access to applications and increased application availability.

In addition, the switches allow IT administrators to create filters and assign URLs to those filters to deny traffic with offending string patterns. This feature is particularly useful for adding protection against viruses such as CODE RED and Nimda and in preventing access to disallowed Web content.

4.4 Application integration

Alteon Application Switches have an architecture that allows for rapid integration of application processors within the chassis. The architecture allows the switches to accommodate compute-intensive applications, simplifying network infrastructure and reducing overall capital and operating expenditure requirements for enterprises. The first two integrated applications supported are SSL acceleration and SSL VPN, based on the capabilities of the Alteon SSL Accelerator, the industry's leading SSL acceleration appliance (Infonetics). These applications are supported in the Alteon 2424-SSL.

4.4.1 SSL acceleration

With its unique ability to set up secure sessions at the application layer between any client and server connected to the Internet, SSL has quickly become the

de facto standard for securing Internet communications. However, application servers bear the increased processor load required to handle the secure session setup as well as the bulk encryption duties required by the SSL protocol. These functions can slow application servers to a crawl if many sessions are initiated at the same time or if a large number of concurrent sessions are required. Instead of adding more servers to increase capacity, the Alteon 2424-SSL can be used to offload this duty from servers, resulting in optimum application performance at a fraction of the cost.

Many secure applications are characterized by a high rate of new session adds that have a limited duration (e.g., customers entering information such as credit card numbers and personal data). The integrated SSL acceleration capability on the Alteon 2424-SSL has been designed for these environments, with support for up to a maximum of 1,000 SSL transactions per second (real-world testing), 16,000 concurrent secure sessions, and the option to seamlessly add up to 255 external Alteon SSL Accelerator appliances to scale performance




Figure 7. Alteon Application Switches protect against application abuse

up to 255,000 transactions per second and 4 million concurrent sessions. When attached in a plug-and-play fashion, an external appliance becomes part of the same security cluster as the on-board SSL acceleration processor, minimizing the IT resources required for configuration. *Figure 8* illustrates integrated SSL acceleration on the Alteon 2424-SSL.

Secure environments can become complex when managing multiple certificates and keys across as few as ten servers. In large server farms, the Alteon 2424-SSL can substantially reduce redundant recurring digital certificate costs by aggregating the certificate installations and management functions on a single device. Consolidating the keys and certificates on the Alteon 2424-SSL improves security by providing better protection for private keys and lowers operations and support costs by streamlining SSL infrastructure and simplifying management.

An extremely important feature of integrated SSL acceleration on the Alteon 2424-SSL is end-to-end encryption. Security-sensitive applications in industries such as financial, healthcare, and government services cannot accept the liability of breaking the client-to-server encrypted path. Many traditional SSL acceleration solutions, however, can only be configured to offload back-end servers by terminating SSL sessions and establishing non-secure, clear text sessions with the back-end servers. This presents a security risk as anyone with access to the back-end infrastructure can sniff packets and pull sensitive information such as credit card numbers and passwords. Installing specialized cryptographic cards in the servers themselves leads to increased capital and management costs, server downtime, and a limited ability to perform load balancing and other con-




Figure 8. Alteon Application Switch with Integrated Security Applications SSL VPN and SSL Acceleration

tent-based services. Less sophisticated back-end encryption schemes establish a second SSL session to the server in response to a request. This approach actually involves two SSL session negotiations and becomes “non-accelerated” back-end encryption. The Nortel Networks solution incorporates session reuse, variable encryption, and connection pooling to allow for secure session proxying that accelerates servers without losing client-to-server encryption.

4.4.2 SSL Virtual Private Networks

The Alteon SSL VPN option for the Alteon 2424-SSL (also available on the Alteon SSL Accelerator) allows the switch to function as a secure remote access gateway. The Alteon SSL VPN is a remote access security solution that extends the reach of enterprise applications to mobile workers, telecommuters, partners, and customers. With SSL

as the underlying security protocol, the Alteon SSL VPN allows for truly unrestricted remote access, using the Internet for remote connectivity and the ubiquitous Web browser as the primary client interface.

A primary benefit of the Alteon SSL VPN is ease of deployment. Because SSL is built into Web browsers, SSL VPNs can be implemented without additional client software. This reduces operations costs by reducing the management associated with client-based deployments. SSL VPNs are also network agnostic, able to traverse firewalls that might otherwise block IPSec VPNs. Finally, the Alteon SSL VPN, through granular user authentication, provides application-specific access. This provides greater application access control than a Layer 2-3 VPN, which provides network-wide access.

The Alteon 2424-SSL has a number of key features that were developed specifically for SSL VPNs:

Client authentication—The Alteon 2424-SSL provides both server and client authentication. In an extranet environment, the switch supports RADIUS authentication with a unique user ID and password, certificate-based authentication, or both. Digital certificates can be generated by the switch and validated by issuing a certificate signing request to a Certificate Authority (CA). Alternatively, an enterprise can act as its own CA and validate with its own digital signature. Because SSL VPNs allow access from any Web browser, the capability can be deployed with two-factor authentication such as RSA's SecurID standard.

Session management—Remote users accessing applications through their browsers are subject to the risk of dropped sessions as some browsers refresh sessions after a period of inactivity. This could result in session interruptions if users are directed to different servers. To solve this issue, the Alteon 2424-SSL uses a second-tier virtual IP address structure that enables SSL session persistence in a distributed environment.

Granular access control—Once remote users are authenticated, access to applications must be controlled. In an SSL VPN installation, the Alteon 2424-SSL acts as a proxy server for available applications and can communicate with access control servers to assign privileges for access to specific applications. This allows enterprise IT departments to allow or deny access to enterprise applications on a per-user basis.

Application tunneling—As part of the application proxy function, the Alteon 2424-SSL maintains SSL sessions to the

backend servers, ensuring that data is not open to attack or snooping at any point along the connection.

Multi-protocol support—The Alteon SSL VPN capability on the Alteon 2424-SSL supports a wide range of applications including secure access to e-mail (SMTP-S, POP3-S, IMAP-S), file sharing (FTP-S), Web-enabled applications, client-server applications, and many others.

Application address translation—One challenge involved with extending internal applications to remote users is the translation of internal IP addresses. The Alteon 2424-SSL creates an external virtual IP address that is mapped to the internal IP address, allowing seamless, secure remote access to the internal applications.

Advanced filtering—SSL traffic on Port 443 is often left open on firewalls, presenting a potential security risk from authenticated users with malicious intent. To overcome this risk, the Alteon 2424-SSL utilizes proven Layer 4-7 filtering capabilities so that IT departments can allow or deny access at the application layer to authenticated users based on IP address, requested URL, or cookie information.

Auditing—The Alteon SSL VPN capability on the Alteon 2424-SSL allows creation of detailed activity reports so that IT administrators can track application usage and user attributes. The information can be exported to popular database tools for analysis.

4.5 Virtual Matrix Architecture

Unlike Layer 2-3 switches that are optimized for forwarding independent data packets based on well-defined MAC and IP protocol fields, application switches are designed to perform session-oriented

traffic management where the definition of a session varies with different applications. Classifying traffic by session is frequently more complex than examining TCP/IP protocol headers because:

- Content is non-deterministic. Content identifiers can be of varying lengths and can occur at unpredictable locations within a request. As a result, scanning through an entire request for a specific string is very processor intensive.
- To parse requests, an application switch temporarily terminates the TCP connection from a client, examines the request, and opens up a connection to an appropriate server. This temporary termination is called delayed binding. During this process, the switch must temporarily buffer the request, which uses system memory.
- With delayed binding, two independent TCP connections span the session—one from the client to the application switch and the second from the switch to the selected server. The application switch must modify the TCP header on every packet that travels between the client and the server for the duration of the session. This function—known as TCP connection splicing—is taxing on a switch, particularly when the switch must process thousands of these sessions simultaneously.
- To ensure that all packets within a session are forwarded to the same server, an application switch must maintain session states on every active session. Depending on session duration, the application switch may need to track tens to hundreds of thousands or more active connections, requiring large amounts of memory.

Because session processing is inherently processor and memory intensive, application switches must have a robust architecture that takes full advantage of all processing and memory resources without introducing bottlenecks. Many application switching vendors use one of two different approaches for the architecture: centralized CPU systems or distributed processing systems.

Centralized architectures are flexible, particularly for topologies where traffic enters from a single port (such as a WAN link), because all processing power and memory can be brought to bear on such traffic. However, when session traffic is heavy or when traffic ingresses from multiple ports, all session processing must pass through the central processor, which easily becomes a bottleneck. As a result, centralized architectures lack scalability and are only suited for sites with low traffic expectations and simple traffic management requirements. As switches incorporate more sophisticated functions, this architecture is clearly dated.

At the other end of the spectrum is a distributed processing architecture. A distributed model is fast and delivers excellent performance, since all functions are dedicated to specified ports. It is optimal for topologies where session traffic ingresses from a large number of ports. However, this architecture also suffers from limitations. Because processing and memory are not shared, one port may not be able to use information on another port to make a decision. In addition, IT administrators must carefully architect their networks to avoid using specific ports that might cause processor overload while other processors sit idle.

Alteon Application Switches use VMA to offer the best of both centralized and distributed architectures. With VMA, Alteon Application Switches have the

resource aggregation and flexibility of centralized models and the performance of distributed processing models. VMA is based on the successful Virtual Matrix Architecture used by Alteon Web Switches, but with accommodations for the new high-performance Alteon Application Switches.

VMA is a fast and flexible architecture that makes efficient use of the entire system's capacity while providing the parallel performance of distributed processing. In this architecture, all processors share load but no single processor must see all traffic. When a packet reaches an Alteon Application Switch, an algorithm is applied to the source IP address that uniquely selects one of the high-performance switch processors on the system as a designated processor. The processor selection algorithm ensures that traffic is evenly shared across all processors on the switch. The algorithm is also deterministic—all packets from the same source IP address will always be handed to the same designated processor. VMA also ensures that pertinent traffic is distributed efficiently to the application processors. In the case of the Alteon 2424-SSL, an application processor would handle decryption and encryption of HTTPS traffic.

Once the ingress port hands the received packet to the designated processor, the processor examines its local session table and makes a forwarding decision which may include parsing content, selecting a server, performing network address translation or TCP connection splicing on the packet, metering bandwidth usage, and so on. With an efficient port selection algorithm and high speed intra-switch communications, the added latency of moving a packet from the ingress port to the designated processor is negligible.

On the return path the algorithm is applied to the destination IP address of the packet, resulting in the same designated processor being selected. In this way, the designated processor sees session traffic in both directions and state information can be kept in the designated processor's memory where it can be accessed more quickly. A simplified view of VMA is illustrated in *Figure 9*.

With VMA, each processor handles a roughly equal subset of traffic going through the switch, resulting in an optimal use of processing and memory resources. However, certain information must be available to each processor. This includes packet filtering and content parsing rules and common information used for every forwarding decision. For example, each processor keeps the access control list entries for all ingress ports, allowing it to filter traffic differently for each ingress port.

With VMA, multiple processors work in parallel to process incoming traffic, increasing total performance of Alteon Application Switches. All memory is utilized, dramatically increasing the amount

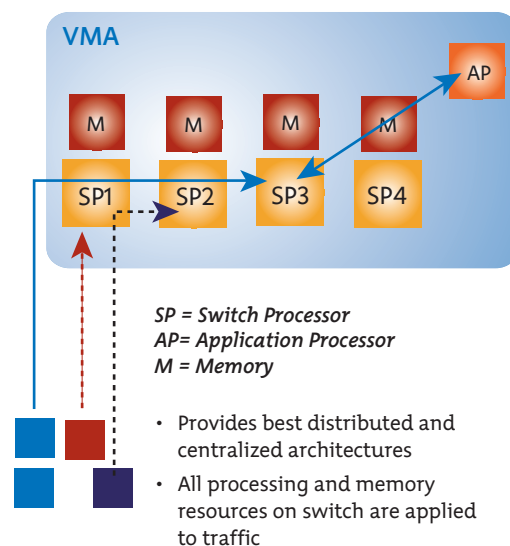


Figure 9. Virtual Matrix Architecture—on Alteon Application Switches

of memory available compared to a traditional distributed model. More memory means more session states stored, allowing session persistence to last longer. It also provides increased buffering for delayed binding and deeper, faster content parsing. Because each processor has access to all content and flow information in the device, topology and connectivity constraints are eliminated.

4.6 Plug-and-play deployment

For IT administrators, the ability to scale networking and server infrastructure and reconfigure the infrastructure to keep up with rapidly changing business requirements is of utmost importance. Alteon Application Switches have a number of key features that enable “plug-and-play deployment,” enabling rapid application set-up, flexibility, and management simplicity. These features include:

- **Multi-application support** enables the use of one Alteon Application Switch for multiple applications—no need for multi-box implementations and the associated management headaches. As an example, one Alteon 2424-SSL switch can support various application-intelligent load balancing applications as well as SSL VPNs to enable secure remote access to business applications.
- **High port density** allows Alteon Application Switches to meet the scalability requirements of growing networks, server farms, and data centers in a small form factor.
- **Virtual IP addresses** allow IT administrators to add capacity without reconfiguring their networks. Alteon Application Switches appear to the network as one or many “virtual servers” each represented by a virtual IP (VIP) address or address range. The switch load balances across

multiple real servers or network devices with real IP (RIP) addresses. By using VIP addresses to represent target real servers or network devices, new servers and devices can be added to a network by simply including the device into the load balance rotation for an existing VIP address. This allows IT administrators to buy what they need, when they need it instead of buying based on fuzzy estimates of future demand. To the users of the network, it’s all transparent. No complicated configuration changes are required.

- **Extensive Layer 2-3 functionality** that allows use of Alteon Application Switches to tie servers into uplink routers, aggregate servers, and network devices, and perform light-weight routing (through support for BGP, OSPF, and RIP).
- **VMA ensures** that all processors on the Alteon Application Switch support all ports. As a result, provisioning is simplified for IT administrators. In effect, the switch provisions itself by distributing load evenly across all processing resources.
- **Sophisticated network management** allows IT administrators to configure and monitor all switch functions via standard Web browsers, SNMP applications, and command line interface (CLI) from the console port or via Telnet. Private MIB and RMON are supported. In addition, a robust port mirroring capability allows administrators to collect detailed data on network performance and usage. The management interface (Alteon Element Management System) is integrated and can be used with Nortel Networks Optivity Network Management System and HP OpenView.

5.0 Maximizing return on IT investment

Alteon Application Switches help enterprises get the most out of every IT dollar they spend, ensuring that networks are optimized for business application performance without breaking the budget. Alteon Application Switches accomplish this by extending the life of existing infrastructure, utilizing assets more efficiently, reducing capital and operating expenses, and supporting multiple applications and functions that contribute significantly to the health and success of the business.

5.1 Reducing capital and operating expenditures

In the recent past, enterprise IT departments used a brute force approach to meet rapidly escalating network requirements—adding more bandwidth to relieve congestion, adding more servers to improve application performance, and buying more equipment than needed to meet projected traffic growth. For many enterprises, this brute force approach has led to overly complex networks that are costly to manage and scale. Alteon Application Switches can bring a new level of cost efficiency to these networks through better utilization of existing infrastructure and more cost-effective scalability for future infrastructure. Alteon Application Switches integrate seamlessly into existing networks to deliver immediate returns by reducing capital expenses and long-term operating expenses. Examples of these savings appear on the following pages.

Improving utilization. Servers are typically deployed on a per-application basis and are often over-provisioned to handle traffic peaks and anticipated future demand. This strategy can lead to poor server utilization, often under 40 percent. Improving server utilization is a

simple step towards improving network efficiency and capturing the value received from existing infrastructure. Alteon Application Switches can increase utilization to more than 75 percent by controlling traffic to and from the servers. Sophisticated filtering and Dynamic Data Path capabilities enable the switch to optimize application-related traffic to the servers while limiting unrelated traffic. As an example, increasing utilization from 40 percent to 70 percent could reduce server requirements and costs by as much as 40 percent. These benefits apply not only to Web and application servers but to firewalls, VPN devices, caches, and other network devices as well.

Deferring capital expenditures.

Moore's Law states that processing power doubles every 18 months. Over three years, that equates to a 300 percent improvement in performance or an equivalent 75 percent reduction in price for a given level of performance. By utilizing Alteon Application Switches for intelligent load balancing, an enterprise can realize significant cost savings by deferring purchases until a lower price point has been reached. Instead of purchasing a large server today that is over-provisioned to meet expected future requirements, an enterprise can purchase smaller, inexpensive servers that cover today's demand and add these smaller servers to a load balanced cluster as requirements dictate. Closer alignment of server provisioning with real-world demand ensures that capital will not be allocated for unnecessary infrastructure, freeing money for other investments, lowering initial capital expenditures, and reducing total cost of ownership over the service period.

Extending network asset life. Alteon Application Switches allow IT departments to optimize network and application performance by using Dynamic Data Path technology to make intelligent decisions based on TCP port numbers, URLs, HTTP headers, and HTTP cookies. By switching based on URL, for example, an Alteon Application Switch makes it possible to host content on servers best suited for the content— heavy duty dynamic applications on high-end servers and static content on low-end servers, which may be older servers that have been re-purposed to handle content for which they are best suited. Alteon Application Switches make it possible to employ servers or other network assets for up to twice the typical life by dynamically directing the most appropriate traffic to the device and by enabling load balancing to improve overall performance using a cluster of devices.

Managing bandwidth. No matter how efficient and well managed the server infrastructure, business application performance is at the mercy of lower level connectivity. Unfortunately, this common infrastructure is shared by many other competing applications, resulting in performance-robbing peaks that can cause lost packets and latency. Instead of over-provisioning costly bandwidth to solve this problem, an enterprise can use an Alteon Application Switch to prioritize traffic based on user or application and throttle back latency tolerant traffic to ensure there is enough capacity for more latency sensitive applications.

Additional detail on the capital and operating expense savings possible from Alteon Application Switches is available in the business case brief entitled "The Alteon Tuned Network, Unlocking Network Value."

5.2 Improving business application performance

In addition to immediate capital and operating expenditure savings, Alteon Application Switches provide benefits that improve business application performance. This translates into higher revenues and reduced costs over time through improved customer satisfaction and employee productivity. Alteon Application Switches improve business application performance by:

- Optimizing network design and reducing operations headaches through support for multiple load balancing, bandwidth management, and security applications from one switch
- Ensuring non-stop access to business applications, thus improving revenues from highly available external applications and increasing employee productivity from greater access to vital internal applications
- Protecting applications from external and internal security threats through substantial multi-layer security features and applications
- Scaling applications efficiently to avoid operations headaches and application downtime
- Allowing IT administrators to adjust network and server infrastructure quickly to ensure that the infrastructure meets rapidly changing business requirements.

5.3 Summary

Through an advanced high-performance architecture and the most comprehensive set of intelligent traffic management features, Alteon Application Switches allow enterprises to optimize business application performance and maximize the return on existing network infrastructure. Alteon Application Switches allow enterprises to take advantage of Dynamic Data Path, sophisticated multi-layer security, and integrated applications to simultaneously accelerate, optimize, and protect vital business applications. The end result is network, server, and application infrastructure that is simple and cost-effective to manage and scale, allowing enterprises to get the most out of existing infrastructure while reducing capital and operating expenditures.

In the United States:

Nortel Networks
35 Davis Drive
Research Triangle Park, NC 27709
USA

In Canada:

Nortel Networks
8200 Dixie Road,
Suite 100
Brampton, Ontario L6T 5P6
Canada

In Caribbean and Latin America:

Nortel Networks
1500 Concorde Terrace
Sunrise, FL 33323
USA

In Europe:

Nortel Networks
Maidenhead Office Park
Westacott Way
Maidenhead Berkshire SL6 3QH
UK

In Asia:

Nortel Networks
6/F Cityplaza 4,
Taikooshing,
12 Taikoo Wan Road,
Hong Kong



Nortel Networks is an industry leader and innovator focused on transforming how the world communicates and exchanges information. The company is supplying its service provider and enterprise customers with communications technology and infrastructure to enable value-added IP data, voice and multimedia services spanning Wireless Networks, Wireline Networks, Enterprise Networks, and Optical Networks. As a global company, Nortel Networks does business in more than 150 countries. More information about Nortel Networks can be found on the Web at:

www.nortelnetworks.com

For more information, contact your Nortel Networks representative, or call 1-800-4 NORTEL or 1-800-466-7835 from anywhere in North America.

*Nortel Networks, the Nortel Networks logo, the globemark design, and Alteon are trademarks of Nortel Networks. All other trademarks are the property of their owners.

Copyright © 2003 Nortel Networks. All rights reserved. Information in this document is subject to change without notice. Nortel Networks assumes no responsibility for any errors that may appear in this document.

GSA Schedule GS-35F-0140L
1-888-GSA-NTEL

NN103321-071403