

Application

Delivery Challenge

MARCH 2006

Produced by
Layland Consulting and Webtorials



Application Delivery Challenge

Introduction
Setting the Stage3



**The Power of Integration:
Secure and Assured Application
Delivery from the Data Center**5



**Managed Application Delivery
Services Reduce the Cost and Complexity
of Application Acceleration**8



**Application Delivery Challenge –
Implementing an Application Front End**11

THE CHALLENGE SERIES

Produced By

Layland Consulting and Webtorials, a venture of Distributed Networking Associates, Inc. Greensboro, N.C. www.webtorials.com

Editor/Publisher

Steven Taylor
taylor@webtorials.com

Design/Layout Artist

Debi Vozikis
dvozikis@rcn.com

Copyright © 2006

Distributed Networking Associates, Inc.

Professional Opinions Disclaimer

All information presented and opinions expressed in this report represent the current opinions of the author(s) based on professional judgment and best available information at the time of the presentation. Consequently, the information is subject to change, and no liability for advice presented is assumed. Ultimate responsibility for choice of appropriate solutions remains with the reader.

Introduction

Setting the Stage

The data center is the heart of the IT infrastructure. It is the repository for enterprise's knowledge. It is the engine that underlies the success or failure of the enterprises' IT strategy. How quickly and securely information flows to and from the data center is the critical factor. This was true in the past, and it is even truer now with data center consolidation. Consolidation makes the data center more important because all the working copies of the enterprise's data now reside there. Security breaches or slow response time has an impact on the entire enterprise. Because of the importance of the data center, network managers must provide the best possible interface between the data center and network.

Network managers need a solution that solves today's challenge. Application Delivery architecture provides the road map for a successful data center and IT implementation. Application delivery architecture moves the focus from switching and routing packets to providing a fast application response time along with security. Network managers must implement an infrastructure that understands the applications. Overtime a full deployed application deliver infrastructure will be present throughout the network but the most important first implementation of the new architecture is in the data center. Fulfilling this critical function is an Application Front End.

An AFE is the next-generation Layer 4-7 switches and Server Load Balancers that links the data center to the network and its users. The AFE builds upon the important features of these earlier solutions of increasing availability and understanding the health and balancing the load between servers and data centers. The AFE takes this tradition and expands the role by increasing security, providing application acceleration and saving money and resources in the data center.

One of the most important new roles for an Application Front End (AFE) is application acceleration. Users are demanding faster response time from applications that increasingly send larger amounts of data for each transaction. An AFE

plays an important role in speeding transactions along using a wide range of techniques.

The solution should help control costs in the data center. One of the main goals of data center consolidation is to control costs. The network's solution needs to contribute by helping control server cost and making sure that server resources are efficiently used. No solution is complete without contributing to security. The solution needs to help defend the data center and server from attacks and new security threats. This is all in addition to providing support for functions the older Layer 4-7 switches / Server Load Balancers (SLB) provided in the past while providing high availability, high throughput and being cost effective. It is a tall order.

The diagram shows an architectural view of the AFE. It sits as the traffic cop between the servers and resources in the data center, the internal networks, and the Internet. The AFE provides four main functions:

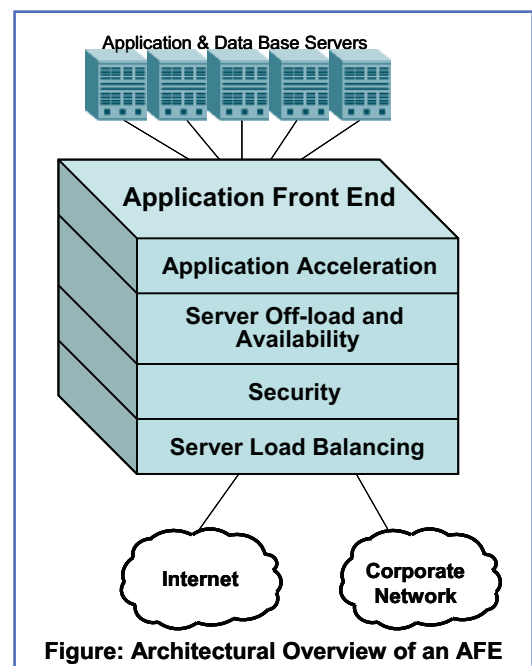


Figure: Architectural Overview of an AFE

Application Acceleration: The AFE provides functions that can help speed up transactions. Examples include caching objects and providing them to the user, eliminating the step to the server; and compressing traffic to reduce both the size and number of packets needed for a transaction, quickly redirecting traffic to the right resource or eliminating steps and signaling to the browsers to use objects in its cache.

Server off-load and availability: AFEs understand the traffic passing through it and the health of the server it is sending the traffic to. This understanding allows it to play a critical role in providing high availability. The AFE can also reduce the overall load on the servers. Servers can be inefficient in processing TCP connections and SSL processing. Additionally, "webified" newer applications use a large number of short-lived TCP connections that consume server resources. The AFE can perform the processing for the server and multiplex the connection over set of long lived connections.

Security: It is becoming evident that every piece of equipment must help with security. Security is a concern for every part of the network and the AFE can play a critical role in the overall security of the data center. The AFE sees all the traffic going to and from the data center, putting it in a unique position to identify threats such as Denial of Service (DoS) attacks and making sure users are authenticated before allowing the transactions to reach the servers. Thus, AFEs are a critical part of the overall security of the data center.

Server Load Balancing: The AFE is the next generation SLB. The role of insuring that no server was over-utilized or under-utilized is still an important one. The AFE takes over the role of the older SLB by bringing their functions into the future. In fact, the AFE expands the role SLB played in the past. It is generally thought that only large networks needed SLBs, but with the increasing role of AFE's outlined above, AFEs are needed in every data center. This will bring the benefits of SLBs to even the smallest of data centers.

The question facing network managers is not whether an AFE is needed but which AFE to use. The Application Delivery Challenge can help answer that question. I have challenged the leading AFE vendors to showcase their AFE solutions and explain why their solution is the best one. While many of the

features of an AFE are common there are significant variations. Different vendors put their research efforts on different factors that make up an AFE. It is important to understand the details and emphasis of each vendor to understand how they will solve your unique network problem. That is the purpose of the Challenge – to allow you to hear from the vendors and for them to explain their solution.

In the pages that follow, three of the top AFE vendors – **Juniper Networks, Netli and Blue Coat Systems** – explain what they do and, most importantly, address the following questions:

- What makes their AFE solution the best?
- What combination of features or technology sets their solution apart from the others?
- What problem or situation is their AFE best designed for? Is it high throughput, superior security, low operating cost, fastest acceleration or richest server features?

What follows are their answers. I have asked them to focus on their most specific and unique advantages. Of course the vendor can do more than what is presented in this brief answer. I encourage you to contact the vendors if their answer intrigues you and ask them to address your particular problem.

If you have any comments about this Challenge you can contact me at Robin@Layland.com.



The Power of Integration:

Secure and Assured Application Delivery from the Data Center

Ironic as it may seem, today's distributed enterprise is all about consolidation. The more businesses extend their operations to locations around the world in an effort to achieve "globalization," the more they desire to centralize critical application resources and IT operations to maintain control, reduce overhead, lower operational costs, and simplify regulatory compliance.


Obviously, this consolidation presents its own set of challenges. One of the most prevalent is maintaining and improving application performance. For many businesses, web-enabling traditional client-server applications has made data center consolidation and server centralization possible by providing universal and uniform application access to virtually anyone, anywhere – remote and branch office employees, telecommuters, mobile users, even partners and customers.

This transition, however, has raised a whole new set of issues. Web-enabled applications consume ten times the bandwidth of their client-server counterparts, monopolizing an inordinate amount of the limited wide-area network (WAN) capacity. Not only does this impact all other applications running over the WAN, it condemns users to a life of agonizingly slow page downloads, frustratingly poor response times, and generally lower productivity.

That's not all. As web-enabled applications grow in popularity, businesses are forced to deploy a variety of niche appliances, as well as more servers, in the data center to deliver scalable, secure and highly available applications. In order to deliver a complete solution, businesses have deployed a variety of independent, one-dimensional point products in their data centers, each focused on a specific task: load balancing, caching, security, SSL acceleration and termination, Authentication, Authorization and Accounting (AAA), compression, and content transformation. This trend defies the original objectives of consolidation – reduce overhead and lower operational costs – by adding complexity, cost and confusion to the data center.



by Michael Banic
Director, Product Marketing
Juniper Networks



The Juniper Networks Philosophy: Simplify the Web Tier

Juniper Networks subscribes to a different philosophy, one that embraces the concept of consolidation by delivering a solution that improves the delivery of web-enabled applications while reducing complexity and cost.

With the data center application acceleration (DX) platform, Juniper dramatically simplifies the "web tier" by consolidating multiple functions currently performed by point solutions into a single, high-performance device. By integrating load balancing, caching, security, SSL acceleration, compression, and content transformation, the DX platform streamlines the application delivery process, enabling organizations to do more with less – a cost-cutting, easy-to-manage solution that honors the spirit of data center consolidation.

Relieving Server Burden

The DX platform integrates multiple features – including TCP connection multiplexing, SSL termination, server load balancing, ultra-fast object caching, adaptive compression, and TCP slow-start mitigation, among others – to accelerate the delivery of web-enabled applications from the data center.

By multiplexing TCP connections, the DX platform reduces thousands of incoming client connections down to just a few, relieving the connection-management burden on back-end servers and allowing them to do what they do best: serve content. By taking over resource-intensive tasks such as session set-up and tear-down and SSL termination, the DX platform frees up considerable CPU cycles on the servers, allowing them to process up to four times the normal number of incoming requests. As a result, transactions are completed much faster, and users receive much better response times.

Thanks to these TCP multiplexing and SSL offload functions, organizations find they can deliver faster application

response times to more users with far fewer servers. They then have the enviable option of either reducing administrative overhead by eliminating some of the servers in the data center, or keeping them to prepare for future growth without requiring any additional infrastructure investment.

Spreading the Wealth

To ensure continuous availability and optimum performance of remaining servers, the DX platform's patented Fewest Outstanding Requests load-balancing algorithm equitably distributes incoming requests among available resources. Unlike legacy TCP-focused load-balancing techniques, the DX platform is the only solution in the industry that distributes requests at the HTTP application layer. As a result, the fastest and most available resources are efficiently utilized without overloading them, ensuring requests are filled quickly and avoiding workflow delays. A global server load balancing (GSLB) feature also enables the DX device to load-balance between data centers in different geographical locations, pro-

viding faster and more efficient performance as well as superior disaster recovery capabilities.

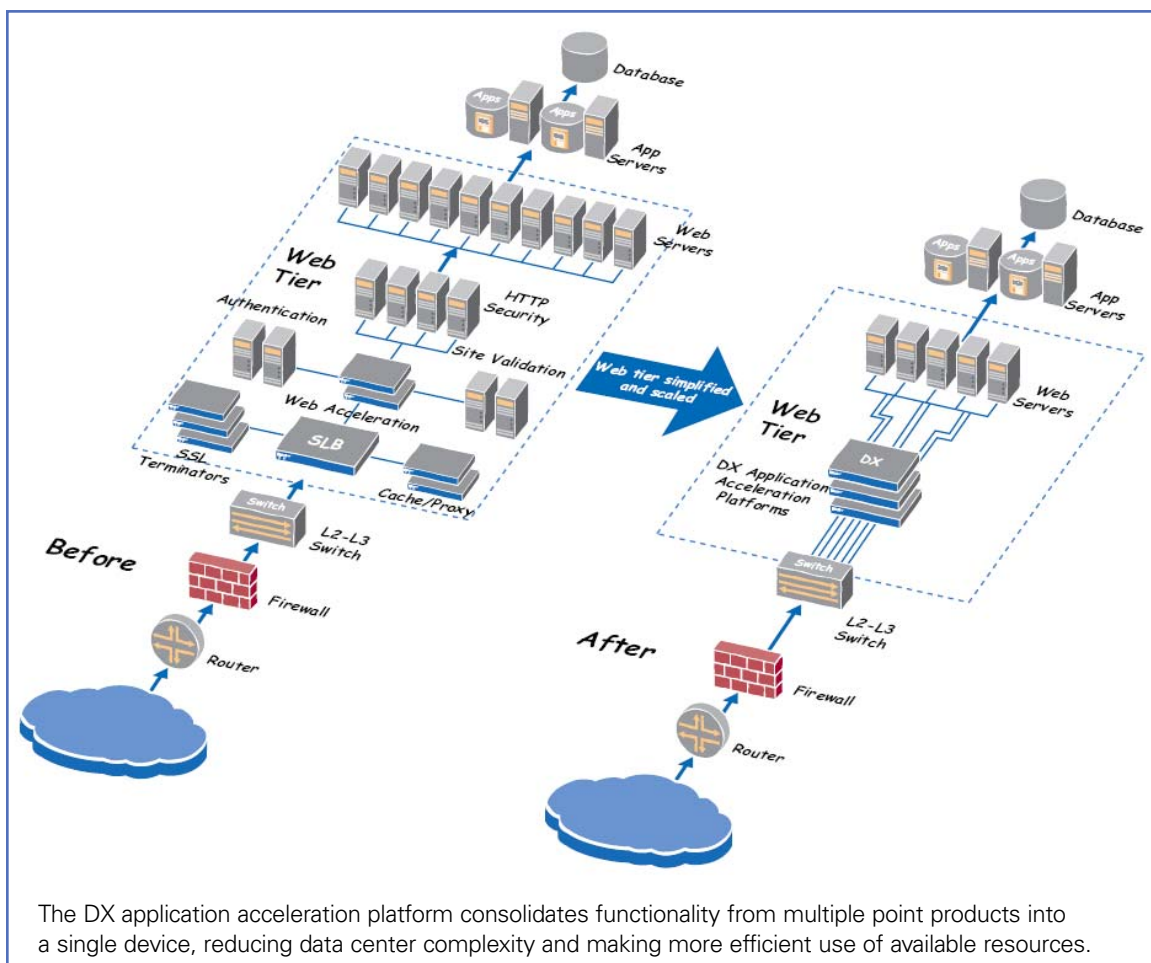
The DX platform can also be deployed to load-balance non-HTTP server farms that handle e-mail, FTP and all TCP and UDP traffic with application-level health checking. As the environment changes, additional functionality can be easily added to meet escalating demands without requiring any additional products or changes to the infrastructure.

Quick Response

Full support for standard GZIP and Deflate compression standards enable the DX platforms to make the most efficient use of available bandwidth, reducing congestion and accelerating content-rich downloads. The DX compresses all application flows, from standard HTTP objects to Microsoft Office documents, and can even detect the type of browser being used and employs the best compression technique for that particular software to further streamline downloads. The DX platform also supports "chunking" – allowing a

browser to display new objects as they are received, rather than waiting for the entire page to download – to speed displays.

The DX further accelerates response times through 3G caching, which stores commonly-requested objects locally on the DX platform's fast DRAM so those requests never reach the server to consume valuable cycles. The DX platform can also force browsers to cache frequently-requested static objects, further accelerating response times for applications that rely on content-rich interfaces.



The DX application acceleration platform consolidates functionality from multiple point products into a single device, reducing data center complexity and making more efficient use of available resources.

To keep operations moving smoothly, the DX platform optimizes TCP by eliminating much of the back-and-forth required to establish connections and complete the “slow-start” process. Brief periods of user idle time can cause a TCP connection to lapse, forcing the browser to initiate a new connection with the user's next mouse-click or key-stroke. As a full proxy, the DX platform keeps all TCP connections alive until the SSL session is terminated, so no time is wasted re-establishing connections and going through slow start.

Secure and Assured Application Delivery

The universal accessibility of web-enabled applications presents certain security risks. The DX platform addresses those risks by acting as a full HTTP proxy that protects back-end resources from external threats.

The DX platform terminates all incoming requests so outside sources never gain direct access to the servers themselves. By providing a buffer between clients and servers, the DX platform validates all requests and responses to ensure only proper HTTP is allowed through. The DX can “cloak” back-end resources with flexible URL translations to proactively prevent threats. Through its easy-to-use and extensible Sentry rules, the DX platform protects against an ever-growing list of common attacks.

Finally, the transaction-based DX platform is “application fluent,” not just “application aware,” enabling the device to transform content to improve application performance, modify workflows, and reduce or eliminate errors. That application fluency is enabled by the DX platform's AppRules™ feature, which gives IT a level of business agility that the network has never delivered in the past.

The Power of Juniper

These features – multiplexing, load balancing, caching and compression – accelerate web-enabled applications for all users, whether they are located in corporate headquarters, a branch office, a hotel room, or even dialing in from home over a 56k modem. All users accessing a web-based application accelerated by a DX platform will notice a dramatic improvement in application performance, transaction response times and page downloads, resulting in a much more satisfying experience and restoring productivity to earlier client-server levels.

Yet as powerful as the standalone DX platform is, its true power is derived from being part of the larger Juniper product portfolio. Working with complementary Juniper products, the DX platform contributes to an even more powerful integrated solution for accelerating – and securing – application delivery from the data center.

Working with the Juniper WX and WXC application acceleration platforms, for instance, the DX platform delivers even greater value to branch-office users accessing centralized applications by accelerating and extending support to client-server-based and legacy applications. The WX and WXC platforms improve application performance across wide-area links through a combination of compression, caching, TCP- and application-specific acceleration technologies, and bandwidth management and Quality of Service capabilities. Deployed in concert, the DX and WX/WXC platforms represent the industry's most complete application acceleration and delivery solution – all from a single source.

Securing the Data Center

Likewise, working with the Juniper SSL/VPN solutions, the DX platform contributes to a complete secure and assured application delivery system.

Deployed together, the DX platform and SSL/VPN enable unprecedented levels of security and access to protect sensitive and confidential data, as well as business-critical resources. The SSL/VPN provides virtualized connectivity for any user, anywhere in the world. Access levels can be defined, applied and enforced for individual users, ranging from very broad down to single files or directories. Access levels can even be assigned based on the type of computer accessing the network, so users on a personal laptop, for instance, would get fewer privileges than someone logging-on using a company-approved and recognized PC.

The Juniper products – the DX and WX/WXC platforms, combined with the SSL VPN and IPsec VPN/Firewall platforms – deliver a complete, compatible, end-to-end application delivery solution, all from a single vendor.

Juniper Networks

<http://www.juniper.net/products/appaccel/dca/>
866.298.6428 (inside the U.S.)
or 978.589.0500 (outside the U.S.)

Managed Application Delivery

Services Reduce the Cost and Complexity of Application Acceleration



by Tim Knudsen
Sr. Director, Product Management
Netli, Inc.



Problem

Today's enterprise Web application deployments face diametrically opposing forces amplifying the risk of failure. First, eCommerce, collaboration and supply-chain activities transcend national boundaries requiring applications to serve an arbitrary, global audience in a fully secure fashion. Second, the underlying HTTP/TCP protocol is poorly suited to serve globally dispersed users from a single point, forcing data center build out. At the same time, profit-maximization pressures are forcing IT to centralize and consolidate IT infrastructure to minimize costs. The result is a recipe for failure – unusable applications, weak adoption, and negative return on significant IT investments.

At the root of the problem is the “Global Delivery Dilemma” – how to ensure the LAN-like performance expectations necessary to ensure business adoption *without* sacrificing security, introducing compliance and visibility holes, implementing and managing application modifications, and having to deploy and manage expensive data center infrastructure?

Challenge

Unfortunately, no application front-end (AFE) appliance, dual-ended appliance or traditional content delivery network (CDN) service supports the complete scope of functions needed to eliminate the dilemma.

Global Sub-second Performance for Arbitrary Users

- Dual-ended solutions can't effectively serve arbitrary users
- Application front-end (AFE) appliance solutions provide some TCP-level enhancements, but can't eliminate the round-trip issue – the core of Internet latency – due to their single-node architecture.

Infrastructure Offload and Scalability

- Using a CDN solution to offload static content provides only partial infrastructure utilization reduction for dynamic Web applications and actually increases overall application management costs and still does not resolve the underlying latency cause

End-to-end Security

- No existing appliance or CDN solution can serve arbitrary users, securely, from end-to-end, for the full range of server and client certificate models

Transparent

- Neither AFE or dual-ended solutions or existing CDN solutions provide transparent (no application, client or server modification required) application acceleration and infrastructure offload
- Breaking apart the application for edge-based computing solutions is not only very costly to deploy and maintain, but introduces significant security risks and visibility and compliance holes

As a result, IT is forced to evaluate, deploy and manage an overlapping conglomeration of appliances and services in an attempt to satisfy as many requirements as possible, but ultimately falling short of eliminating the *Global Delivery Dilemma*. The end-result is a costly infrastructure that partially mitigates the root-cause of the problem, while introducing significant complexity and unnecessary risks.

Solution – Global, Secure and Transparent Managed Service

The solution is to take a different approach – a global, secure and transparent Web application delivery service that eliminates the *Global Delivery Dilemma* for both sides of the Web application stack – end-user and machine-to-machine Web service communications.

The Netli One™ Platform

The Netli One platform is an overlay network infrastructure that supports the Netli® Application Delivery Network (ADN) services, which include NetLightning, NetliOffload™, NetliContinuity™, and NetliView™.

NetLightning securely accelerates dynamic, enterprise Web applications to assure LAN-like response times worldwide without requiring application modifications. With zero cache warming and management required, NetliOffload offloads static content delivery to shift bandwidth, computing, and storage requirements from the enterprise data center to the on-demand Netli network infrastructure. NetliContinuity provides global traffic management to ensure optimal performance and availability across data centers. NetliView continuously monitors of business-critical, Netli-optimized applications to provide visibility into application response time and health.

Netli One core components include a series of globally distributed points of presence, a high-performance global DNS redirection and IP address mapping system, and high-performance protocol and content optimization software that is tightly integrated to eliminate the *Global Delivery Dilemma* in the following ways:

Global Sub-second Performance for Arbitrary Users

- Global bi-nodal architecture serves any arbitrary end-user, from any region in the world
- Netli protocol enhances TCP to optimize application transport, reducing response times by enabling all requests to be delivered in no more than two to three total round trips
- Predictive object delivery intelligently anticipates object requests to accelerate total page delivery
- Optimized loss and congestion management maximizes transaction availability
- Global redirection and traffic management ensures optimal application availability with intelligent site selection based on real-time network and application performance data

Infrastructure Offload and Scalability

- Transparent static content offload reduces origin server utilization and creates additional scalability for existing Web application infrastructure
- Connection management between the AAP and origin server eliminates latency resulting from connection creation and tear-down

- SSL offload and acceleration recovers valuable server resources, creating additional utilization headroom for unexpected traffic spikes

End-to-end Security

- Netli's security framework secures transactions from end-to-end, while providing full freedom of security models by supporting both server and client certificate models
- Netli's security framework improves security by enhancing encryption levels under acceleration
- Connection and URL filtering combined with a bi-nodal architecture provides 24x7x 365 intrusion and attack prevention

Transparent

- No application modifications required to accelerate and optimize applications
- All dynamic requests served from the origin, enabling a complete audit trail of business-critical transactions one source of truth
- End-to-end monitoring and robust reporting provides complete, real-time visibility into application response time and availability from the perspective that matters most – the end-user

How it Works

The Netli One platform consists of a series of globally located virtual data centers (VDCs), which are distributed points of presence around the world that serve the Internet's end users and application access points (AAPs), that reside in the data center (or near proximity) and provide the same benefits as the AFE Layer 4 through 7 switching appliances (see Figure 1).

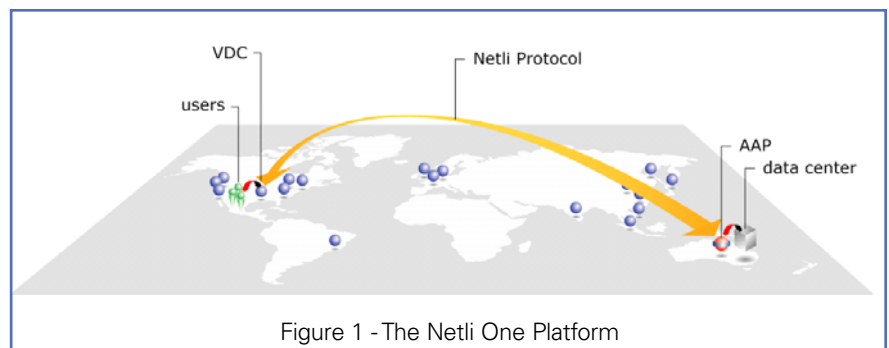


Figure 1 - The Netli One Platform

The Netli architecture essentially replaces TCP and HTTP while still utilizing the standard IP transport of the Internet. As a result, end users that are distant from the web server enjoy the same level of performance that local users receive, even in the event of network congestion and packet loss. It supports HTTP and HTTPS applications and is completely managed by Netli, requiring minimal customer IT resources for installation, support, and maintenance. There are no application software, content, or client changes and no capital expenditures with Netli services. Companies subscribe to a Netli for one or more Web-enabled applications and delegate the DNS processing for those applications to Netli.

When end users access any of the optimized Web applications, Netli's system transparently optimizes the application delivery. With Netli, Web-enabled applications transparently appear to be globally replicated even while being served from a centralized datacenter. The result is a global infrastructure that enterprises can leverage to deliver sub-second response times, increased availability, more efficient asset utilization and greater visibility into performance worldwide-with no server or client changes and no capital expenditure. Additionally, Netli can provision the services in a matter of hours with little more than a simple DNS change.

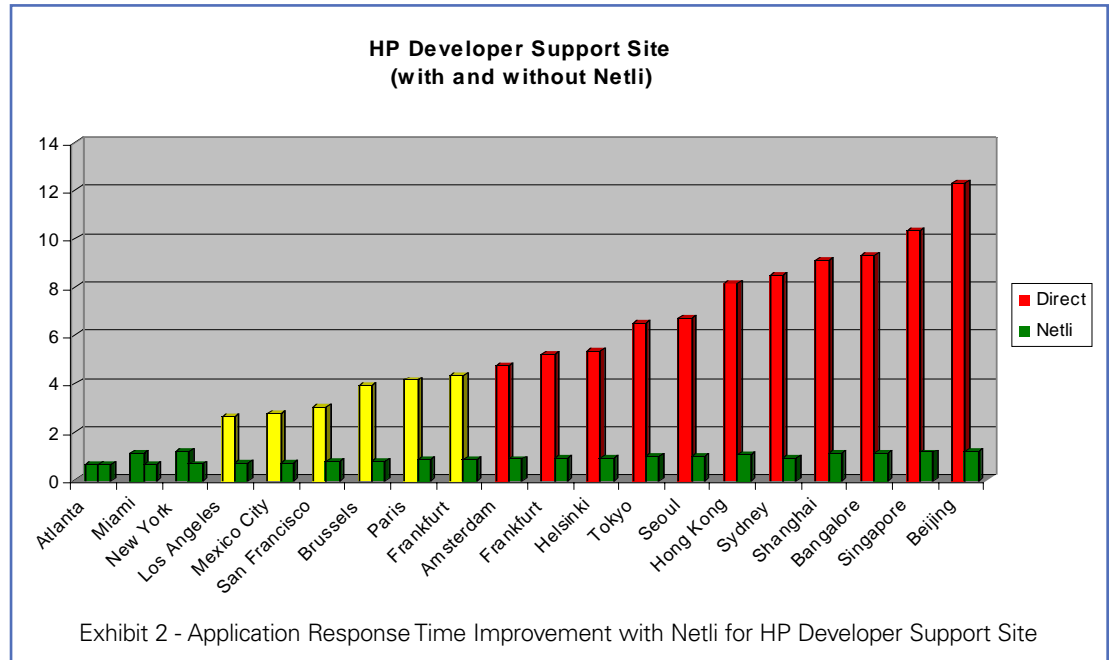
The results are real. Figure 2 illustrates how Netli's solutions improve application response time performance regardless of end-user location.

Summary – Reduce Cost, Complexity, Risk and Inefficiency

Eliminating the *Global Delivery Dilemma* is more than an IT issue. It is a mission-critical business problem where the risks of implementing a partial solution are more than displeased users, but significant business costs in the form of lost revenue and productivity, and sky rocketing IT costs. Solving the problem with a multi-layer, 'build it yourself'

approach is costly, risky, and does not ultimately solve the problem. Solving the problem requires a new approach – a global, secure, and transparent managed service.

The Netli One Platform and ADN services solve the problem by providing global sub-second response time for arbitrary users and end-to-end security in a fully transparent



delivery model. Netli optimizes the delivery of applications for the world's three biggest mobile phone suppliers, the largest technology reseller, the largest beauty products company, the two largest computer manufacturers, and the top two import auto manufacturers – processing more than one billion enterprise transactions across its network service platform every day. With Netli, industry-leading companies solve the *Global Delivery Dilemma*, resulting in increased application adoption, revenue, and competitive advantage.


Netli, Inc.

www.netli.com
3980 Fabian Way
Palo Alto, CA 94303
Phone: 1-866-26-NETLI

Application Delivery Challenge – Implementing an Application Front End



by R. 'Lynn' Nye Jr.
Enterprise System Architect
Blue Coat Systems Inc.



The Application Front End (AFE) has deep roots in the history of networked applications and was the cornerstone of IBM's System Network Architecture (SNA). Known as the Front End Processor (FEP), it was the device that provided boundary services for mainframe applications. It primarily offloaded the mainframe from redundant tasks and provided the foundation for scaling services. While the FEP did a lot of the "heavy lifting" it was a component in a larger system architecture, where every node had a role in assuring reliable application delivery service levels. All of the components weaved together as an architecture defines an *application delivery infrastructure*.

As a solution requirement the industry has been evolving its understanding of what the requirements are and more importantly, what they are not limited to. Since the build-out of the *packet delivery infrastructure*, which provides universal IP connectivity, the system challenges have centered around the relationship (supply / demand) between applications and networks. The initial 'promised land' was based on various network Quality of Service (QoS) approaches which through 'deep packet inspection' would determine the identity of a particular application flow on a network and apply network priority. It has evolved to include a number of WAN optimization and content caching techniques that not only address bandwidth contention issues, but also improve end-user performance.

While on a parallel track, there was the emerging challenge associated with managing the relationship users have with applications, which was unleashed with the 'web model' of browsers and system transparency. Once application access wasn't limited to proprietary client / server relationships, users have few if any inhibitions about accessing anything and everything. While initial energy was applied to various authentication and firewall techniques this market requirement has evolved to incorporate web content filtering, Spyware filtering, etc.

Blue Coat Systems Inc. has had the unique company history starting with application performance and building security on

top of that platform. Blue Coat supports both reverse proxy (AFE) and forward proxy configurations, (deployed in thousands of enterprise networks), and has continued to expand application and protocol support. Today a Blue Coat proxy can protect, control and accelerate HTTP(S), as well as Instant Messaging, Streaming Media, P2P, FTP, CIFS and MAPI.

All of these services are centrally managed and through the Blue Coat Policy Center, from which an organization can apply very granular policy to manage services. With strong authentication integration, the Blue Coat proxy understands the context of the user-to-application (session) relationship, which is the foundation for both business and technical policy administration.

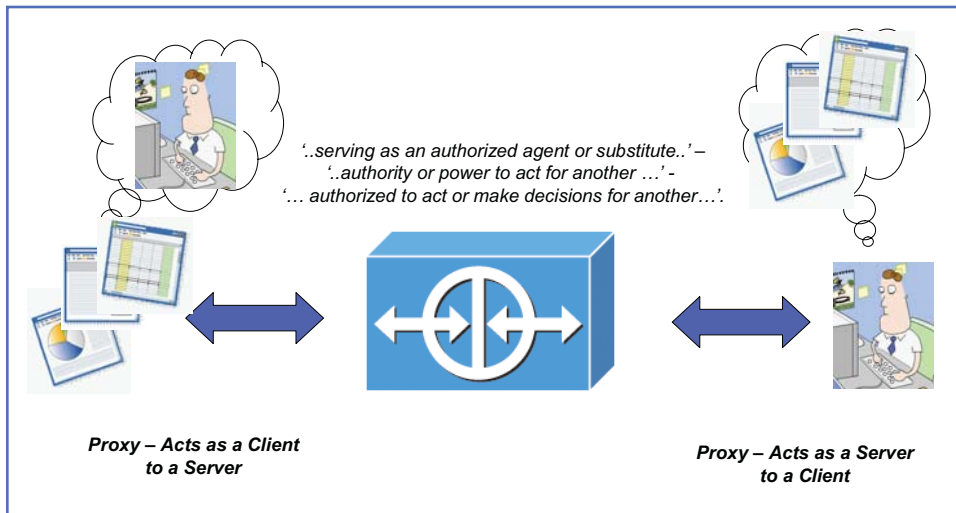
Power of the Proxy (Protect - Control - Accelerate)

From Blue Coat's perspective a proxy has a number of inherent responsibilities when managing sessions - far more than just inspecting packets moving across a network. Since applications are primarily made available to users, a next generation solution must be architected to understand the context of a user-to-application relationship.

prox-y (prŏk'sē) n., pl. -ies.

1. A person authorized to act for another; an agent or substitute.
2. The authority to act for another.
3. The written authorization to act in place of another.

As indicated in the definition lifted from the dictionary, a proxy acts on behalf of another. In the context of networked applications the proxy acts like a client to the server and a server to a client. Being in this position, the proxy can take literally do anything at layer 3 through 7 of the OSI stack. Taking



advantage of that capability is driven by granular and global policy, which essentially ‘tailors’ the proxy for any environment.

The differentiating quality of a proxy in contrast to network-level deep packet inspection or flow control is context. As an example by terminating both sides of a session, the proxy knows the address resolved by DNS, the authentication process used by the user and the breadth of content required to complete a transaction. So rather than attempting to associate a number of packets or flows by IP address or other schemes commonly employed by network level devices, a proxy inherently understands and can manage all aspects of a session.

The Blue Coat Difference

To address the questions posed in this challenge, there are a number of items that quickly begin to differentiate Blue Coat.

- Complete system solution: One vendor / One OS / One management structure / One support framework
- Application Proxy: HTTP / HTTPS / IM / P2P / Streaming / FTP / CIFS / MAPI
- Price Performance: Platform Breadth / Special Purpose OS
- Integrated Session Security: Content Filtering / Anti Virus Scanning / Authentication

Therefore Blue Coat represents not only an AFE solution for our customers, but a complete proxy architecture infrastructure for a breadth of critical application delivery requirements. The ‘power of the proxy’ is based on the heritage of Blue Coat, which was caching Web content during the days the company was known as CacheFlow. This background is unique in the industry and represents several years of pro-

viding not only acceleration of Web applications, but also inherently providing a policy structure that delivers application-level control. The control associated with an application-level proxy is unique and can’t be matched by any other ‘packet level’ approach.

However, the breadth of requirements for assuring protection, providing session control and delivering on performance isn’t isolated to the data center. Web-based applications are used ubiquitously for employees, customers and business partners, which dictate that

the proxy needs to support traffic between the internal network and the Internet as well as to address the challenges associated with protecting, controlling and delivering performance to users in remote locations. There are also application services such as File and Exchange, as well other legacy TCP applications that benefit from a ‘front-end’ and optimized delivery to remote locations.

Therefore ‘best of breed’ products end up being features on platforms that eventually become system infrastructure, since dealing with multiple platforms / vendors drives the TCO (Total Cost of Ownership) out of justifiable range quickly. Therefore a system solution not only addresses one or two applications, but rather is a platform for an application delivery infrastructure, for today’s challenges and tomorrows requirements.

Most of the energy behind AFEs and how they are evaluated revolves around the core services:

- Offload redundant protocol services
 - Cache
 - TCP (connection set-up / tear down and multiplexing)
 - SSL Encryption
 - HTTP Compression
- Extending Application Services
 - Content Switching
 - Content Transformation

While all of these are important features, they are rapidly becoming the commoditized. The things that are beginning to differentiate vendors, and what enterprise architects are

looking at, are related to the user relationships or sessions with those applications, which incorporate the logical / physical access points of the Internet and branch offices. In other words, who is accessing what? Is the context of this session not only authorized, but is it being conducted in a standardized manner or are there variables that raise concerns?

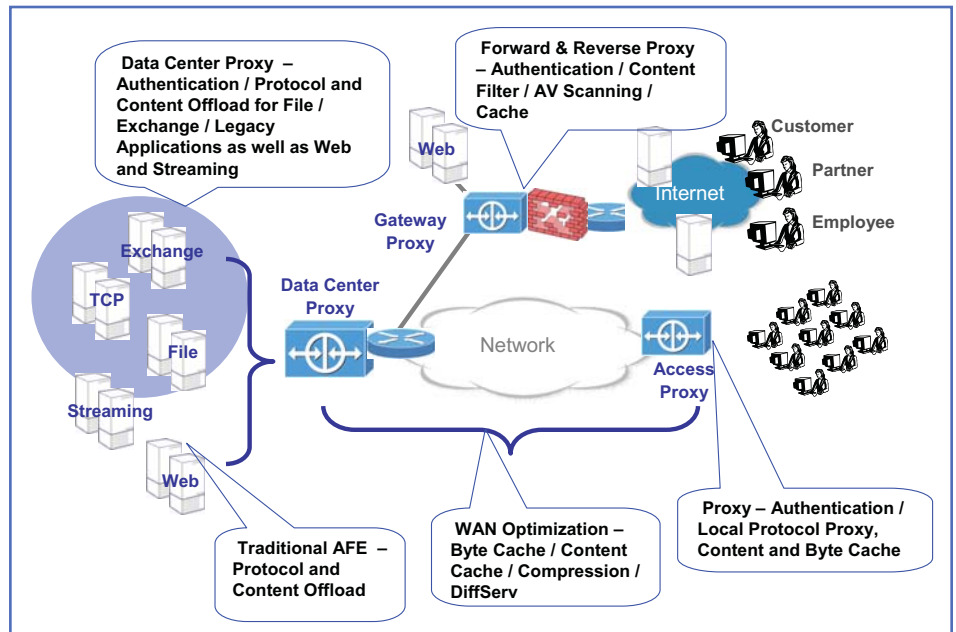
Blue Coat Delivery

To meet those requirements, Blue Coat has continued to leverage the proxy to include a broader set of application services. Because a proxy understands the context of a session, which is inclusive of the user and the application, the fundamentals of application acceleration are well understood. Blue Coat's framework delivers content to users in an application orientation whenever possible, but assures application performance for all packet-based applications through 'byte caching'.

Building on the initial web application Forward and Reverse Proxy cache, Blue Coat has continued to add full application proxy support including:

- HTTP
- HTTPS
- FTP
- IM - AOL / YAHOO / MSN
- Streaming - Real / Media Player / Quicktime
- P2P - bitTorrent / eDonkey / Kazaa
- File - CIFS
- Exchange - MAPI

Those application level proxy services provide significant application performance improvements by caching application content. It is complimented by a number of WAN optimization services such as compression and a differential compression technique (Byte Cache). To assure end-to-end delivery, the proxy also provides granular DiffServ packet marking which leverages investments enterprises have made in their *packet delivery infrastructure*.



All of these application and WAN services are complimented by a broad set of integrated security services, including authentication, *Anti-Virus / Malware / Spyware* protection and as well as *content filtering*.

Summary

The AFE is a clear example of how a proxy is becoming the de facto platform for accelerating, controlling and protecting networked application services. It represents a component of a much wider requirement of what enterprise architects are looking at to address their system requirements.

Blue Coat has been involved with the evolution of the Web infrastructure since its commercial viability was established. The Web application environment has continued to evolve to where the definition of a browser is getting more and more challenging, but it also doesn't stand alone in the world of critical applications within the enterprise.

A proxy can uniquely address the breadth of application services and provides an application delivery infrastructure for today's requirements and those not yet defined in the future. This is the foundation of the application delivery infrastructure.

Blue Coat Systems Inc.

<http://www.bluecoat.com>

North America: 1.408.220.2199

Europe, Middle East & Africa: 0800.085.2356 Toll Free

Australia & New Zealand : +61.416.382.645

Asia Pacific: +852.2166.8121