

WebSphere Application Accelerator for Public Networks

*An IBM product offering created in partnership
with Akamai.*



Contents

- 2 The Akamai Ready initiative and WebSphere Application Accelerator for Public Networks
- 3 IBM WebSphere DataPower appliance with Application Optimization
 - 3 Technology overview
 - 4 How it works
 - 4 Value: Performance, availability, scale, security
- 5 Akamai Public Network Acceleration
 - 5 Technology overview
 - 7 How it works
 - 7 Value: Performance, availability, scale, security
- 9 Application Optimization Solutions with WebSphere Application Accelerator for Public Networks and WebSphere DataPower
 - 9 Smarter caching brings applications closer to the user
 - 10 Adaptive load optimization responds to changes in traffic demands
 - 10 Tighter cloud security: Defense in depth
 - 11 Additional benefits and capabilities
- 11 Summary

Executive summary

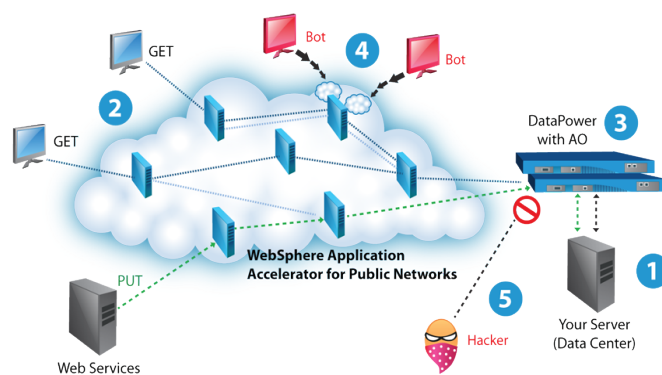
From the browser to the cloud or enterprise server, Akamai and IBM® are partnering to provide unmatched performance and security for Internet-based applications and cloud services. Akamai is the leader in accelerated and secure content delivery over the Internet. IBM is the leader in web application middleware and enterprise- and cloud-based data services. Add both together, and the value is greater than the sum of the parts.

The Akamai Ready initiative and WebSphere Application Accelerator for Public Networks

The *Akamai Ready* initiative is a collaboration between IBM and Akamai Technologies. It combines Internet and enterprise delivery optimization strategies that can help it provide value greater than the sum of its parts. *WebSphere Application*

Accelerator for Public Networks is the first IBM product offering from the Akamai Ready initiative. *WebSphere Application Accelerator for Public Networks* provides capabilities to help optimize, accelerate, secure, enhance and manage enterprise applications to Internet users. This can provide dramatic advantages for many usage scenarios (see Figure 1) such as:

- Client browsers accessing high-traffic enterprise websites
- Business partner access to enterprise web services
- Security and threat protection



- 1 Simple to deploy - no application changes needed
- 2 Consistently reach customers regardless of location with enhanced performance that offloads your server
- 3 Reduce response time even further with Elastic caching inside your data center
- 4 DDOS (malicious attack) protection amplified
- 5 Enhance Security inside your firewall

Figure 1.

WebSphere Application Accelerator for Public Networks combines many capabilities of the Akamai Web Application Accelerator and is best used with IBM WebSphere® DataPower® service-

oriented architecture (SOA) appliances to deliver a configurable end-to-end solution for web-based applications. IBM and Akamai have collaborated to produce a solution that delivers performance, security and availability beyond what was previously available by traditional application delivery infrastructures.

This white paper provides technical and implementation details of *WebSphere Application Accelerator for Public Networks* as it pertains to web application acceleration with WebSphere DataPower appliances. It provides a background overview of IBM's DataPower SOA appliance portfolio and Akamai's WebSphere Application Accelerator technology followed by an in-depth analysis of the benefits of implementing *WebSphere Application Accelerator for Public Networks*.

IBM WebSphere DataPower appliance with Application Optimization

Technology overview

IBM WebSphere DataPower provides a line of purpose-built, easy-to-deploy network appliances to simplify, secure and accelerate enterprise and cloud computing scenarios. The DataPower Security Gateway XS40 and Integration Appliance XI50 are purpose-built devices well-known for their capabilities in web services and XML. They provide accelerated validation, transformation, encryption/decryption, threat protection, routing and more, to offload processor-intensive operations from overloaded internal servers. The XS40 focuses on security, and is typically deployed to improve Internet-facing applications and services in the DMZ. The XI50 combines the capabilities of the XS40 with enterprise service bus (ESB) integration features within the enterprise.

IBM recently introduced the Application Optimization (AO) feature pack option that expands the XS40 and XI501 to fully cover web application delivery and security. AO provides mission-critical performance and resiliency for web applications and Web 2.0 technologies with features such as self-balancing and intelligent load distribution.

As shown in Figure 2, a typical usage model for the XS40 with AO is deployment in the DMZ. In this scenario, the XS40 acts as a reverse proxy, accepting incoming requests for a wide variety of network protocols. The appliance can then perform security verifications, acceleration, transformation, logging and other services on the request before it is dynamically routed to the appropriate server resource(s). Responses return back to the client through the DataPower appliance, where the same, or other services, can be applied.




Figure 2.

The AO option adds additional scaling, performance and availability capabilities to provide the premier application delivery solution for mission-critical web applications.

Examples include:

- Self-balancing for virtually unlimited scalability
- Intelligent load distribution, which leverages information from the server tier to optimally route incoming requests
- Application versioning, allowing versions of mission-critical web applications to be managed with minimal disruption.
- New versions can be partially rolled into and out of production for testing and verification, including full quiesce capabilities

The mission of the DataPower XS40 appliance with AO is to provide a consolidation solution in the DMZ. It is not uncommon to find a number of software- and appliance-based

servers deployed in the DMZ, each providing one or more key functionalities. The DataPower XS40 appliance with AO provides a governed, highly manageable, secure and scalable integration solution for delivery of web services and web applications in a single appliance.

How it works

The DataPower XS40 Security Gateway is a highly specialized, rack-mountable network appliance. However, unlike most network appliances, the XS40 processes information at the application protocol layer, not the network protocol layer. The XS40 is fluent in XML, web services, WSDL and WS-*, SOAP, Messaging, HTTP, HTTPS, web applications, and many more application protocol standards. The XS40 correlates request and response messages, and optimizes throughput by buffering or streaming network I/O as appropriate. This is a significant advantage over traditional network security gateways that only inspect the individual packets on an IP network. A traditional edge network device cannot differentiate message content that is taken from encrypted web services, XML or other common application-level payloads. The XS40's application-level processing allows it to perform operations, validations and transformations on the message as a whole or at the field level, leading to greater accuracy, protection and much richer runtime capabilities.

As a network device, the XS40 focuses on performance, scalability and security. DataPower appliances achieve this through purpose-built hardware and firmware. At the core of the DataPower appliance is advanced technology that compiles the XML transformation language, XSL, into optimized machine code. On top of this is acceleration for XML processing and processor-intensive encryption/decryption operations. With the AO option, self-balancing provides scalability, while the appliance can leverage information from backend application servers for dynamic and optimized routing. For security, the XS40 provides authentication, authorization, signature validation and other standard security features. The appliance itself is also designed for security. The case is built with innovative tamper-proof technology that disables the appliance when tripped. This is coupled with

controlled external access ports for full auditing and governance, and an optional hardware security module (HSM) for storage of sensitive cryptographic material.

Value: Performance, availability, scale, security

WebSphere DataPower SOA Appliances are a key element in IBM's holistic approach to SOA. The DataPower XS40 with AO offers a complete hardware platform for validating, accelerating, and securing web services and applications in a highly manageable and scalable way.

Performance

WebSphere DataPower's innovative XML hardware delivers wire speed performance that isn't possible with any general software-based solutions, employing the following capabilities:

- Next-generation XML hardware acceleration
- Hardware accelerated encryption/decryption
- Advanced compilation of XSL into machine code instructions
- Gigabit Ethernet

Scale

WebSphere DataPower contributes to improved infrastructure scalability through:

- Intelligent back-end application workload balancing with session affinity (AO)
 - DataPower appliances provide intelligent distribution of application traffic to server members or groups using feedback from the application cluster, health checks or application session affinity. This feedback allows for dynamic adjustments that direct traffic to servers such that the server utilization, response time and throughput are optimized.
- Self balancing (AO)
 - Two or more DataPower appliances will appropriately and automatically self-balance client requests amongst themselves. The operation is transparent as appliances join or leave the self-balance group.
- Service level management

Availability

WebSphere DataPower helps ensure applications and websites are always available, by leveraging:

- Standards-based, centralized governance and security
- Intelligent back-end application workload balancing with failover (AO)
- Application versioning (AO)
 - DataPower appliances act as a dynamic router, providing “application version”-based routing and non-disruptive version rollout and quiesce. This feature requires the advanced features of WebSphere Virtual Edition (WVE) application server. DataPower will retrieve application version information from a WVE deployment and use it for intelligent routing based on the chosen policies.
- Bridges to Web 2.0 technologies with JSON filtering and validation, support for REST verbs, and converting and bridging of REST and web services

Security

WebSphere DataPower enhances origin security with:

- Mature message-level security and access control
 - Messages can be filtered, validated, encrypted and signed, helping to provide more secure enablement of high-value applications. Supported technologies include WS-Security, WS-Policy, WS-SecurityPolicy, WS-ReliableMessaging, WS-SecureConversation, WS-Trust, SAML and LDAP. The XS40 is a sealed network device for high reliability and increased security assurance.
- Web application firewall to protect against cross-site scripting, SQL injection and a wide variety of XML and content-borne threats
- Control access to applications, services and data based on customizable roles and rights
- Advanced hardware security to protect sensitive information
- Standards-based, centralized governance and security

The WebSphere DataPower platform is the obvious choice for enterprise customers looking to enhance the edge of their enterprise application delivery platform with unsurpassed performance, scalability, reliability and security. With a broad range of features, WebSphere DataPower appliances consolidate many of the application front-end functions that are spread across traditional software and servers today, and promote a data center consolidation strategy at the edge of the enterprise network.

Akamai Public Network Acceleration Technology overview

Akamai employs a massive network of tens of thousands of servers in more than 1,000 networks and over 70 countries across the globe, controlled by network intelligence systems that route requests, balance load and ensure extreme network uptime. These servers at the Internet’s edge, called EdgeServers, are the foundation of the Akamai EdgePlatform, placing Akamai within a single network hop of 90 percent of the world’s Internet users. By applying a sophisticated set of techniques to accelerate dynamic content, Akamai goes well beyond static caching to optimize application delivery bottlenecks for fully dynamic, enterprise applications.

The widely distributed nature of the EdgePlatform places an EdgeServer region in close proximity to every Internet user and corresponding centralized application, regardless of their location. Akamai uses an intelligent dynamic mapping system to direct each end-user and origin location to an optimal EdgeServer, in essence serving as an on-ramp and off-ramp to the Akamai network. The mapped Akamai EdgeServers form a direct binodal network between centralized applications and users across the edge of the Internet. By having servers as close as possible to end users, latency and packet loss are minimized.

Akamai uses a variety of optimizations, collectively called the Akamai Protocol, to address routing, transport and application layer bottlenecks inherent to the Internet.

- *Routing.* Akamai SureRoute technology eliminates the inefficiencies of Border Gateway Protocol (BGP), which seeks to minimize the number of hops Internet traffic takes, but may direct traffic to congested or completely blocked Internet routes. SureRoute leverages Akamai's network of EdgeServers and proprietary algorithms to provide a real-time weather map of the Internet. At any given time, for each independent user, SureRoute determines the highest performing and most available path to communicate between two Akamai EdgeServers, dynamically optimizing round-trip time between each end user and the application server and reducing packet loss. This optimizes the round-trip time and availability of the Internet, ensuring end-user requests reach the application server regardless of Internet bottlenecks such as service provider blackouts, brownouts, de-peering, network outages and earthquakes. Real-time routing decisions become increasingly important as web applications become more time sensitive and interactive with new technologies such as AJAX, while integrating functionality like live-chat and VoIP.
 - *Transport.* The Akamai Protocol removes the inefficiencies of TCP, eliminating the chattiness within the core Internet protocols by substituting more efficient protocols to communicate between origin and end-user EdgeServers. Akamai eliminates the three-way handshake for connection establishment and teardown by establishing a set of long-lived persistent communication connections between the Akamai EdgeServers. It also eliminates slow start, because EdgeServers maintain detailed knowledge of network conditions, such as bandwidth and latency, allowing them to communicate immediately at an optimal window size and avoid TCP's inherently sluggish mechanism. In addition, WebSphere Application Accelerator employs pipelining that allows multiple HTTP requests to be multiplexed over a single connection without waiting for a response, and a more intelligent retransmit methodology than the TCP timeout parameter.
 - *Application.* The Akamai Protocol employs intelligent prefetching, compression and caching to eliminate HTTP inefficiencies that require multiple TCP connections be established and torn down to deliver a page.
 - *Intelligent prefetching.* When a client requests the base page from the browser, Akamai simultaneously parses the HTML base page and immediately issues requests for the remaining elements of the page, before the browser even requests them, so they are already waiting at the edge to be delivered, as if the origin server were only a few milliseconds away.
 - *Compression.* Data is dynamically compressed and decompressed, reducing bandwidth usage and the total packets required to deliver data between nodes.
 - *Caching.* Cacheable content is stored on EdgeServers across the globe, close to end users. Serving requests from the cache at the Internet's edge eliminates the latency from traversing the globe. Time to live can be defined within metadata, to refresh the cache with the latest content. The frequency can be set to check with the origin every time before serving, to seconds, minutes, hours or days, depending on the nature of the information. Data may also be identified as dynamic, eliminating the caching component, but still benefiting from the optimized routing and transport protocols to deliver information directly from the origin to global end users with high performance and availability.
- Akamai also addresses security threats in the cloud. The distributed nature of the Akamai EdgePlatform makes it difficult to successfully launch a DDoS attack against an origin that's accelerated by Akamai. The EdgePlatform acts as a security perimeter in the cloud, absorbing traffic at the edge, before it ever reaches the origin, routing traffic appropriately and ensuring the application is available, even when under heavy load.

The following steps and Figure 3 depict how Akamai managed services work:

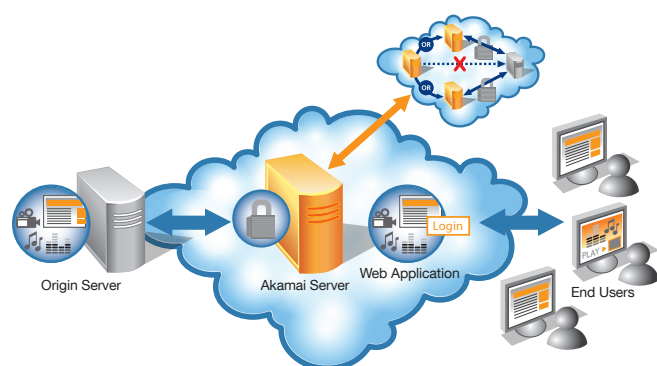


Figure 3.

How it works:

1. Akamai's dynamic mapping system directs user requests for secure application content to an optimal Akamai server.
2. Route-optimization technology identifies the fastest and most reliable path back to the origin infrastructure to retrieve dynamic application content.
3. A high-performance transport protocol transparently optimizes communications between the Akamai server and the origin, improving performance and reliability.
4. The Akamai server retrieves the requested application content and returns it to the user over secure optimized connections.

Value: Performance, availability, scale, security

Enterprises that deliver web-based applications on the Akamai network help offload bandwidth from the origin, reducing infrastructure costs and facilitating data center consolidation initiatives. Connection and route optimization techniques dynamically avoid Internet problem spots and maximize

throughput. The result for each application is superior performance, higher availability, increased scalability and an added security perimeter in the cloud. This, in turn, leads to increased usage, adoption and improved productivity.

Performance

The transport optimizations result in a dramatic reduction in effective round trips taken over the Internet—further improved by SureRoute optimization whenever a trip to the origin must be taken, yielding the following benefits:

- Global users experience local response times, regardless of their distance from infrastructure
- Higher application adoption
- Improved end-user productivity
- Promotion of data center consolidation
- End-user response time improvements of up to five times

Figure 4 illustrates performance benefits of Akamai acceleration for a Customer Support Portal's four-step dynamic transaction, showing how Akamai helps “flatten the world” by making users feel as if they are close to a data center, regardless of their location.

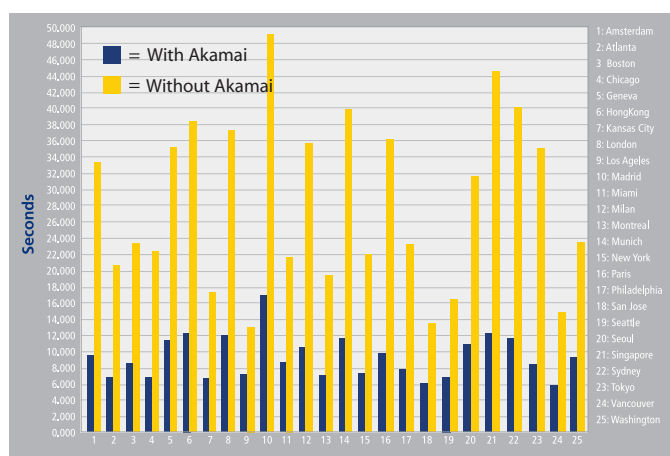


Figure 4: Global performance with and without Akamia

Scale

Akamai also greatly improves server scale within the data center by addressing bottlenecks outside of it:

- Static content can be offloaded out of the data center through caching.
- By offloading content and storage in the cloud, WAA reduces the bandwidth needed to support your application.

Figure 5 shows 75 percent origin offload. This and other key performance metrics are conveniently monitored using Akamai’s customer portal. When combining the offload capabilities with the Akamai Protocol, the result is a dramatic reduction in server hits, freeing up precious server resources while optimizing power and rack-space within the data center.

This graph shows the extent to which Akamai Edge Servers reduced the number of hits on your origin.

Origin Hits Reduction: 74.6%




Figure 5. Origin Offload of Hits

Security

With a growing number of Distributed Denial of Service (DDoS) attacks—websites and applications are frequently at risk of attack. Akamai improves security by eliminating the public entry points to corporate infrastructure and taking the initial hit at the Akamai edge, outside of the data center.

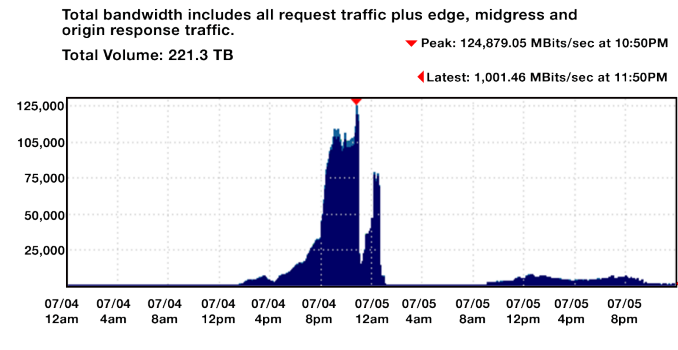


Figure 6. Akamai takes the first hit at the edge, reducing attack risks

Figure 6 shows what can happen to an origin under a DDoS attack. In this example, almost 125 Gbps of attack traffic was aimed at the origin. Akamai’s highly distributed EdgePlatform was able to absorb the traffic without causing origin failure.

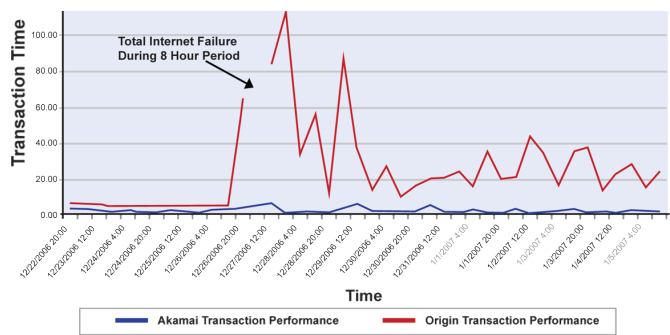


Figure 7: Akamai avoids application availability issues

The availability benefits of SureRoute include:

- Determination of a path that avoids Internet bottlenecks
- Elimination of end-user denial of service caused by Internet issues outside of your control
- Transparency to end users, who are simply happy that your application or website is accessible, compared to other applications or sites, not using Akamai, which are not available

The graph in Figure 7 shows the average time taken to run a simple web transaction—equivalent to a user viewing a web page—around the time of the December 26, 2006 Taiwanese earthquake. The tests were run from three locations in China: Beijing, Hong Kong and Shanghai. The web server was located in the U.S. state of Pennsylvania. Users viewing the page without Akamai would have seen total failures, or would have waited for close to two minutes. A week later, these users still had an average 30-second delay. Users viewing the page when it was delivered by Akamai noticed very little issue accessing content.

Application Optimization Solutions with WebSphere Application Accelerator for Public Networks and WebSphere DataPower

WebSphere Application Accelerator for Public Networks together with WebSphere DataPower combines the best-in-class Internet application delivery technology from Akamai with best-in-class enterprise application infrastructure technologies from IBM. This provides end-to-end optimization, acceleration, security and management from the user through the Internet to enterprise services and back. The next few sections describe some of the capabilities that are possible with this application optimization solution.

Smarter caching brings applications closer to the user

The Internet itself is often the root cause for application delivery problems to distributed users. To overcome this problem, the IBM *WebSphere Application Accelerator for Public Networks* uses Akamai's global overlay network to physically bring application resources closer to the end user by caching them at the edge of the network.

The WebSphere DataPower appliance—with *WebSphere Application Accelerator for Public Networks* capabilities configured—leverages Akamai's advanced distributed caching engine to cache resources that were previously not cacheable.

One example of this is the caching of RESTful service components. Many Web 2.0-style applications today use a RESTful pattern from the web browser to retrieve information integral to dynamically rendering the web application, creating a rich user experience. In the quintessential RESTful pattern, an HTTP "GET" is used to retrieve resources, while an HTTP "PUT" is used to update. The standard Akamai technology would not be able to cache these resources at the edge because they are dynamically generated by the application and not identified as cacheable content. However, when the *WebSphere Application Accelerator for Public Networks* product is implemented with a WebSphere DataPower appliance, the appliance can recognize the complex patterns of cacheable RESTful components and provide an appropriate injection of meta-data to make these cacheable. In addition, the WebSphere DataPower appliance can keep track of HTTP PUT requests that update the dynamic resource and work with *WebSphere Application Accelerator for Public Networks* to invalidate the corresponding cache entry.

The application intelligence provided by the WebSphere DataPower appliance and *WebSphere Application Accelerator for Public Networks* in this scenario provides capabilities that were previously not possible. Users can experience dramatically improved response times for both static and dynamic content, while significant processing is offloaded from the backend web applications at the same time.

The graph in Figure 8 illustrates response times for a distributed user base in three different architectures. The top line uses neither DataPower nor Akamai technology; the second line improves response time using default DataPower and *WebSphere Application Accelerator for Public Networks* settings; and the third shows dramatic improvements with DataPower and *WebSphere Application Accelerator for Public Networks* configured to work together.²

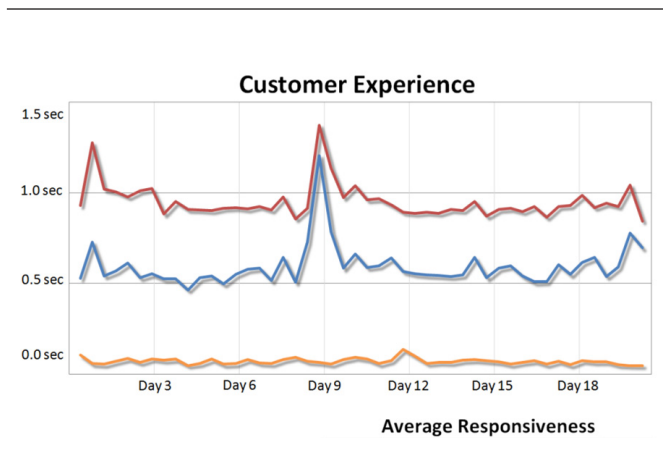


Figure 8.

Adaptive load optimization responds to changes in traffic demands

Another significant capability is the ability to adapt to changing load conditions when *WebSphere Application Accelerator for Public Networks* leverages DataPower's knowledge of the application state in the backend. When enabled with the Application Optimization option, DataPower intelligently senses and reacts to application workload conditions, not only balancing traffic across a WebSphere Network Deployment infrastructure, but also dynamically modifying Web Application Accelerator cache settings in the cloud to adapt and throttle traffic reaching the data center. During spikes of activity, users will still enjoy snappy response time with information fed from the Akamai cache. As the backend workload decreases, the caching duration of dynamic data is decreased. This improves application availability and responsiveness while still providing the freshest data to users. This is illustrated in Figure 8, where peak loads caused performance degradation, but with *WebSphere Application Accelerator for Public Networks*, performance remained fast and consistent.

Tighter cloud security: Defense in depth

The WebSphere DataPower appliance with *WebSphere Application Accelerator for Public Networks* can provide end-to-end protection of application infrastructure. By leveraging the adaptive load optimization described above, *WebSphere Application Accelerator for Public Networks* is able to provide improved protection against DDoS attacks. Since the Akamai network is built to absorb DDoS at the edge, DataPower can dynamically increase cache setting timeouts for all content—including content which is normally dynamic—thereby enabling Akamai protection of the origin infrastructure by absorbing those attacks at the edge, while continuing to provide fast and reliable service to end users. When the attack subsides, DataPower returns cache settings to their normal state.

The WebSphere DataPower appliance together with *WebSphere Application Accelerator for Public Networks* provides a second layer of security protection with a security-enhanced handshake between the Akamai network and DataPower, which helps ensure that only traffic carried over the Akamai network reaches application infrastructure—thereby extending the security perimeter into the cloud and protecting the origin from external threats. (See Figure 9.)




Figure 9.

The pervasive nature of the Internet creates tremendous opportunity for businesses to expose critical applications directly to end users on a global scale. After investing heavily to develop and promote increasingly complex online applications, enterprises are realizing the same traits that make the Internet universal can also undermine performance, reliability and security.

Additional benefits and capabilities

In addition to the performance, availability and security benefits, the Akamai Ready solution can also provide improved time to value through a streamlined implementation process. Many of the *WebSphere Application Accelerator for Public Networks* features can be deployed quickly, with little to no customization. An example of this is the SureRoute service provided by DataPower. The *WebSphere Application Accelerator for Public Networks* requires a SureRoute service at the web application origin to which constant measurements are taken to find the fastest and most reliable routes. Typically, this service would need to be integrated into a customer's web application or web server. The Akamai Ready Starter Toolkit provides customizable templates to minimize the time necessary to customize and deploy these cutting-edge capabilities comes with the SureRoute service ready to run.

The Akamai Ready Starter Tool Kit is available at no charge from IBM.com. It is strongly recommended that customers utilize the standard professional services offering from IBM to implement their *WebSphere Application Accelerator for Public Networks* solution. These service professionals work with you to identify opportunities to maximize time to value and return on investment.

Summary

WebSphere DataPower and *WebSphere Application Accelerator for Public Networks* work together to solve problems and amplify application delivery from your enterprise, through the cloud, to global end users. The Akamai Ready WebSphere

DataPower appliance acts as a smart intermediary, balancing, securing and optimizing online traffic from within your data center, while the *WebSphere Application Accelerator for Public Networks* acts as a high-speed Internet overlay, providing improved performance and reliability to global end users, with an added layer of security. When combined, these solutions amplify performance, reliability, scalability and security, delivering greater value than the sum of their parts. In this case, one plus one does in fact equal three.

For more information

To learn more about the IBM *WebSphere Application Accelerator for Public Networks* please contact your IBM sales representative or IBM Business Partner, or visit the following website: ibm.com/websphere/datapower

Additionally, financing solutions from IBM Global Financing can enable effective cash management, protection from technology obsolescence, improved total cost of ownership and return on investment. Also, our Global Asset Recovery Services help address environmental concerns with new, more energy-efficient solutions. For more information on IBM Global Financing, visit: ibm.com/financing

*The WebSphere Application Accelerator for Public Networks Toolkit has a non-warranted license and a license agreement must be accepted in order to download the toolkit.



© Copyright IBM Corporation 2011

IBM Software Group
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
April 2011
All Rights Reserved

IBM, the IBM logo, ibm.com and WebSphere are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at ibm.com/legal/copytrade.shtml

Other company, product and service names may be trademarks or service marks of others.

References in this publication to IBM products and services do not imply that IBM intends to make them available in all countries in which IBM operates.



Please Recycle
