The Impact of Web Services on the Network

Overview

For years, enterprises have struggled to find reliable, cost-effective ways to integrate and automate critical processes between different application packages. Web services technology has the potential to answer an enterprise's needs, providing the ability to integrate different systems and application types regardless of their platform, operating system, programming language, or location.

The deployment of Web services can take many different approaches through the use of middleware via a centralized broker or bus to a distributed approach, treating Web services as independent entities on the network. Resembling different network topologies, Web services could be deployed in the traditional, centralized or brokered client/server approach or they could be peer to peer, being distributed and driving control to the network itself. In a peer-to-peer structure, Web services could act as a server, providing functionality to a requestor or act as a client, receiving functionality from any service. However, in this type of structure, every node must be responsible for its own security, availability, performance, and management. Regardless of which topology you use, Web services are definitely going to tax your network infrastructure.

In the rush toward realizing the benefits of a Service Oriented Architecture via Web services, enterprise architects and developers must factor in the role the network plays in the successful delivery of Web services. Successful, from a network perspective, means achieving the core IT objectives of high availability, performance, and security for Web services. Ignoring the network will mean the failure of nearly any deployment due to any number of reasons as shown in the following figure:

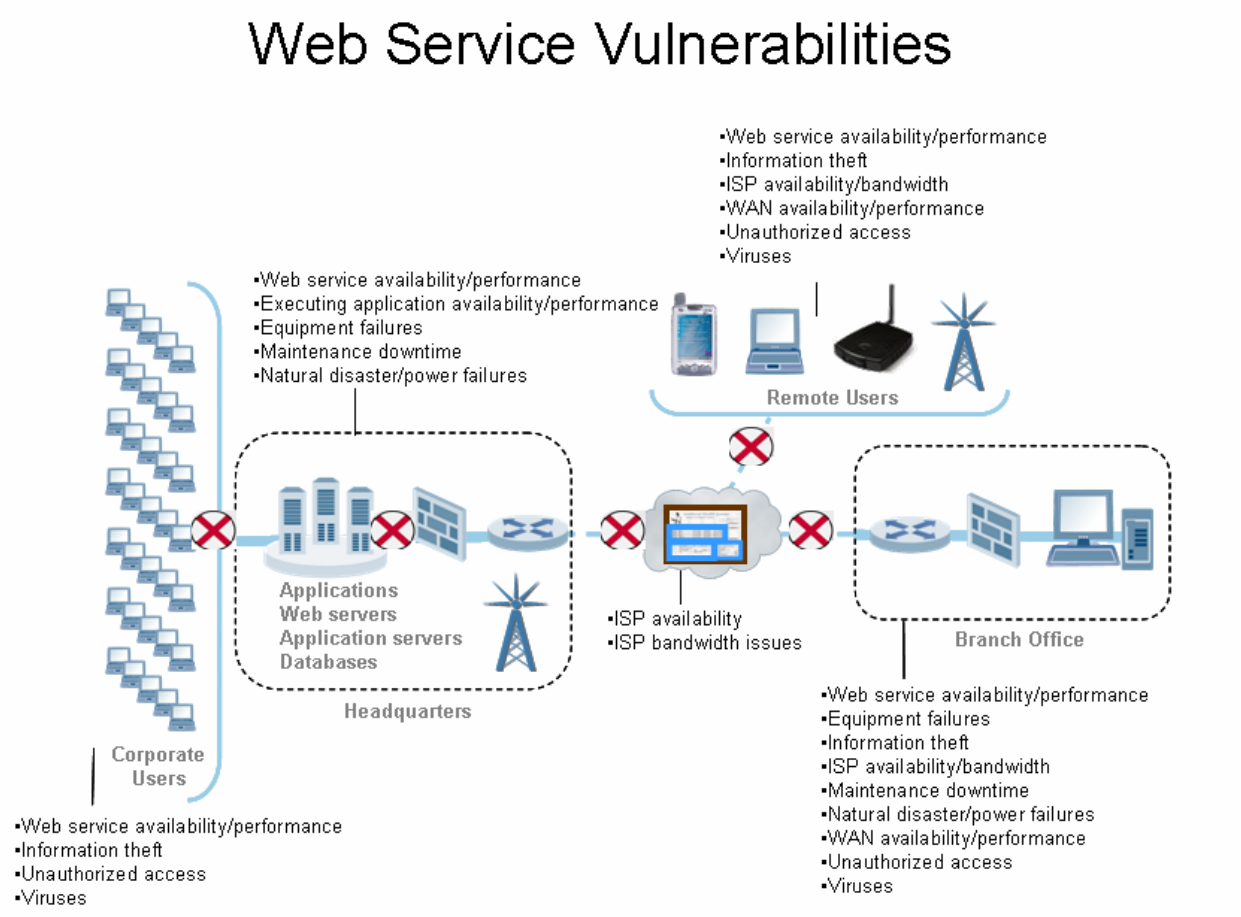

Challenges

Not understanding the network's capability means that the architects or developers must either design high availability, scalability, security, and performance optimization into every application or Web service itself or choose to ignore these objectives. The problem is that most enterprises have dozens if not hundreds of existing applications. Early research indicates that the number of Web services will be at least twice the number of base applications from which they are derived. So, in addition to the several dozen or hundreds of applications that organizations must support, they will also have to support hundreds if not thousands of discreet Web services.

Imagine having to design, implement, and maintain a different high availability or security scheme for each Web service. The cost, complexity, and security exposure would doom any organization to failure. However, success is achievable if architects and developers understand how to identify robust networking solutions to leverage their high availability, scalability, security, and performance optimization capabilities when deploying Web services.

To flourish, businesses must implement highly intelligent network products that can quickly process *any* application or Web service, ensuring quick response times, reliable sessions, easy scalability, and application-level security – all through a single network device. Additionally, the products they choose must be flexible, with an architecture that can easily accommodate future applications and protocols while protecting and improving the performance of their current applications.

The key challenges to keeping Web services up and running well, include:

- **Reliability –** The distributed nature of Web service applications demand a stable and reliable network environment and server infrastructure. With different components scattered across geographically-dispersed networks, reliable communication and application performance becomes paramount to deployment success.
- **Quality of Service –** In addition to communication reliability, organizations will need to prioritize requests. Requests will need to be intercepted, analyzed, and directed to the proper resource to provide quality of service based on an organization's business policies.
- **High Availability –** As the demand for Web services increase, the availability of each component within the service and the applications that process the requests will be critical. Key systems and devices that ensure Web service availability and reliability will be required to direct requests to healthy resources.
- **Scalability –** Flexible deployment scenarios will be required as demand for Web services increases. Organizations will need to act quickly to add resources to service requests without interrupting business.
- **Performance –** The quick adoption rate and ease of deployment of Web services will place huge demands on network infrastructures. Traffic generated by Web services can be significant – a single request can easily trigger 4 to 8 other requests. For example, requesting a stock quote might initiate as many as 8 components to perform functions to serve that user's request. In total, the entire transaction could generate up to 40 related requests.
- **Application Security** – The need to secure applications without sacrificing performance is extremely important. Enterprises must offload intensive SSL processing from application servers, allowing them to handle the performance demands required in a Web services environment.
- **Network Security –** When designing their Web services infrastructure, organizations are challenged with finding traditional network devices and tools that not only increase reliability, but also provide an extra layer of security. Network devices need

flexible, comprehensive, and secure feature sets to increase an enterprise's control over network traffic and protect the organization from existing and future attacks.

Solution

F5 divides Web service challenges into three categories:

- Availability
- Security
- Performance

The following sections describe the Web service challenges for each category and how F5 provides solutions for each challenge.

Web Services Availability

Web services are loosely coupled, self contained units of code that are exchanged via messages to notify each other of events, request information, or demand that an action be done on their behalf. Developers can write new Web services or they can use products that expose existing code as a Web service, which especially prevalent in financial services. Most often, these Web services are combined along with new code for use in a Web application. With this scenario, there are two areas of concern:

- Maintaining the availability of the Web application
- Maintaining the availability of each individual Web service in the Web application

The following figure shows a Web application that consists of three different Web services with the actual transactions being executed at the originating server, which are located in three different, geographically-dispered data centers.

If any one of these data centers or servers within a single data center goes down, the Web service and therefore, the entire Web application goes down. Also consider if the ISP link to any data center fails, the Web application goes down as shown in the following figure:

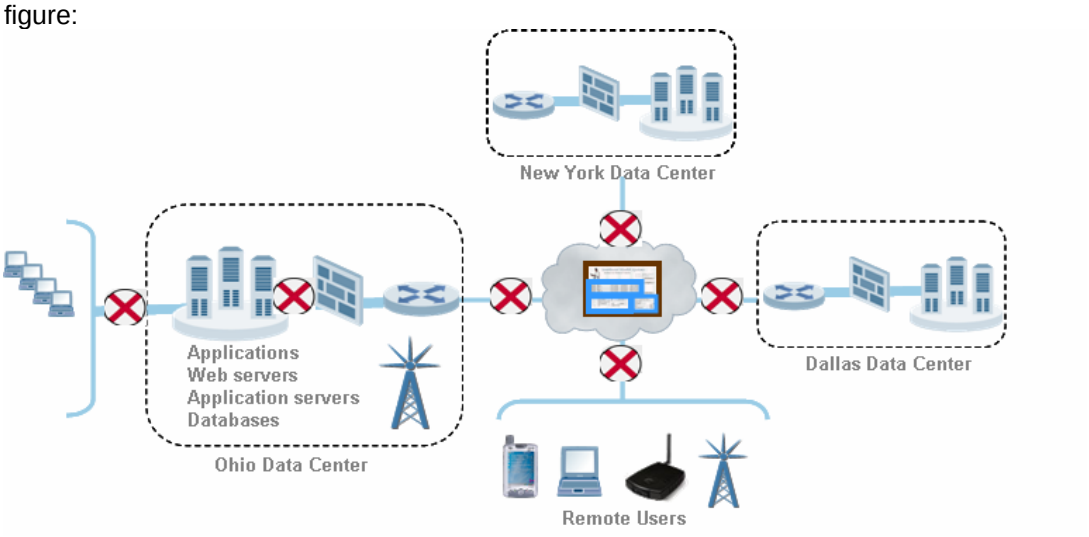

F5 takes a modular approach to making sure your Web services and Web applications remain up and running:

- **Multiple data centers** – ***BIG-IP® Global Traffic Manager*** maintains Web services availability that span more than one data center
- **ISP Links** – ***BIG-IP Link Controller*** maintains Web services availability over ISP links

- **LAN** – ***BIG-IP Local Traffic Manager*** maintains Web services availability within the data center

Web Services Availability across Multiple Data Centers

The BIG-IP Global Traffic Manager provides high availability for applications running across multiple and globally-dispersed data centers. The BIG-IP Global Traffic Manager distributes end-user application requests according to business policies, and data center and network conditions to ensure the highest possible availability. Using topology-based routing, the BIG-IP Global Traffic Manager can also route traffic over multiple links based on the user's location, choosing the best-performing link for the best end-user experience, while solving ISP Peering disputes.

By collecting performance and availability metrics from each data center, ISP connections, servers, caches, and even end-user content, the BIG-IP Global Traffic Manager ensures high availability and adequate capacity *prior* to directing traffic to a site. This gives organizations an intelligent way to manage application availability by holistically monitoring the health of Web applications and factoring in the dependencies between the application's Web services to automate the failover process when necessary. When the BIG-IP Global Traffic Manager detects a problem with the availability of a Web service in an application, it automatically and transparently reroutes users to a functional application with all the services running properly.

Another key component of availability is persistence, ensuring that users are directed to the right resources with the right content without breaking sessions or losing data. The BIG-IP Global Traffic Manager routes users to the same site regardless of their entry point to maintain consistency for applications or transactions such as e-trading, e-commerce, and financial services. Persistence information is also propagated to the local DNS servers to reduce the frequency of synchronizing back-end databases. In the case of worldwide availability, the BIG-IP Global Traffic Manager can resolve IP addresses down to the country, increasing topological control for managing global traffic and ensuring users get the information in their own language.

Web Services Availability over ISP Links

As organizations increase their use of the Internet to deliver Web applications, maintaining only one link to the public network exposes a single point of failure, which poses a serious network vulnerability. The BIG-IP Link Controller monitors the availability and performance of multiple WAN ISP connections to intelligently manage bi-directional traffic flows to a site, providing fault tolerant and optimized Internet access.

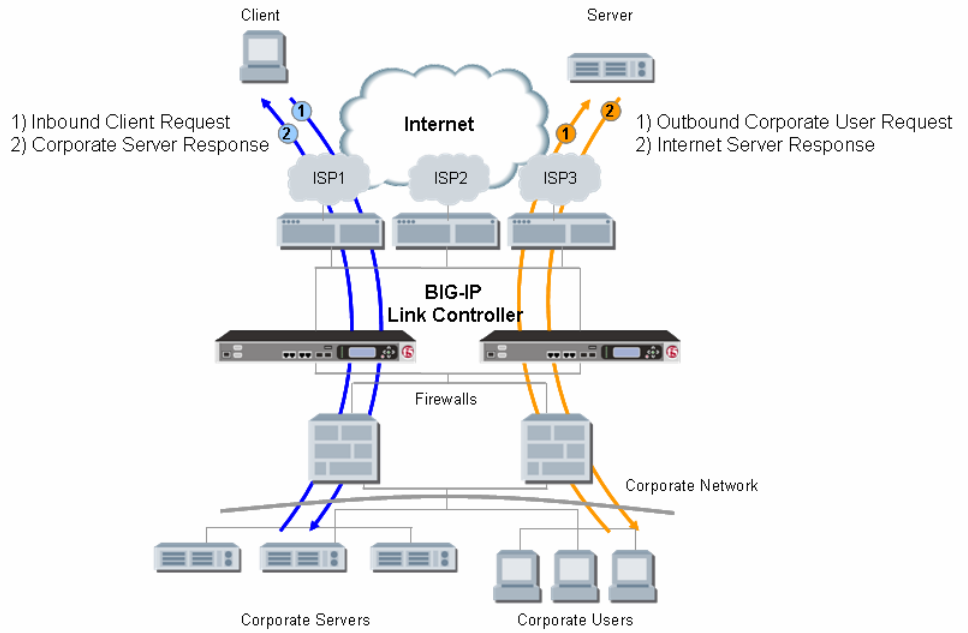

The BIG-IP Link Controller uses sophisticated monitors to detect errors across an entire link to provide end-to-end, reliable WAN connectivity. It monitors the health and availability of each connection, detecting outages to a link or ISP. In the event of a failure, traffic is transparently directed across other available links so end users stay connected.

The BIG-IP Link Controller simplifies multi-homed deployments so you no longer need ISP cooperation, large bandwidth connections, designated IP address blocks, ASNs, or high-end routers to protect your network from ISP failures. Using DNS-based technology that removes the dependency on BGP to provide failover capabilities, the BIG-IP Link Controller eliminates multi-homed problems such as latency, high update overhead, and inferior traffic management. You can also aggregate inexpensive links, with more control over which link to use based on performance, costs, and business policies.

Web Services Availability in the LAN

In a LAN environment, the BIG-IP Local Traffic Manager delivers sub-second failover and connection mirroring to maintain the availability of applications regardless of system, server, or application failure. Sophisticated monitors check device, application, and content availability. Rich static and dynamic load balancing methods track server performance levels to select the best resources. This solution also switches and persists on information unique to a specific vendor's application, server, or custom values for mobile and wireless applications.

Web Services Security

In the 2005 FBI Computer Crime Survey, of the 2,066 organizations with more than $1M in annual revenue that responded, 87% reported a computer security break within the last year, with spyware and viruses as the most common problem, followed by port scan, sabotage of data or networks, and adult pornography. In fact, spyware tripled in 2005 according to a recent Yankee Group report (http://www.cio-today.com/story.xhtml?story_id=02100000IOIO) And, over 50% of hacking attempts came from the US and China.

Today, IT has a variety of tools to enable security. Firewalls control information entering and leaving the enterprise while providing employees and partners with secure access to corporate resources. There are all kinds of point products to protect your networks from viruses, worms, Trojan horses, blended threats, and other unwanted content. But it's not just about stopping attacks – it's also about simultaneously serving legitimate users. F5 delivers the best of both worlds, providing a suite of security services that play a significant role in bolstering network and application security.

Web Services Security across Data Centers

Organizations are increasingly being exploited at the DNS level with DoS attacks that compromise the security of their site. Difficulty in differentiating between legitimate DNS requests and attacks is also a very real concern. The BIG-IP Global Traffic Manager's unmatched DNS performance can tolerate high levels of DNS attacks, protecting organizations while still maintaining maximum and continuous availability for applications and services.

Running on the latest version of BIND, the BIG-IP Global Traffic Manager inherits all the security protection against DNS cache poisoning that cause Pharming attacks. Packet filtering limits or denies access to and from Websites based on monitoring the traffic source, destination, or port. A hardened device, the BIG-IP Global Traffic Manager is designed to resist common attacks, including teardrop, ICMP, does not run SMTPd, FTPd, Telnetd daemons, zone file tampering, dynamic DNS-based attacks, implement DNS security policies via iRules, DNS attacks.

The BIG-IP Global Traffic Manager strengthens site security and diffuses attacks before they can start. iRules can help you create policies that inspect data packets to block DNS requests from rogue sites or known sources of attacks before they can do damage.

Web Services Security over ISP Links

To protect your network, the BIG-IP Link Controller Module ignores directed subnet broadcasts and does not respond to broadcast ICMP echoes used to initiate Smurf and Fraggle attacks. A connection table matches existing connections so that spoofed connections, such as in a LAN attack, are not passed on to the servers. The BIG-IP Link Controller checks for proper frame alignment to protect against common fragmentation attacks such as Teardrop, Boink, Bonk and Nestea. Other threats, such as WinNuke, Sub7, and Back Orifice are denied through the default blockage of ports. With the ability to reassemble overlapping TCP segments and IP fragments, organizations can thwart a new class of unknown attacks that are becoming more prevalent.

The BIG-IP Link Controller protects against a wide variety of attacks including common threats such as SYN and ACK Floods and also provides specialized detection features such as the SYN Check and Dynamic Reaping that detect and filter connection attacks from back-end servers.

Web Services Security in the WAN

With the influx of online financial services and e-commerce, people are regularly sending their most private, most personal data through the Internet. Although most Web sites work diligently to protect personal information, unscrupulous hackers can exploit unsecure Web sites to steal Social Security numbers and credit card numbers, which could be used to steal identities. Hackers can also uncover personal information while it's in transit to a Web site by capturing text in fixed patterns (xxx-xx-xxxx) from the payload.

Web services take data supplied to them via HTTP requests and either process transactions locally or act as an aggregation point for multiple external services residing on other back-end systems.

F5's WANJet protects data while it's in transit by providing site-to-site encryption so that all TCP traffic between WANJet appliances is fully encrypted. (WANJet encryption isn't required for sites with SSH (Secure Shell) between a client and server or if the site is using Secure HTTP (https://) since both SSH and HTTPS use SSL encryption.

WANJet encryption is based on the Transport Layer Security v1 protocol and supports:

- RSA is used for the shared (public and private) keys
- RC4_128 encryption algorithm (128 bit encryption)
- MD5 for the hash or message digest algorithm

Each WANJet ships with a self-signed certificate to ensure that a WANJet appliance will only establish an SSL tunnel with another WANJet appliance.

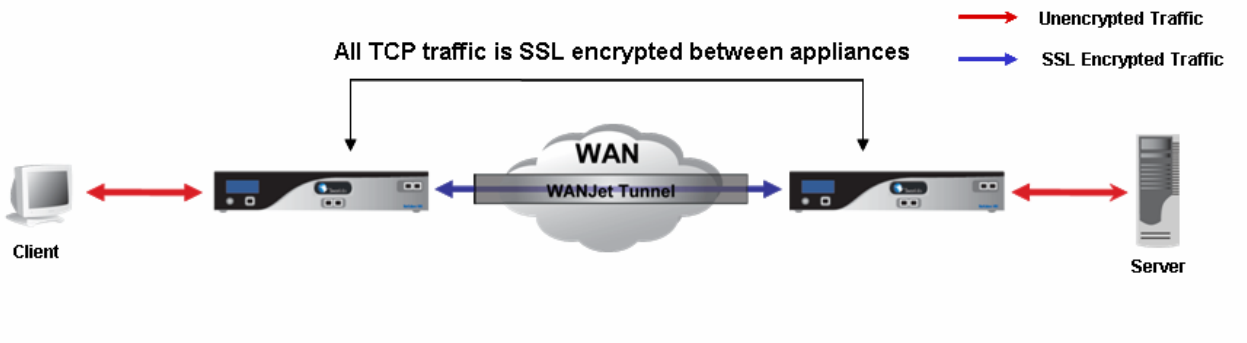

You can deploy WANJets at either end of the SSL tunnel to maintain security for insecure traffic over the Internet. WANJet also gives you the option of selectively encrypting certain types of traffic.

Web Services Security in the LAN

The BIG-IP Local Traffic Manager addresses security concerns using:

- An XML Gateway
- Authentication
- Encryption Services
- Application Protection

XML Gateway

Most security practices focus on protecting the LAN's resources, including servers, mainframes, application servers, Web servers, and databases. When you consider Web services, it's all about protecting the contents of the data packet as it travels from the Web application to the server that will execute the transaction with the data it received from the client.

The BIG-IP Local Traffic Manager uses an XML gateway for authentication and packet inspection to block, deny, or mask the contents of a packet. iRules, a component of the TMOS architecture for the BIG-IP Local Traffic Manager, has a built in XML parser that you can use to make traffic management and security decisions based on the content of the XML payload. This enables you to allow or deny specific Web service methods based on:

- The method name
- Required headers
- Method parameters

In addition to the built in support for SOAP messages, support for forward XPath expressions enables you to determine whether a given XML document adheres to a required format.

Employing the reverse proxy features of F5's TMOS and bi-directional TCP control, the BIG-IP Local Traffic Manager coupled with iRules can inspect information coming into your Web servers through HTTP requests, and going out to users through HTTP responses. This enables the BIG-IP Local Traffic Manager to analyze all data flowing through an HTTP session, down to the:

- Content
- Variable and value data
- Session management information
- Basic server response data

Many states are now forcing companies to notify their customer when their personal information has been lost or stolen. The BIG-IP Local Traffic Manager with iRules gives you the ability to search the header in a SOAP message and scrub (remove or replace) authentication information such as social security numbers, credit card numbers, account numbers, or any other identifiable text patterns from HTTP responses. This not only protects sensitive customer information from disclosure, but also protects your company's reputation.

The following iRule fragment replaces credit card numbers with Xs.

```
# Find ALL the possible credit card numbers in one pass
set card_indices [regexp -all -inline -indices {(?:3[4-
7]\d{13})|(?:4\d{15})|(?:5[1-5]\d{14})|(?:6011\d{12})} [HTTP::payload]]

foreach card_idx $card_indices {
  set card_start [lindex $card_idx 0]
  set card_end [lindex $card_idx 1]
  set card_len [expr {$card_end - $card_start + 1}]
  set card_number [string range [HTTP::payload] $card_start $card_end]

  set double [expr {$card_len & 1}]
  set chksum 0
  set isCard invalid

  # Calculate MOD10
  for { set i 0 } { $i < $card_len } { incr i } {
     set c [string index $card_number $i]
     if {($i & 1) == $double} {
        if {[incr c $c] >= 10} {incr c -9}
     }
     incr chksum $c
  }

  # Determine Card Type
  switch [string index $card_number 0] {
     3 { set type AmericanExpress }
     4 { set type Visa }
     5 { set type MasterCard }
     6 { set type Discover }
     default { set type Unknown }
  }

  # If valid card number, then mask out numbers with X's
  if { ($chksum % 10) == 0 } {
     set isCard valid
     HTTP::payload replace $card_start $card_len [string repeat "X" $card_len]
  }
}
```

Authentication

By acting as an authentication proxy for various traffic types, organizations can provide top level authentication for applications on the BIG-IP Local Traffic Manager. This allows you to push your security perimeter one level further down in the network (away from the applications), providing greater protection for your Web and application tiers.

The BIG-IP provides centralized AAA capabilities to limit, grant, and log access to specific information based on user credentials or SSL certificates. SSL involves the key exchange and then encrypting or decrypting traffic. Most SSL acceleration technologies only accelerate the key exchange with special hardware. F5's uses a hardware ASIC to accelerate both the key exchange and the encryption and decryption process so you don't have to trade performance for authentication.

If FTP or HTTP traffic comes to the device, the BIG-IP can take login information and interface with a 3rd-party authorization device, just like SSL works today. Authentication policies can be customized for various traffic types and are applied using rules or an authorization profile. Five common authentication profile types are supported:

Obtain credentials using HTTP basic authentication:

- LDAP
- RADIUS
- TACACS+

Obtain credentials using SSL client certificates:

- Client Certificate-based LDAP
- OCSP

Developers can also use iRules for custom authentication.

Encryption Services

The BIG-IP Local Traffic Manager provides a secure reverse proxy for SSL acceleration, termination, and re-encryption to Web servers. This enables the device to holistically, partially, or conditionally encrypt data using the Advanced Encryption Standard (AES) algorithms to provide the most secure SSL encryption available on the market at no additional cost. When coupled with an iRule, you can inspect the payload and if the user's request is not using at least 128 bits of encryption, the BIG-IP returns an informational error, preventing that user from accessing a secure Webpage as shown in the following iRule fragment:

```
when HTTP_REQUEST {
  # check for at least 128 bits of encryption
  if { [SSL::cipher bits] < 128 }{
    # when browser cannot do at least 128 bits of encryption
    # redirect to a un-encrypted page with an informational error
    HTTP::redirect http://10.10.10.10/error/sslerr.html
  }
}
```

With the BIG-IP, you can also specify policies that specify limited access for certain IP addresses, VLAN, users, etc.

Application Protection

The BIG-IP Local Traffic Manager protects your applications, data, and network resources, using the following features:

- **Resource Cloaking** – The BIG-IP virtualizes and hides all application, server error codes, and real URL references that may provide hackers with clues about infrastructure, services, and their associated vulnerabilities. For example, the following iRule fragment removes headers from server responses:

```
when HTTP_RESPONSE {
  #
  # Remove all but the given headers.
  #
  HTTP::header sanitize "ETag" "Content-Type" "Connection"
}
```

- **Content Protection** – The BIG-IP acts as a security proxy that is designed to protect the entire network against DoS and DDoS attacks, SYN Floods, and other network-based attacks. Combined with Dynamic Reaping, an adaptive method for reaping idle connections, the BIG-IP provides robust security to filter out the heaviest attacks while simultaneously delivering uninterrupted service for legitimate connections. SYNCheck™ provides comprehensive SYN Flood protection for the servers that sit behind the BIG-IP device.
- **Customized Application Attack Filtering** – BIG-IP's packet inspection and event-based iRules greatly enhance your ability to search for and apply rules to block known L7 attacks while also defining policies for blocking access or disallowing commands to be run.
- **Protocol Attacks** – The BIG-IP uses Protocol Sanitization and a Full TCP Termination point that independently manages client-side and server-side connections, protecting all back-end systems and applications from malicious attacks.
- **Application Firewall** – The BIG-IP integrates a control point to define and enforce L4-based filtering rules (based on PCAP that are similar to network firewalls) to improve network protection.
- **Identity Theft** – Coupled with iRules, the BIG-IP can replace or remove customer social security numbers, credit card numbers, account numbers or any other identifiable text pattern from servers, protecting against disclosure, fines, and damage to your business reputation.
- **Hardened Appliances** – Protects servers from attacks and ensures that only valid responses get through. Even the most sophisticated attacks can be efficiently identified, isolated, and eliminated without producing any negative effect on the site's performance and without harming legitimate application transactions.
- **Random Attacks** – Application-layer packet inspection and behavioral logic protect against counterfeit application activity, providing precise attack mitigation and granular blocking against script kiddies, known worms and vulnerabilities, requests for restricted object and file types, and other known exploits.
- **Web Server Protection** – Hides your Web infrastructure so that hackers can't tell what Web servers you're running. The BIG-IP strips out identifying OS and Web server information from message headers, conceals HTTP error messages from users, and removes application error messages from Webpages sent to users while checking to make sure no server code leaks out onto Webpages.

Web Services Performance

Zona Research reported that over $25 billion dollars are lost every year due to poor Web performance. This could mean unhappy customers that result in loss revenue and unhappy employees that result in lost productivity.

XML traffic via Web services is more bulky, so it consumes more bandwidth. More bandwidth usually means buying bigger pipes, but that's not always economically feasible not does it always solve the problem. Why not use your existing resources more efficiently? F5 offers a number of options, depending on the end user's location:

- Data centers and branch offices
- Mobile workers
- ISP Links
- LAN

LAN-Like Performance over the WAN

WANJet™ is an appliance-based solution that delivers LAN-like application performance over the WAN. It does this by accelerating file transfer, email, client/server applications, data replication, and other operations, resulting in predictable, fast performance for all WAN users. WANJet solutions work seamlessly across all wide-area networks including dedicated links, IP VPNs, frame relay, and even satellite connections.

Operating at Layer 5 of the OSI reference model, WANJet has full application knowledge and network awareness. Using adaptive protocol optimization, site-to-site encryption, and quality of service applied to application streams, WANJet delivers significantly more bandwidth for applications, effectively expanding WAN capacity. Unique technologies such as Transparent Data Reduction, ensure high-speed application performance and reduce the amount of data transferred over the WAN by up to 95%.

Web Application Performance for Mobile Workers

Mobile workers access enterprise applications from coffee shops, airports and offices. These workers expect their Web applications to perform well in all locations. If any part of the application delivery system falters, end-to-end performance degrades and productivity suffers.

F5 WebAccelerator™ is an advanced application delivery solution that provides superior Web application performance for mobile workers. WebAccelerator accelerates Hyperion™, Peoplesoft™, Plumtree™, SAP™, Siebel™ and other Web applications, increasing interactive performance by up to 500%.

Installed single-ended in the data center or dual-ended in branch locations, WebAccelerator enhances interactive performance from any location, improves download times for static and dynamic data, utilizes bandwidth more efficiently, and reduces the cost of delivering Web-enabled applications to the mobile workforce.

Web Services Performance over ISP Links

Compressing traffic without control over specific types of users (broadband users, dial-up users, etc.) can adversely affect application performance as well as the client experience. Using Round Trip Time and line quality calculations, the BIG-IP Link Controller dynamically calculates user latency and bandwidth throughput, devoting more compression power to those users who will benefit most.

TCP protocol inefficiencies can also cause unnecessary chattiness that adversely affects bandwidth utilization of the link. The BIG-IP Link Controller leverages TCP Express, a highly optimized TCP stack, to overcome TCP protocol inefficiencies, delivering efficient

bandwidth utilization of the WAN link by completely filling the pipe, prioritizing bandwidth for mission-critical applications, and improving end-to-end performance for dial-up and broadband clients over the WAN.

Using its topology database, the BIG-IP Link Controller can accurately determine the location of users and route traffic over the desired link based on pre-defined policies. This lets you choose the best performing link to deliver a superior end user experience based on location, while avoiding inter-ISP routing issues that can result in high latency and poor performance.

Web Services Performance in the LAN

The BIG-IP Local Traffic Manager's TMOS architecture includes a highly optimized TCP stack (TCP Express) to reduce TCP inefficiencies, delivering up to 80% performance gain for users and up to a 4x improvement in bandwidth efficiency. Using the latest RFCs and F5-unique improvements, F5's optimized TCP stack reduces latency due to retransmissions of lost packets, distance and network congestion, and improves end-to-end performance by providing a significant increase in bandwidth. And since TMOS is a fast application proxy, the BIG-IP can isolate client-side flows from server-side flows to independently optimize performance for each connecting device, translating communications between systems for improved application performance.

Additional BIG-IP Local Traffic Manager optimization techniques include:

- **Intelligent Compression** – By asymmetrically offloading HTTP compression from servers, the BIG-IP increases server capacity to decrease your TCO by 65%. A patent-pending approach determines connection latency, decreasing bandwidth usage by up to 80% while improving end-user response times by over 200%.
- **Cache Client Requests** – Aggregates millions of client application requests into hundreds of server connections, improving server capacity by up to 60%.
- **Prioritize Bandwidth for Critical Applications** – The BIG-IP uses Rate Shaping to control available bandwidth so that mission-critical and latency-sensitive applications get the bandwidth they need to perform well. This capability enables you to define traffic and application limits, control the rate at which those resources are allowed to spike or burst, provide queuing to prioritize traffic types, and define relationships where certain traffic types can borrow from other traffic types. Paired with iRules, you can isolate different traffic types and assign an unlimited number of policies to control bandwidth usage for each traffic type.
- **Multi-store Caching** – Manages distinct cache repositories per application or per department to intelligently control priority applications, boosting server capacity by 9x.
- **Offload SSL Processing** – Delivers up to 2 Gbps of sustained SSL throughput to offload SSL processing by accelerating the encryption of communications using more secure ciphers without a performance penalty. The BIG-IP sets the performance bar with best-in-market SSL TPS, bulk encryption, and the highest concurrent SSL connections available today.

Conclusion

The ability to provide interoperability between various applications running on disparate platforms and repurposing existing code makes Web services extremely versatile. Development toolkits from all the major vendors are making it easier for application developers to jump into the fray. But your job doesn't end once you post your Web application. Ignore the network impact on the deployment and delivery of your applications that consist of Web services and all your hard work may be at risk. Since a

one-size-fits-all approach rarely works as advertised, F5 offers a variety of solutions to keep your Web services up and running.

About F5

F5 Networks is the global leader in Application Delivery Networking. F5 provides solutions that make applications secure, fast, and available for everyone, helping organizations get the most out of their investment. By adding intelligence and manageability into the network to offload applications, F5 optimizes applications and allows them to work faster and consume fewer resources. F5's extensible architecture intelligently integrates application optimization, protection for the application and the network, and delivers application reliability – all on one universal platform. Over 10,000 organizations and service providers worldwide trust F5 to keep their applications running. The company is headquartered in Seattle, Washington with offices worldwide. For more information, go to www.f5.com.