

Network Infrastructure for Web Services

- Won't talk about applications, software, tools or platforms...
- Web services are also about networks
- Talk about network infrastructure:
 - Network equipment
 - Network services
- New protocols create new pressures and demands on existing networks

Eugene Kuznetsov, eugene@datapower.com

DataPower Technology, Inc. <http://www.datapower.com>



Why might WS not take off?

- Too difficult/expensive to develop or deploy
- Lack of sound business model / ROI
- Security fiascos
- Performance problems
- Reliability/availability
- Some of these also plagued the Web
- Both lessons and infrastructure from Web build-out

Quick History of HTTP LB

- Performance and reliability needs escalated well past those offered by one server
- Server load balancing (SLB) provided virtualization of web servers
- SLB quickly moved from software into content-aware network hardware (F5, ArrowPoint, Alteon)
- Complexity of load balancing algorithms quickly grew to meet application needs: URL switching, cookies...
- App. developers aware of SLB's, sometimes influences app architecture

Quick History of HTTP CDN

- Performance and reliability needs escalated well past those offered by one server cluster
- Global reach drove bandwidth costs
- Distributed edge caches moved content closer to user, reduce latency, load on origin servers, bw \$
- Network service and shared infrastructure (e.g., 10000+ Akamai servers)
- Interesting combinations with global load balancing, caching appliances within enterprise WANs, etc.
- Commoditization, consolidation and standardization

Overview of Network Security Infrastructure

- IP Firewall blocks incoming and outgoing traffic by IP address, IP port and often content
- Web proxies (transparent or not) process, possibly log and filter web requests
- Email scanners censor content and scan for viruses
- SSL accelerators offload SSL from servers
- VPN devices provide telecommuter and partner connectivity via virtualized networks
- Intrusion detection devices monitor & raise alarms
- DoS protection devices & services
- No knowledge of XML today

XML Family of New Protocols

- XML -- data-encoding for protocols
- SOAP & surrounding specs
- UDDI, WSDL
- BXXP
- XSLT/XPath
- XML DSIG/ENC
- SAML

WS Challenges

- Performance
 - XSLT bottlenecks in real projects
 - XML/SOAP
 - “it was fast enough on my machine”
 - success breeds performance problems
- Security
 - “concerns” limit WS deployment
 - can’t expose endpoints without guarantees
- Management
 - “where are your web services?”
 - deployment headaches

Network Needs


- Acceleration
 - bandwidth & processing intensive protocols
 - based on app. requirements
 - reserve of capacity or plan on success
- Security
 - programmatic access to ent. Backend
 - standalone, sw-independent security
- Management
 - more WS nodes → centralized management
 - RAS: reliability, availability & scalability

XML-Aware Network Devices

- New type of content-aware networking equipment
 - capable of XML-processing, SOAP support
 - New name: “WS-aware”, or “XML-aware”, or “XML-router”..
- Network hardware capable of parsing and processing XML data streams
- SOAP load balancer
- XML firewall
- XML accelerator/off-loader
- WS billing / payment processor

SOAP Load Balancer

- Receives requests and relays to pool of servers
- Server IP's hidden from client
- Load balancing scheduling
 - round robin, target URL/endpoint, server load
 - request, client, session
- Virtualizes server resources
- Hides server failure, overload
- Performs application-aware routing based on XML data

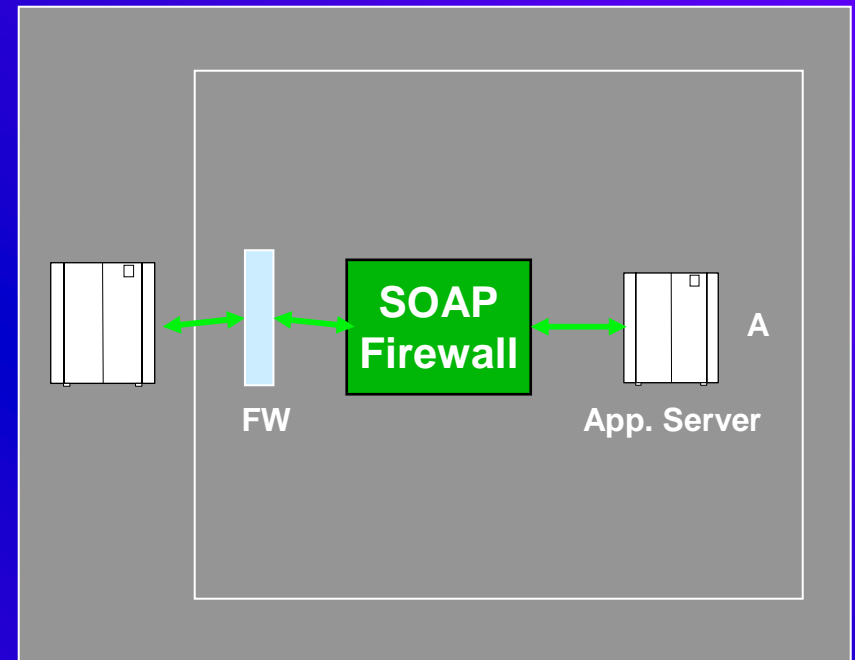


Question of Session State

- TCP Connection
- HTTP-layer bindings (cookies, URL)
- SOAP envelope parameter
- SSL session ID (?)
- Problem of long-running requests
- Asynchronous messages
- Endpoint exposure is a network infrastructure problem

XML (or SOAP) Firewall

- Intercepts incoming & outgoing WS requests
- Augments or incorporates IP-layer firewall
- Checks for well-formed XML
- Performs XML DoS checks
- Validates against standard (e.g. SOAP) or custom user schemas
- Applies XPath or similar rules
- Integrates with access control solution



Some Technical Challenges

- Using existing string matching / deep packet filtering
 - Example, look for “<foo>” in first 200 bytes
 - Problem #1: often hard limit on # of bytes
 - Problem #2: unaware of XML tree structure
 - Problem #3: no single canonical form (<foo/> or Unicode)
- Caching for performance
 - Caching technology to avoid XSLT delays
 - Caching for web services ... ?
- Performance challenges
 - Much more complex processing than previously attempted
 - Wirespeed is the network standard

More Technical Challenges


- Security via same engine/system as endpoint
 - Example, XML filtering
 - Problem #1: same vulnerability will affect FW as host
 - Problem #2: different requirements for technology
- Tightly coupled systems
 - Example, platform code in XML processing
 - Problem #1: not web services way!
 - Problem #2: makes migration of function into network difficult

XML Acceleration

- Offload resource-intensive tasks to purpose-built network hardware
- XSLT transformation
- Protocol gateways
- Schema validation
- Parsing/object marshalling

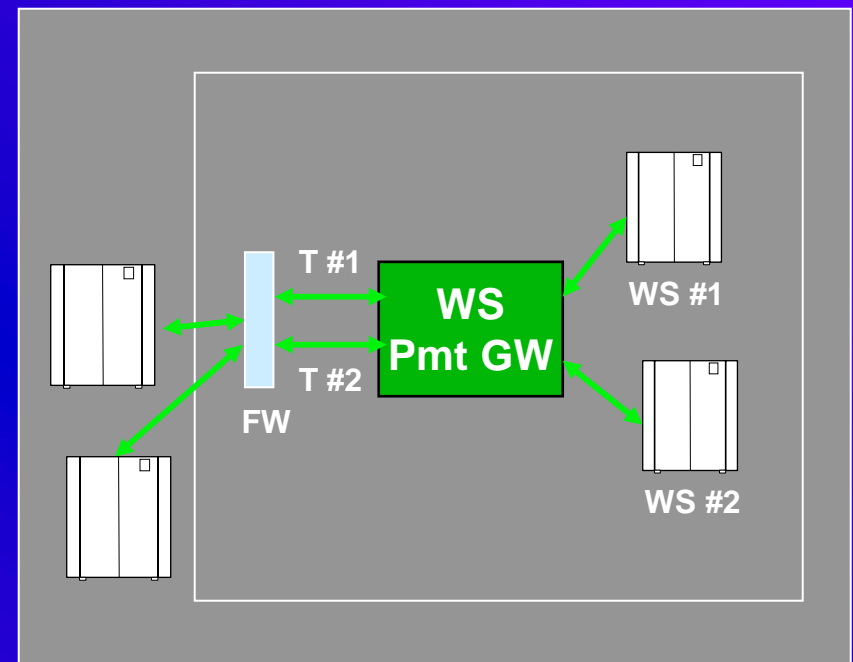
XSLT Accelerator Example

- Server returns XML content
- XSLT turned off on server
- Proxy passes HTTP requests to the server
- Intercepts XML server replies
- Converts XML to HTML for end user
- Non-XML server replies are unaffected
- XML doc refs XSL by URL
- Stylesheet retrieved from web server
- Easy setup, great performance



WS Payment/Billing GW

- Token-based web services payments
- Tokens as authorizations for specific web service use:
T#1-> WS1, T#2->WS2
- Based on XML-DSIG, PKI, XKMS & crypto → difficult to deploy across all servers
- GW provides a centralized access point to all web service end nodes
- Dedicated, central point of access control and management: PnP revenue



Network Issues Today

- Some practical considerations for network infrastructure today
- Uncertain scalability requirements
 - latency & throughput calculations
 - testing with projected loads is revealing
 - plan for off-loading to accelerators
 - existing load balancing h/w
- Security
 - no XML-aware firewalls yet (so be very careful)
 - existing SSL accelerator if using SSL
 - existing IP firewall (reconfigured)
 - performance impact
- Network Management
 - SOAP traffic on specific port
 - basic network statistics

XML-Aware Network Services

- New “in-the-network” services for WS
- Directories of web services, UDDI, etc.
- Edge processing and acceleration
- Guaranteed delivery of XML docs & transactions
- Cryptographic tokens and PKI certificates
- Managed security service providers for WS
- Fully outsourced deployment infrastructure

Network a year from now

- What kind of network equipment is likely to be associated with a web services project a year from now? Predictions.
- Acceleration
 - XSLT/XML accelerator
 - SOAP load balancer
 - Probably SOAP intermediary router
- Security
 - Custom configuration of existing IP firewall
 - Existing or integrated SSL accelerator
 - SOAP/WS/XML firewall (standalone or integrated)
 - Access server with integrated WS access control, XML
- Network Management
 - Deployment server management (existing)
 - Possibly web-services specific network management

Summary

- WS are also about XML-aware network infrastructure
- Apply lessons of the past
- XML performance is key
- Today's software functions → network functions