

Mozilla Infrastructure

Paul Querna

12/16/2004

Currently...

- Deploy new servers for specific tasks, one at a time.
- Add caching when performance degrades
- Difficulty meeting peak demands

New Architecture

- Goals:
 - Easy Monitoring and Maintenance
 - Performance
 - Scalability
 - Redundancy
 - Cost Effective

Load Directors




Content & Cache Servers



Application Servers

Databases



Load Directors

- Distribute Traffic to other Servers
- Public IP addresses
- Allow back-end servers to be shutdown
- Provide Fall-over for other LDs

No Load Directors?

- LDs are generally... idle.
- Cheaper, but:
 - Less Control over balancing (DNS Only)
 - Hosts are on or off...
- LD Advantages:
 - Easier to Manage Change
 - Fewer Public IPs

Linux Virtual Server

- Automatic IP address Fall-over
- Forwarding:
 - NAT
 - Direct Routing
 - LDs just change ethernet frame

Alternatives

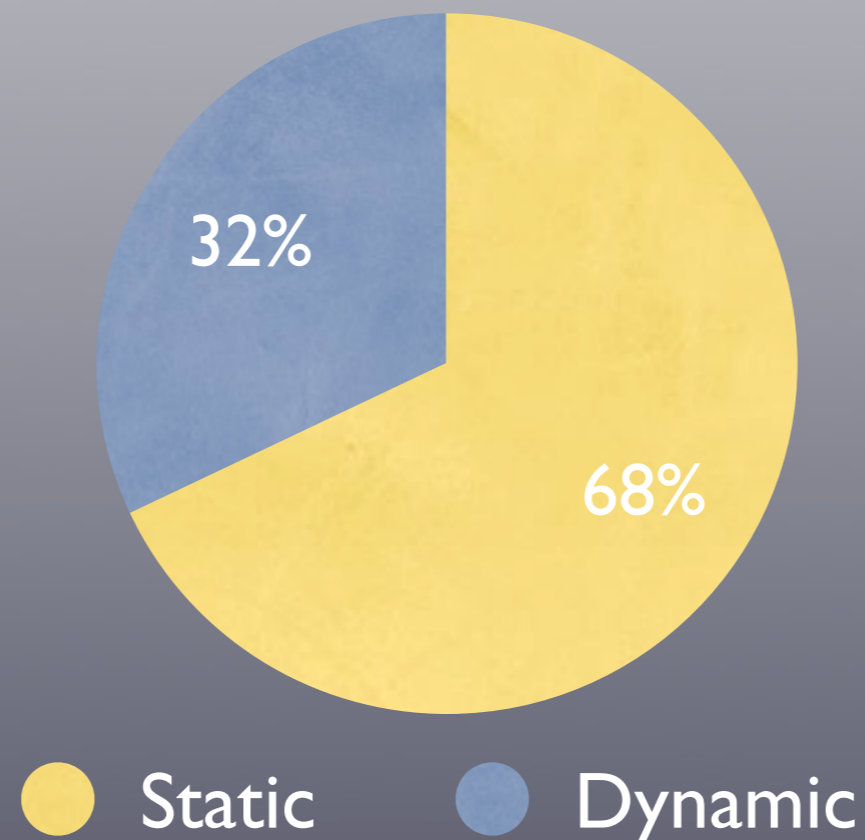
- OpenBSD/FreeBSD
 - CARP
- UCARP
- Cisco SLBs
- No Load Directors

Content and Cache Servers

- Cache Dynamic Content
- Serve Static Content
- Balance Load to Application Servers
- SSL Decoding

Caching

- Most Content can be cached!
- update.mozilla.org:



More on Caching

- About 50% of dynamic requests are cached (November Squid Stats)
- Need to work with Mozilla Devs:
 - HTTP Headers
 - Higher Cache Hit rates
 - 5 minute timeouts?

Caching

- Apache 2
 - mod_cache
 - mod_proxy
 - Load Balancing
 - Sticky Sessions
 - Pooled Connections
 - Event & Worker MPMs

mod_proxy_balancer

- Round Robin
- Weighted
- Disable down servers
- Today: Scripted
- Future: mod_proxy_backhand?

Static Content

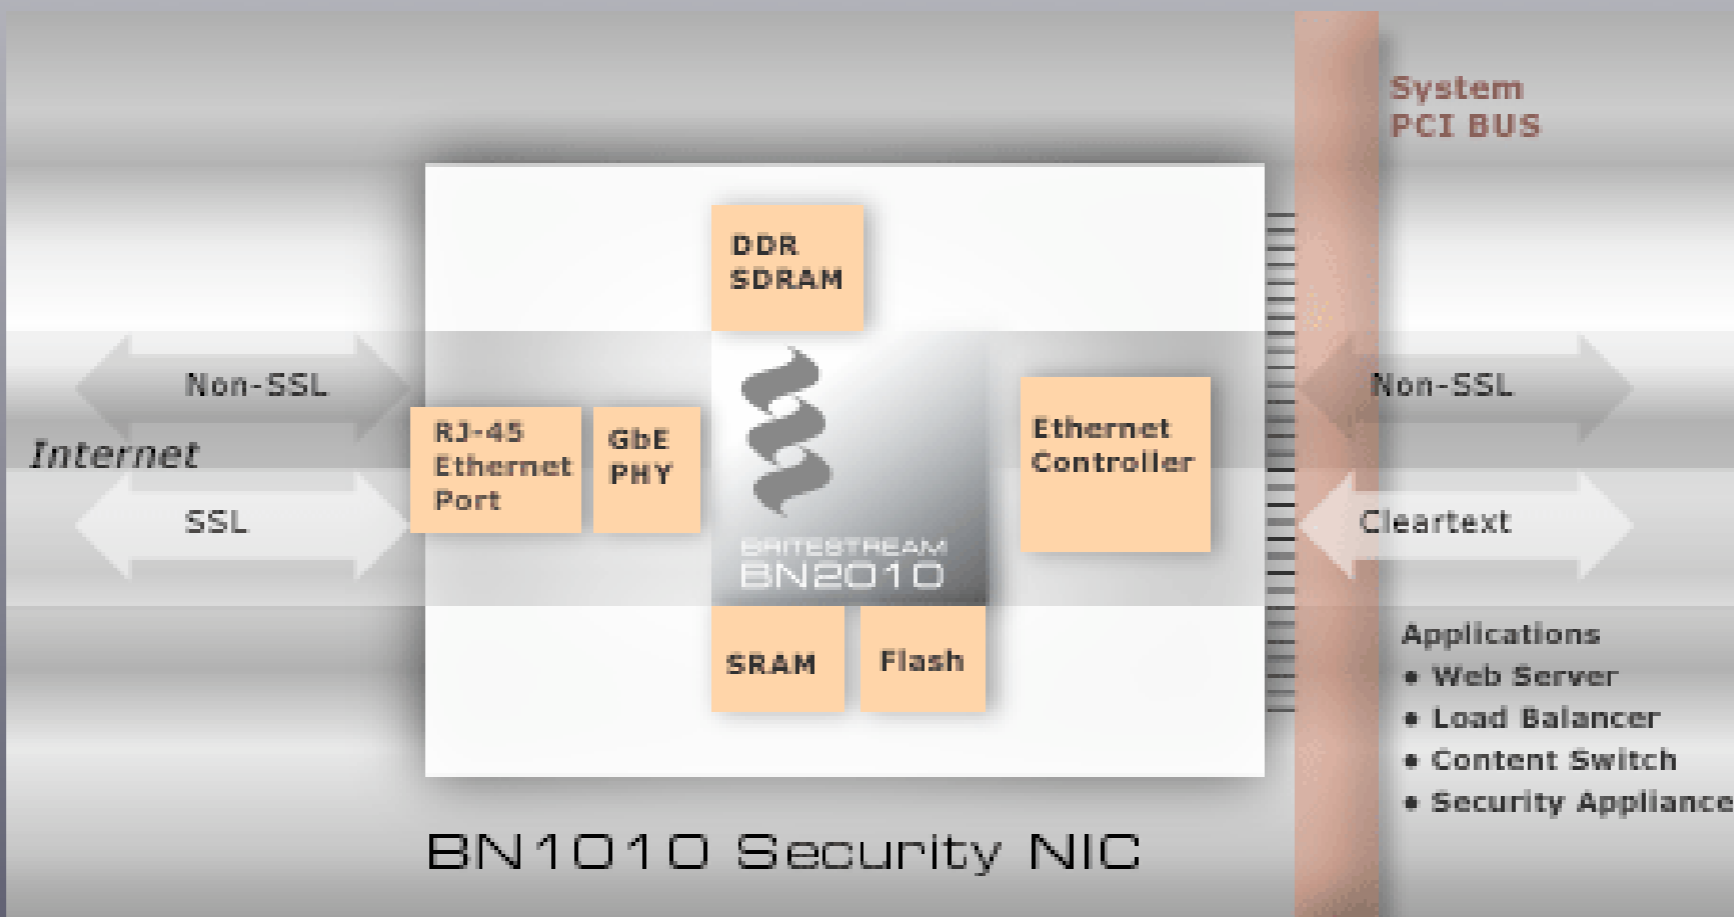
- Why here?
 - Flexibility
 - Application Layer Failures
 - www.mozilla.org is completely static!

SSL Decoding

- Why?
- Recommend Purchase:
 - Traditional SSL Accelerators ~ \$1500
 - PCI Card does Crypto Operations
 - Ethernet based products ~ ?
 - Zoom Zoom
 - SSL is 'free'



BN1010 Security NIC



BN1010 Security NIC

Alternatives

- Squid
 - Issues:
 - Static Content
 - Auth, mod_rewrite, etc..
 - Another Piece of Software to Manage
- No Caching?
 - More Hardware at Application Level

Application Servers






- Anything that serves HTTP could work.
- Keep it Simple:
 - Apache 2 for all PHP based sites
 - Everything is PHP?
 - Coordinate with Mozilla Devs

Applications need Data

- MySQL?
- Work with the Mozilla Devs:
 - Graceful Failure
 - Application Servers
 - Read Only Mirrors for common DBs
 - memcached?

MySQL

Application Servers



Load Directors




Content & Cache Servers



Application Servers

Databases



Monitoring is required.

- Near Real Time Stats at all levels
 - Apache: mod_log_spread
- Use Data to:
 - Short term, Change Balancer Weights
 - Long term, Adjust Number of Servers

Backups!

- No Backups:
 - Load Directors
 - Cache/Content Servers
 - Application Servers
- Database Servers
- Deployment Server

Deployment

- Must be Automated.
- Everything is in Version Control.
- Easy to add machines, or change their roles
- Local Packages for Apache, PHP, others.
- rsync Doc Roots
- Linux... But which one?

Starting a War..

- Gentoo
 - Gentoo Installer is not ready.
 - Can be done with plenty of custom work
 - Future Looks Promising.....

Network Booting..

- Debian
 - Fully Automated Install
 - Local Apt Repos
 - Recommend: Testing/Sarge
- Redhat/Fedora
 - Kickstart
 - Local Apt Repos

Scaling Out.

- Network Boot Server
- Fetches Data from Deployment Server
- Installs all Packages for Role
- Reboots.
- Starts Serving Requests
- 5-10 Minutes for a new Server

Real Numbers...

- Load Directors: 3 Machines
 - 2x NICs
 - Public
 - Private
 - CPU/Disk/RAM less important
 - Software: LVS

Middle Tier

- Cache/Content Servers: 3-5 Machines
 - 2x NICs
 - Public
 - Private
 - Max out RAM
 - SSL Accelerators
 - Software: Apache 2, mod_{proxy,cache}

Application Tier

- Application Servers: 3-5 Machines
 - 1x NIC
 - Private
 - Software: Apache 2, PHP, MySQL

Database Tier

- Database Servers: 2-3 Machines
 - 1 Master, 1-2 'hot' slaves
 - RAID 10 or 5
 - 1x NICs
 - Private

Administrative Machines

- Two Machines:
 - Backups/Monitoring
 - Build/Deployment/Monitoring
- RAID 10 or 5 on both

Questions?

- Slides Online:
 - <http://www.apache.org/~pquerna/osl.pdf>