

Apache Web Server Performance Tuning

Jaqui Lynch
Mainline Information Systems
Email – jaqui.lynch@mainline.com

Agenda

- Anatomy of a Web Transaction
- Web Server Parameters
- Security
- Network & System Parameters
- Scalability
- Web Page Coding & Other Tips
- Questions

Anatomy of a Web Transaction

What's in a URL?

- http://www.cmg.org:80/conference/index.html
- Http:// gopher, ftp, http, https - protocol
- www.cmg.org:80/
- Request template
 - conference/index.html
- Methods
 - Get and post
- URL means Universal Record Locator

Anatomy of a Web Transaction

- User enters URL (server waiting for reqs)
- Client does dns lookup to find out server ip
- Client connects TCP/IP to server
- Client sends http request
- Server analyses the request, reads file from local storage, sends mime header (i.e. text/html) and then the text
- Server may do a dns lookup
- Server makes a log entry of the get
- Client formats according to mime header
- Each gif, animation, etc is a separate get
- Each get is a separate request

Web Server Parameters 1

- Indexing no
- ServerType standalone
 - Standalone or inetd
- MinSpareServers 5
- Max SpareServers 10
 - HARD_SERVER_LIMIT in http.h is 256
- StartServers 5
- MaxClients 150
 - max concurrent threads/children
- MaxRequestsPerChild 30

Web Server Parameters 2

- Timeout 300
 - Wait on get, post, put, acks
- Log Types
 - Referrer
 - Agent
 - Access/transfer
 - Error
 - Combined
 - Scriptlog
- Logs can now go to separate logs, or to syslog
- Put logs in their own filesystem and on their own disks in a busy system
- Only keep logs you actually will use

Web Server Parameters 5

- HostnameLookups
 - Off, On or Double
- If setting this on on a busy server consider putting a local DNS server on the system
- Can turn hostname lookups off inside of server_status and other directives as follows:
- HostnameLookups off
 - <Files ~ "\.(html|cgi)\$">
 - HostnameLookups on
 - </Files>
- Leave off and use logresolve to postprocess the logs

Web Server Parameters 6

- `http://www.cmg.org` vs `http://www.cmg.org/`
- FollowSymLinks and SymLinksIfOwnerMatch
 - If not on it causes multiple lsstat commands if use short and long URLs
 - `http://www.abc.com`
 - `http://www.abc.com/htdocs/`
 - `http://www.abc.com/htdocs/index.html`
 - lsstat results are not cached
 - Use FollowSymLinks only
- AllowOverride
 - Set to none or Apache will try to open the default password file (usually `.htaccess`) every time you try to access this directory

Security

- No security
- Basic
 - Password & data clear text
- Digest
 - MD5 – only data clear text
- SSL
- Indexing
- Robots.txt and spiders/web crawlers
- Time Synchronization
 - Use NTP or similar

Network Parameters (AIX)

- Values set by no
 - Specifically TCP send and receive spaces
- Watch mbufs
- Watch transmit and receive queue sizes for the adapters
- 100mb or faster is better
- lsattr -E -l ent0
- lsattr -E -l css0

no Settings

Starting point for AIX with 100mb network

Put in /etc/rc.tune or rc.local and add to inittab

On SP this is tuning.cust

- thewall 262096 kb
 - sb_max 1310720 bytes
 - rfc1323 1
 - tcp_sendspace 655360
 - tcp_recvspace 655360
 - udp_recvspace 655360
 - udp_sendspace 1310720
- Make sure sb_max and pools are in 4096 byte multiples

System Parameters

(AIX starting point)

- maxuproc lsattr -E -l sys0
 - Increase drastically from 40 to 1000 or 2000
 - Server forks/spawns subprocesses for each connection

- minperm 10%
- maxperm 50%

- maxrandwrt 32 x 4kb pages

Other

- Run only one log analyzer
- Move logs elsewhere and postprocess
- Schedule backups carefully
- Use I/O pacing when running backups
 - chdev -l sys0 -a maxpout =17 -a minpout=12
 - Run the backup
 - chdev -l sys0 -a minpout=0
 - chdev -l sys0 -a maxpout =0

Scalability – Redirectors

DNS

- Round-robin DNS
- xyz.com has several ip addresses
- No idea if server is up or overloaded
- DNS changes propagate slowly so must use TTL=0
- All servers treated equally
- Can apply weights to the entries if using MX records
- Max of 32 entries per record

Scalability – Redirectors

Hardware

- Uses a black box to redirect sessions
- DNS points to the black box
- Redirector analyzes request against its table of mappings
- Maintains info on
 - Server response time
 - Network utilization
 - Application mix
- Detects failed servers and marks as unavailable
- Allows transparent deletes/adds of servers and services
- Some use MD5 or ACLs to reduce DOS
- Hardware changes may be needed to add protocols
- IP headers rewritten using NAT

Scalability – Redirectors

Software

- Uses software on one or more systems to redirect sessions
- DNS points to the systems
- Redirector analyzes request against its table of mappings
- Maintains info on
 - Server response time, Network utilization, Application mix
 - Metrics – CPU, network traffic, memory, I/O
 - Load balancing at the IP or URL level
- Detects failed servers and marks as unavailable
- Servers can be chosen using Ping Triangulation rather than current load – I.e. can balance by proximity or load
- Allows transparent deletes/adds of servers & services

Scalability – Redirectors

Software

- All control packets authenticated
- Uses UDP for performance between agents and controller
- Protocol changes do not need hardware changes
- Supports TCP and UDP applications including NNTP, pop3, SMTP, telnet, http, SSL, FTP
- Packets are not modified
- Only inbound client-to-server requests are monitored
- Outbound traffic goes directly from server to client
- Copes well with heterogeneous systems
- Allows dynamic reallocation of services plus a very quick way to increase server capacity

Other points to remember

- High speed private network to backend servers
 - Application servers
 - Reporting servers
 - Authentication and/or Authorization servers
 - Business Logic servers
 - Database servers
- Backups
 - Timing
 - Don't forget to plan for capacity on the backup server

Program Design

- Optimization & compiler options
- Languages
 - Compiled - C and C++
 - Interpreted - perl, Java & Javascript
- File Locking
- Where are the resources used
 - Javascript - on the client
 - C, perl, Java - on the server

Images

- Specify width & height
 - ` `
 - Reserves the space on the page and loads the picture later - otherwise you wait
- Provide text only options for dial users (alt=" ")
- Keep images small (<25k)
- Thumbnails for large images
 - Don't use the same physical image with a different width and height
- Flag thumbnails with the real image size and use them to represent large images
- Save as gif89a - limits size (jpgs are better for photo quality)
- Have only a small number per page

Special Effects 1/2

- Never on the home page
- Use cycles on your server
- Don't use cycles on the users system
 - Can also perform badly here and make your system look bad
 - You have no control over the client and thus performance
- Generally Java based - may not work properly on all browsers
- People turn off Java & Javascript so make them optional

Special Effects 2/2

- Tend to distract people from your message
- Blink annoys people
- Dissolves, fades, etc look different depending on the machine and the browser
- Pages load slower with animations and will just sit there blank
- Animations don't let you scroll while they are loading
- Pages don't load properly so the red loading light doesn't go out
- Animations are the closest thing the web has to viruses
- **If it is important enough to merit space on your server then let them read it in peace!**

Other things that really eat bandwidth

- Audio
 - Realaudio & realplayer
 - Netphone
 - Napster and other mp3 sharing programs
- Video
 - Flag the video with the size
 - Realvideo, CU-seeme
- Error checking
- Turn off debug options & open of stderr



Summary

- Above all use commonsense!!
- Plan ahead for the timing of system tasks
- Keep it simple
- Basic tweaking of system parameters can make a huge difference

- Any questions?
- Email – jaqui.lynch@mainline.com

