


Basic Apache Configuration and Tuning

Ken Coar


*coar@Apache.Org,
The Apache Group*





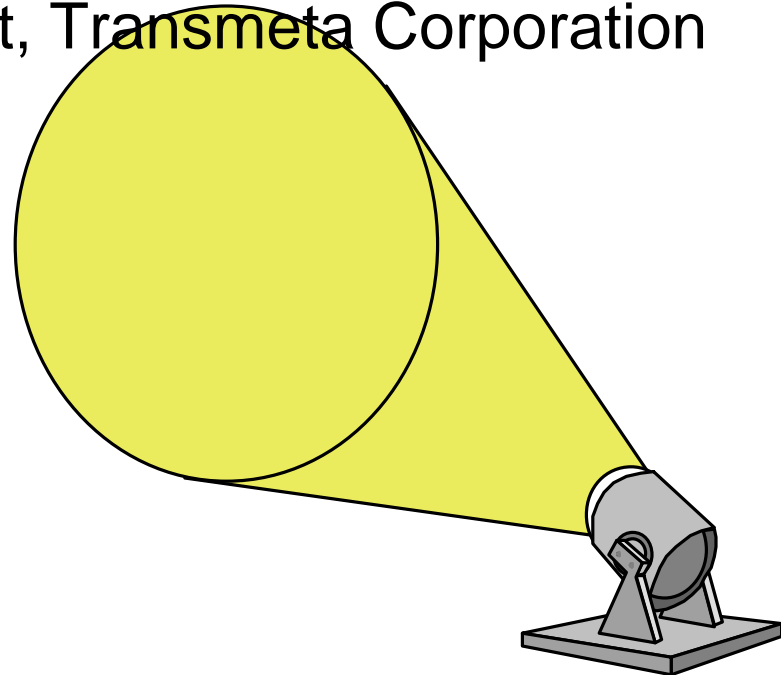
Disclaimers


- UNIX is a registered trademark of X/Open Company Ltd
- PostScript is a registered trademark of Adobe Systems, Inc
- All other trademarks are the property of their respective owners, and no misuse is intended



Acknowledgements


- I would like to gratefully acknowledge the contributions of the following:
 - Neil Gregory, Paul Scherrer Institut
 - Dean Gaudet, Transmeta Corporation






Copies of This Presentation

- *WILL BE* available online at
<http://Web.GoLux.Com/coar/>




Contents

- What *IS* Apache, Anyway?
- Getting Apache
- Installing the server
- Configuring Apache
- Logfiles
- Improving Performance
- Troubleshooting




What *IS* Apache, Anyway?

- Open-Source Web server originally based on NCSA server
- Available on over 160 varieties of Unix -- and Windows NT
- Over 56% of Internet Web servers run Apache or an Apache derivative



Getting Apache

- You can pull the distribution down from the Apache Web site
`<http://www.apache.org/dist/>`
- Some added-value vendors (e.g., RedHat) provide pre-built packaged versions
- Some companies offer support (for a price, of course!)
- The Apache site is mirrored all over the world, so there's probably a site nearer to you



Installing the Server on Windows NT


- Just run the installer
- If upgrading, be sure to save your content and configuration files, and **uninstall** the current version before installing the new one
- You can uninstall the current version by choosing

Start->

Settings->


Control Panel->

Add/Remove Programs




Installing the Server on Unix

- If you have a pre-built package
 - Install it
- Otherwise,
 - Unpack the tarball
 - Choose your modules
 - Compile the server
 - Modify configuration files to reflect your environment



Configuring Apache


- Choosing functionality
 - Apache functionality is available through modules which are either built into or loaded into the server
- Server instructions
 - Apache reads its run-time configuration instructions from text files
 - No GUI available
 - 182 configuration directives in base package



Configuring Apache


(continued)

- When used with `-d` option, server reads `httpd.conf`; `-f` allows use of a different name
- After `httpd.conf`, server reads `srm.conf`, then `access.conf` (unless the latter two have been renamed by `ResourceConfig` and `AccessConfig` directives, respectively)
- Consider combining these into a single file for simplicity




Pre-Forked or Multi-Threaded

- On Unix, Apache uses a pre-forking model, which means that several processes are created to implement the server and handle requests
 - Control *via* the parent process; don't touch the children!
- On Windows NT, there are only two processes to handle multiple request threads




Virtual Hosts

- A single Apache server installation can actually support hundreds of different Web sites
- NSPs and ISPs do this a lot
- Allows additional Web presence without accompanying hardware or software investment required




Content Negotiation

- If you have multiple language-specific versions of documents, Apache can automatically choose the correct one according to the client's language preferences
- Server can choose among different graphics formats; if it has both GIF and JPEG versions, it can send the one for which the client has expressed a preference




Scripts, SSIs, Servlets, and SSL

- Apache can handle run-time execution of external applications (scripts)
- Server-side includes (SSIs) allow the server to modify pages at run-time before sending them
- Java servlet modules are available, but not yet part of the base package (still under development)
- SSL-enabled versions of Apache are available, both from commercial vendors and under open-source licence




Logfiles

- Two basic logfiles
 - Access log -- who's been visiting your server and what they wanted
 - Error log -- problems the server has encountered and things it has noticed
- Can be configured for each virtual host, or for entire server
- Access log format can be customised



Improving Performance


- Have enough physical memory
 - Apache works best when it doesn't need to swap
 - Use `MaxClients` to limit concurrent requests to a rational and supportable value
- Have fast enough hardware
 - CPU, disk, network interface
- Don't use NFS! And definitely not for your `ServerRoot` or logfile directories



Improving Performance

(continued)


- Reduce filesystem lookups
 - Looking for and opening `.htaccess` files is expensive; apply `AllowOverrides None` liberally
 - Organise document tree to minimise cache thrash
 - Don't enable symbolic link tracing (omit `FollowSymLinks` and `SymLinksIfOwnerMatch` from `Options` directives)



Improving Performance

(continued)


- Disable DNS lookups with `HostNameLookups Off` except where needed (in `ScriptAlias` directory, `<Files .cgi>`, *et cetera*)
- Restrict content negotiation; only enable `Options MultiViews` where needed, and use `DirectoryIndex x y z` with the most common choices listed first
- Consider increasing the value of `MaxRequestsPerChild` unless base child memory consumption is large or there is danger of memory leaks



Improving Performance


(continued)

- Use `mod_status` to monitor server children; tune `StartServers`, `MinSpareServers`, and `MaxSpareServers` accordingly -- but IFF indicated
- Consider adjusting `KeepAliveTimeout` if `keepalive` connexions aren't likely (e.g., simple pages without inlined images)
- Don't use SSIs except where needed



Troubleshooting

- First and foremost, *check the error log!*
- Check the Apache FAQ
`<http://www.apache.org/docs/misc/FAQ.html>`
- Check the Apache bug database to see if it's been reported yet
`<http://bugs.apache.org/>`
- Check & ask on USENET
`comp.infosystems.www.servers.unix`
- **If all else fails**, submit a bug report
`<http://www.apache.org/bug_report.html>`



Going Further

- The main Apache site
<http://www.apache.org/>
- HTML Syntax Validator
<http://validator.w3.org/>
- CGI Specification:
<http://www.w3.org/pub/WWW/CGI/>
- CGI RFC project:
<http://Web.Golux.Com/coar/cgi/>
- USENET News:
comp.infosystems.www.servers.unix

© 1998 by Ken A L Coar

IBM Corporation,
MeepZor Consulting, and
The Apache Group