

Upgrading NCSA httpd

httpd 1.2 -> httpd 1.3

httpd 1.3 is a maintenance release designed to fix some of the bugs which were introduced in httpd 1.2.

Changes from httpd 1.2

- Now compiles cleanly under A/UX and Solaris.
- Directory indexing bugs repaired, new options to suppress last modified, size, and description columns (thanks Tanmoy.)
- Wildcard based access control should work this time.
- Fixed core dump related to recursive parsed document include.
- Fixed bug by which httpd would not follow a symbolic link it should have.
- Fixed file typing (no longer case sensitive like in 1.2.)
- Parent directory escaped in indexes.
- IdentityCheck should work again.
- LAST_MODIFIED and #flastmod are now local zone.
- Time related functions work with NeXT now.
- Fixed spelling error in 500 error page.
- Server will log proper timeout message for non-DNS hosts.
- Added x-bit hack for people to use the x-bit of an HTML file to determine if it is parsed or not.
- Added compile defines so people can customize the amount of DNS the server performs (for speed).

httpd 1.1 -> httpd 1.2

httpd 1.2 makes some significant changes over httpd 1.1. These changes make its "look and feel" a bit different than 1.1.

New features

- The CGI interface is now revised to version 1.1. See the changes that were made.
- A new logfile format was agreed upon which should simplify the lives of the authors of logfile analyzers. Quickly, the format is:

```
host rfc931 authuser [DD/Mon/YYYY:hh:mm:ss [+/-]HHMM] "request" status bytes
```

Host = hostname

rfc931 = RFC931 user name if IdentityCheck active, – if not.

authuser = HTTP/1.0 authenticated user, – if none.

DD/Mon/YYYY:hh:mm:ss [+/-]HHMM: Local time of the request with timezone offset from GMT at end.

"request" = the request as sent by the client

status = the HTTP/1.0 status code from this transaction

bytes = the count of the bytes sent in this transaction, not including the header. If not applicable, this will be a – character.

- CGI scripts are now allowed anywhere. See this tutorial on the new setup (and how it relates to the old setup), and how to disable this feature in places you don't want it.
- The server side includes interface has been completely rewritten. See this tutorial to learn how to convert your old INC SRV documents, and to learn how to take advantage of the new features.
- There is a new access control option to disable symbolic links only if the owner of the pointer is not the same as the owner of that which is pointed to. This means your users can have symbolic links to things they own, but not to dangerous things like /etc.
- A new access control method called mutual-failure has been added. This method is a bit unorthodox, but allows you to allow hosts from one domain while excluding certain hosts (such as public access machines) from that domain.
- Wildcard expressions are now allowed in various areas of server

configuration, to allow patterns to be specified. This is most useful in the Directory directive.

- Directory indexing has been revamped. It looks much different, and I've written a short tutorial on how to set it up.
- Access Control Files now allow the indexing directives as well as the DefaultType directive.
- require user now allows quotes for PGP usernames with spaces.
- Server now explicitly kills CGI scripts when the client aborts.
- Server now verifies the DNS hostname it gets from the IP number to prevent PTR spoofs.
- Support for 304 and If-modified-since.

Bug fixes

All of the known bugs in 1.1 have now been fixed. Now it's time to find the ones I introduce with 1.2.

- Fixed problem running scripts in ServerMode inetd under IRIX.
- Fixed bad port problem under OSF/1.
- Inserted missing return statement for the IdentityCheck directive.
- Fixed problem whereby errors would stop being logged after a restart.
- Fixed 256 character limitation on CGI URLs.

httpd 1.0 → httpd 1.1

httpd 1.1 should plug right in if you already have httpd 1.0.

- Introduced experimental PEM/PGP based encrypted user authentication. See this overview to read about trying it out.
- Improved directory indexing. See the new directives in srm.conf to use it.
- Cleaned up error output and fixed horrible output when server-side include error occurred
- Fixed slight bug in buffering code
- Directory indexing for user-supported directories fixed
- Retrieval of user supported directory with no trailing slash issues redirect

again

- Now supports RFC931 identd for logging purposes. See the IdentityCheck directive.
 - stderr for scripts and server side includes now sent to error_log
 - Fixed bug in NCSA POST script code which would cause Location: to be ignored for local files
 - Removed misfeature wherein Location: url's were being escaped by the server
 - Args to INC SRVURL escaped to avoid unpleasant surprises
 - Location: /cgi-bin/foo?arg now works
 - HEAD only for CGI scripts now ignores body put out by stupid scripts
-

httpd 1.0a5 → httpd1.0

httpd 1.0 is fairly similar to httpd 1.0a5.

Changes which directly require configuration changes

- ScriptAlias now used for CGI scripts only, access your NCSA scripts with OldScriptAlias
- Fixed AddType directive
- The NCSA scripts are no longer included, CGI replacements are now included

Other changes

- NCSA POST scripts now have REMOTE_HOST set
 - Added AddEncoding directive to srm.conf and .htaccess
 - error_log now logs all failed accesses and the reason for failure
 - Time headers are now RFC822 compliant
 - You can now alias or disable user-supported directories
 - Added a require directive to allow you to allow any valid user
 - I/O is now buffered into chunks for speed
 - Redirect no longer supported from .htaccess files.
 - Added -v and -f command line flags
-

httpd@ncsa.uiuc.edu