WHAT'S NEW IN HTTPD 2.4

Rich Bowen
rbowen@apache.org
@rbowen

Slides are at: **tm3.org/newin24**

TIMELINE

- 1.3 - June1998
- 2.0 - April 2002
- The earth cooled ...
- ApacheCon 2005, Stuttgart, my "Why I Hate Apache" lightning talk
- 2.2 - December 2005
- Feb 2010, 1.3 declared "end of life"
- Mid-2012, 2.0 EOL discussion. No formal action taken
- 2.4 - August 2012
- Current: 2.4.3

IMPORTANT EVENTS

- 1.3 ran more than 75% of the world's websites
- 2.0 introduced the idea of MPMs
- 2.2 added much new awesomeness
- 2.4 adds so much more awesomeness it makes 2.2 look positively neolithic

SO, WHAT'S NEW?

Well, let's get started ...

LOADABLE MPMS

MPMS

- Multi-Processing Modules
- Control how httpd deals with concurrent requests
- Worker, Prefork, Event, WinNT, etc

BEFORE ...

- If you wanted to change MPMs, you had to rebuild (or reinstall a package)

NOW ...

- Build MPMs as shared objects:

```
./configure  --enable-mpms-shared=all
```

NOW ...

- Load your desired MPM

```
LoadModule mpm_event_module \
modules/mod_mpm_event.so
```

BE CAREFUL ...

- Since some directives are MPM-specific, be sure that you test your config before putting it into production

WHY THIS IS AWESOME

- Easily test performance with other MPMs without having to rebuild
- Quickly revert when you discover it doesn't work

ERRORLOGFORMAT

BEFORE

- Error log format not configurable

NOW

- ErrorLogFormat has a similar syntax to LogFormat, and allows for configuration of the error log format

```
ErrorLogFormat "[%t] [%l] [pid %P] %F: %E: [client %a] %M"
```

SYNTAX

- Syntax is the same as LogFormat, with a different available variables
- Adds the %L variable, which allows you to correlate an error log entry with an access log entry

TYPICAL ENTRY

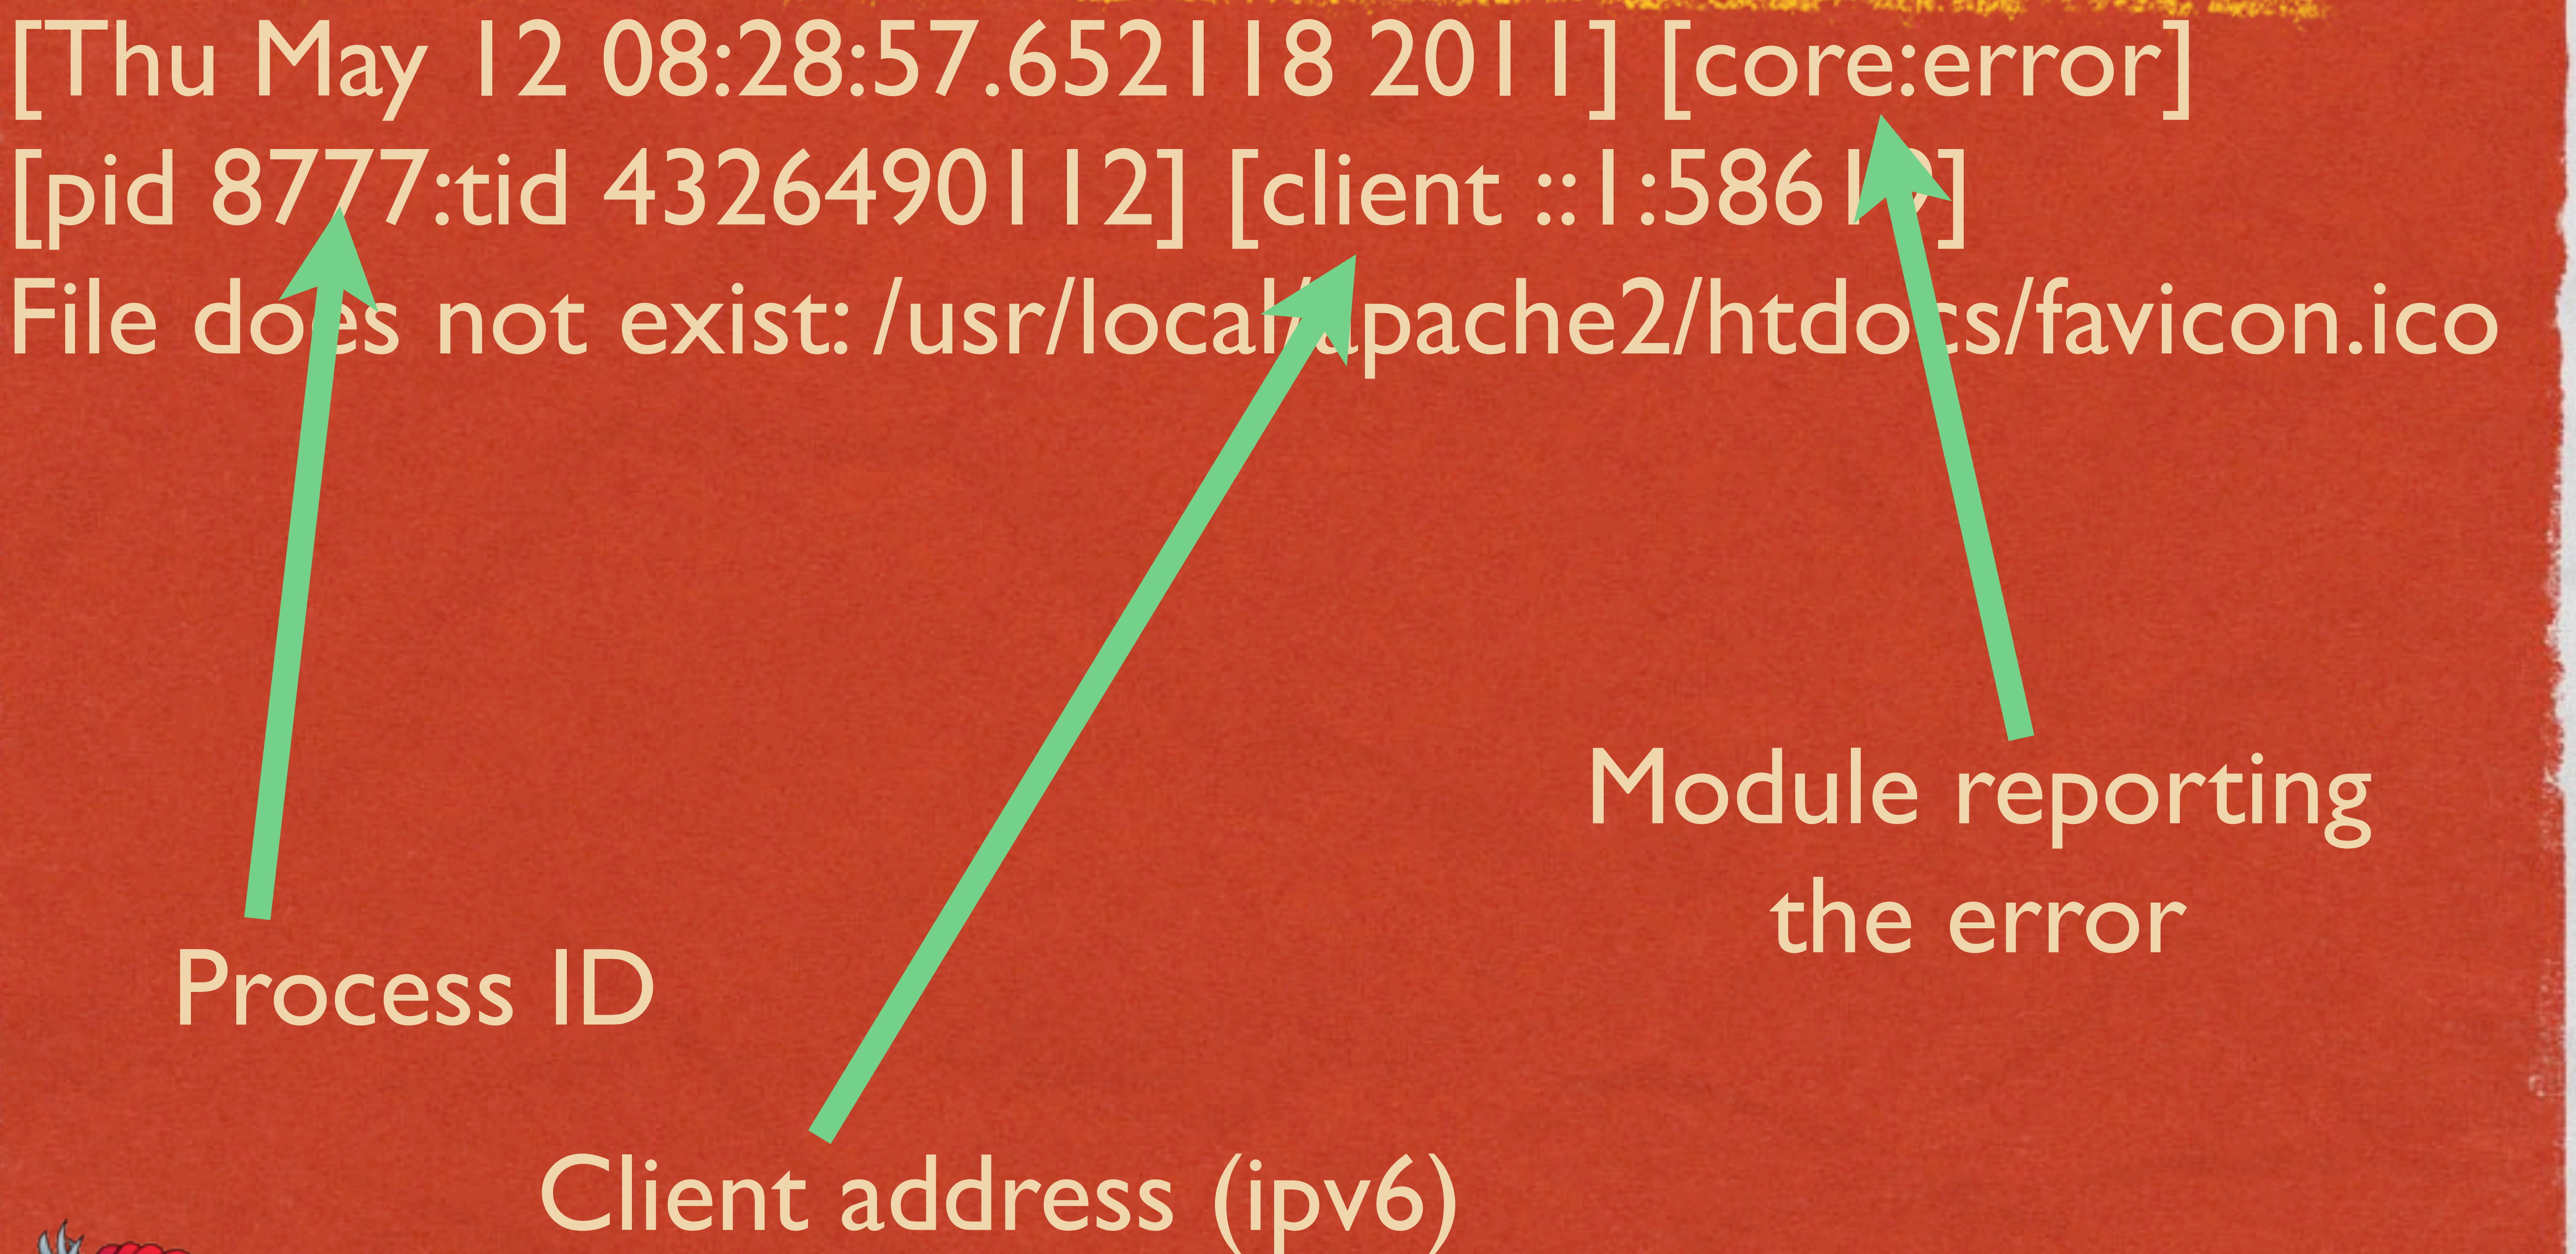

```
[Thu May 12 08:28:57.652118 2011] [core:error]
[pid 8777:tid 4326490112] [client ::1:58612]
File does not exist: /usr/local/apache2/htdocs/favicon.ico
```

TELLS YOU WHAT MODULE IS RESPONSIBLE

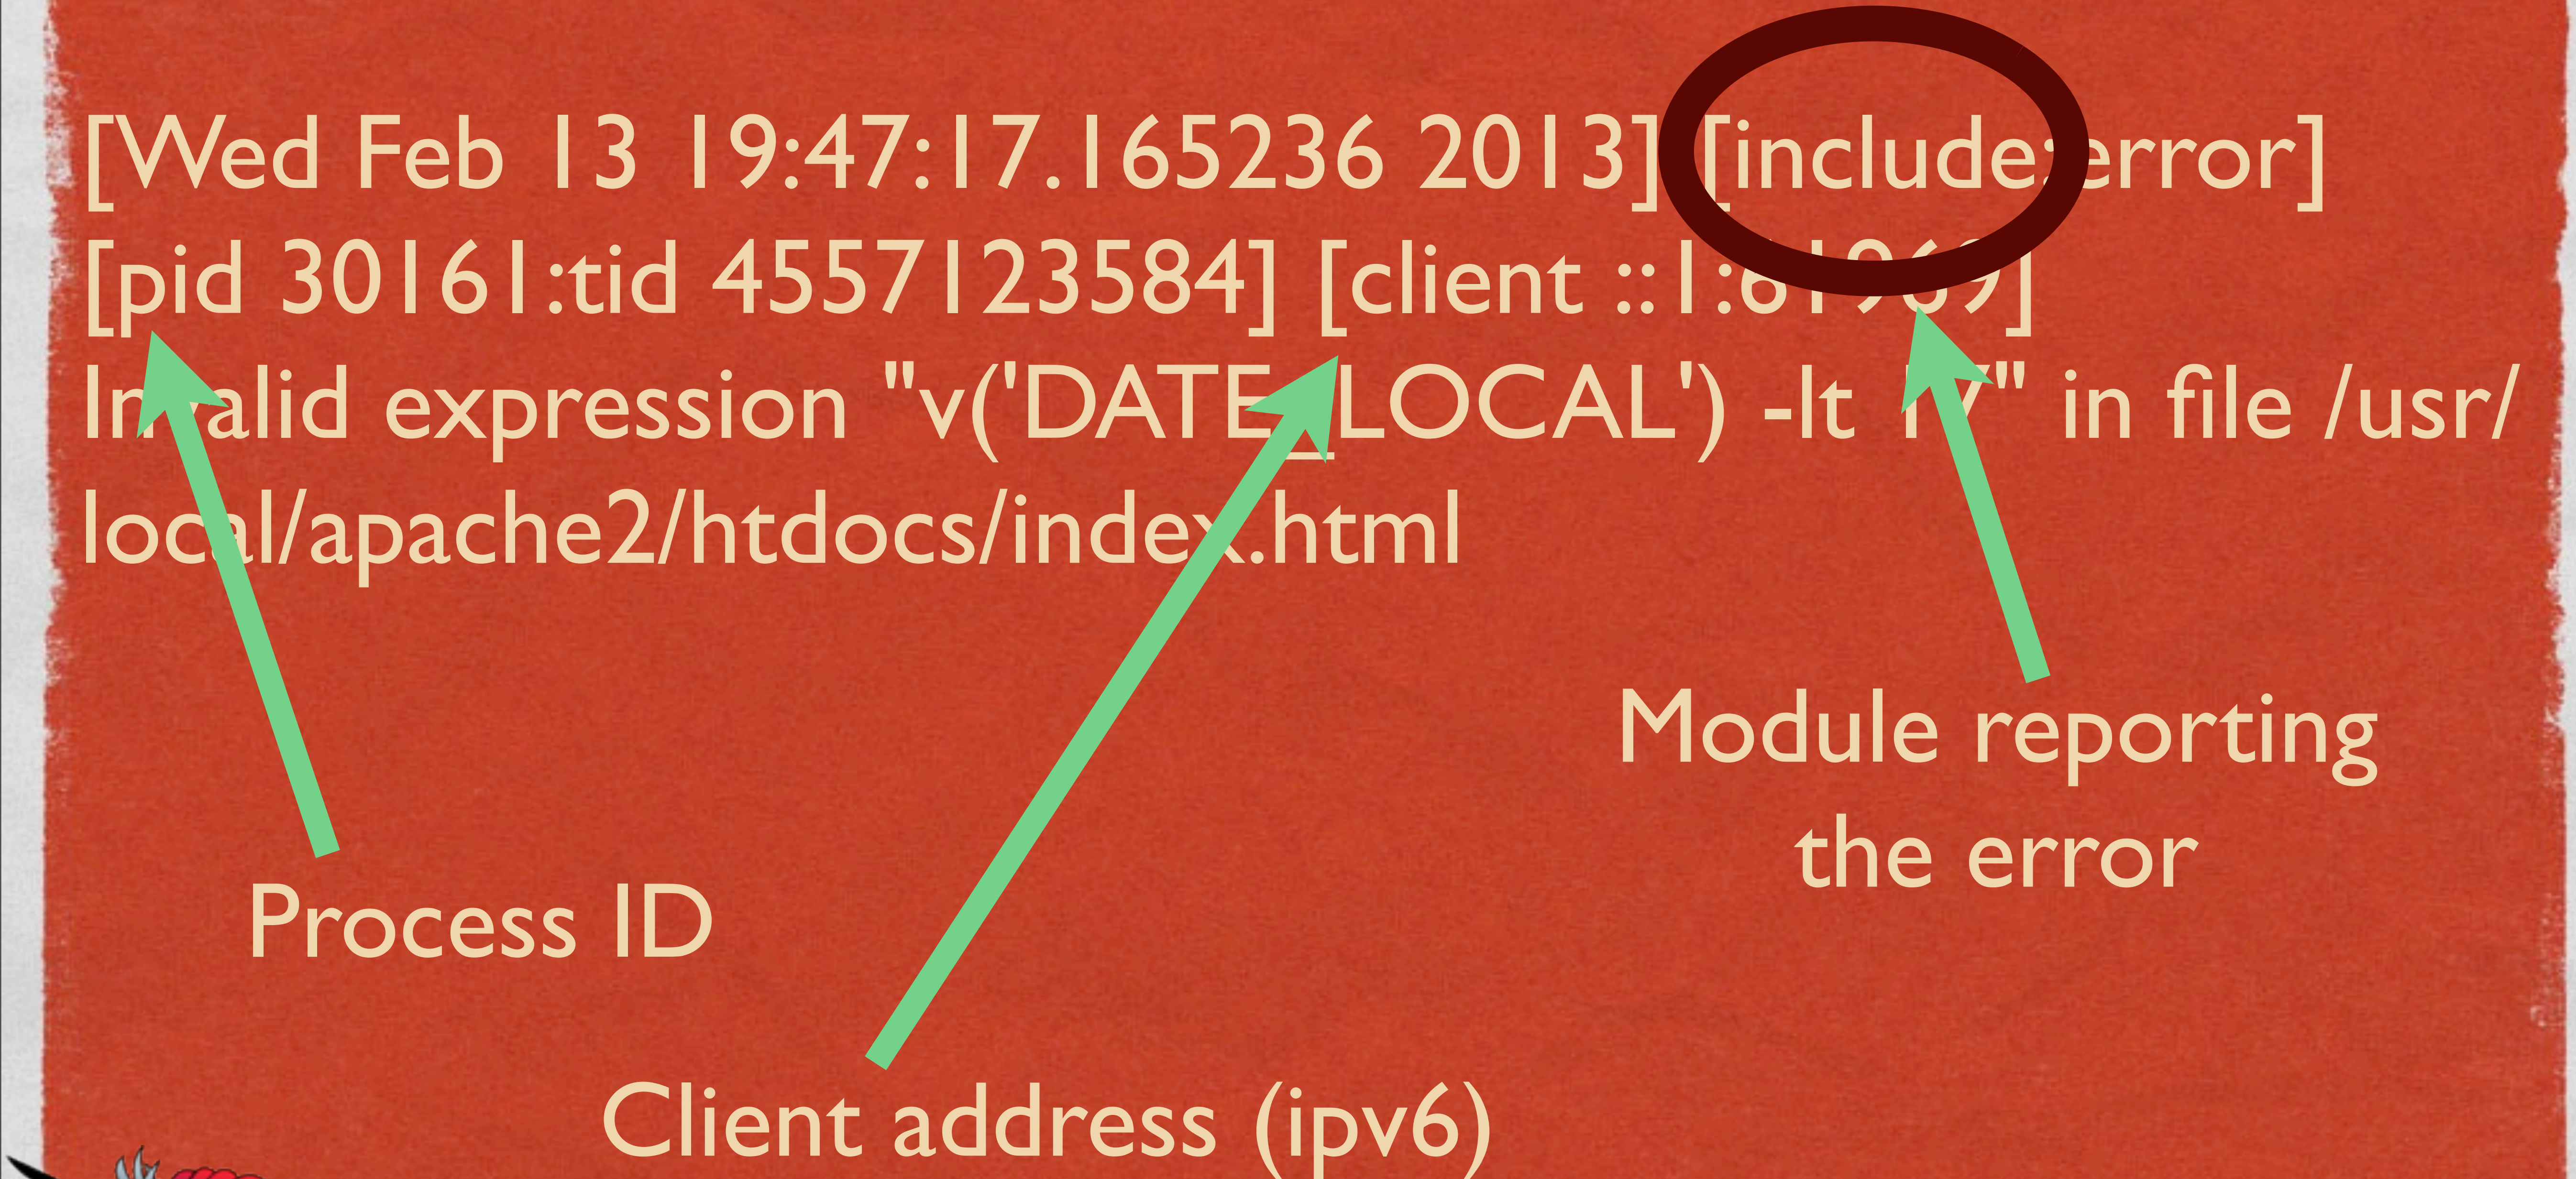

```
[Wed Feb 13 19:47:17.165236 2013] [include:error]
[pid 30161:tid 4557123584] [client ::1:61969]
Invalid expression "v('DATE_LOCAL') -lt 17" in file /usr/
local/apache2/htdocs/index.html
```

Module reporting the error

Process ID

Client address (ipv6)

PER-MODULE/DIRECTORY LOGLEVEL CONFIGURATION

LOGLEVEL

- Configures the verbosity of error logging

```
LogLevel crit
```

BEFORE ...

- 2.2 and earlier, LogLevel can be set globally or per virtual host
- LogLevel setting affects all modules, so one module's log messages might overwhelm another's.

NOW ...

- LogLevel may be placed in a <Directory> block to restrict the scope of more verbose logging

```
LogLevel crit

<Directory /var/www/cgi-bin>
   LogLevel debug
</Directory>
```

NOW ...

- LogLevel may be set per module, to get just the logging you're actually interested in

```
LogLevel info ssl:warn
```

TRACE

- In addition to the loglevel values you already know about, trace1 ... trace8 have been added.
- These add additional detailed debugging information

REWRITELOG

- Side-effect - the RewriteLog directive goes away and is replaced with:

```
LogLevel warn rewrite:trace5
```

REWRITELOG

- Get those log entries from the error log with

```
tail -f error_log | grep '[rewrite:'
```

- Which brings us to the next topic …

<IF>, <ELSEIF>, AND <ELSE>

THE KILLER FEATURE

- This is *the* reason to upgrade today
- This is what you've been asking for as long as I've been around httpd

```
# Compare the host name to example.com
#   and redirect to www.example.com if it matches
<If "%{HTTP_HOST} == 'example.com'">
    Redirect permanent / http://www.example.com/
</If>
```

Isn't that better than RewriteRules?

Or, the other way around

```
# Compare the host name to example.com
#   and redirect to www.example.com if it matches
<If "%{HTTP_HOST} != 'example.com'">
    Redirect permanent / http://example.com/
</If>
```

IF .. ELSEIF .. ELSE

```
<If "-R '10.1.0.0/16'">
  #...
</If>
<ElseIf "-R '10.0.0.0/8'">
  #...
</ElseIf>
<Else>
  #...
</Else>
```

REWRITERULE KILLER

- Everything that you've been using mod_rewrite for ...
- Guess I need to write a new edition of my book

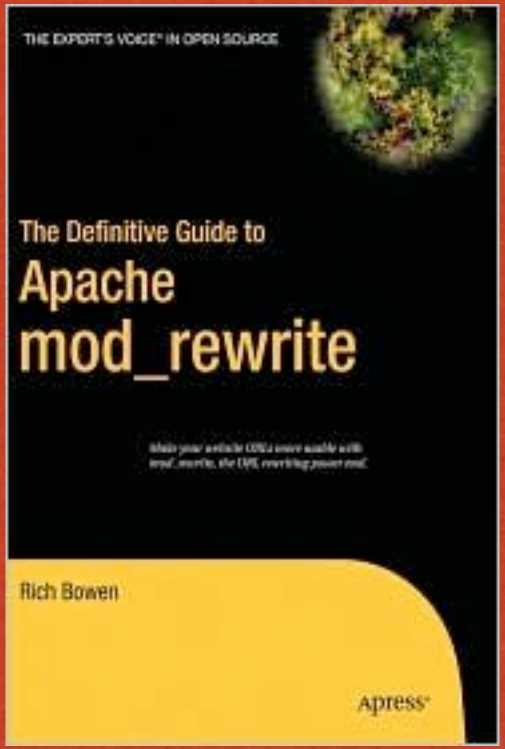

EXPRESSIONS

- <If> can use any expression that httpd can evaluate
- Which brings us to ...

EXPRESSION PARSER

THE OTHER KILLER FEATURE

- This is the other killer feature
- Note: Overuse will result in performance degradation

BEFORE

- Several different syntaxes for expressions
- Regex
- glob (shell file match syntax)

NOW

- A single expression engine that any module can call
- Extremely flexible, and can access a wide variety of variables and conditions

DOCS

- http://httpd.apache.org/docs/2.4/expr.html
- If you're into documentation, we could use some help writing a "recipes" style doc to augment the formal doc

VARIABLES

- request headers
- response headers
- environment variables
- user-settable variables

EXAMPLES

```
# Only allow access to this content
#      during business hours
<Directory "/foo/bar/business">
    Require expr %{TIME_HOUR} -gt 9 && %
                             {TIME_HOUR} -lt 17
</Directory>
```

See new Require syntax, later

NAMEVIRTUALHOST

BEFORE

```
NameVirtualHost *:80

<VirtualHost *:80>
  ServerName example.com
  # ...
</VirtualHost>

<VirtualHost *:80>
  ServerName example.co.de
  # ...
</VirtualHost>
```

BEFORE

- Omitting the NameVirtualHost line would result in an error message that wasn't very helpfule

NOW

- So we figured, why not do what we knew you meant?

NOW

```
<VirtualHost *:80>
  ServerName example.com
  # ...
</VirtualHost>

<VirtualHost *:80>
  ServerName example.co.de
  # ...
</VirtualHost>
```

NOW

- NameVirtualHost isn't required, because we'll do what you meant
- Just create the <VirtualHost> blocks, and we'll do the rest
- If there's overlap between vhosts, we'll tell you that, too

OVERRIDE CONFIGURATION

ALLOWOVERRIDE

- AllowOverride is too broad
- It allows categories of directives to be used in .htaccess files
- What if you want more control than that?

NOW

- AllowOverride None is now the default, as it should always have been
- Yay

ALSO ...

- AllowOverrideList lets you tweak what's permitted

```
AllowOverride AuthConfig
AllowOverrideList Redirect RedirectMatch
```

CONFIG FILE VARIABLES

OH, LOOK, ANOTHER KILLER FEATURE

- Did I say *the* reason to upgrade?
- Here's another
- You've been asking for this since 1997, too

DEFINE

- The Define directive lets you define a variable
- You can then use that variable in other directives, or in <IfDefine> blocks

EXAMPLE

```
# Invoke with -DTEST
Define TEST 1

<IfDefine TEST>
  Define servername test.example.com
</IfDefine>
<IfDefine !TEST>
  Define servername www.example.com
  Define SSL
</IfDefine>

DocumentRoot /var/www/${servername}/htdocs
```

EXAMPLE

```
# Don't repeat yourself ...
Define DOCROOT /var/www/htdocs

DocumentRoot ${DOCROOT}
<Directory ${DOCROOT}>
    Require all granted
</Directory>
```

MOD_HEARTMONITOR

JIM'S TALK

- You went to Jim's talk, right?
- Monitor a back-end server to see if it's alive, and swap it out of the proxy rotation accordingly

MOD_LUA

DANIEL'S TALK

- Embeds Lua in your httpd
- You should go to Daniel's talk, tomorrow
- He'll be showing some very cool stuff

MOD_RATELIMIT

DOWNLOAD LIKE IT'S 1999

- Limit the bandwidth, per Location or Directory

```
<Location /downloads>
    SetOutputFilter RATE_LIMIT
    SetEnv rate-limit 400
    # rate in Kb/s
</Location>
```

MOD_AUTHZ_CORE AND REQUIRE CHANGES

REQUIRE

- There's a whole talk on this this afternoon
- Require now makes sense
- None of that allow/deny/order/satisfy crap that everybody got wrong

```
<RequireAny>
    Require method GET POST OPTIONS
    Require valid-user
</RequireAny>
```

```
<Directory /www/mydocs>
    <RequireAll>
        <RequireAny>
            Require user superadmin
            <RequireAll>
                Require group admins
                Require ldap-group cn=Administrators,o=Airius
                <RequireAny>
                    Require group sales
                    Require ldap-attribute dept="sales"
                </RequireAny>
            </RequireAll>
        </RequireAny>
        <RequireNone>
            Require group temps
            Require ldap-group cn=Temporary Employees,o=Airius
        </RequireNone>
    </RequireAll>
</Directory>
```

4PM

- Be there

MOD_REWRITE ENHANCEMENTS

QSD

- New [QSD] flag (QueryStringDiscard) to remove query string as part of a rewrite
- New [END] flag to do what you always thought [L] did
- RewriteCond can use expressions as discussed above
- RewriteMap can make SQL queries
- You should have attended my tutorial on Monday

REWRITEMAP

```
RewriteMap myquery \
    "fastdbd:SELECT destination FROM rewrite WHERE source = %s"
```

ROTATELOGS ENHANCEMENTS

AUTOMATIC LOG ROTATION

- RewriteLogs rotates your log files without restarting the httpd process
- Rotates by time or file size
- Operates as a piped log handler

NOW

- -L : Can create a symlink to the current log
- Previously, you never knew for sure what your log file was called at any given moment

```
CustomLog
   "|bin/rotatelogs /var/logs/logfile 86400
           -L /var/log/access_log"
                 common
# That's all one line, by the way ...
```

```
tail -F /var/logs/access_log
```

POST-ROTATE

- -p : Invoke something when the log file is rotated
- Log parsing, for example

```
CustomLog
   "|bin/rotatelogs /var/logs/logfile 5M
           -p /usr/bin/process_logfile"
                 common
# That's all one line, by the way ...
```

MOD_MACRO

BONUS!

- mod_macro has just recently been added as an included module
- And there was much rejoicing

MOD_MACRO

- Insert macros into your httpd.conf that can be executed at startup
- Eliminate any repeated config sections
- Was a third-party module until a few months ago
- Not yet in 2.4. Stay tuned.

```
<Macro DirGroup $dir $group>
  <Directory $dir>
    require group $group
  </Directory>
</Macro>

Use DirGroup /www/apache/private private
Use DirGroup /www/apache/server  admin
# ... and so on N times
```

WHICH IS THE SAME AS ...

```
<Directory /www/apache/private>
  require group private
</Directory>

<Directory /www/apache/server>
  require group admin
</Directory>

# ... etc N times
```

FIN

FIN

Rich Bowen
rbowen@apache.org
@rbowen

Slides are at: **tm3.org/newin24**