

# 20 Things You Didn't Know You Could Do With Apache (aka httpd power tools)

Rich Bowen  
[rbowen@apache.org](mailto:rbowen@apache.org)  
ApacheCon US 2006





1

Check Spelling



# CheckSpelling



CheckSpelling On

CheckCaseOnly On



# Good for ...

- Windows users
- Simple typos
- Stupid people who just can't get stuff right on brochures



# Bad for ...

- Performance

- Every 404 results in a directory scan to see if something was missed



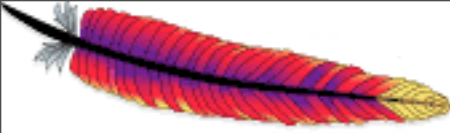
# mod\_speling

- Yes, mod\_speling
- See, it's misspelled. Funny, hmm?
- No, I didn't think so either



# 2 Caching





# Sort-of-dynamic

- Your blog that you update once every three months
- The front page of your corporate site with database-driven press releases





# Caching dynamic content

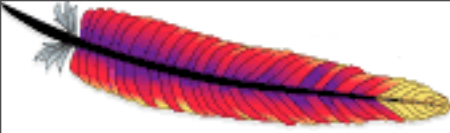
```
CacheEnable disk /  
CacheRoot /var/www/cache  
CacheDefaultExpire 3600  
# Only in 2.3+  
CacheMinExpire 3600
```

```
# Or, in 2.2-  
CacheIgnoreCacheControl On  
CacheDefaultExpire 600
```



# Caching

- CacheDir needs to be writeable by the Apache user
- Also great for proxying content from a back-end server
- Can also cache to memory, if you have a lot of RAM



# Even if I say not to ...

CacheIgnoreCacheControl On  
CacheStoreNoStore On  
CacheStorePrivate On

Resources requiring authentication  
are NEVER cached!



# Don't Cache That

CacheIgnoreHeaders Set-Cookie

CacheDisable /images



3  
mod\_deflate

# Deflate



```
AddOutputFilterByType DEFLATE \  
text/html text/plain text/xml
```



# Compressed

- All content gzipped on the way out to the client
- Browsers decompress it as it arrives
- 2-3 times performance improvement
- Can log compression percentage. See docs.



# Caveats

- Works best for sites that are text-heavy
- Don't try to deflate non-textual data. IE gets angry.
- You **can** cache the output, but probably don't want to





4

ext\_filter



# ExtFilterDefine

```
ExtFilterDefine fixtext mode=output \  
  intype=text/html \  
  cmd="/bin/sed s/verdana/arial/g"
```

```
<Location />  
  SetOutputFilter fixtext  
</Location>
```

# Silly

- Yes, it's a somewhat silly example
- And very slow
- But good for rapid prototyping and quick fixes





# Ok, better example

```
ExtFilterDefine c-to-html mode=output \  
  intype=text/c outtype=text/html \  
  cmd="/usr/bin/enscript --color -W html -Ec -o - -"  
  
<Directory "/www/htdocs/c">  
  SetOutputFilter c-to-html  
  AddType text/c .c  
  
  # Extra logging  
  ExtFilterOptions DebugLevel=1  
</Directory>
```



# Enscript

- Pretty-prints C source code
- Color syntax coding
- Declines to process files that aren't "text/c"



# 5

syslog

# Syslog

```
ErrorLog syslog:local1
```

```
local1.* @sysloghost:601
```





# Syslog

- Many servers logging to one location
- Log critical errors to somewhere you can't miss

```
local1.crit /var/log/panic
```





# Access log

- Sending the access log to syslog is harder
- Can be done with a piped log handler
- `mod_log_spread` might be a better way to handle this



# Syslog

CustomLog |/usr/local/apache/bin/apache\_syslog combined

```
#!/usr/bin/perl
use Sys::Syslog qw( :DEFAULT setlogsock );

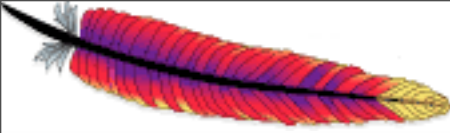
setlogsock('unix');
openlog('apache', 'cons', 'pid', 'local2');

while ($log = <STDIN>) {
    syslog('notice', $log);
}
closelog
```



6

IndexStyleSheet



# Index of /icons

Name	Last modified	Size	Description
Parent Directory		-	
HEADER.html	17-Nov-2005 16:18	85	
README	21-Nov-2004 09:35	4.8K	
README.html	17-Nov-2005 16:24	43	
a.gif	21-Nov-2004 09:35	246	
a.png	21-Nov-2004 09:35	293	
alert.black.gif	21-Nov-2004 09:35	242	
alert.black.png	21-Nov-2004 09:35	279	
alert.red.gif	21-Nov-2004 09:35	247	
alert.red.png	21-Nov-2004 09:35	298	
apache_pb.gif	21-Nov-2004 09:35	2.3K	
apache_pb.png	21-Nov-2004 09:35	1.4K	
apache_pb2.gif	21-Nov-2004 09:35	2.4K	
apache_pb2.png	21-Nov-2004 09:35	1.4K	
apache_pb2_ani.gif	21-Nov-2004 09:35	2.1K	
back.gif	21-Nov-2004 09:35	216	
back.png	21-Nov-2004 09:35	284	
ball.gray.gif	21-Nov-2004 09:35	233	
ball.gray.png	21-Nov-2004 09:35	277	
ball.red.gif	21-Nov-2004 09:35	205	
ball.red.png	21-Nov-2004 09:35	265	
binary.gif	21-Nov-2004 09:35	246	
binary.png	21-Nov-2004 09:35	296	
binhex.gif	21-Nov-2004 09:35	246	
binhex.png	21-Nov-2004 09:35	304	
blank.gif	21-Nov-2004 09:35	148	
blank.png	21-Nov-2004 09:35	195	
bomb.gif	21-Nov-2004 09:35	308	
bomb.png	21-Nov-2004 09:35	356	
box1.gif	21-Nov-2004 09:35	251	
box1.png	21-Nov-2004 09:35	308	
box2.gif	21-Nov-2004 09:35	268	
box2.png	21-Nov-2004 09:35	322	
broken.gif	21-Nov-2004 09:35	247	
broken.png	21-Nov-2004 09:35	305	
burst.gif	21-Nov-2004 09:35	235	
burst.png	21-Nov-2004 09:35	314	



# IndexStyleSheet

IndexStyleSheet "/css/style.css"



# CSS

```
#indexlist{font-family:"Lucida Sans Typewriter","Lucida Console","Luxi Mono","Andale Mono",Monaco,"Bitstream Vera Mono",monospace;margin:auto;}
```

```
.indexhead{background-color:gray;}
```

```
.indexbreakrow{visibility:hidden;}
```

```
.indexcolicon,.indexcollastmod,.indexcolname,.indexcolsize,.indexcoldesc  
{padding:0 1em 0 1em;}
```













```
.indexhead .indexcolicon, .indexhead .indexcoldesc {visibility: hidden;}
```

```
.indexeven{background-color:#666;}
```

```
.indexodd{background-color:#333;}
```

i am code. what are you?

## Index of /projects/apache

	Name	Last modified	Size	
	Parent Directory		-	
	cgi-win.awk	2006-06-19 01:29	13K	AWK Script
	httpd-2.2.2-14jun06.exe	2006-06-14 13:18	17M	Executable File
	httpd-2.2.2-25may06.exe	2006-05-25 19:59	17M	Executable File
	isapi-handler.2.0.diff	2006-06-26 08:22	3.5K	Diff Patch
	isapi-handler.2.2.diff	2006-06-26 08:24	3.7K	Diff Patch
	isapi-handler.trunk.diff	2006-06-26 08:29	3.2K	Diff Patch
	martin.patch	2006-06-26 09:11	1.6K	Diff Patch
	mod_isapi-response-attempt2.diff	2006-05-31 01:37	1.3K	Diff Patch
	mod_isapi-response-attempt3.diff	2006-06-01 19:37	3.7K	Diff Patch
	mod_isapi-response.diff	2006-05-30 22:18	515	Diff Patch
	mod_isapi.patched.zip	2006-06-01 22:15	70K	Diff Patch



# Documentation

- [http://httpd.apache.org/docs/2.2/mod/mod\\_autoindex.html#indexstylesheet](http://httpd.apache.org/docs/2.2/mod/mod_autoindex.html#indexstylesheet)
- 2.2 only
- Not yet documented. Will be soon.





# And ...

- With HeaderName and ReadmeName, wrap all that content in whatever shell you want.

```
HeaderName /style/header.html  
ReadmeName /style/footer.html  
IndexStyleSheet /style/style.css
```

```
IndexOptions FancyIndexing \  
  SuppressHTMLPreamble \  
  SuppressLastModified
```



7

mod\_authnz\_ldap



# Auth against LDAP

```
LoadModule authn_alias_module modules/mod_authn_alias.so
```

```
AuthLDAPBindDN cn=youruser,o=ctx  
AuthLDAPBindPassword yourpassword  
AuthLDAPURL ldap://ldap.host/o=ctx
```

```
<Directory /webpages/secure>
```

```
Order deny,allow
```

```
Allow from all
```

```
AuthBasicProvider ldap
```

```
AuthType Basic
```

```
AuthName LDAP_Protected_Place
```

```
AuthzLDAPAuthoritative off
```

```
require valid-user
```

```
</Directory>
```



# LDAP auth

- Configuration is a bit scary
- Presumably that meant something to folks who grok LDAP



# Active Directory auth

- Unfortunately, some of us work in MS-centric environments
- I have had this working, briefly
- I didn't understand it




WebDAV



# WebDav

- FTP is insecure
- SCP or SFTP is too complicated for your manager
- Maybe DAV is the right solution



```
DavLock /var/dav/davlock  
<Directory /www/upload>  
  Dav On  
</Directory>
```





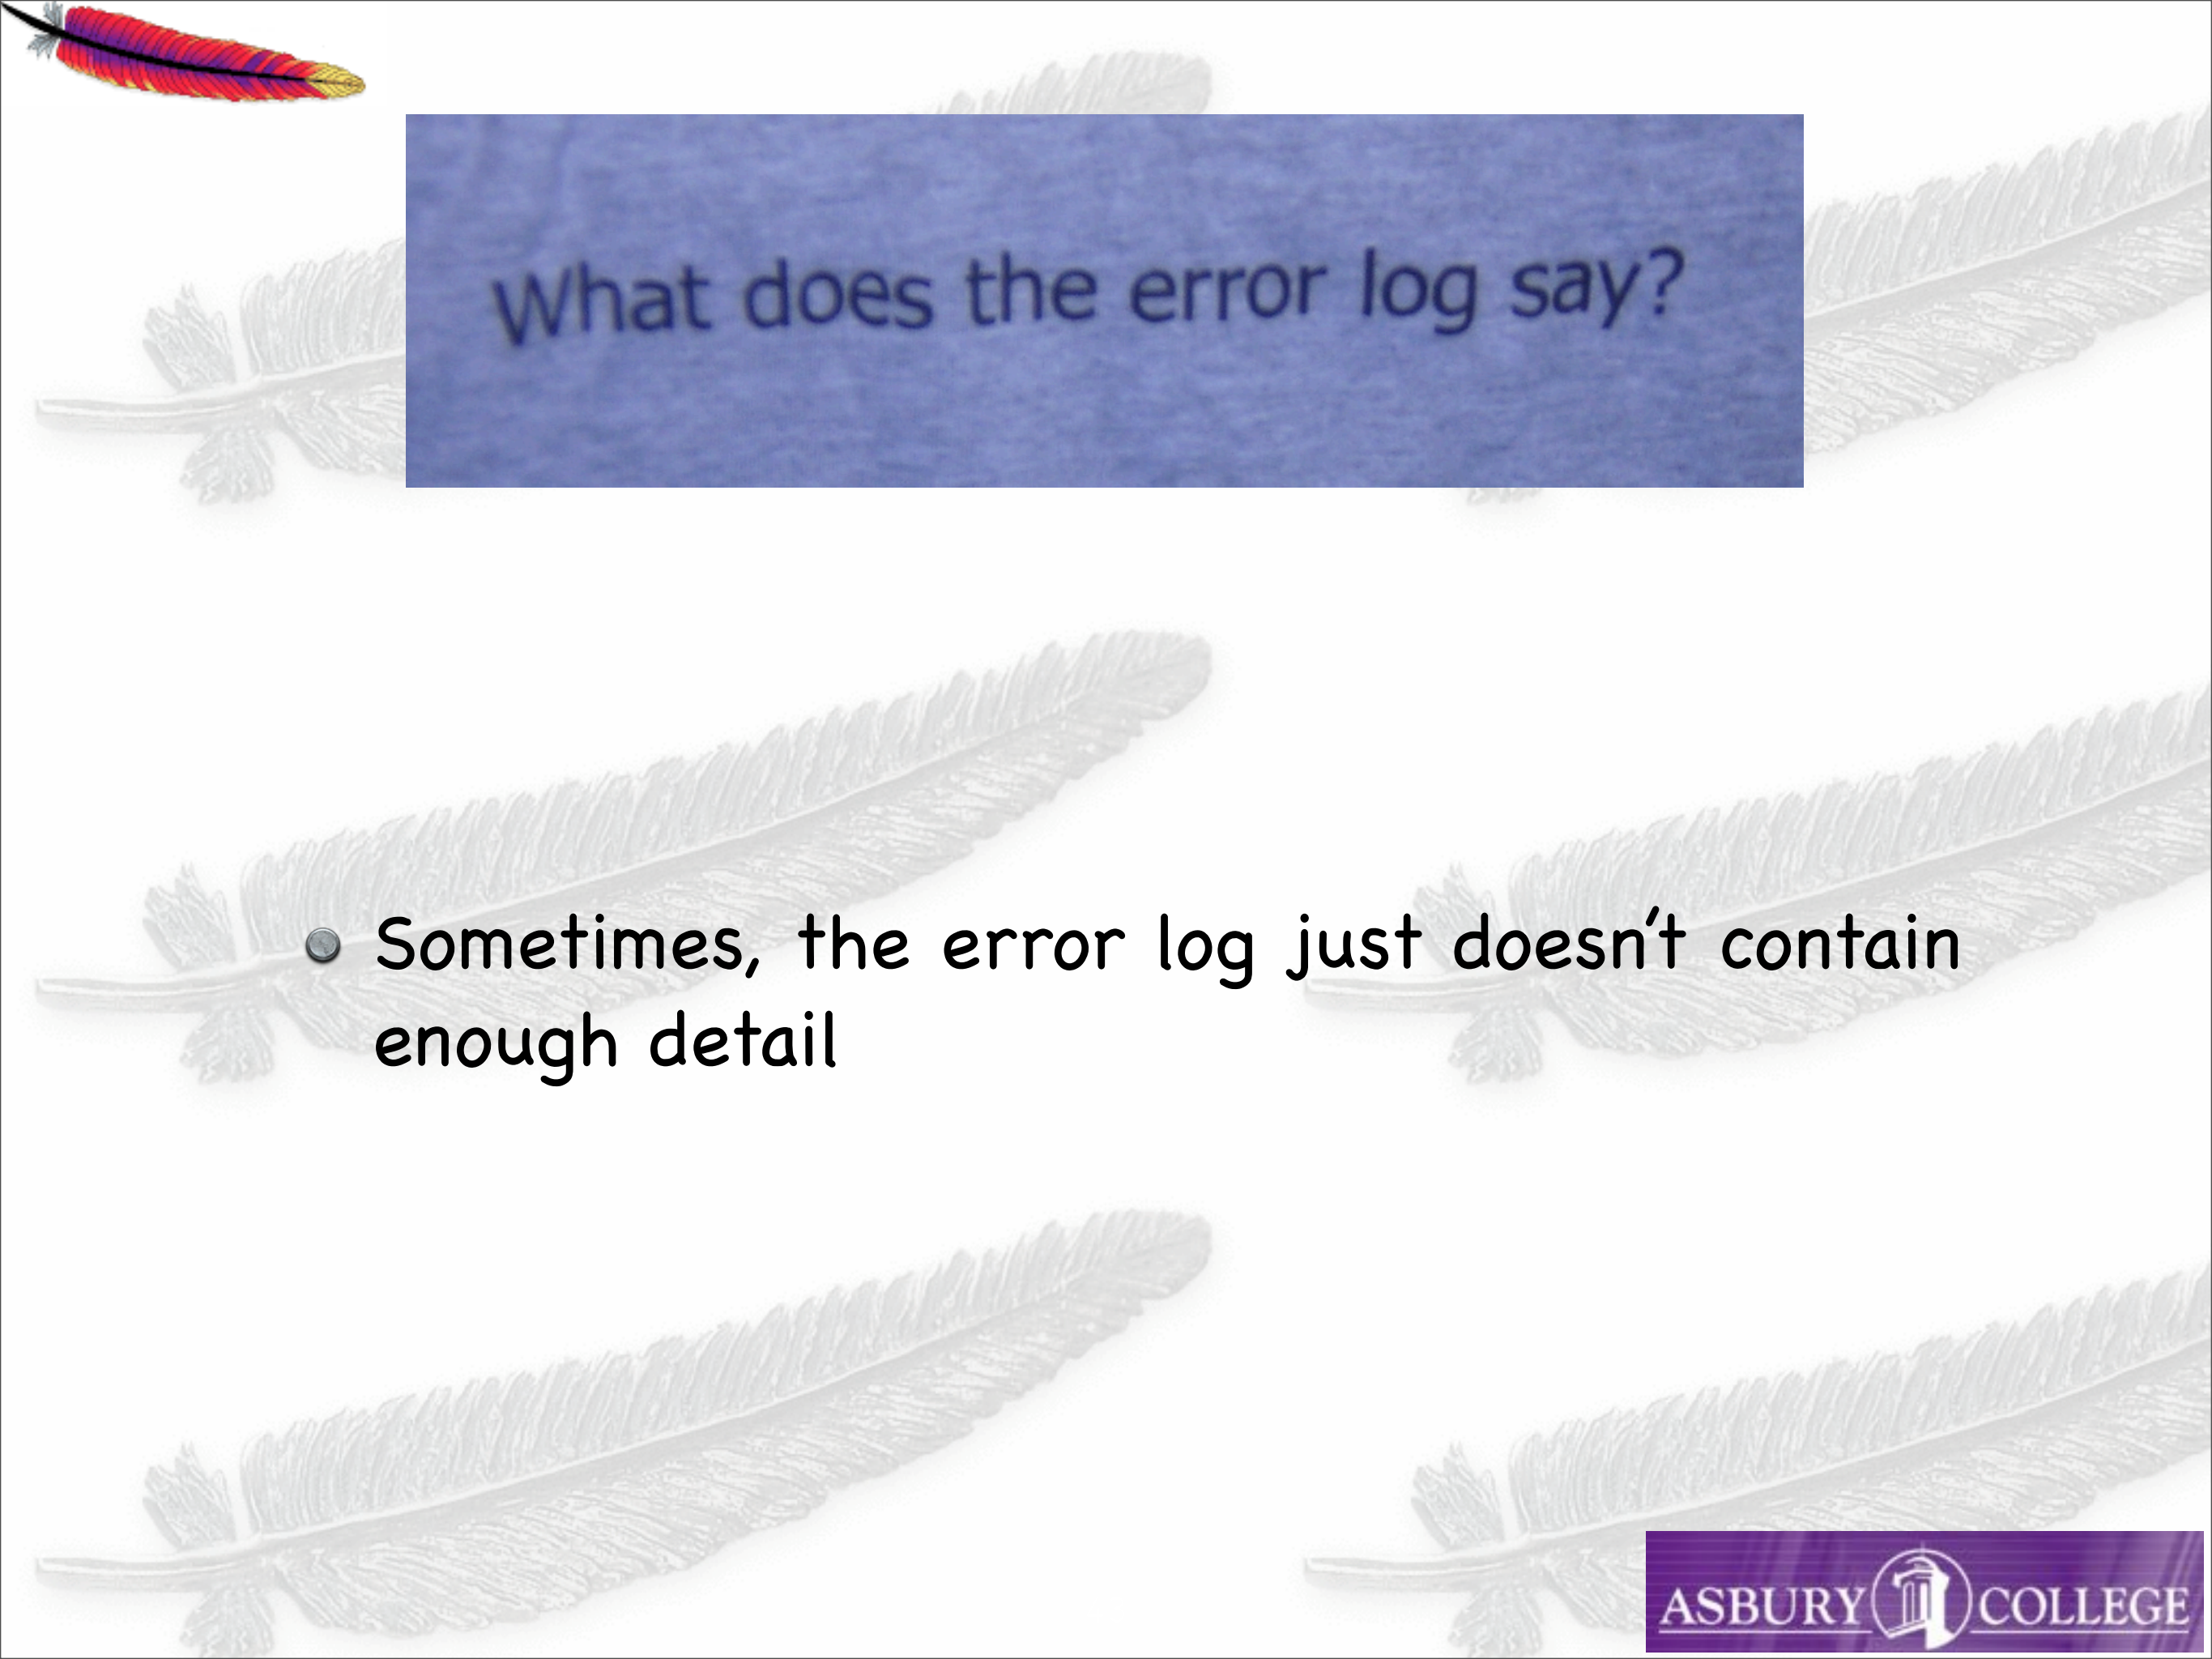
# Mount from any OS

- Well, any \*real\* operating system
- Windows doesn't play nice, but can usually be persuaded to



9

Debug log modules



What does the error log say?

- Sometimes, the error log just doesn't contain enough detail



# mod\_dumpio

DumpIOInput On  
DumpIOOutput On

# BIG log





# mod\_log\_forensic

- Logs the beginning and end of a request
- check-forensic script complains if a particular request wasn't created
- Logs the details of the request
- Relies on unique\_id in 2.0, and can use it in 2.2 if you want



# 10

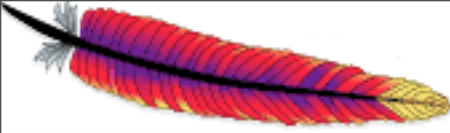
Proxy Balancer

# mod\_proxy\_balancer

- One front end server running mod\_proxy and mod\_proxy\_balancer
- Multiple backend servers running HTTP, FTP, or AJP protocol servers

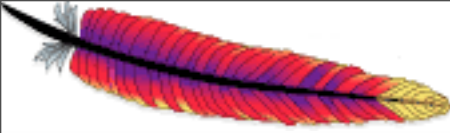






# Not that it's related ...





# He's dead, Jim

- Does heartbeat, so that dead servers are removed
- Weighted traffic distribution
- Does sticky sessions (so that an established session always goes back to the same server)



# mod\_proxy\_ajp



- Replaces mod\_jk and all the other icky Tomcat proxy thingies



# Balancer manager

- Relies on mod\_proxy
- Gives a webby management interface for the proxy balancer

```
<Location /balancer-manager>  
  SetHandler balancer-manager  
</Location>
```



# Example 1

```
ProxyPass /special-area http://special.example.com/ smax=5 max=10  
ProxyPass / balancer://mycluster/
```

```
<Proxy balancer://mycluster>
```

```
  BalancerMember http://1.2.3.4:8009
```

```
  BalancerMember http://1.2.3.5:8009 smax=10
```

```
  # Less powerful server, don't send as many requests there
```

```
  BalancerMember http://1.2.3.6:8009 smax=1 loadfactor=20
```

```
</Proxy>
```

# Example 2

```
ProxyPass / balancer://hotcluster/  
<Proxy balancer://hotcluster>  
  BalancerMember http://1.2.3.4:8009 loadfactor=1  
  BalancerMember http://1.2.3.5:8009 loadfactor=2  
  # The below is the hot standby  
  BalancerMember http://1.2.3.6:8009 status=+H  
  ProxySet lbmethod=bytraffic  
</Proxy>
```



# 11

mod\_logio



# Normal logging

```
129.101.54.111 - - [13/Sep/2006:19:45:35  
-0400] "GET /images/Kenya_world_map.jpg  
HTTP/1.0" 200 42739
```

- 42729 is the bytes downloaded, NOT including headers, and does NOT include the request.
- Total transaction size is more than this





# mod\_logio

```
LogFormat "%h %l %u %t \"%r\" %>s %b \"%  
{Referer}i\" \"%{User-Agent}i\" %I %O"
```

combinedio

- %I is total input
- %O is total output



# 12

PATH\_INFO



# PATH\_INFO

<http://www.foo.com/example.php/one/two/>

This bit





# Granny says ...

- The most important thing about ~~mod\_rewrite~~ magic is knowing when not to use it.

(With apologies to Terry Pratchett)



# “Pretty” URLs

- mod\_rewrite frequently used to rewrite query strings into path information:

RewriteEngine On

RewriteRule ^/book/([^/]+)/([^/]+)/? \

/cgi-bin/book.php?author=\$1&topic=\$2 [PT,L]



# Which gives you

<http://www.foo.com/book/bowen/apache>

- ... instead of ...

<http://www.foo.com/cgi-bin/book.php?author=bowen&topic=apache>




# Bad news

- This doesn't actually result in higher search engine ranking
- It does, however, make it easier to read the URL over the phone



# Using PATH\_INFO instead

- Remove the unnecessary additional processing
- Faster
- Much shinier
- You get to be smug with people who do it the old-and-crufty way







# Step One

- 'book' needs to be a php file
- Without the .php file extension

```
<Files book>  
  SetHandler application/x-httpd-php  
</Files>
```

# Step Two

```
<?php
```

```
list($zero, $one, $two) =  
    explode("/", $_SERVER['PATH_INFO']);
```

```
?>
```



# Switch

```
$args = split('/', $_SERVER['PATH_INFO']);  
switch( $args[1] ) {  
    case `foo` :  
        // stuff here  
        break;  
  
    case `bar` :  
        // stuff here  
        break;  
  
    default :  
        // Other stuff  
        break;  
}
```




13  
graceful-stop



# graceful-stop

- 2.2

- Stop the server, but allow connections to complete what they are doing



You are graceful

Good for:  
Airline Hostess,  
Ballet Dancer,  
Model and  
Actress

`apachectl graceful-stop`



14  
authn\_alias



# mod\_authn\_alias

- Allows you to create aliases of various auth sources, so that you can check multiple sources of the same type



```
LoadModule authn_alias_module modules/mod_authn_alias.so
```

```
<AuthnProviderAlias ldap ldap-alias1>  
  AuthLDAPBindDN cn=youruser,o=ctx  
  AuthLDAPBindPassword yourpassword  
  AuthLDAPURL ldap://ldap.host/o=ctx  
</AuthnProviderAlias>
```

```
<AuthnProviderAlias ldap ldap-other-alias>  
  AuthLDAPBindDN cn=yourotheruser,o=dev  
  AuthLDAPBindPassword yourotherpassword  
  AuthLDAPURL ldap://other.ldap.host/o=dev?cn  
</AuthnProviderAlias>
```

```
Alias /secure /webpages/secure
```

```
<Directory /webpages/secure>
```

```
  Order deny,allow
```

```
  Allow from all
```

```
  AuthBasicProvider ldap-other-alias ldap-alias1
```

```
  AuthType Basic
```

```
  AuthName LDAP_Protected_Place
```

```
  AuthzLDAPAuthoritative off
```

```
  require valid-user
```

```
</Directory>
```





Or

```
<AuthProviderAlias file file1>  
  AuthUserFile /tmp/user1  
</AuthProviderAlias>
```

```
<AuthProviderAlias file file2>  
  AuthUserFile /tmp/user2  
</AuthProviderAlias>
```

```
<Location /x>  
  AuthType Basic  
  AuthBasicProvider file1 file2  
  AuthName "Wooga"
```

```
  Require valid-user  
</Location>
```



# 15

Protocol modules



# mod\_pop3



- That's right! It's a pop3 server!
- [http://svn.apache.org/viewcvcs.cgi/httpd/mod\\_pop3/](http://svn.apache.org/viewcvcs.cgi/httpd/mod_pop3/)



# See also ...

- `mod_echo`
- `mod_smtpd`
- `mod_spamd`



# Protocol modules

- Apache httpd is a generic protocol server
- Protocol modules can be any protocol



# Why?

- Common configuration format
- Common authentication architecture
- Common rewriting/mapping capabilities



16  
mod\_dbd



# mod\_dbd

- Manages database connection pooling
- Not actually much use by itself
- Great for other modules that need to talk to databases
- Such as ...





# mod\_authn\_dbd

- Auth against any SQL database
- Well, “any” might be a little too optimistic



DBDriver pgsql

#Connection string: database name and login credentials

DBDParams "dbname=htpasswd user=apache password=xxxxxx"

#Parameters for Connection Pool Management

DBDMin 1

DBDKeep 2

DBDMax 10

DBDExptime 60

<Directory /usr/www/myhost/private>

AuthType Basic

AuthName "My Server"

AuthBasicProvider dbd

Require valid-user

AuthDBDUserPWQuery "select password from authn where username = %s"

</Directory>



17

server-info?config



# server-info?config

```
26: ForceLanguagePriority Prefer Fallback  
: </Directory>
```

**In file: /usr/local/apache/conf/extra/httpd-dav.conf**

```
15: DavLockDB "/usr/local/apache/var/DavLock"  
17: Alias /uploads "/usr/local/apache/uploads"  
19: <Directory "/usr/local/apache/uploads">  
20:   Dav On  
34:   order allow,deny  
35:   allow from all  
37:   Options Indexes  
   : </Directory>  
46: BrowserMatch "Microsoft Data Access Internet Publishing Provider" redirect-caref  
47: BrowserMatch "MS FrontPage" redirect-carefully  
48: BrowserMatch "^WebDrive" redirect-carefully  
49: BrowserMatch "^WebDAVFS/1.[0123]" redirect-carefully  
50: BrowserMatch "^gnome-vfs/1.0" redirect-carefully  
51: BrowserMatch "^XML Spy" redirect-carefully  
52: BrowserMatch "^Dreamweaver-WebDAV-SCM1" redirect-carefully
```

**In file: /usr/local/apache/conf/httpd.conf**

```
472: SSLRandomSeed startup builtin  
473: SSLRandomSeed connect builtin  
476: <AuthnProviderAlias file file1>  
477:   AuthUserFile /tmp/user1
```



# What's in there?

- All your configuration settings
- What file they live in
- What order they are applied



# 18

mod\_imagemap



# Server-side image maps

- Yes, I miss 1995
- Things made sense back then
- And server-side image maps were cool
- Netscape 1.1 ruined **everything**





# Enough nostalgia

AddHandler imap-file map

- And then ...

```
#Comments are printed in a 'formatted' or 'semiformatted' menu.  
#And can contain html tags. <hr>  
base referer  
poly map "Could I have a menu, please?" 0,0 0,10 10,10 10,0  
rect .. 0,0 77,27 "the directory of the referer"  
circle http://www.inetnebr.com/lincoln/feedback/ 195,0 305,27  
rect another_file "in same directory as referer" 306,0 419,27  
point http://www.zyzyva.com/ 100,100  
point http://www.tripod.com/ 200,200  
rect mailto:nate@tripod.com 100,150 200,0 "Bugs?"
```





# In your HTML ...

```
<a href="/maps/imagemap1.map">  
    
</a>
```

- A blink tag, and an “under construction” image or two might help the ambience



# 19

mod\_negotiation



# File extensions are *\*so\**

## Windows 3.1

- <http://www.example.com/index> is much cooler than <http://www.example.com/index.php>
- And lets you change to .html or .cfm later on without changing your URLs



# It's a feature

## Options +MultiViews

- A request for the resource called index is served by the variant called index.php
- index.php.en and index.php.fr are also variants of this resource
- Accept-Language (in your browser prefs)



20  
svn auto-commit



Ever wish life had an  
undo button?





# Revision control everything!

```
<Location /dav>
```

```
  DAV svn
```

```
  SvnPath /path/to/svn/repos
```

```
  SVNAutoversioning On
```

```
</Location>
```



# R.C. for everyone

- Your users just mount this like a regular DAV share
- All file modifications generate an SVN commit
- Can always roll back to a previous version





# Caveats


- MANY times more drive space required
- No commit messages
- Diffs not possible on binary files (usually)
- Can revert by date/time if you know it



21  
mod\_pony



# mod\_pony





# Acme::Pony

- It's on CPAN
- Add it to your existing code and delight your code reviewers



# Before

```
use Acme::Pony;  
print "Hello, World\n";
```

# After

```
use Acme::Pony;
    buf
    fYBUffY
    buFFYbUff
    YbUFfyBUff
    YBuffYbUFFyb
    uffYBuffYbuf
    FybuffYbuF
    fYbUffYBuffY          BuFF
y    buFFyBUfFYBUfFYbUFFYbUffyb
uffYbuFFYbUfFYBUFFyBUffYbuFFYbu
fFY    bUFfybUffYBUffYBUfFYbUFFyBUf
fyBU    FfYbuFFYbUFfyBUffYbuFFyBUff
        ybUfF ybuffYbUffybUffYbu
        ffYB          UFFyBUFFYBUf
        FyB          uFfy bUffY
        BUf          fYbufFybuF
        fY          BUF FyB uFFY
        BU          fFy  XX
        XXX          XX
```

And, yes, it still runs



# So, naturally ...

- We thought, why not make this an Apache module, so that the whole world can enjoy it?



# The bad news ...

- Causes a segfault on every page load
- But ... then again ... maybe that's a feature?
- Matt Lewandowski is the author of this fine module





```
buf
fYBUffY
buFFYbUff
YbUFfyBUff
YBUffYbUFFyb
uffyBUffYbuf
FybuFFYbuF
fYbUffYBUffY          BuFF
y  buFFyBUffYBUffYbUffYbUffyb
uffybuFFYbUffYBUffYBUffYbuFFYbu
fFY  bUffybUffYBUffYBUffYbUffYBUf
fyBU  FfYbuFFYbUffYBUffYbuFFYBUff
      ybUff  ybuFFYbUffYbuFfYbu
      ffYB          UFFyBUffYBUf
      FyB          uFfy  buffY
      BUf          fYbuFFybuF
      fY          BUF  FyB  uFFY
      BU          fFy  XX
      XXX          XX
      XXX          XX
```



# #apache

- For anything else you don't know, #apache is the source of all wisdom
- [irc.freenode.net](https://irc.freenode.net)
- Support Freenode
- [rbowen@apache.org](mailto:rbowen@apache.org)