

Apache Cookbook



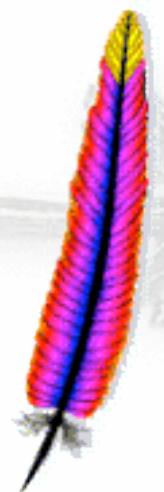
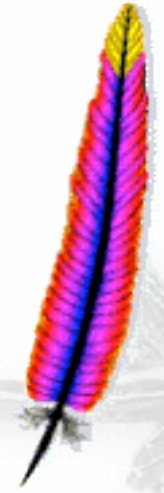
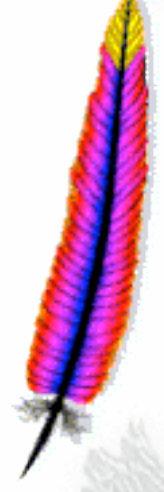
ApacheCon EU
2008
Amsterdam

Rich Bowen - Asbury College
rbowen@apache.org





Table of Contents

- 
- 
- 
- SSL vhosts
 - Rewrite based on query string
 - Preventing "image theft"
 - Logging more information
 - Logging to syslog
 - WebDAV
 - Preventing malicious requests with mod_security
 - Enabling PHP
 - Mass virtual hosting
 - Customized error messages
 - URL handler ("rewrite everything")
 - Fancy directory listings
 - Caching dynamic content
 - /server-info goodness
 - /server-status goodness
 - UserDir without the ~



Recipes

- At http://people.apache.org/~rbowen/presentations/apache_cookbook_recipes.tar.gz



Caveat: Versions

- 2.2 is the current version of Apache
- If you are running 1.3, you really should upgrade
- Some, not all, of these recipes will work in 2.0



SSL vhosts



- Multiple SSL hosts, one IP address



Problem



- One SSL cert per IP address
- Certificate is negotiated before the HOST: header is sent



Solution



- Three options:
 - Wildcard certificate
 - Get more IP addresses
 - Ignore the error messages



Wildcard certificate



- Costs twice regular certs
- Works for *.domain.tld
- Cannot span multiple domains
- Set up name-based vhosts the normal way

Wildcard certificate

```
NameVirtualHost *:443

# Wildcard certificate for *.domain.com
SSLCertificateFile /var/www/conf/server.crt
SSLCertificateKeyFile /var/www/conf/server.key

<VirtualHost *:443>
  ServerName one.domain.com
  DocumentRoot /var/www/one/htdocs
  SSLEngine On
</VirtualHost>

<VirtualHost *:443>
  ServerName two.domain.com
  DocumentRoot /var/www/two/htdocs
  SSLEngine On
</VirtualHost>
```





Multiple IP addresses



- This is the best solution
- Not always an option

Multiple IP addresses

```
<VirtualHost 172.20.4.10:443>
  ServerName one.domain.com
  DocumentRoot /var/www/one/htdocs

  SSLCertificateFile /var/www/conf/one.crt
  SSLCertificateKeyFile /var/www/conf/one.key

  SSLEngine On
</VirtualHost>

<VirtualHost 172.20.4.11:443>
  ServerName two.domain.com
  DocumentRoot /var/www/two/htdocs

  SSLCertificateFile /var/www/conf/two.crt
  SSLCertificateKeyFile /var/www/conf/two.key

  SSLEngine On
</VirtualHost>
```



Ignore errors



- SSL cert will be valid for only one hostname
- Other named vhosts will be encrypted
- Browser will report that the cert doesn't match the hostname
- SSL is encryption + validation. You're losing the validation.

Ignore the errors



```
NameVirtualHost *:443

# Certificate for one.domain.com
SSLCertificateFile /var/www/conf/one.crt
SSLCertificateKeyFile /var/www/conf/one.key

<VirtualHost *:443>
  ServerName one.domain.com
  DocumentRoot /var/www/one/htdocs
  SSLEngine On
</VirtualHost>

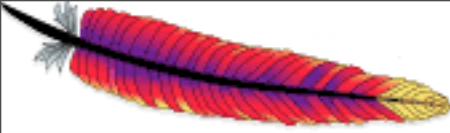
# Will be secure, but will generate errors
<VirtualHost *:443>
  ServerName two.domain.com
  DocumentRoot /var/www/two/htdocs
  SSLEngine On
</VirtualHost>
```



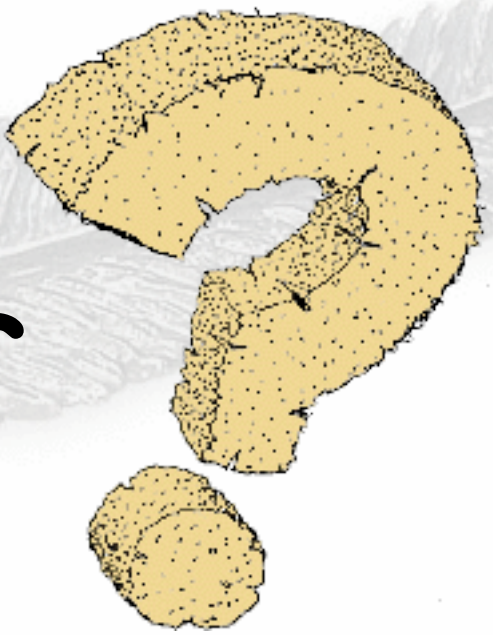
Other options



- Efforts are underway to escape this limitation
- Browser support is the big hurdle



Rewrite based on QUERY_STRING or PATH_INFO



- Sometimes what gets asked is:

“I want to forbid access if the
QUERY_STRING doesn't contain
foo=bar”



Rewrite by QUERY_STRING



- The sensible solution would be to handle this in your script/handler/program
- But, if that's not an option, `mod_rewrite` might be a good choice



Problem



- RewriteRule **doesn't** have access to the `QUERY_STRING`
- Only the URI - the bit after `http://` `hostname.com` and before the `?` - is **accessible** to RewriteRule



Solution

- RewriteCond has access to the entire requested URL, and any other server variables



RewriteCond %{VARIABLE} regex



RewriteCond

- Does the QUERY_STRING contain foo=bar

```
RewriteCond %{QUERY_STRING} foo=bar  
RewriteRule ^ - [F]
```





^ rather than .*

- ^ means "starts with"
- All strings start, even empty strings.
- Thus, all strings match ^
- ^ is more efficient than .*





Backreferences

- Or, you can do a rewrite based on the value of the QUERY_STRING

```
RewriteCond %{QUERY_STRING} user=(.+)\b  
RewriteRule (.*) /home/%1/www$1
```





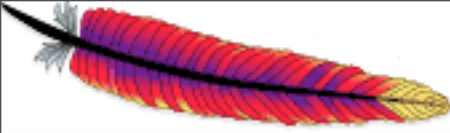
More frequently ...

- People want to map <http://example.com/one/two/three> to <http://example.com/something.php?a=one&b=two&c=three>



See also

- Upcoming recipe “URL Handler”
- Not quite the same, but many similar techniques



PATH_INFO

- Everything after the final / is the path info
- “Final /” refers to the / following an actual file or resource

<http://example.com/index.php/one/two/three>



PATH_INFO

- The trick is to figure out which bit is a valid resource, and which bit is PATH_INFO
- Two approaches



URL Prefix

- <http://example.com/prefix/one/two/three>
- You know that only URLs starting with **prefix** need special attention

```
RewriteRule ^/prefix(.*) \  
/handler.php?args=$1
```



File existence

- Check to see if the requested file exists
- If not, rewrite
- May interfere with other rewrite matches

```
RewriteCond %{REQUEST_FILENAME} !-f  
RewriteCond %{REQUEST_FILENAME} !-d  
RewriteRule (.*) /handler.php?args=$1
```



Caveats

- May need to prepend a directory path

```
RewriteCond \  
/var/www%{REQUEST_FILENAME} !-f
```

- Still need to do something useful with the value of \$1, if you want it to be split into args.

The full recipe

```
RewriteRule ^/prefix/([^/]+)/([^/]+) \  
/handler.php?one=$1&two=$2 [PT,L]
```



Caveats

- Exactly two arguments
- No more, no less
- Perhaps you want this to be more flexible?



More flexible

```
RewriteRule ^/prefix/([^/]+)?/?([^/]+)? \
/handler.php?one=$1&two=$2 [PT,L]
```

- Matches are now optional
- Arguments will be passed null - just ignore them in handler.php, or check for null values and take appropriate measures



More arguments

- This technique can be repeated for up to 9 arguments.
- \$1 - \$9
- \$10 is not available



Preventing image theft

- “Image theft” is the term used for other sites embedding your images in their pages.
- Ideally, you want to forbid having your images in any pages but your own
- There are several ways to accomplish this



SetEnvIf



- SetEnvIf is provided by mod_setenvif
- Sets environment variables if certain conditions are met

SetEnvIf

```
SetEnvIf Referer "^http://myhost\.com" localref=1  
<FilesMatch "\.(gif|jpg|png)">  
  Order Deny,Allow  
  Deny from all  
  Allow from env=localref  
</FilesMatch>
```



10_image_theft

Problem

- Some browsers don't set the Referer value

```
SetEnvIf Referer "^http://myhost\.com" localref=1  
SetEnvIf Referer "^$" localref=1
```

```
<FilesMatch "\.(gif|jpg|png)">  
  Order Deny,Allow  
  Deny from all  
  Allow from env=localref  
</FilesMatch>
```





mod_rewrite

- Or, you could do it with a RewriteRule

RewriteEngine on

RewriteCond %{HTTP_REFERER} !=""

RewriteCond %{HTTP_REFERER} !example\.com [NC]

RewriteRule \.(jpe?g|gif|png)\$ - [F,NC]



11_image_theft



But, more usefully

- If you're just going to fail the request, use SetEnvIf. It's more efficient
- But if you wanted to do something more interesting ...



Redirect the request



```
RewriteEngine on
RewriteCond %{HTTP_REFERER} !=""
RewriteCond %{HTTP_REFERER} !example\.com [NC]
RewriteCond %{REQUEST_URI} !go_away.png
RewriteRule \.(jpe?glgif|png)$ /images/go_away.png [NC,L]
```

Or ...



```
RewriteEngine on
RewriteCond %{HTTP_REFERER} !=""
RewriteCond %{HTTP_REFERER} !example\.com [NC]
RewriteRule \.(jpe?g|gif|png)$ \
  http://othersite.com/images/unsavory.jpg [NC,R]
```




Logging more information

- The standard log file is sometimes not sufficient.
- This recipe shows you how to get a little more information





mod_log_config



- Variables available for other values
- Always use 'combined' rather than 'common'



combined



```
LogFormat "%h %l %u %t \"%r\" %>s %b \"%{Referer}i\" \"%{User-Agent}i\"" \
    combined
CustomLog logs/access_log combined
```



Additional variables



- http://httpd.apache.org/docs/2.2/mod/mod_log_config.html#formats
- Most of the actual useful variables are already in 'combined'
- Most log analysis packages understand the 'combined' format



Important variables



- `%{something}C` - the value of the 'something' cookie
- `%{something}i` - the 'something' request (input) header
- `%{something}o` - the 'something' response (output) header
- `%q` - The query string
- and ...



mod_logio

- %b gives the size of the response in bytes
- Does not include headers
- Does not include the request
- mod_logio gives both of these





mod_logio



- %I - total size of request (Input) in bytes
- %O - total size of response (Output) in bytes
- Includes headers in each case.



mod_dumpio



- http://httpd.apache.org/docs/2.2/mod/mod_dumpio.html
- Dumps all input and output to the error log

```
# DumpIOLogLevel notice (2.3)
DumpIOInput On
DumpIOOutput On
```




mod_log_forensic



- http://httpd.apache.org/docs/2.2/mod/mod_log_forensic.html
- Logs at the start, end of a request
- Uses unique IDs to match the two
- check_forensic script alerts you to requests that did not complete



LogLevel



- LogLevel changes the level at which error messages are emitted
- Can increase/decrease the volume of your error_log
- In practice, this seldom adds useful information



RewriteLog



- Should always turn on the RewriteLog when RewriteRules aren't doing what you expect them to do
- Can only be turned on in main config, not in .htaccess files

RewriteLog



RewriteLog logs/rewrite_log
RewriteLogLevel 9



Other logs



- suexec
- SSL



Logging to syslog

- “Offsite” logs, in the event of catastrophe
- Multiple servers logging to the same place





ErrorLog

ErrorLog syslog
...
ErrorLog syslog:local0





Then, in /etc/syslog.conf


```
local0.* /var/log/error_log  
...  
local1.* @192.168.1.22:32376
```





access_log

- mod_log_config doesn't log to syslog
- Have to use piped log handlers





Solution

CustomLog |/usr/bin/apache_syslog combined

- Where the script looks like:

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

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

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



18_perl_syslog



...

- Sys::Syslog is a standard Perl module, so you already have it installed
- Piped logging is a standard feature
- Script is started at server startup and remains running for the life of the server





Or ...

- A simpler solution:

```
CustomLog "|/usr/bin/logger -t apache -i -p  
local6.notice" combined
```





WebDAV

- Network filesystem over HTTP (or HTTPS)
- Manage your web content
- Access your files from anywhere
- Impress your friends



DAV

- Distributed
- Authoring
- Versioning



Modules

- mod_dav
- mod_dav_fs

```
./configure --enable-modules=most \  
  --enable-mods-shared=all \  
  --enable-dav --enable-dav-fs
```



Recipe

```
DavLockDb dav/davlock
```

```
Alias /dav /var/www/dav
```

```
<Directory /var/www/dav>
```

```
    Dav On
```

```
</Directory>
```




Accessing

- <http://servername.com/dav/>



Client applications

- Most modern operating systems
- cadaver – Simple command-line application
- NetDrive – Windows
- DavExplorer – Java



For More Information

- Thursday morning
- Bill Rowe
- <http://www.eu.apachecon.com/program/talk/39>



Caveat

- Files must be writeable by the Apache user
- This makes most of us VERY uncomfortable



Solution

- Run two Apache instances, with different permissions:
- Instance 1, runs as `apache.apache`, content owned by `dav.dav`
- Instance 2, runs as `dav.dav`, has access to these directories
- Instance 2 runs over SSL, and is authenticated

Like ...

1

```
User apache  
Group apache  
DocumentRoot /var/www
```

2

```
User dav  
Group dav  
DocumentRoot /var/www  
<Directory /var/www>  
    Dav On  
</Directory>
```

```
/var/www> ls -lad .  
drwxrwxr-x  9 dav  dav  306 Mar 23 22:42 .
```



Preventing malicious requests with mod_security



- modsecurity.org
- Apache module to do request filtering

New syntax



- Syntax has changed considerably in mod_security 2, so some of these recipes might not work quite as expected, depending on what version you're using.

Core rules



- Download the core rules from <http://modsecurity.org/download/index.html>
- Try to understand before using - this will avoid blocking desirable traffic

Basic Configs



- Turn on the engine
- Enable scanning of request body

```
# Basic configuration options
SecRuleEngine On
SecRequestBodyAccess On
SecResponseBodyAccess Off
```

Trivial example



```
# Trivial SQL blocking rule
```

```
SecDefaultAction \
```

```
log,auditlog,deny,status:403,phase:2,t:lowercase
```

```
SecRule REQUEST_URI|QUERY_STRING insert
```

phase:2 indicates that this runs after URL mapping. phase:1 runs before URL mapping. t:lowercase lowercases the variable before comparison is applied

SecRule



SecRule VARIABLES OPERATOR [ACTIONS]

- ACTIONS is optional - SecDefaultAction will be used
- Use multiple variables like REQUEST_URI|ARGS|QUERY_STRING
- OPERATOR is a regex match, by default

More complex example



```
# file injection
SecRule REQUEST_FILENAME|ARGS|
ARGS_NAMES|REQUEST_HEADERS "(?:\b(?:\.
(?:ht(?:access|passwd|group)|www_?acl)|global
\.asa|httpd\.conf|boot\.ini)\b|\/etc\/)" \
    "capture,ctl:auditLogParts=
+E,deny,log,auditlog,status:501,msg:'Remote File
Access Attempt.',severity:'2'"
```

Note:



- mod_security is extremely powerful
- mod_security 2 adds a huge amount of new functionality and flexibility
- I'm just beginning to learn it, so you should go to the mailing lists with your questions
- <http://modsecurity.org/>



Enabling PHP



- There's a certain amount of disagreement about the Right Way to do this
- So, if there's any confusion, you should keep in mind one important rule of thumb



● Rich is Right





Now that we've got that
out of the way ...



AddType



- AddType associates a MIME type with a file extension
- It tells the browser how to display a particular type of content
- e.g. image/gif files should use the GIF rendering engine, and application/pdf files should use Adobe Acrobat



AddType



```
AddType image/gif .gif
```



AddHandler



- AddHandler tells the server how to process a certain type of file
- Calls a Handler which does something to the file before passing it along to the client



AddHandler



AddHandler cgi-script .cgi



PHP



- PHP is a handler
- However, PHP predates the AddHandler directive, and so uses the AddType directive
- This is a grotty hack, and should be shunned



The right way:



php

```
AddHandler application/x-httpd-php .php
```

The other way

The PHP logo, consisting of the lowercase letters 'php' in a white, italicized font inside a blue oval, which is itself inside a blue square with a thin black border.

```
AddType application/x-httpd-php .php
```




Multiple file extensions

- In either case, multiple file extensions can cause problems.
- foo.php.txt
- With php as a handler, it will still be executed
- With php as a mime type, it will lose its text/plain attribute



Discussion



- They both work
- Since it's a handler, I recommend using AddHandler
- Rasmus disagrees



LoadModule



- Must also ensure that the php module is loaded:

```
LoadModule php5_module modules/libphp5.so
```



Testing

```
<?php  
    phpinfo();  
?>
```



PHP Version 5.2.0



System	Darwin Rocinante.local 8.9.1 Darwin Kernel Version 8.9.1: Thu Feb 22 20:55:00 PST 2007; root:xnu-792.18.15~1/RELEASE_i386 i386
Build Date	Jan 2 2007 10:06:07
Configure Command	'./configure' '--with-apxs2=/usr/local/apache/bin/apxs' '--with-mysql' '--with-zlib-dir=/sw'
Server API	Apache 2.0 Handler
Virtual Directory Support	disabled
Configuration File (php.ini) Path	/usr/local/lib
PHP API	20041225
PHP Extension	20060613
Zend Extension	220060519
Debug Build	no
Thread Safety	disabled
Zend Memory Manager	enabled
IPv6 Support	enabled
Registered PHP Streams	php, file, data, http, ftp, compress.zlib
Registered Stream Socket Transports	tcp, udp, unix, udg
Registered Stream Filters	string.rot13, string.toupper, string.tolower, string.strip_tags, convert.*, consumed, convert.iconv.*, zlib.*

This program makes use of the Zend Scripting Language Engine:
 Zend Engine v2.2.0, Copyright (c) 1998-2006 Zend Technologies

Powered By




PHP Credits

Configuration



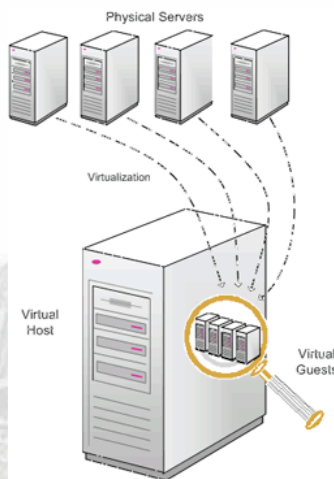
Mass Virtual Hosting

- Several ways to do it
- Most of them are icky
- Don't do this unless you really need to



When?

- When you have LOTS of vhosts
- Most of us don't have that many vhosts
- Most of us are better off just making <VirtualHost> blocks




Include

- Put each vhost in its own file
- Include them

```
Include conf/vhosts/*.conf
```


- 000Default.conf
- ZZZWildcard.conf





mod_vhost_alias


- Comes with Apache
- Very well documented
- Rather limiting





mod_vhost_alias

- Substitutes bits of the hostname into the directory path, using templates, like ...




%0 gives you the entire hostname:

VirtualDocumentRoot /var/www/%0

www.example.com maps to

/var/www/www.example.com




%1 gives you the first part of the hostname:

VirtualDocumentRoot /var/www/%1

www.example.com maps to

/var/www/www




%2 gives you the second part of
the hostname:

VirtualDocumentRoot /var/www/%2

www.example.com maps to


/var/www/example



%3 gives you the third part of
the hostname:

VirtualDocumentRoot /var/www/%3

www.example.com maps to
/var/www/com




And so ...

VirtualDocumentRoot /var/www/%1/%2/%3


www.example.com maps to

/var/www/www/example/com




-1, -2, -3 counts from the right
VirtualDocumentRoot /var/www/**%-1**/**%-2**

www.example.com maps to
/var/www/**com**/**example**




m.n lets you choose particular letters
VirtualDocumentRoot \
 /var/www/**%-2.1**/**%-2.2**/**%-2.3**+

www.example.com maps to
/var/www/**e**/**x**/ample



```
# likewise ...
VirtualScriptAlias \
    /var/www/%-2.1/%-2.2/%-2.3+/cgi


# /cgi-bin maps to the directory
# /var/www/e/x/ample/cgi
# for www.example.com
```






Advantages

- Don't have to restart to add a new vhost
- All your vhosts are identical and predictable




Caveats

- All your vhosts must be identical
- You can't intermix vhost_alias vhosts and regular vhosts on the same IP address
- mod_alias and mod_userdir always override vhost_alias directives




Vhosts with mod_rewrite

```
RewriteEngine On  
RewriteCond %{HTTP_HOST} \  
                ^([\^.]\.example\.com  
RewriteRule (.*) /var/www/%1$1
```



Disadvantages

- May cause interactions with other RewriteRules (like in .htaccess files) that may cause breakage.





Customized Error Messages



- Override the default boring error responses
- Less jarring to the user
- Give them useful information or links



ErrorDocument



ErrorDocument 404 /errors/404.html



Not always an error message

- Can be used as a “default document” when something is not found
- ErrorDocument 404 /index.html
- ErrorDocument 401 /register.html



Embedded logic



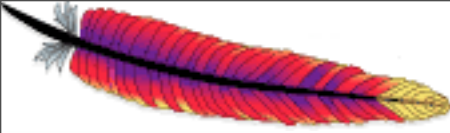
- Can contain basic embedded logic using SSI
- See `extras/httpd-multilang-error.doc.conf` for extended example



Embedded logic



```
Alias /error /www/error
<Directory /www/error>
  Options IncludesNoExec
  AddOutputFilter Includes html
</Directory>
ErrorDocument 404 /error/404.html
```



Then 404.html is ...

```
<html>
```

```
<head><title>Not Found</title></head>
```

```
<body>
```

```
<!--#if expr="HTTP_REFERER" -->
```

The link from <!--#echo var="HTTP_REFERER" --> appears to be bad.

```
<!--#else -->
```

The URL you entered could not be found here.

```
<!--#endif -->
```

```
</body>
```

```
</html>
```





URL Handler (“Rewrite Everything”)

- One content handler for all requests



Recipe

```
RewriteEngine On  
RewriteCond %{REQUEST_FILENAME} !-f  
RewriteCond %{REQUEST_FILENAME} !-d  
RewriteCond $1 !=/handler.php  
RewriteRule (.*) /handler.php [PT]
```



handler.php

- handler.php would know what was actually requested by looking at `$_SERVER['REQUEST_URI']`
- Other files (images, css, static files) are served as normal, due to the `-f` test.














ErrorDocument

- Can also be done with an ErrorDocument
- ErrorDocument 404 /handler.php
- HOWEVER, ErrorDocuments can't receive POST data, so this is a rather limited solution



Fancy Directory Listings

Index of /icons

Name	Last modified	Size	Description
 Parent Directory		-	
 a.gif	20-Nov-2004 15:16	246	
 a.png	20-Nov-2004 15:16	293	
 alert.black.gif	20-Nov-2004 15:16	242	
 alert.black.png	20-Nov-2004 15:16	279	
 alert.red.gif	20-Nov-2004 15:16	247	
 alert.red.png	20-Nov-2004 15:16	298	
 apache_pb.gif	20-Nov-2004 15:16	2.3K	
 apache_pb.png	20-Nov-2004 15:16	1.4K	
 apache_pb2.gif	20-Nov-2004 15:16	2.4K	
 apache_pb2.png	20-Nov-2004 15:16	1.4K	

- Auto directory listings are ugly
- It would be nice to have more control over them

Suppress unwanted columns

Index of /icons

Name	Last modified	Size	Description
Parent Directory		-	
a.gif	20-Nov-2004 15:16	246	
a.png	20-Nov-2004 15:16	293	
alert.black.gif	20-Nov-2004 15:16	242	
alert.black.png	20-Nov-2004 15:16	279	
alert.red.gif	20-Nov-2004 15:16	247	
alert.red.png	20-Nov-2004 15:16	298	
apache_pb.gif	20-Nov-2004 15:16	2.3K	
apache_pb.png	20-Nov-2004 15:16	1.4K	
apache_pb2.gif	20-Nov-2004 15:16	2.4K	
apache_pb2.png	20-Nov-2004 15:16	1.4K	

IndexOptions SuppressLastModified \
SuppressDescription

Insert "wrapper" htn

Index of /icons

Name	Last modified	Size	Description
Parent Directory	-	-	-
a.gif	20-Nov-2004 15:16	246	
a.png	20-Nov-2004 15:16	293	
alert.black.gif	20-Nov-2004 15:16	242	
alert.black.png	20-Nov-2004 15:16	279	
alert.red.gif	20-Nov-2004 15:16	247	
alert.red.png	20-Nov-2004 15:16	298	
apache_pb.gif	20-Nov-2004 15:16	2.3K	
apache_pb.png	20-Nov-2004 15:16	1.4K	
apache_pb2.gif	20-Nov-2004 15:16	2.4K	
apache_pb2.png	20-Nov-2004 15:16	1.4K	

```
IndexOptions SuppressHTMLPreamble
HeaderName /style/header.html
ReadmeName /style/footer.html
```



Wrapper

Index of /icons

Name	Last modified	Size	Description
Parent Directory		-	
a.gif	20-Nov-2004 15:16	246	
a.png	20-Nov-2004 15:16	293	
alert.black.gif	20-Nov-2004 15:16	242	
alert.black.png	20-Nov-2004 15:16	279	
alert.red.gif	20-Nov-2004 15:16	247	
alert.red.png	20-Nov-2004 15:16	298	
apache_pb.gif	20-Nov-2004 15:16	2.3K	
apache_pb.png	20-Nov-2004 15:16	1.4K	
apache_pb2.gif	20-Nov-2004 15:16	2.4K	
apache_pb2.png	20-Nov-2004 15:16	1.4K	

```
<html>
<head><title>Directory Listing</title>
</head>
<body>
```

38_header

... Listing goes here ...

```
</body>
</html>
```

39_footer



CSS

Index of /icons

Name	Last modified	Size	Description
Parent Directory		-	
a.gif	20-Nov-2004 15:16	246	
a.png	20-Nov-2004 15:16	293	
alert.black.gif	20-Nov-2004 15:16	242	
alert.black.png	20-Nov-2004 15:16	279	
alert.red.gif	20-Nov-2004 15:16	247	
alert.red.png	20-Nov-2004 15:16	298	
apache_pb.gif	20-Nov-2004 15:16	2.3K	
apache_pb.png	20-Nov-2004 15:16	1.4K	
apache_pb2.gif	20-Nov-2004 15:16	2.4K	
apache_pb2.png	20-Nov-2004 15:16	1.4K	

IndexStyleSheet "/css/style.css"



Caching Dynamic Content

- Much of your 'dynamic' content doesn't change very often
- Cache it to improve performance



Warning

- Caching dynamic content, by definition, causes stale content to be served
- Note that “private” content will not (usually) be cached



Cache for 10 minutes

```
CacheRoot /usr/local/apache/cache  
CacheEnable disk /  
CacheDirLevels 5  
CacheDirLength 3  
  
# Cache stuff for 10 minutes  
CacheDefaultExpire 600  
CacheIgnoreCacheControl On
```




Cleaning the cache

- There are two ways to clear the cache
- Depending on how much you care ...



htcacheclean

- Cleans up your cache periodically
- Can specify an upper limit on size
- -t deletes empty directories (in the cache)

```
htcacheclean -d 10 \  
-p /var/cache/apache \  
-l 50M \  
-t
```



htacheclean

- Runs every 10 minutes (or whatever you specify)
- Keeps cache below 50M (or whatever ...)
- Purges older content first



rm -rf

- If you don't care about gradually expiring content, just delete everything in the cache directory
- Faster - if you need to quickly purge the cache



/server-info goodness



- mod_info gives useful information about your server configuration



Configuration

```
<Location /server-info>  
    SetHandler server-info  
    # Order deny,allow  
    # deny from all  
    # allow from 192.168  
</Location>
```



Security considerations

- Should protect this resource
- Don't give crackers additional information



/server-info



Apache Server Information

Subpages:

[Configuration Files](#), [Server Settings](#), [Module List](#), [Active Hooks](#)

Sections:

[Server Settings](#), [Startup Hooks](#), [Request Hooks](#)

Loaded Modules:

[mod_php5.c](#), [mod_rewrite.c](#), [mod_alias.c](#), [mod_speling.c](#), [mod_actions.c](#),
[mod_imagemap.c](#), [mod_dir.c](#), [mod_negotiation.c](#), [mod_dav_fs.c](#),
[mod_cgi.c](#), [mod_info.c](#), [mod_autoindex.c](#), [mod_status.c](#), [mod_dav.c](#),
[mod_mime.c](#), [mod_proxy_balancer.c](#), [mod_proxy_http.c](#), [mod_proxy.c](#),
[mod_version.c](#), [mod_setenvif.c](#), [mod_headers.c](#), [mod_expires.c](#),
[mod_mime_magic.c](#), [mod_env.c](#), [mod_logio.c](#), [mod_log_config.c](#),
[mod_deflate.c](#), [mod_filter.c](#), [mod_include.c](#), [mod_dumpio.c](#),
[mod_auth_digest.c](#), [mod_auth_basic.c](#), [mod_authz_default.c](#),
[mod_authz_user.c](#), [mod_authz_groupfile.c](#), [mod_authz_host.c](#),



/server-info?config



Configuration:

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

```
40: Listen 80
119: User daemon
120: Group daemon
140: ServerAdmin you@example.com
149: ServerName rocinante.rcbowen.com
156: DocumentRoot "/usr/local/apache/htdocs"
166: <Directory />
167:     Options FollowSymLinks
168:     AllowOverride None
169:     Order deny,allow
170:     Deny from all
    : </Directory>
183: <Directory "/usr/local/apache/htdocs">
196:     Options Indexes FollowSymLinks
203:     AllowOverride None
208:     Order allow,deny
209:     Allow from all
    : </Directory>
218: DirectoryIndex index.html
225: <FilesMatch "^\.ht">
226:     Order allow,deny
227:     Deny from all
228:     Satisfy All
    : </FilesMatch>
238: ErrorLog logs/error_log
```

?config



- Includes Include'd files
- Shows line numbers, file names

```
26: ForceLanguagePriority Prefer Fallback
   : </Directory>
In file: /usr/local/apache/conf/cookbook/balancer.conf
3: <Proxy balancer://mycluster>
4:   BalancerMember http://192.168.1.50:80
5:   BalancerMember http://192.168.1.51:80
   : </Proxy>
7: ProxyPass /test balancer://mycluster/
10: <Location /balancer-manager>
11:   SetHandler balancer-manager
   : </Location>
In file: /usr/local/apache/conf/cookbook/php.conf
5: AddHandler application/x-httpd-php .php
```



?config



- Particularly useful on third-party distros of Apache with unfamiliar config file layout
- Locate overlapping or conflicting configuration settings

?server



Server Settings

Server Version: Apache/2.2.3 (Unix) DAV/2 PHP/5.2.0

Server Built: Mar 23 2007 22:06:40

Module Magic Number: 20051115:3

Hostname/port: localhost:80

Timeouts: connection: 300 keep-alive: 300

MPM Name: Prefork

MPM Information: Max Daemons: 256 Threaded: no Forked: yes

Server Architecture: 32-bit

Server Root: /usr/local/apache

Config File: /usr/local/apache/conf/httpd.conf

Server Built With:

- D APACHE_MPM_DIR="server/mpm/prefork"
- D APR_HAS_MMAP
- D APR_HAVE_IPV6 (IPv4-mapped addresses enabled)
- D APR_USE_SYSVSEM_SERIALIZE
- D SINGLE_LISTEN_UNSERIALIZED_ACCEPT
- D APR_HAS_OTHER_CHILD
- D AP_HAVE_RELIABLE_PIPED_LOGS
- D HTTPD_ROOT="/usr/local/apache"
- D SUEXEC_BIN="/usr/local/apache/bin/suexec"
- D DEFAULT_ERRORLOG="logs/error_log"
- D AP_TYPES_CONFIG_FILE="conf/mime.types"
- D SERVER_CONFIG_FILE="conf/httpd.conf"



?server



- Equivalent to `httpd -v`

?list



Server Module List

mod_php5.c
mod_rewrite.c
mod_alias.c
mod_speling.c
mod_actions.c
mod_imagemap.c
mod_dir.c
mod_negotiation.c
mod_dav_fs.c
mod_cgi.c
mod_info.c
mod_autoindex.c
mod_status.c
mod_dav.c
mod_mime.c
mod_proxy_balancer.c
mod_proxy_http.c
mod_proxy.c
mod_version.c
mod_setenvif.c
mod_headers.c
mod_expires.c
mod_mime_magic.c
mod_env.c

mod_env.c
mod_logio.c
mod_log_config.c
mod_deflate.c
mod_filter.c
mod_include.c
mod_dumpio.c
mod_auth_digest.c
mod_auth_basic.c
mod_authz_default.c
mod_authz_user.c
mod_authz_groupfile.c
mod_authz_host.c
mod_authn_default.c
mod_authn_file.c
mod_so.c
http_core.c
prefork.c
core.c

?hooks



Startup Hooks

Pre-Config:

```
-10 mod_log_config.c
-10 mod_logio.c
-10 prefork.c
10 mod_php5.c
10 mod_headers.c
10 mod_proxy.c
10 mod_rewrite.c
```

Test Configuration:

```
10 mod_so.c
```

Post Configuration:

```
-10 mod_include.c
-10 mod_cgi.c
-10 core.c
00 mod_mime_magic.c
10 mod_mime.c
10 mod_dav.c
10 mod_status.c
10 mod_rewrite.c
10 mod_php5.c
10 mod_auth_digest.c
10 mod_headers.c
10 mod_proxy.c
```

Open Logs:

```
10 prefork.c
-10 core.c
10 mod_log_config.c
```

Child Init:

```
10 mod_proxy.c
10 mod_auth_digest.c
```

And if you select one ...



Module Name: `mod_log_config.c`

Content handlers: *none*

Configuration Phase Participation: Create Server Config, Merge Server Configs

Request Phase Participation: Logging

Module Directives:

CustomLog - a file name, a custom log format string or format name, and an optional "env=" (see docs)

TransferLog - the filename of the access log

LogFormat - a log format string (see docs) and an optional format name

CookieLog - the filename of the cookie log

BufferedLogs - Enable Buffered Logging (experimental)

Current Configuration:

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

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

```
253: LogFormat "%h %l %u %t \"%r\" %>s %b" common
```

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

```
267: CustomLog logs/access_log common
```

- That's `?mod_log_config.c`



/server-status goodness



- Displays the current status of the server
- Also some basic statistical reports

Configuration



```
<Location /server-status>  
  SetHandler server-status  
  # Order deny,allow  
  # deny from all  
  # allow from 192.168  
</Location>
```



Security



- As with /server-info, protect
- Also, reveals what users are looking at what content

ExtendedStatus



- ExtendedStatus On
- Gives more information



/server-status



Apache Server Status for rocinante.rcbowen.com

Server Version: Apache/2.2.3 (Unix) DAV/2 PHP/5.2.0

Server Built: Mar 23 2007 22:06:40

Current Time: Saturday, 21-Apr-2007 10:46:45 EDT

Restart Time: Saturday, 21-Apr-2007 10:45:09 EDT

Parent Server Generation: 7

Server uptime: 1 minute 35 seconds

Total accesses: 1 - Total Traffic: 2 kB

CPU Usage: u0 s0 cu0 cs0

.0105 requests/sec - 21 B/second - 2048 B/request

1 requests currently being processed, 5 idle workers

W.....
.....
.....
.....

Scoreboard Key:

"_" Waiting for Connection, "s" Starting up, "R" Reading Request,

"w" Sending Reply, "k" Keepalive (read), "D" DNS Lookup,

"c" Closing connection, "L" Logging, "G" Gracefully finishing,

"I" Idle cleanup of worker, "." Open slot with no current process

Srv	PID	Acc	M	CPU	SS	Req	Conn	Child	Slot	Client	VHost	Request
0-7	21321	0/0/0	W	0.00	0	0	0.0	0.00	0.00	127.0.0.1	rocinante.rcbowen.com	GET /server-status HTTP/1.1
1-7	21322	0/1/1	_	0.00	91	3	0.0	0.00	0.00	127.0.0.1	rocinante.rcbowen.com	GET /server-status HTTP/1.1

Srv Child Server number - generation



ExtendedStatus



Client	VHost	Request
121.233.25.240	archive.apache.org	GET /dist/struts/binaries/jakarta-struts-1.2.2.zip HTTP/1.1
218.241.91.2	struts.apache.org	GET /dtds/struts-config_1_2.dtd HTTP/1.1
85.57.130.253	ibatis.apache.org	GET /dtd/sql-map-2.dtd HTTP/1.1
62.143.135.190	www.apache.org	GET /images/redarrow.gif HTTP/1.1
58.213.10.75	www.apache.org	GET /dist/httpd/httpd-2.2.4.tar.bz2 HTTP/1.0
85.207.73.206	ant.apache.org	GET /manual/toc.html HTTP/1.1
74.140.181.181	httpd.apache.org	GET /server-status HTTP/1.1
74.6.74.35	maven.apache.org	GET /maven-1.x/plugins/ejb/announcements/announcement-1.7.txt HTTP/1.1
85.95.80.125	tomcat.apache.org	GET /robots.txt HTTP/1.1
218.225.203.163	tapestry.apache.org	GET /favicon.ico HTTP/1.1
66.131.172.154	harmony.apache.org	GET /site.css HTTP/1.1
74.6.86.152	mail-archives.apache.org	GET /mod_mbox/jakarta-oro-user/200206.mbox/date HTTP/1.0
88.191.13.246	tomcat.apache.org	GET /tomcat-5.0-doc/images/printer.gif HTTP/1.1
84.113.197.3	ant.apache.org	GET /images/tabSel-left.gif HTTP/1.1
209.47.162.29	jakarta.apache.org	GET /struts/dtds/struts-config_1_1.dtd HTTP/1.1
67.82.58.195	ant.apache.org	GET /manual/coretasklist.html HTTP/1.0
76.21.243.223	maven.apache.org	GET /maven-1.x/images/icon_error_sml.gif HTTP/1.1
88.100.207.162	jakarta.apache.org	GET /poi/skin/images/valid-html401.png HTTP/1.1
88.100.207.162	jakarta.apache.org	GET /poi/resources/images/group-logo.gif HTTP/1.1
60.208.254.184	www.apache.org	GET /style/style.css HTTP/1.1
83.24.219.35	httpd.apache.org	GET /docs/2.2/configuring.html HTTP/1.1
60.208.254.184	www.apache.org	GET /foundation/preFAQ.html HTTP/1.1
208.246.78.240	portals.apache.org	GET /jetspeed-1/channels/turbine.rss HTTP/1.1
65.54.188.11	mail-archives.apache.org	GET /mod_mbox/struts-user/200304.mbox/%3cSea1-F149eurQQtHT1D000
209.249.11.3	wiki.apache.org	GET /cocoon/WebDAVCMS HTTP/1.0
84.113.197.3	ant.apache.org	GET /images/menu-right.gif HTTP/1.1

/server-status?auto



- Machine-readable
- Useful for things like mrtg

```
Total Accesses: 128900979
Total kBytes: 20000683955
CPULoad: 1.25654
Uptime: 1887566
ReqPerSec: 68.2895
BytesPerSec: 10850300
BytesPerReq: 158887
BusyWorkers: 167
IdleWorkers: 89
Scoreboard: ___WVC_KWW__WK_KW_KW__K_KKCKKK_KW_RWK_R_WK__C_RKKKRKWWR_K___KKK.....
```


Example mrtg script




```
#!/usr/bin/perl
use LWP::Simple;
$content = get("http://localhost/server-status?auto");

$content =~ m/BusyWorkers: (\d+)/s;
print $1 . "\n";

$content =~ m/IdleWorkers: (\d+)/s;
print $1 . "\n";
```



/server-status?refresh=4



- Automatically refreshes every N seconds
- Or, combine them:

<http://rocinante.rcbowen.com/server-status?auto&refresh=2>

- Not sure what that's useful for ...



UserDir without the ~

- Using `mod_rewrite` to create a per-user URL, without the ~



Problem

2

- We want:

<http://example.com/username/foo>

- To work the same as:

<http://example.com/~username/foo>



But

2

- Somehow, ...

<http://example.com/not-a-username/foo>

- still needs to work properly



-d

2

RewriteEngine On

If that home directory exists ...

RewriteCond /home/\$1 -d

RewriteRule ^/([^/]+)/(.*) /home/\$1/www/\$2

How this works

2

```
RewriteEngine On  
RewriteCond /home/$1 -d  
RewriteRule ^/([^/]+)/(.*) /home/$1/www/$2
```

- That's right, **\$1** is used in the RewriteCond **before** it is defined in the RewriteRule
- Pretty cool, hmm?



Huh?

2

- RewriteRules are always evaluated before the corresponding RewriteConds
- You can watch this in the RewriteLog:

<http://example.com/rbowen/index.html>

- (3) applying pattern '^/([^/]+)/(.*)' to uri '/rbowen/index.html'
- (4) RewriteCond: input='/home/rbowen' pattern='-d' => matched
- (2) rewrite '/rbowen/index.html' -> '/home/rbowen/www/index.html'
- (2) local path result: /home/rbowen/www/index.html
- (1) go-ahead with /home/rbowen/www/index.html [OK]



That's all, folks

- rbowen@apache.org
- <http://people.apache.org/~rbowen/>



- ProxyPass / http://somewhere
- RewriteRule (.*) http://somewhere\$1 [P]