

Tutorial Slides: mod_perl 2.0 By Example

by Stas Bekman
<http://stason.org/>
<stas@stason.org>
TicketMaster

ApacheCon US 2002
Monday, Nov 18 2002
Las Vegas, Nevada, USA

This talk is available from:
<http://stason.org/talks/>

Last modified Sun Nov 17 05:26:13 2002 GMT

1 Getting Your Feet Wet with mod_perl 2.0

1.1 About

- Prerequisites
- Installation
- Configuration
- Server Launch and Shutdown
- Registry Scripts
- Handler Modules

1.2 Prerequisites

- Apache 2.0 is required.
- Perl, depending on MPM:
 - prefork: 5.6.0, better 5.6.1
 - threaded: 5.8.0 + ithreads
- The installation details are in the handouts

1.3 Installation

```
% lwp-download \  
  http://perl.apache.org/dist/mod_perl-2.0.01.tar.gz  
% tar -xvzf mod_perl-2.x.xx.tar.gz  
% cd modperl-2.0  
% perl Makefile.PL MP_AP_PREFIX=$HOME/httpd/prefork \  
  MP_INST_APACHE2=1  
% make && make test && make install
```

- `MP_AP_PREFIX` == Apache installation prefix

1.4 Configuration

- Enable DSO:

```
LoadModule perl_module modules/mod_perl.so
```

- Find 2.0 modules

```
PerlModule Apache2
```

- Enable the 1.0 compatibility layer

```
PerlModule Apache::compat
```

1.5 Server Launch and Shutdown

- Start:

```
% $HOME/httpd/prefork/bin/apachectl start
```

```
% tail -f $HOME/httpd/prefork/logs/error_log
```

```
[Tue Sep 03 12:34:57 2002] [notice] Apache/2.0.41-dev (Unix)
```

```
mod_perl/1.99_05-dev Perl/v5.8.0 mod_ssl/2.0.41-dev OpenSSL/0.9.6d
```

```
DAV/2 configured -- resuming normal operations
```

- Stop:

```
% $HOME/httpd/prefork/bin/apachectl stop
```

1.6 Registry Scripts

- Configuration

```
Alias /perl/ /home/httpd/httpd-2.0/perl/  
<Location /perl/>  
    SetHandler perl-script  
    PerlResponseHandler ModPerl::Registry  
    PerlOptions +ParseHeaders  
    Options +ExecCGI  
</Location>
```

- Restart the server

- A simple script:

```
#!/usr/bin/perl
print "Content-type: text/plain\n\n";
print "mod_perl 2.0 rocks!\n";
```

- Make it executable and readable:

```
% chmod a+rx /home/httpd/httpd-2.0/perl/rock.pl
```

- Request:

```
% lwp-request http://localhost/perl/rock.pl
mod_perl 2.0 rocks!
```

1.7 Handler Modules

- /home/httpd/httpd-2.0/perl/MyApache/Rocks.pm

```
package MyApache::Rocks;
```

```
use strict;
```

```
use warnings;
```

```
use Apache::RequestRec ();
```

```
use Apache::RequestIO ();
```

```
use Apache::Const -compile => qw(OK);
```

```
sub handler {  
    my $r = shift;  
  
    $r->content_type('text/plain');  
    print "mod_perl 2.0 rocks!\n";  
  
    return Apache::OK;  
}  
1;
```

- Adjust @INC via a startup file

```
use lib qw(/home/httpd/httpd-2.0/perl);
```

- loaded from *httpd.conf*:

```
PerlRequire /home/httpd/httpd-2.0/perl/startup.pl
```

- Configure:

```
<Location /rocks>  
    SetHandler perl-script  
    PerlResponseHandler MyApache::Rocks  
</Location>
```

- Restart server

- Issue request:

```
% lwp-request http://localhost/rocks  
mod_perl 2.0 rocks!
```


2 Introducing mod_perl Handlers

2.1 About

- Handler Anatomy
- mod_perl Handlers Categories
- Bucket Brigades
- Single Phase's Multiple Handlers Behavior

2.2 Handler Anatomy

- Apache distinguishes between numerous phases
- Each phase provides a hook: *ap_hook_<phase_name>*
- These hooks are used by modules to alter the default Apache behavior
- Hooks are usually referred to as handlers or callbacks
- Naming convention: `PerlFooHandler`
- e.g. `PerlResponseHandler` configures the response callback.

```
package MyApache::CurrentTime;

use strict;
use warnings;

use Apache::RequestRec ();
use Apache::RequestIO ();

use Apache::Const -compile => qw(OK);

sub handler {
    my $r = shift;

    $r->content_type('text/plain');
    $r->print("Now is: " . scalar(localtime) . "\n");

    return Apache::OK;
}
1;
```

- Configuring the response handler:

```
PerlModule MyApache::CurrentTime
<Location /time>
    SetHandler modperl
    PerlResponseHandler MyApache::CurrentTime
</Location>
```

- A request to *http://localhost/time* returns the current time

2.3 mod_perl Handlers Categories

The mod_perl handlers can be divided by their application scope in several categories:

- **Server life cycle**
 - PerlOpenLogsHandler
 - PerlPostConfigHandler
 - PerlChildInitHandler
 - PerlChildExitHandler
- **Protocols**
 - PerlPreConnectionHandler
 - PerlProcessConnectionHandler

- **Filters**
 - `PerlInputFilterHandler`
 - `PerlOutputFilterHandler`
- **HTTP Protocol**
 - `PerlPostReadRequestHandler`
 - `PerlTransHandler`
 - `PerlInitHandler`
 - `PerlHeaderParserHandler`
 - `PerlAccessHandler`
 - `PerlAuthenHandler`
 - `PerlAuthzHandler`
 - `PerlTypeHandler`
 - `PerlFixupHandler`
 - `PerlResponseHandler`
 - `PerlLogHandler`
 - `PerlCleanupHandler`

2.4 Bucket Brigades

- Apache 2.0 implements request and response data flow filtering
- Modules can filter each other's output
- No need to modify Apache to accommodate SSL, compressions, transformation filters.
- The *Bucket Brigades* technology was introduced to make IO filtering efficient and avoid unnecessary copying.


Buckets:

- A bucket represents a chunk of data.
- Buckets linked together comprise a brigade.
- Each bucket in a brigade can be modified, removed and replaced with another bucket.
- Bucket types: files, data blocks, flush and end of stream indicators, pools, etc.
- To manipulate a bucket one doesn't need to know its internal representation.

Bucket Brigades:

- A stream of data is represented by bucket brigades.
- Filters manipulate brigades one by one (adding/modifying/removing buckets)
- ... and pass the brigade to the next filter on the stack

- Here's an imaginary bucket brigade after it has passed through several filters.
- Some buckets were removed, some modified and some added.



- More about BBs when we talk about protocols and filters

2.5 Single Phase's Multiple Handlers Behavior

- Phases' behavior varies when there is more than one handler registered to run for the same phase.
- The following table specifies each handler's behavior in this situation:

| Directive | Type |
|---------------------------------|----------------|
| PerlOpenLogsHandler | RUN_ALL |
| PerlPostConfigHandler | RUN_ALL |
| PerlChildInitHandler | VOID |
| PerlPreConnectionHandler | RUN_ALL |

| | |
|---|------------------------|
| <code>PerlProcessConnectionHandler</code> | <code>RUN_FIRST</code> |
| <code>PerlPostReadRequestHandler</code> | <code>RUN_ALL</code> |
| <code>PerlTransHandler</code> | <code>RUN_FIRST</code> |
| <code>PerlInitHandler</code> | <code>RUN_ALL</code> |
| <code>PerlHeaderParserHandler</code> | <code>RUN_ALL</code> |
| <code>PerlAccessHandler</code> | <code>RUN_ALL</code> |
| <code>PerlAuthenHandler</code> | <code>RUN_FIRST</code> |
| <code>PerlAuthzHandler</code> | <code>RUN_FIRST</code> |
| <code>PerlTypeHandler</code> | <code>RUN_FIRST</code> |
| <code>PerlFixupHandler</code> | <code>RUN_ALL</code> |
| <code>PerlResponseHandler</code> | <code>RUN_FIRST</code> |
| <code>PerlLogHandler</code> | <code>RUN_ALL</code> |
| <code>PerlCleanupHandler</code> | <code>RUN_ALL</code> |
| <code>PerlInputFilterHandler</code> | <code>VOID</code> |
| <code>PerlOutputFilterHandler</code> | <code>VOID</code> |

The types:

- VOID
 - Executed in the order they have been registered disregarding their return values.
 - Though in `mod_perl` they are expected to return `Apache::OK`.
- RUN_ALL
 - Executed in the order they have been registered until the first handler that returns something other than `Apache::OK` or `Apache::DECLINED`.

- RUN_FIRST

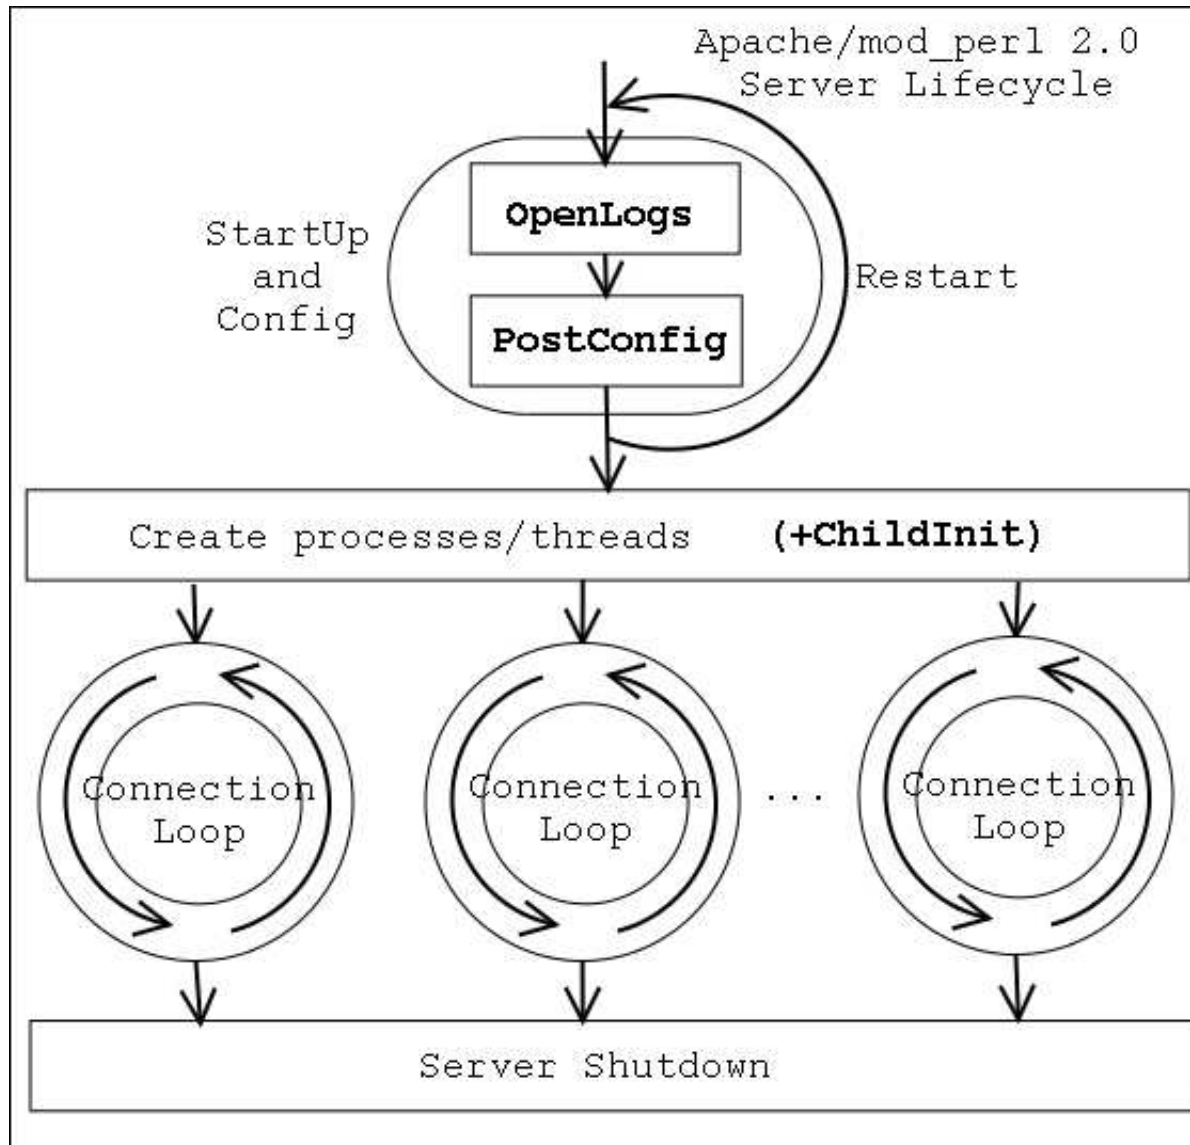
- Executed in the order they have been registered until the first handler that returns something other than `Apache::DECLINED`.
- If the return value is `Apache::DECLINED`, the next handler in the chain will be run.
- If the return value is `Apache::OK` the next phase will start.
- In all other cases the execution will be aborted.

3 Server Life Cycle Handlers

3.1 About

- Server Life Cycle
- Startup Phases Demonstration Module
- PerlOpenLogsHandler
- PerlPostConfigHandler
- PerlChildInitHandler

3.2 Apache 2.0 Server Life Cycle



1. Parse the configuration file, `open_logs`, `post_config`
2. Restart to test graceful restarts
3. Parse the configuration file, `open_logs`, `post_config`
4. Spawn workers: `procs`, `threads`, `mix`
 - Run `PerlChildInit` for each spawned process
5. Process requests over connections
6. Shutdown

3.2.1 Startup Phases Demonstration Module

```
package MyApache::StartupLog;
```

```
use strict;
```

```
use warnings;
```

```
use Apache::Log ();
```

```
use Apache::ServerUtil ();
```

```
use File::Spec::Functions;
```

```
use Apache::Const -compile => 'OK';

my $log_file = catfile "logs", "startup_log";
my $log_fh;

sub open_logs {
    my($conf_pool, $log_pool, $temp_pool, $s) = @_;
    my $log_path = Apache::server_root_relative($conf_pool, $log_file);

    $s->warn("opening the log file: $log_path");
    open $log_fh, ">>$log_path" or die "can't open $log_path: $!";
    my $oldfh = select($log_fh); $| = 1; select($oldfh);

    say("process $$ is born to reproduce");
    return Apache::OK;
}

sub post_config {
    my($conf_pool, $log_pool, $temp_pool, $s) = @_;
    say("configuration is completed");
    return Apache::OK;
}

sub child_init {
```

```
my($child_pool, $s) = @_;  
say("process $$ is born to serve");  
return Apache::OK;  
}  
  
sub say {  
    my($caller) = (caller(1))[3] =~ /([^\:]+)$/;  
    printf $log_fh "[%s] - %-11s: %s\n", scalar(localtime), $caller, $_[0];  
}  
  
END {  
    say("process $$ is shutdown\n");  
}  
  
1;
```

- And the *httpd.conf* configuration section:

```
PerlModule                MyApache::StartupLog
PerlOpenLogsHandler      MyApache::StartupLog::open_logs
PerlPostConfigHandler    MyApache::StartupLog::post_config
PerlChildInitHandler     MyApache::StartupLog::child_init
```

```
% bin/apachectl start && bin/apachectl stop
```

- Checking *logs/startup_log*:

```
[Thu Aug 22 15:57:08 2002] - open_logs : process 21823 is born to reproduce  
[Thu Aug 22 15:57:08 2002] - post_config: configuration is completed  
[Thu Aug 22 15:57:09 2002] - END : process 21823 is shutdown
```

```
[Thu Aug 22 15:57:10 2002] - open_logs : process 21825 is born to reproduce  
[Thu Aug 22 15:57:10 2002] - post_config: configuration is completed  
[Thu Aug 22 15:57:11 2002] - child_init : process 21830 is born to serve  
[Thu Aug 22 15:57:11 2002] - child_init : process 21831 is born to serve  
[Thu Aug 22 15:57:11 2002] - child_init : process 21832 is born to serve  
[Thu Aug 22 15:57:11 2002] - child_init : process 21833 is born to serve  
[Thu Aug 22 15:57:12 2002] - END : process 21825 is shutdown
```

- Apache restarts itself
- Only the parent process runs the `END` block on shutdown

pre_config phase:

- There is also the *pre_config* phase, which is executed before the configuration files are parsed
- ... but this is of no use to `mod_perl`, because `mod_perl` is loaded only during the configuration phase.

3.2.2 *PerlOpenLogsHandler*

- The *open_logs* phase happens just before the *post_config* phase.
- `PerlOpenLogsHandler` handlers used for opening module-specific log files.
- At this stage `STDERR` goes to the console, not *error_log*
- This phase is of type `RUN_ALL`.
- The handler's configuration scope is `SRV`.

The *open_logs* handler from our example:

```
sub open_logs {
    my($conf_pool, $log_pool, $temp_pool, $s) = @_;
    my $log_path = Apache::server_root_relative($conf_pool, $log_file);

    $s->warn("opening the log file: $log_path");
    open $log_fh, ">>$log_path" or die "can't open $log_path: $!";
    my $oldfh = select($log_fh); $| = 1; select($oldfh);

    say("process $$ is born to reproduce");
    return Apache::OK;
}
```

- The *open_logs* phase handlers accept four arguments:
 - the configuration pool,
 - the logging streams pool,
 - the temporary pool
 - and the server object
- The configuration:

PerlOpenLogsHandler MyApache::StartupLog::open_logs

3.2.3 *PerlPostConfigHandler*

- The *post_config* phase happens right after Apache has processed the configuration files, before any child processes were spawned.
- used for initializing things to be shared between all child processes.
- You can do the same in the startup file, but in the *post_config* phase you have an access to a complete configuration tree.
- This phase is of type `RUN_ALL`.
- The handler's configuration scope is `SRV`.

- The handler from the example:

```
sub post_config {  
    my($conf_pool, $log_pool, $temp_pool, $s) = @_;  
    say("configuration is completed");  
    return Apache::OK;  
}
```

- Args are identical to the *open_logs* phase's handler.
- Configuration:

```
PerlOpenLogsHandler MyApache::StartupLog::post_config
```

3.2.4 *PerlChildInitHandler*

- The *child_init* phase happens immediately after the child process is spawned.
- Each child process (not a thread!) runs the hooks of this phase only once in their life-time.
- In the prefork MPM this phase is useful for initializing any data structures which should be private to each process. e.g.:
 - `Apache::DBI` pre-opens database connections during this phase
 - `Apache::Resource` sets the process' resources limits.

- This phase is of type `VOID`.
- The handler's configuration scope is `SRV`.

- The handler from the example:

```
sub child_init {  
    my($child_pool, $s) = @_;  
    say("process $$ is born to serve");  
    return Apache::OK;  
}
```

- The *child_init()* handler accepts two arguments:
 - the child process pool
 - and the server object.

- Configuration:

```
PerlOpenLogsHandler MyApache::StartupLog::child_init
```


4 Protocol Handlers

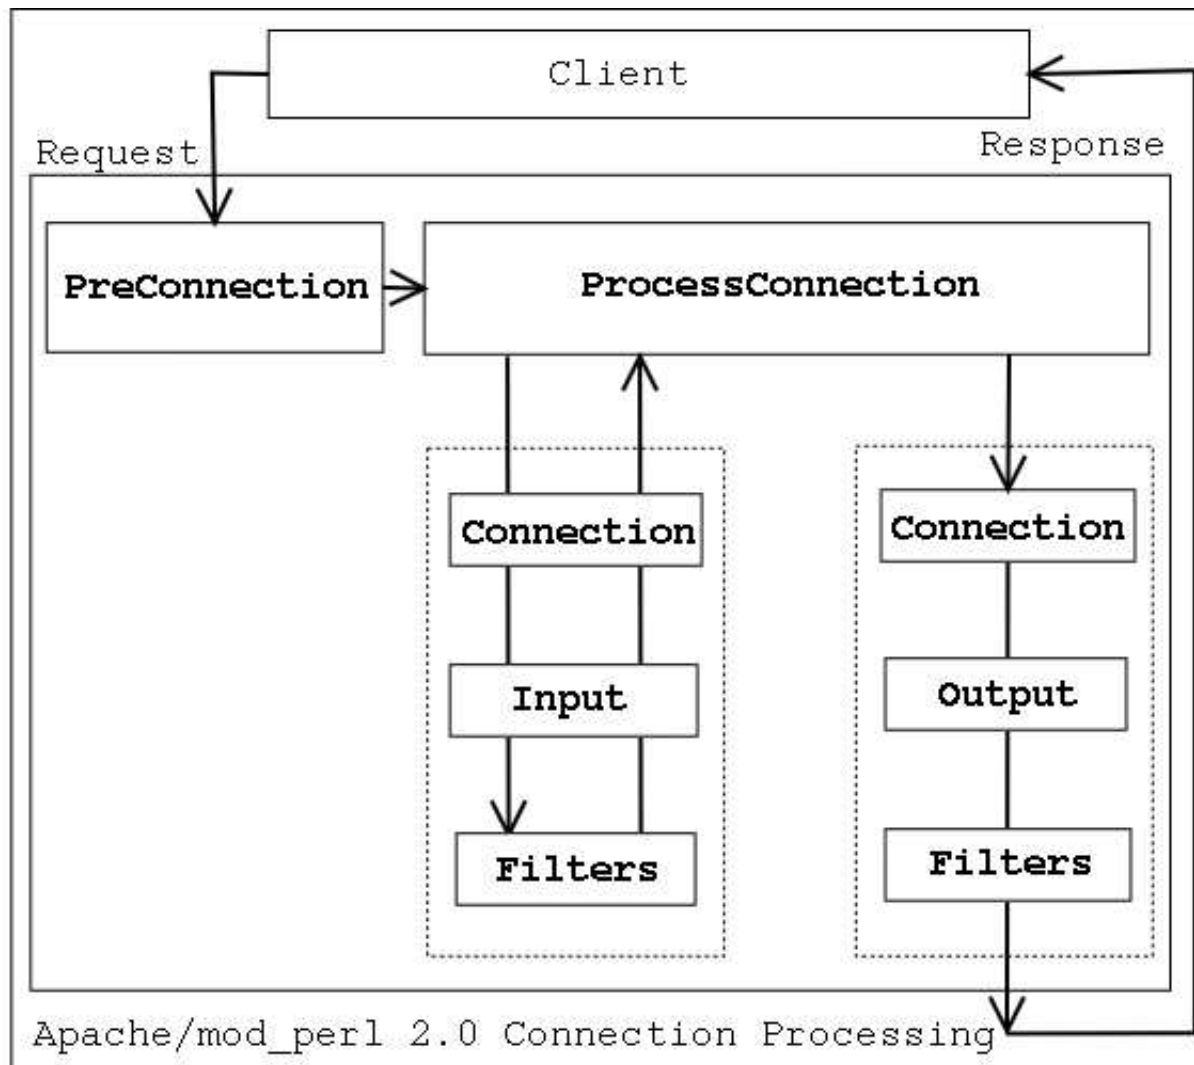
4.1 About

- Connection Cycle Phases
- PerlPreConnectionHandler
- PerlProcessConnectionHandler
- Socket-based Protocol Module
- Bucket Brigades-based Protocol Module

4.2 Connection Cycle Phases

- Each child server is engaged in processing connections.
- Each connection may be served by different connection protocols,
 - e.g., HTTP, POP3, SMTP, etc.
- Each connection may include more than one request,
 - e.g., several HTTP requests can be served over a single connection, when a response includes several images.

Connection Life Cycle diagram:



4.2.1 *PerlPreConnectionHandler*

- The *pre_connection* phase happens just after the server accepts the connection, but before it is handed off to a protocol module to be served.
- It gives modules an opportunity to modify the connection as soon as possible and insert filters if needed.
- The core server uses this phase to setup the connection record based on the type of connection that is being used.
- `mod_perl` itself uses this phase to register the connection input and output filters.

Apache::Reload:

- In mod_perl 1.0 `Apache::Reload` was used to automatically reload modified since the last request Perl modules.
- It was invoked during *post_read_request*, the first HTTP request's phase.
- In mod_perl 2.0 *pre_connection* is the earliest general phase
- So now we can invoke the `Apache::Reload` handler during the *pre_connection* phase if the interpreter's scope is:

PerlInterpScope connection

- Though this is not good for production server,
- where there are several requests coming on the same connection
- and only one handled by mod_perl
- and the others by the default images handler
- the Perl interpreter won't be available to other threads while the images are being served.

- This phase is of type `RUN_ALL`.
- The handler's configuration scope is `SRV`, because it's not known yet which resource the request will be mapped to.
- **META:** As of this writing `PerlPreConnectionHandler` is not being executed by `mod_perl`. Stay tuned.

Skeleton:

```
sub handler {  
    my ($c, $socket) = @_;  
    # ...  
    return Apache::OK;  
}
```

- A *pre_connection* handler accepts two arguments:
 - connection record
 - and socket objects

4.2.2 *PerlProcessConnectionHandler*

- The *process_connection* phase is used to process incoming connections.
- Only protocol modules should assign handlers for this phase, as it gives them an opportunity to replace the standard HTTP processing with processing for some other protocols (e.g., POP3, FTP, etc.).
- This phase is of type `RUN_FIRST`.
- The handler's configuration scope is `SRV`.
 - Therefore the only way to run protocol servers different than the core HTTP is inside dedicated virtual hosts.

Skeleton:

```
sub handler {  
    my ($c) = @_;  
    my $socket = $c->client_socket;  
    # ...  
    return Apache::OK;  
}
```

- A *process_connection* handler accepts a connection record object as its only argument,
- a socket object can be retrieved from the connection record object.

- Let's look at the following two examples of connection handlers:
 1. Using the connection socket to read and write the data.
 2. Using bucket brigades to accomplish the same and allow for connection filters to do their work.

4.2.2.1 Socket-based Protocol Module

- The `MyApache::EchoSocket` module simply echoes the data read back to the client.
- This module's implementation works directly with the connection socket and therefore bypasses connection filters if any.

Configuration:

```
Listen 8010
<VirtualHost _default_:8010>
    PerlModule                               MyApache::EchoSocket
    PerlProcessConnectionHandler MyApache::EchoSocket
</VirtualHost>
```

- use the `Listen` and `<VirtualHost>` directives to bind to the non-standard port **8010**:

Demo:

```
panic% httpd
panic% telnet localhost 8010
Trying 127.0.0.1...
Connected to localhost (127.0.0.1).
Escape character is '^]'.
Hello
Hello

fOo BaR
fOo BaR

Connection closed by foreign host.
```

The code:

```
package MyApache::EchoSocket;

use strict;
use warnings FATAL => 'all';

use Apache::Connection ();
use APR::Socket ();

use Apache::Const -compile => 'OK';

use constant BUFF_LEN => 1024;

sub handler {
    my $c = shift;
    my $socket = $c->client_socket;
```

```
my $buff;
while (1) {
    my $rlen = BUFF_LEN;
    $socket->recv($buff, $rlen);

    last if $rlen <= 0 or $buff =~ /^[\r\n]+$/;

    my $wlen = $rlen;
    $socket->send($buff, $wlen);

    last if $wlen != $rlen;
}

Apache::OK;
}
1;
```

4.2.2.2 Bucket Brigades-based Protocol Module

- The same module, but this time implemented by manipulating bucket brigades,
- and which runs its output through a connection output filter that turns all uppercase characters into their lowercase equivalents.

Configuration:

```
Listen 8011
<VirtualHost _default_:8011>
    PerlModule MyApache::EchoBB
    PerlProcessConnectionHandler MyApache::EchoBB
    PerlOutputFilterHandler MyApache::EchoBB::lowercase_filter
</VirtualHost>
```

Demo:

```
panic% httpd
panic% telnet localhost 8011
Trying 127.0.0.1...
Connected to localhost (127.0.0.1).
Escape character is '^]'.
Hello
hello

fOo BaR
foo bar
```

- As you can see the response now was all in lower case, because of the output filter.

The code:

```
package MyApache::EchoBB;

use strict;
use warnings FATAL => 'all';

use Apache::Connection ();
use APR::Bucket ();
use APR::Brigade ();
use APR::Util ();

use APR::Const -compile => qw(SUCCESS EOF);
use Apache::Const -compile => qw(OK MODE_GETLINE);
```

```
sub handler {
    my $c = shift;

    my $bb_in  = APR::Brigade->new($c->pool, $c->bucket_alloc);
    my $bb_out = APR::Brigade->new($c->pool, $c->bucket_alloc);
    my $last = 0;

    while (1) {
        my $rv = $c->input_filters->get_brigade($bb_in,
                                                Apache::MODE_GETLINE);

        if ($rv != APR::SUCCESS or $bb_in->empty) {
            my $error = APR::strerror($rv);
            unless ($rv == APR::EOF) {
                warn "get_brigade: $error\n";
            }
            $bb_in->destroy;
            last;
        }
    }
}
```

```
while (!$bb_in->empty) {
    my $bucket = $bb_in->first;

    $bucket->remove;

    if ($bucket->is_eos) {
        $bb_out->insert_tail($bucket);
        last;
    }

    my $data;
    my $status = $bucket->read($data);
    return $status unless $status == APR::SUCCESS;

    if ($data) {
        $last++ if $data =~ /^[\r\n]+$/;
        # could do something with the data here
        $bucket = APR::Bucket->new($data);
    }
}
```

```
    }  
  
    $bb_out->insert_tail($bucket);  
}  
  
my $b = APR::Bucket::flush_create($c->bucket_alloc);  
$bb_out->insert_tail($b);  
$c->output_filters->pass_brigade($bb_out);  
last if $last;  
}  
  
Apache::OK;  
}
```

```
use base qw(Apache::Filter);
use constant BUFF_LEN => 1024;

sub lowercase_filter : FilterConnectionHandler {
    my $filter = shift;

    while ($filter->read(my $buffer, BUFF_LEN)) {
        $filter->print(lc $buffer);
    }

    return Apache::OK;
}

1;
```

- A simplified pseufo-code handler:

```
while ($bb_in = get_brigade()) {
    while ($bucket_in = $bb_in->get_bucket()) {
        my $data = $bucket_in->read();
        # do something with $data
        $bucket_out = new_bucket($data);

        $bb_out->insert_tail($bucket_out);
    }
    $bb_out->insert_tail($flush_bucket);
    pass_brigade($bb_out);
}
```

- This handler could be much simpler, since we don't modify the data.
- We could simply pass the whole brigade unmodified without even looking at the buckets.
- But from this example you can see how to write a connection handler where you actually want to read and/or modify the data.
- We will discuss the filters next

5 Input and Output Filters

5.1 About

- I/O Filtering
- PerlInputFilterHandler
- PerlOutputFilterHandler
- Connection vs. HTTP Request Filters
- All-in-One Filter

- Input Filters
 - Connection Input Filters
 - HTTP Request Input Filters
- Output Filters
 - Connection Output Filters
 - HTTP Request Output Filters
 - Stream-oriented Output Filter
 - Bucket Brigade-based Output Filters

5.2 I/O Filtering

- Apache 2.0 considers all incoming and outgoing data as chunks of information,
- disregarding their kind and source or storage methods.
- These data chunks are stored in *buckets*, which form *bucket brigades*.
- Both input and output filters, filter the data in bucket brigades.

5.2.1 PerlInputFilterHandler

- The `PerlInputFilterHandler` handler registers a filter for input filtering.
- This handler is of type `VOID`.
- The handler's configuration scope is `DIR`

5.2.2 *PerlOutputFilterHandler*

- The `PerlOutputFilterHandler` handler registers and configures output filters.
- This handler is of type `VOID`.
- The handler's configuration scope is `DIR`

5.2.3 Connection vs. HTTP Request Filters

- `mod_perl` can do connection and HTTP request filtering.
- (Apache supports several other types)
- method attributes set the filter type:

```
sub handler : FilterRequestHandler    { ... }  
sub handler : FilterConnectionHandler { ... }
```

- HTTP Request filter handlers are declared using the `FilterRequestHandler` attribute.
- HTTP request input and output filters skeleton:

```
package MyApache::FilterRequestFoo;
use base qw(Apache::Filter);

sub input : FilterRequestHandler {
    my($filter, $bb, $mode, $block, $readbytes) = @_;
    #...
}

sub output : FilterRequestHandler {
    my($filter, $bb) = @_;
    #...
}
1;
```

- If the attribute is not specified, the default `FilterRequestHandler` attribute is assumed.
- Filters specifying subroutine attributes must subclass `Apache::Filter`, others only need to:

```
use Apache::Filter ();
```

- Request filters are usually configured in the `<Location>` or equivalent sections:

```
PerlModule MyApache::FilterRequestFoo
PerlModule MyApache::NiceResponse
<Location /filter_foo>
    SetHandler modperl
    PerlResponseHandler      MyApache::NiceResponse
    PerlInputFilterHandler   MyApache::FilterRequestFoo::input
    PerlOutputFilterHandler  MyApache::FilterRequestFoo::output
</Location>
```

- Connection filter handlers are declared using the `FilterConnectionHandler` attribute.
- Connection input and output filters skeleton:

```
package MyApache::FilterConnectionBar;
use base qw(Apache::Filter);

sub input : FilterConnectionHandler {
    my($filter, $bb, $mode, $block, $readbytes) = @_;
    #...
}

sub output : FilterConnectionHandler {
    my($filter, $bb) = @_;
    #...
}
1;
```

- This time the configuration must be done outside the `<Location>` or equivalent sections,
- usually within the `<VirtualHost>` or the global server configuration:

```
Listen 8005
```

```
<VirtualHost _default_:8005>
```

```
    PerlModule MyApache::FilterConnectionBar
```

```
    PerlModule MyApache::NiceResponse
```

```
    PerlInputFilterHandler MyApache::FilterConnectionBar::input
```

```
    PerlOutputFilterHandler MyApache::FilterConnectionBar::output
```

```
<Location />
```

```
    SetHandler modperl
```

```
    PerlResponseHandler MyApache::NiceResponse
```

```
</Location>
```

```
</VirtualHost>
```

- For HTTP requests the only difference between connection filters and request filters is that the former see everything: the headers and the body, whereas the latter see only the body.
- mod_perl provides two interfaces to filtering:
 - a direct mapping to buckets and bucket brigades
 - and a simpler, stream-oriented interface
- META: as of this writing the latter is available only for the output filtering
- Following examples will explain the difference

5.3 All-in-One Filter

- The `MyApache::FilterSnoop` handler silently snoops on the data that goes through request and connection filters, in the input and output modes.
- First let's develop a simple response handler that simply dumps the request's *args* and *content* as strings:

```
package MyApache::Dump;  
  
use strict;  
use warnings;  
  
use Apache::RequestRec ();  
use Apache::RequestIO ();  
use APR::Table ();
```

```
use Apache::Const -compile => qw(OK M_POST);

sub handler {
    my $r = shift;
    $r->content_type('text/plain');

    $r->print("args:\n", $r->args, "\n");

    if ($r->method_number == Apache::M_POST) {
        my $data = content($r);
        $r->print("content:\n$data\n");
    }

    return Apache::OK;
}
```

```
sub content {
    my $r = shift;

    $r->setup_client_block;

    return '' unless $r->should_client_block;

    my $len = $r->headers_in->get('content-length');
    my $buf;
    $r->get_client_block($buf, $len);

    return $buf;
}

1;
```

- which is configured as:

```
PerlModule MyApache::Dump
<Location /dump>
    SetHandler modperl
    PerlResponseHandler MyApache::Dump
</Location>
```

- If we issue the following request:

```
% echo "mod_perl rules" | POST 'http://localhost:8002/dump?foo=1&bar=2'
```

- the response will be:

args:

foo=1&bar=2

content:

mod_perl rules

- As you can see it simply dumped the query string and the posted data.

- Now let's write the snooping filter:

```
package MyApache::FilterSnoop;
```

```
use strict;  
use warnings;
```

```
use base qw(Apache::Filter);  
use Apache::FilterRec ();  
use APR::Brigade ();  
use APR::Bucket ();
```

```
use Apache::Const -compile => qw(OK DECLINED);  
use APR::Const -compile => ':common';
```

```
sub connection : FilterConnectionHandler { snoop("connection", @_) }  
sub request    : FilterRequestHandler   { snoop("request",    @_) }
```

```

sub snoop {
    my $type = shift;
    my($filter, $bb, $mode, $block, $readbytes) = @_;

    # $mode, $block, $readbytes are passed only for input filters
    my $stream = defined $mode ? "input" : "output";

    # read the data and pass-through the bucket brigades unchanged
    my $ra_data = '';
    if (defined $mode) {
        # input filter
        my $rv = $filter->next->get_brigade($bb, $mode,
                                           $block, $readbytes);

        return $rv unless $rv == APR::SUCCESS;
        $ra_data = bb_sniff($bb);
    }
    else {
        # output filter
        $ra_data = bb_sniff($bb);
        my $rv = $filter->next->pass_brigade($bb);
        return $rv unless $rv == APR::SUCCESS;
    }
}

```

```
# send the sniffed info to STDERR so not to interfere with
# normal output
my $direction = $stream eq 'output' ? ">>>" : "<<<";
print STDERR "\n$direction $type $stream filter\n";
my $c = 1;
while (my($btype, $data) = splice @$ra_data, 0, 2) {
    print STDERR "    o bucket $c: $btype\n";
    print STDERR "[$data]\n";
    $c++;
}

return Apache::OK;
}
```

```
sub bb_sniff {  
  my $bb = shift;  
  my @data;  
  for (my $b = $bb->first; $b; $b = $bb->next($b)) {  
    $b->read(my $bdata);  
    $bdata = '' unless defined $bdata;  
    push @data, $b->type->name, $bdata;  
  }  
  return \@data;  
}  
  
1;
```

- We want to use this somewhat complicated filter to visualize how various kind of filters work.
- Since this module combines several kind of filters it's the best to try to understand its implementation after you understand standalone filters' implementation.
- At this stage what's important is to see it at work.

- Let's snoop on connection and request filter levels in both directions by applying the following configuration:

```
Listen 8008
```

```
<VirtualHost _default_:8008>
```

```
    PerlModule MyApache::FilterSnoop
```

```
    PerlModule MyApache::Dump
```

```
    # Connection filters
```

```
    PerlInputFilterHandler    MyApache::FilterSnoop::connection
```

```
    PerlOutputFilterHandler  MyApache::FilterSnoop::connection
```

```
    <Location /dump>
```

```
        SetHandler modperl
```

```
        PerlResponseHandler MyApache::Dump
```

```
        # Request filters
```

```
        PerlInputFilterHandler    MyApache::FilterSnoop::request
```

```
        PerlOutputFilterHandler  MyApache::FilterSnoop::request
```

```
    </Location>
```

```
</VirtualHost>
```

- If we issue the following request:

```
% echo "mod_perl rules" | POST 'http://localhost:8008/dump?foo=1&bar=2'
```

- the response doesn't change our snooping filter
- though we can see all the diagnostics in *error_log*

- First we can see the connection input filter at work, as it processes the HTTP headers.
- We can see that for this request each header is put into a separate brigade with a single bucket.
- The data is conveniently enclosed by [] so you can see the new line characters as well.

```
<<< connection input filter
  o bucket 1: HEAP
[POST /dump?foo=1&bar=2 HTTP/1.1
]
```

```
<<< connection input filter
  o bucket 1: HEAP
[TE: deflate,gzip;q=0.3
]
```

```
<<< connection input filter
```

```
o bucket 1: HEAP
```

```
[Connection: TE, close
```

```
]
```

```
<<< connection input filter
```

```
o bucket 1: HEAP
```

```
[Host: localhost:8008
```

```
]
```

```
<<< connection input filter
```

```
o bucket 1: HEAP
```

```
[User-Agent: lwp-request/2.01
```

```
]
```

```
<<< connection input filter
```

```
o bucket 1: HEAP
```

```
[Content-Length: 14  
]
```

```
<<< connection input filter  
  o bucket 1: HEAP
```

```
[Content-Type: application/x-www-form-urlencoded  
]
```

```
<<< connection input filter  
  o bucket 1: HEAP
```

```
[  
]
```

- Here the HTTP header has been terminated by a double new line.

- So far all the buckets were of the *HEAP* type, meaning that they were allocated from the heap memory.
- Notice that the HTTP request input filters will never see the bucket brigade with HTTP headers, as it has been consumed by the last core connection filter.

- The following two entries are generated when `MyApache::Dump::handler` reads the POSTed content:

```
<<< connection input filter
  o bucket 1: HEAP
[mod_perl rules]
```

```
<<< request input filter
  o bucket 1: HEAP
[mod_perl rules]
  o bucket 2: EOS
[ ]
```

- The connection input filter is run before the request input filter.

- Both filters see the same data, since they don't modify it
- The bucket of type *EOS* indicates the end of stream.

- Next we can see that `MyApache::Dump::handler` has generated its response.
- However only the request output filter is filtering it at this point:

```
>>> request output filter
      o bucket 1: TRANSIENT
[args:
foo=1&bar=2
content:
mod_perl rules
]
```

- This happens because Apache hasn't sent yet the response HTTP headers to the client.

- Apache postpones the header sending so it can calculate and set the `Content-Length` header.
- This *TRANSIENT* bucket is allocated from the stack memory

- When the content handler returns Apache sends the HTTP headers through connection output filters:

```
>>> connection output filter
  o bucket 1: HEAP
[HTTP/1.1 200 OK
Date: Wed, 14 Aug 2002 07:31:53 GMT
Server: Apache/2.0.41-dev (Unix) mod_perl/1.99_05-dev
Perl/v5.8.0 mod_ssl/2.0.41-dev OpenSSL/0.9.6d DAV/2
Content-Length: 42
Connection: close
Content-Type: text/plain; charset=ISO-8859-1

]
```

- Now the response body in the bucket of type *HEAP* is passed through the connection output filter,
- followed by the *EOS* bucket to mark the end of the request:

```
>>> connection output filter
      o bucket 1: HEAP
[args:
foo=1&bar=2
content:
mod_perl rules
]
      o bucket 2: EOS
[]
```

- Finally the output is flushed,
- to make sure that any buffered output is sent to the client:

```
>>> connection output filter
      o bucket 1: FLUSH
[]
```

- This module helps to understand that each filter handler can be called many time during each request and connection.
- It's called for each bucket brigade.
- the HTTP request input filter is called only if there is some POSTed data to read
- If you run the same request without POSTing any data or simply running a GET request, the request input filter won't be called.

5.4 Input Filters

- mod_perl supports 2 kinds of input filters:
 1. Connection input filters
 2. HTTP Request input filters

5.4.1 Connection Input Filters

- Let's write a poor man's s/GET/HEAD/ rewrite handler
- The handler looks for the data like:

```
GET /perl/test.pl HTTP/1.1
```

- and turns it into:

```
HEAD /perl/test.pl HTTP/1.1
```

```
package MyApache::InputFilterGET2HEAD;

use strict;
use warnings;

use base qw(Apache::Filter);

use Apache::Connection ();
use Apache::ServerUtil ();
use APR::Brigade ();
use APR::Bucket ();

use Apache::Const -compile => 'OK';
use APR::Const -compile => ':common';
```

```
sub handler : FilterConnectionHandler {
  my($filter, $bb, $mode, $block, $readbytes) = @_;

  my $c = $filter->c;
  my $ctx_bb = APR::Brigade->new($c->pool, $c->bucket_alloc);
  my $rv = $filter->next->get_brigade($ctx_bb, $mode,
                                     $block, $readbytes);
  return $rv unless $rv == APR::SUCCESS;

  while (!$ctx_bb->empty) {
    my $bucket = $ctx_bb->first;

    $bucket->remove;

    if ($bucket->is_eos) {
      $bb->insert_tail($bucket);
      last;
    }

    my $data;
    my $status = $bucket->read($data);
    return $status unless $status == APR::SUCCESS;
  }
}
```

```
    if ($data and $data =~ s|^GET|HEAD|) {  
        $bucket = APR::Bucket->new($data);  
    }  
  
    $bb->insert_tail($bucket);  
}  
  
Apache::OK;  
}  
1;
```

- For example, consider the following response handler:

```
package MyApache::RequestType;

use strict;
use warnings;

use Apache::RequestIO ();
use Apache::RequestRec ();
use Apache::Response ();

use Apache::Const -compile => 'OK';

sub handler {
    my $r = shift;

    $r->content_type('text/plain');
    my $response = "the request type was " . $r->method;
    $r->set_content_length(length $response);
    $r->print($response);
}
```

```
    Apache::OK;  
}  
  
1;
```

- which returns to the client the request type it has issued.
- In the case of the HEAD request Apache will discard the response body, but it'll still set the correct Content-Length header, which will be 24 in case of the GET request and 25 for HEAD.

- Therefore if this response handler is configured as:

```
Listen 8005
<VirtualHost _default_:8005>
  <Location />
    SetHandler modperl
    PerlResponseHandler +MyApache::RequestType
  </Location>
</VirtualHost>
```

and a GET request is issued to /:

```
panic% perl -MLWP::UserAgent -le \
'$r = LWP::UserAgent->new()->get("http://localhost:8005/"); \
print $r->headers->content_length . ": ". $r->content'
24: the request type was GET
```

- where the response's body is:
the request type was GET
- And the Content-Length header is set to 24.

- However if we enable the `MyApache::InputFilterGET2HEAD` input connection filter:

```
Listen 8005
```

```
<VirtualHost _default_:8005>
```

```
    PerlInputFilterHandler +MyApache::InputFilterGET2HEAD
```

```
    <Location />
```

```
        SetHandler modperl
```

```
        PerlResponseHandler +MyApache::RequestType
```

```
    </Location>
```

```
</VirtualHost>
```

- And issue the same `GET` request, we get only:

25:

- which means that the body was discarded by Apache, because our filter turned the GET request into a HEAD request.
- If Apache wasn't discarding the body on HEAD, the response would be:

the request type was HEAD

- that's why the content length is reported as 25 and not 24 as in the real GET request.

5.4.2 HTTP Request Input Filters

- Request filters are similar to connection filters, but have an access to a request object and they don't see the headers.

- Here is the request input filter that lowercases the request's body `MyApache::InputRequestFilterLC`:

```
package MyApache::InputRequestFilterLC;
```

```
use strict;
```

```
use warnings;
```

```
use base qw(Apache::Filter);
```

```
use APR::Brigade ();
```

```
use APR::Bucket ();
```

```
use Apache::Const -compile => 'OK';
```

```
use APR::Const -compile => ':common';
```

```
sub handler : FilterRequestHandler {
  my($filter, $bb, $mode, $block, $readbytes) = @_;

  my $c = $filter->c;
  my $bb_ctx = APR::Brigade->new($c->pool, $c->bucket_alloc);
  my $rv = $filter->next->get_brigade($bb_ctx, $mode,
                                     $block, $readbytes);
  return $rv unless $rv == APR::SUCCESS;

  while (!$bb_ctx->empty) {
    my $b = $bb_ctx->first;

    $b->remove;

    if ($b->is_eos) {
      $bb->insert_tail($b);
      last;
    }
  }
}
```

```
    my $data;  
    my $status = $b->read($data);  
    return $status unless $status == APR::SUCCESS;  
  
    $b = APR::Bucket->new(lc $data) if $data;  
  
    $bb->insert_tail($b);  
}  
  
Apache::OK;  
}  
  
1;
```

- the `MyApache::Dump` response handler dumps the query string and the content body as a response
- Using the following configuration:

```
<Location /lc_input>  
    SetHandler modperl  
    PerlResponseHandler      +MyApache::Dump  
    PerlInputFilterHandler   +MyApache::InputRequestFilterLC  
</Location>
```

- When issuing a POST request:

```
% echo "mOd_pEr1 RuLeS" | POST 'http://localhost:8002/lc_input?FoO=1&BAR=2'
```

- we get a response:

args:

FoO=1&BAR=2

content:

mod_perl rules

- Indeed we can see that our filter has lowercased the POSTed body, before the content handler received it.
- The query string didn't change, since the filter didn't operate on headers.

5.5 Output Filters

- mod_perl supports 2 kinds of output filters:
 1. Connection input filters
 2. HTTP Request input filters

5.5.1 Connection Output Filters

- HTTP request input filters don't see the outgoing headers.
- Whereas connection input filters see **all** the outgoing data
- Connection and HTTP request input filters implementation is similar.
- We concentrate on HTTP request filters.

5.5.2 *HTTP Request Output Filters*

- Output filters can be written using the bucket brigades manipulation or the simplified stream-oriented interface.
- Here is a response handler that send two lines of output in a single string:

1. numerals: 1234567890

2. the English alphabet

```
package MyApache::SendAlphaNum;
```

```
use strict;
```

```
use warnings;
```

```
use Apache::RequestRec ();
use Apache::RequestIO ();

use Apache::Const -compile => qw(OK);

sub handler {
    my $r = shift;

    $r->content_type('text/plain');

    $r->print(1..9, "0\n");
    $r->print('a'..'z', "\n");

    Apache::OK;
}
1;
```

- The purpose of our request output filter is to reverse every line of the response, preserving the new line characters in their places.
- Since we want to reverse characters only in the response body we will use the HTTP request output filter.

5.5.2.1 Stream-oriented Output Filter

- The stream-oriented filtering API implementation:

```
package MyApache::FilterReverse1;

use strict;
use warnings;

use Apache::Filter ();

use Apache::Const -compile => qw(OK);

use constant BUFF_LEN => 1024;
```

```
sub handler : FilterRequestHandler {
  my $filter = shift;

  while ($filter->read(my $buffer, BUFF_LEN)) {
    for (split "\n", $buffer) {
      $filter->print(scalar reverse $_);
      $filter->print("\n");
    }
  }

  Apache::OK;
}
1;
```

- Add the following configuration to *httpd.conf*:

```
PerlModule MyApache::FilterReverse1
PerlModule MyApache::SendAlphaNum
<Location /reverse1>
    SetHandler modperl
    PerlResponseHandler      MyApache::SendAlphaNum
    PerlOutputFilterHandler  MyApache::FilterReverse1
</Location>
```

- Now when a request to `/reverse1` is made, the response handler `MyApache::SendAlphaNum::handler()` sends:

1234567890

abcdefghijklmnopqrstuvwxy

- as a response and the output filter handler `MyApache::FilterReverse1::handler` reverses the lines, so the client gets:

0987654321

zyxwvutsrqponmlkjihgfedcba

5.5.2.2 Bucket Brigade-based Output Filters

- The bucket brigades filtering API implementation:

```
package MyApache::FilterReverse2;

use strict;
use warnings;

use Apache::Filter;

use APR::Brigade ();
use APR::Bucket ();

use Apache::Const -compile => 'OK';
use APR::Const -compile => ':common';
```

```
sub handler : FilterRequestHandler {
  my($filter, $bb) = @_;

  my $c = $filter->c;
  my $bb_ctx = APR::Brigade->new($c->pool, $c->bucket_alloc);

  while (!$bb->empty) {
    my $bucket = $bb->first;
    $bucket->remove;

    if ($bucket->is_eos) {
      $bb_ctx->insert_tail($bucket);
      last;
    }

    my $data;
    my $status = $bucket->read($data);
    return $status unless $status == APR::SUCCESS;

    if ($data) {
      $data = join "",
        map {scalar(reverse $_), "\n"} split "\n", $data;
      $bucket = APR::Bucket->new($data);
    }
  }
}
```

```
        $bb_ctx->insert_tail($bucket);
    }

    my $rv = $filter->next->pass_brigade($bb_ctx);
    return $rv unless $rv == APR::SUCCESS;

    Apache::OK;
}
1;
```

- And the corresponding configuration:

```
PerlModule MyApache::FilterReverse2
PerlModule MyApache::SendAlphaNum
<Location /reverse2>
    SetHandler modperl
    PerlResponseHandler      MyApache::SendAlphaNum
    PerlOutputFilterHandler  MyApache::FilterReverse2
</Location>
```

- Now when a request to `/reverse2` is made, the client gets:

0987654321

zyxwvutsrqponmlkjihgfedcba

- as expected.

6 HTTP Handlers

6.1 About


- HTTP Request Cycle Phases
- PerlPostReadRequestHandler
- PerlTransHandler
- PerlHeaderParserHandler
- PerlInitHandler
- PerlAccessHandler
- PerlAuthenHandler

- PerlAuthzHandler
- PerlTypeHandler
- PerlFixupHandler
- PerlResponseHandler
- PerlLogHandler
- PerlCleanupHandler

6.2 HTTP Request Cycle Phases

- Those familiar with `mod_perl 1.0` will find the HTTP request cycle in `mod_perl 2.0` to be almost identical to the `mod_perl 1.0`'s model.
- The only difference is in the *response* phase which now includes filtering.
- Also the `PerlHandler` directive has been renamed to `PerlResponseHandler` to better match the corresponding Apache phase name (*response*).
- Here is the HTTP request life cycle in `mod_perl 2.0`:

Apache/mod_perl 2.0
HTTP Request
Lifecycle



- From the diagram it can be seen that an HTTP request is processed by 11 phases, executed in the following order:
 1. PerlPostReadRequestHandler (PerlInitHandler)
 2. PerlTransHandler
 3. PerlHeaderParserHandler (PerlInitHandler)
 4. PerlAccessHandler
 5. PerlAuthenHandler
 6. PerlAuthzHandler
 7. PerlTypeHandler

8. PerlFixupHandler
9. PerlResponseHandler
10. PerlLogHandler
11. PerlCleanupHandler

6.2.1 *PerlPostReadRequestHandler*

- The *post_read_request* phase is the first request phase and happens immediately after the request has been read and HTTP headers were parsed.
- This phase is usually used to do processing that must happen once per request.
- For example `Apache::Reload` is usually invoked at this phase to reload modified Perl modules.
- This phase is of type `RUN_ALL`.
- The handler's configuration scope is `SRV`

- The following `ModPerl::Registry` script prints when the last time *httpd.conf* has been modified, compared to the start of the request process time:

```
use strict;
use warnings;

use Apache::ServerUtil ();
use File::Spec::Functions qw(catfile);

my $r = shift;
$r->content_type('text/plain');

my $conf_file = catfile Apache::server_root_relative($r->pool, 'conf'),
    "httpd.conf";

printf "$conf_file is %0.2f minutes old", 60*24*(-M $conf_file);
```

- The script reports incorrect time most of the time

- Because `-M` reports the difference between file's modification time and `$^T`
- Under `mod_perl` `$^T` is set when the process starts and doesn't change after that.
- Solution: Reset `$^T` on each request

```
package MyApache::TimeReset;

use strict;
use warnings;

use Apache::RequestRec ();

use Apache::Const -compile => 'OK';

sub handler {
    my $r = shift;
    $^T = $r->request_time;
    return Apache::OK;
}
1;
```

- We could do:

```
$^T = time();
```

- But `$r->request_time` already stores the request's start time, so we get it without performing an additional system call.

- To enable it just add to *httpd.conf*:

PerlPostReadRequestHandler MyApache::TimeReset

either to the global section, or to the `<VirtualHost>` section if you want this handler to be run only for a specific virtual host.

6.2.2 *PerlTransHandler*

- The *translate* phase is used to perform the translation of a request's URI into an corresponding filename.
- If no custom handler is provided, the server's standard translation rules (e.g., `Alias` directives, `mod_rewrite`, etc.) will continue to be used.
- Also used to modify the URI itself and the request method.
- This is also a good place to register new handlers for the following phases based on the URI.
- This phase is of type `RUN_FIRST`.

- The handler's configuration scope is SRV.

Let's look at a handler that rewrites request URIs, ala `mod_rewrite`.

For example if the previously static pages are now autogenerated, and

`http://example.com/news/20021031/09/index.html`

is now handled by:

`http://example.com/perl/news.pl?date=20021031&id=09&page=index.html`

the following handler can do the rewriting work transparent to *news.pl*, so you can still use the former URI mapping:

```
package MyApache::RewriteURI;

use strict;
use warnings;

use Apache::RequestRec ();

use Apache::Const -compile => qw(DECLINED);

sub handler {
    my $r = shift;

    my ($date, $id, $page) = $r->uri =~ m|^/news/(\d+)/(\d+)/(.*)|;
    $r->uri("/perl/news.pl");
    $r->args("date=$date&id=$id&page=$page");

    return Apache::DECLINED;
}
1;
```

- To configure this module simply add to *httpd.conf*:

PerlTransHandler +MyApache::RewriteURI

6.2.3 *PerlHeaderParserHandler*

- The *header_parser* phase is the first phase to happen after the request has been mapped to its `<Location>` (or an equivalent container).
- At this phase the handler can examine the request headers and to take a special action based on these.
- e.g., this phase can be used to block evil clients targeting certain resources, while few resources were wasted so far.
- This phase is of type `RUN_ALL`.
- The handler's configuration scope is `DIR`.

- Apache handles the HEAD, GET, POST and several other HTTP methods.
- But you can invent your own methods and make Apache accept them.
- e.g., emails are very similar to HTTP messages: they have a set of headers and a body, sometimes a multi-part body.
- `MyApache::SendEmail` extends HTTP by adding a support for the EMAIL method.
- It enables the new extension and pushes the real content handler during the `PerlHeaderParserHandler` phase:

```
<Location /email>
```

```
    PerlHeaderParserHandler MyApache::SendEmail
```

```
</Location>
```

```
package MyApache::SendEmail;

use strict;
use warnings;

use Apache::RequestRec ();
use Apache::RequestIO ();
use Apache::RequestUtil ();

use Apache::Const -compile => qw(DECLINED OK);

use constant METHOD          => 'EMAIL';
use constant SMTP_HOSTNAME => "localhost";

sub handler {
    my $r = shift;

    return Apache::DECLINED unless $r->method eq METHOD;

    Apache::method_register($r->pool, METHOD);
    $r->handler("perl-script");
    $r->push_handlers(PerlHandler => \&send_email_handler);
}
```

```
    return Apache::OK;
}

sub send_email_handler {
    my $r = shift;

    my %headers = map {$_ => $r->headers_in->get($_)}
        qw(To From Subject);
    my $content = content($r);

    my $status = send_email(\%headers, \$content);

    $r->content_type('text/plain');
    $r->print($status ? "ACK" : "NACK");
    return Apache::OK;
}
```

```

sub content {
    my $r = shift;

    $r->setup_client_block;
    return '' unless $r->should_client_block;
    my $len = $r->headers_in->get('content-length');
    my $buf;
    $r->get_client_block($buf, $len);

    return $buf;
}

sub send_email {
    my($rh_headers, $r_body) = @_;

    require MIME::Lite;
    MIME::Lite->send("smtp", SMTP_HOSTNAME, Timeout => 60);
    my $msg = MIME::Lite->new(%$rh_headers, Data => $$r_body);

    $msg->send;
}
1;

```

- when you extend an HTTP protocol you need to have a client that knows how to use the extension.
- So here is a simple client that uses `LWP::UserAgent` to issue an `EMAIL` method request over HTTP protocol:

```
file:send_http_email.pl
-----
#!/usr/bin/perl

use strict;
use warnings;

require LWP::UserAgent;

my $url = "http://localhost:8000/email/";

my %headers = (
    From      => 'example@example.com',
    To        => 'example@example.com',
    Subject   => '3 weeks in Tibet',
);
```

```
my $content = <<EOI;
I didn't have an email software,
but could use HTTP so I'm sending it over HTTP
EOI
```

```
my $headers = HTTP::Headers->new(%headers);
my $req = HTTP::Request->new("EMAIL", $url, $headers, $content);
my $res = LWP::UserAgent->new->request($req);
print $res->is_success ? $res->content : "failed";
```

6.2.4 *PerlInitHandler*

- When configured inside any container directive, except `<VirtualHost>`, this handler is an alias for `PerlHeaderParserHandler` described earlier.
- Otherwise it acts as an alias for `PerlPostReadRequestHandler` described earlier.
- It is the first handler to be invoked when serving a request.
- This phase is of type `RUN_ALL`.
- The best example here would be to use `Apache::Reload` which takes the benefit of this directive.

- Usually `Apache::Reload` is configured as:

```
PerlInitHandler Apache::Reload
```

```
PerlSetVar ReloadAll Off
```

```
PerlSetVar ReloadModules "MyApache::*"
```

- which during the current HTTP request will monitor and reload all `MyApache::*` modules that have been modified since the last HTTP request.
- However if we move the global configuration into a `<Location>` container:

```
<Location /devel>
```

```
PerlInitHandler Apache::Reload
```

```
PerlSetVar ReloadAll Off
```

```
PerlSetVar ReloadModules "MyApache::*"
```

```
    SetHandler perl-script
    PerlHandler ModPerl::Registry
    Options +ExecCGI
</Location>
```

- Apache::Reload will reload the modified modules, only when a request to the */devel* namespace is issued, because PerlInitHandler plays the role of PerlHeaderParserHandler here.

6.2.5 *PerlAccessHandler*

- The *access_checker* phase is the first of three handlers that are involved in what's known as AAA: Authentication and Authorization, and Access control.
- This phase can be used to restrict access from a certain IP address, time of the day or any other rule not connected to the user's identity.
- This phase is of type `RUN_ALL`.
- The handler's configuration scope is `DIR`.
- The concept behind access checker handler is very simple:

- `return Apache::FORBIDDEN` if the access is not allowed,
- otherwise `return Apache::OK`.

- This handler denies requests made from IPs on the blacklist:

```
package MyApache::BlockByIP;

use strict;
use warnings;

use Apache::RequestRec ();
use Apache::Connection ();

use Apache::Const -compile => qw(FORBIDDEN OK);

my %bad_ips = map {$_ => 1} qw(127.0.0.1 10.0.0.4);

sub handler {
    my $r = shift;
    return exists $bad_ips{$r->connection->remote_ip}
```

```
    ? Apache::FORBIDDEN  
    : Apache::OK;  
}  
1;
```

- To enable the handler simply add it to the container that needs to be protected.
- For example to protect an access to the registry scripts executed from the base location */perl* add:

```
<Location /perl/>  
    SetHandler perl-script  
    PerlResponseHandler ModPerl::Registry  
    PerlAccessHandler MyApache::BlockByIP  
    Options +ExecCGI  
</Location>
```

6.2.6 *PerlAuthenHandler*

- The *check_user_id* (*authen*) phase is called whenever the requested file or directory is password protected.
- This, in turn, requires that the directory be associated with *AuthName*, *AuthType* and at least one *require* directive.
- This phase is usually used to verify a user's identification credentials.
- If the credentials are verified to be correct, the handler should return `Apache::OK`.
- Otherwise the handler returns `Apache::AUTH_REQUIRED` to indicate that the user has not authenticated successfully.

- When Apache sends the HTTP header with this code, the browser will normally pop up a dialog box that prompts the user for login information.
- This phase is of type `RUN_FIRST`.
- The handler's configuration scope is `DIR`.
- The following example handler verifies that:

```
SECRET_LENGTH == length join " ", $username, $password;
```

```
package MyApache::SecretLengthAuth;

use strict;
use warnings;

use Apache::Access ();
use Apache::RequestUtil ();

use Apache::Const -compile => qw(OK DECLINED AUTH_REQUIRED);

use constant SECRET_LENGTH => 14;

sub handler {
    my $r = shift;

    my ($status, $password) = $r->get_basic_auth_pw;
    return $status unless $status == Apache::OK;

    return Apache::OK
        if SECRET_LENGTH == length join " ", $r->user, $password;
```

```
$r->note_basic_auth_failure;  
return Apache::AUTH_REQUIRED;  
}  
1;
```

- To enable this handler you have to tell Apache:
 - what authentication scheme to use (`AuthType: Basic` or `Digest`)
 - what's the authentication realm (`AuthName: any string`)
 - the `Require` directive is needed to specify which usernames are allowed to authenticate. If you set it to `valid-user` any username will do.

```
<Location /perl/>  
  SetHandler perl-script  
  PerlResponseHandler ModPerl::Registry  
  PerlAuthenHandler MyApache::SecretLengthAuth  
  Options +ExecCGI
```

```
AuthType Basic
AuthName "The Gate"
Require valid-user
</Location>
```

6.2.7 *PerlAuthzHandler*

- The *auth_checker* (*authz*) phase is used for authorization control.
- This phase requires a successful authentication from the previous phase, because a username is needed in order to decide whether a user is authorized to access the requested resource.
- As this phase is tightly connected to the authentication phase, the handlers registered for this phase are only called when the requested resource is password protected, similar to the *auth* phase.

- The handler is expected to return:
 - `Apache::DECLINED` to defer the decision,
 - `Apache::OK` to indicate its acceptance of the user's authorization,
 - or `Apache::AUTH_REQUIRED` to indicate that the user is not authorized to access the requested document.
- This phase is of type `RUN_FIRST`.
- The handler's configuration scope is `DIR`.

- The `MyApache::SecretResourceAuthz` handler grants access to certain resources only to certain users who have already properly authenticated:

```
package MyApache::SecretResourceAuthz;

use strict;
use warnings;

use Apache::Access ();
use Apache::RequestUtil ();

use Apache::Const -compile => qw(OK AUTH_REQUIRED);

my %protected = (
    'admin' => ['stas'],
    'report' => [qw(stas boss)],
);
```

```
sub handler {
    my $r = shift;

    my $user = $r->user;
    if ($user) {
        my($section) = $r->uri =~ m|^/company/(\w+)/|;
        if (defined $section && exists $protected{$section}) {
            my $users = $protected{$section};
            return Apache::OK if grep { $_ eq $user } @$users;
        }
        else {
            return Apache::OK;
        }
    }

    $r->note_basic_auth_failure;
    return Apache::AUTH_REQUIRED;
}

1;
```

- The configuration is similar to PerlAuthenHandler, this time we just add the PerlAuthzHandler setting. The rest doesn't change.

```
Alias /company/ /home/httpd/httpd-2.0/perl/  
<Location /company/>  
    SetHandler perl-script  
    PerlResponseHandler ModPerl::Registry  
    PerlAuthenHandler MyApache::SecretLengthAuth  
    PerlAuthzHandler MyApache::SecretResourceAuthz  
    Options +ExecCGI  
  
    AuthType Basic  
    AuthName "The Secret Gate"  
    Require valid-user  
</Location>
```

6.2.8 *PerlTypeHandler*

- The *type_checker* phase is used to set the response MIME type (`Content-type`) and sometimes other bits of document type information like the document language.
- e.g., `mod_autoindex`, which performs automatic directory indexing, uses this phase to map the filename extensions to the corresponding icons which will be later used in the listing of files.
- Of course later phases may override the mime type set in this phase.
- This phase is of type `RUN_FIRST`.

- The handler's configuration scope is `DIR`.
- When overriding the default *type_checker* handler, which is usually the `mod_mime` handler, don't forget to set the response handler:

```
$r->handler('perl-script'); # or $r->handler('modperl');  
$r->set_handlers(PerlResponseHandler => \&handler);
```

- It's the easiest to leave this stage alone and do any desired settings in the *fixups* phase.

6.2.9 *PerlFixupHandler*

- The *fixups* phase is happening just before the content handling phase.
- It gives the last chance to do things before the response is generated.
- For example in this phase `mod_env` populates the environment with variables configured with *SetEnv* and *PassEnv* directives.
- This phase is of type `RUN_ALL`.
- The handler's configuration scope is `DIR`.

- This fixup handler tells Apache at run time which handler and callback should be used to process the request based on the file extension of the request's URI.

```
package MyApache::FileExtDispatch;

use strict;
use warnings;

use Apache::RequestIO ();
use Apache::RequestRec ();

use Apache::Const -compile => 'OK';

use constant HANDLER => 0;
use constant CALLBACK => 1;

my %exts = (
    cgi => ['perl-script',      \&cgi_handler],
    pl  => ['modperl',          \&pl_handler ],
    tt  => ['perl-script',      \&tt_handler ],
    txt => ['default-handler',  undef          ],
);
```

```

sub handler {
    my $r = shift;

    my($ext) = $r->uri =~ /\.(\\w+)$/;
    $ext = 'txt' unless defined $ext and exists $exts{$ext};

    $r->handler($exts{$ext}->[HANDLER]);

    if (defined $exts{$ext}->[CALLBACK]) {
        $r->set_handlers(PerlHandler => $exts{$ext}->[CALLBACK]);
    }

    return Apache::OK;
}

sub cgi_handler { content_handler($_[0], 'cgi') }
sub pl_handler   { content_handler($_[0], 'pl')  }
sub tt_handler   { content_handler($_[0], 'tt')  }

```

```
sub content_handler {  
    my($r, $type) = @_;  
  
    $r->content_type('text/plain');  
    $r->print("A handler of type '$type' was called");  
  
    return Apache::OK;  
}  
  
1;
```

- Here is how this handler is configured:

```
Alias /dispatch/ /home/httpd/httpd-2.0/htdocs/  
<Location /dispatch/>  
    PerlFixupHandler MyApache::FileExtDispatch  
</Location>
```

- Notice that there is no need to specify anything, but the fixup handler.
- It applies the rest of the settings dynamically at run-time.

6.2.10 *PerlResponseHandler*

- The *handler (response)* phase is used for generating the response.
- This is arguably the most important phase and most of the existing Apache modules do most of their work at this phase.
- This is the only phase that requires two directives under `mod_perl`:

```
<Location /perl>  
    SetHandler perl-script  
    PerlResponseHandler ModPerl::Registry  
</Location>
```

- `SetHandler` set to `perl-script` or `modperl` tells Apache that `mod_perl` is going to handle the response generation.
- `PerlResponseHandler` tells `mod_perl` which callback is going to do the job.
- This phase is of type `RUN_FIRST`.
- The handler's configuration scope is `DIR`.

- This handler that prints itself:

```
package MyApache::Deparse;

use strict;
use warnings;

use Apache::RequestRec ();
use Apache::RequestIO ();
use B::Deparse ();

use Apache::Const -compile => 'OK';

sub handler {
    my $r = shift;

    $r->content_type('text/plain');
    $r->print('sub handler ', B::Deparse->new->coderef2text(\&handler));

    return Apache::OK;
}
1;
```

- To enable this handler add to *httpd.conf*:

```
<Location /deparse>  
    SetHandler modperl  
    PerlResponseHandler MyApache::Deparse  
</Location>
```

- Now when the server is restarted and we issue a request to *http://localhost/deparsed* we get the following response:

```
sub handler {  
    package MyApache::Deparsed;  
    my $r = shift @_;  
    $r->content_type('text/plain');  
    $r->print('sub handler ', 'B::Deparsed'->new->coderef2text(\&handler));  
    return 0;  
}
```

- If you compare it to the source code, it's pretty much the same code. `B::Deparsed` is fun to play with!

6.2.11 *PerlLogHandler*

- The *log_transaction* phase is always executed.
- Even if the previous phases were aborted.
- At this phase log handlers usually log various information about the request and the response.
- This phase is of type `RUN_ALL`.
- The handler's configuration scope is `DIR`.

- The following handler logs the request data into user-specific files.
- We assume that all URIs include the username in the form of: */users/username/*

```
package MyApache::LogPerUser;
```

```
use strict;
```

```
use warnings;
```

```
use Apache::RequestRec ();
```

```
use Apache::Connection ();
```

```
use Fcntl qw(:flock);
```

```
use Apache::Const -compile => qw(OK DECLINED);
```

```
sub handler {
    my $r = shift;

    my($username) = $r->uri =~ m|^/users/([^\|+)|;
    return Apache::DECLINED unless defined $username;

    my $entry = sprintf qq(%s [%s] "%s" %d %d\n),
        $r->connection->remote_ip, scalar(localtime),
        $r->uri, $r->status, $r->bytes_sent;

    my $log_path = Apache::server_root_relative($r->pool,
        "logs/$username.log");
    open my $fh, ">>$log_path" or die "can't open $log_path: $!";
    flock $fh, LOCK_EX;
    print $fh $entry;
    close $fh;

    return Apache::OK;
}
1;
```

- Configuration:

```
<Location /users/>  
    SetHandler perl-script  
    PerlResponseHandler ModPerl::Registry  
    PerlLogHandler MyApache::LogPerUser  
    Options +ExecCGI  
</Location>
```

- After restarting the server and issuing requests to the following URIs:

```
http://localhost/users/stas/test.pl
```

```
http://localhost/users/eric/test.pl
```

```
http://localhost/users/stas/date.pl
```

- The `MyApache::LogPerUser` handler will append to *logs/stas.log*:

```
127.0.0.1 [Sat Aug 31 01:50:38 2002] "/users/stas/test.pl" 200 8  
127.0.0.1 [Sat Aug 31 01:50:40 2002] "/users/stas/date.pl" 200 44
```

- and to *logs/eric.log*:

```
127.0.0.1 [Sat Aug 31 01:50:39 2002] "/users/eric/test.pl" 200 8
```

6.2.12 PerlCleanupHandler

META: not implemented yet

7 mod_perl 1.0 to mod_perl 2.0 Migration

7.1 About

- Configuration Directives Changes
- mod_perl API Changes
- mod_perl Modules Changes

7.2 Migrating from mod_perl 1.0 to mod_perl 2.0

- Some configuration directives were renamed or removed
- Several APIs have changed, renamed, removed, or moved to new packages.
- ...
- But there is a back-compatibility layer!

7.3 The Shortest Migration Path from 1.0

```
use Apache2;  
use Apache::compat;
```

- Certain Configuration directives and APIs have changed
- See <http://perl.apache.org/docs/2.0/user/compat/compat.html>

7.4 Migrating Configuration Files

- Several configuration directives are deprecated in 2.0
- but still available for backwards compatibility
- Otherwise, consider using the directives that have replaced them.

7.4.1 PerlHandler

`PerlHandler => PerlResponseHandler`

7.4.2 PerlSendHeader

PerlSendHeader On => PerlOptions +ParseHeaders

PerlSendHeader Off => PerlOptions -ParseHeaders

7.4.3 PerlSetupEnv

PerlSetupEnv On => PerlOptions +SetupEnv

PerlSetupEnv Off => PerlOptions -SetupEnv

7.4.4 PerlTaintCheck

- The taint mode now can be turned on with:

PerlSwitches -T

- It's disabled by default.
- You cannot disable it once it's enabled.

7.4.5 *PerlWarn*

- Warnings now can be enabled globally with:

PerlSwitches -w

7.4.6 *PerlFreshRestart*

- `PerlFreshRestart` is a `mod_perl 1.0` legacy
- In `mod_perl 2.0` a full tear-down and startup of interpreters is done on restart.
- To reuse the same *httpd.conf* for 1.0 and 2.0:

```
<IfDefine !MODPERL2>  
    PerlFreshRestart  
</IfDefine>
```

7.5 Code Porting

- Certain APIs couldn't be preserved in 2.0, hence:

```
use Apache::compat;
```

- which is slower, because it's implemented in Perl.
- Consider getting rid of using the compat layer

- Remember the installation?

```
% perl Makefile.PL ... MP_INST_APACHE2=1
```

- Now mod_perl 1.0 and 2.0 can co-exist on the same machine
- mod_perl 2.0 modules go into the dir *Apache2/*

```
use Apache2;
```

- looks for path(s) with *Apache2/* using @INC
- prepends them to @INC

7.6 ModPerl::Registry Family

- s/Apache::Registry/ModPerl::Registry/

```
Alias /perl/ /home/httpd/perl/  
<Location /perl>  
    SetHandler perl-script  
    PerlResponseHandler ModPerl::Registry  
    Options +ExecCGI  
    PerlOptions +ParseHeaders  
</Location>
```

- Issues with chdir(), which affects all threads
- MPM doesn't matter, should work the same everywhere

- Cook your own registry with `ModPerl::RegistryCooker`
 - is-a: Subclass one of the classes
 - has-a: Build your own using aliasing

7.7 Method Handlers

- `$$` doesn't work since some callbacks accept more than 2 args

```
package Bird;
@ISA = qw(Eagle);

sub handler : method {
    my($class, $r) = @_;
    ...;
}
```

- See the *attributes* manpage.

7.8 Apache::StatINC Replacement

- Replace:

```
PerlInitHandler Apache::StatINC
```

- with:

```
PerlInitHandler Apache::Reload
```

- Apache::Reload provides an extra functionality
- See the module's manpage.

8 That's all folks!

8.1 Thanks

Thanks to TicketMaster for sponsoring my work on mod_perl



8.2 References

- All the information can be found at:


`http://perl.apache.org/docs/`

- Further Questions?
 - Grab me at the corridor and demand answers
 - Ask at `modperl@perl.apache.org`

8.3 A Shameless Plug

Programming, Administration, Performance Tips

Practical mod_perl



O'REILLY®

Stas Bekman & Eric Cbolet

9 Bonus Chapter: What's new under the sun?

9.1 About

- What's new in Apache 2.0
- What's new in Perl 5.6.0 - 5.8.0
- What's new in mod_perl 2.0

9.2 Versioning Convention

- To make things simple here and in the new docs:
- mod_perl:
 - mod_perl 1.0 (not mod_perl 1.xx)
 - mod_perl 2.0 (not mod_perl 2.0.xx)
- Apache:
 - Apache 1.3
 - Apache 2.0

9.3 Why the 2.0 Rewrite?

- Too patchy (6 years!), backward compatibility with:
 - Apache 1.3.0 - 1.3.27
 - Perl 5.003 - 5.8.0
- mod_perl 2.0 starts afresh with:
 - Apache 2.0 (incompatible with Apache 1.3)
 - Perl 5.6.0 (has semi-thread-safe Perl Interpreters)
 - Threaded mpms: 5.8.0 (really thread-safe)

- A new build system autogenerates the code used to
 - autogenerates the code that is used to ...
 - ...
 - which generates the final code ...
 - ... and it all works
- Automatically supports new Apache APIs

9.3.1 *The Apache::Test Framework*

- The core:

`All tests successful.`

`Files=85, Tests=568, 75 wallclock secs...`

- Any Perl module needing mod_perl 1.0 or 2.0
- Any Apache module (both 1.3 and 2.0), PHP, Python, C...
- Already used by httpd-test to test Apache!!!

9.4 New in Apache 2.0

- Apache Portable Runtime
- Multi Processing Model modules (MPMs).
 - processes: prefork
 - threads: worker, leader, perchild...
 - os: mpmt_os2, netware, winnt, beos...
- Protocol Modules (HTTP, POP3, SMTP...)

- I/O Filtering
- Bucket Brigades
- Parsed Configuration Tree
- New Hook Scheme (Flexible, Order-able)
- Optional Functions

9.5 New in Perl m/5\.(6|8)\.ld/

- Thread-safe Interpreter (5.8.0) via `perl_clone()`
- Subroutine attributes:

```
sub handler : FilterRequestHandler { ... }
```

- `CORE::GLOBAL::` subs overriding `CORE::`
- PerlIO layers => `APR::PerlIO:`

```
open my $fh, "<:APR", $file, $r;
```

- I18n: Unicode, UTF...

- Safe signal handling (5.8.0)
- Large file support (files > 2 gigabytes)
- XSLoader, a lighter alternative to DynaLoader
- fine tuned warnings control:

```
use warnings FATAL => 'all';
```

- Lots of performance enhancements. Though threads slow things down.
- Numerous memory leaks and bugs were fixed

9.6 New in mod_perl 2.0

- All the new Apache 2.0 and Perl 5.6.0+ features
- Plus its own new features

9.6.1 *Threads Support*

- Thread Interpreters Pool
- `scalar @perl_interpreters != scalar @apache_threads`
 - no need for front-end/back-end separation

- Two classes of interpreters: *parent* and *clone*
- parent: preload modules and *perl_clone()* clones
- clones: do the real work
 - mutable data is copied by the clone
 - read-only data such as the syntax tree is shared
 - clone pools are FIFO => memory re-use

9.6.2 *Thread-environment Issues*

- Manipulating Perl data is thread-safe (5.8.0)

`push()`, `map()`, `chomp()`, ...

- The rest, depends on the underlying implementation

`localtime()`, `readdir()`, `srand()`, ...

- Thread-safe but Process-scoped

`chdir()`, `umask()`, `chroot()`, ...

- See `perlthrtut(3)`

9.6.3 Perl interface to the APR and Apache APIs

- Apache:: API, which handles issues specific to the web server.
- APR:: API, which implements a portable and efficient API to handle generically work with files, threads, processes, shared memory, etc.
- APR:: API, can be used outside of Apache as well.

9.6.4 *Other New Features*

- Protocol modules
- Simplified stream-oriented filtering API
- etc...

9.6.5 *Optimizations*

- "Compiled" Perl*Handlers
- Inlined Apache::*.xs calls
- Use of Apache Pools for memory allocations

10 Bonus Chapter: mod_perl 2.0 Server Configuration

10.1 About

- `mod_perl` configuration directives
- Enabling `mod_perl`
- Accessing the `mod_perl` 2.0 Modules
- Startup File
- Perl's Command Line Switches
- `mod_perl` 2.0 Handlers
- `PerlOptions` Directive

- **Handlers Directives**
- **Threads Mode Specific Directives**
- **mod_perl Directives Argument Types and Allowed Location**
- **Server Startup Options Retrieval**
- **Perl Interface to the Apache Configuration Tree**

10.2 Configuring mod_perl 2.0

- DSO:

```
LoadModule perl_module modules/mod_perl.so
```

- Static: nada

10.2.1 Accessing the Modules

- `mod_perl 2.0` Perl libs go to *Apache2/*
- Adjust `@INC`:

```
use Apache2 ();
```

```
# @INC before:
```

```
/usr/lib/perl5/site_perl/5.8.0/i686-linux-thread-multi
```

```
# @INC after:
```

```
/usr/lib/perl5/site_perl/5.8.0/i686-linux-thread-multi/Apache2
```

10.2.2 *PerlRequire'd Startup File*

```
use Apache2 ();  
# use Apache::compat (); # 1.0 compat  
  
use lib qw(/home/httpd/perl);  
  
use ModPerl::Util (); #for CORE::GLOBAL::exit  
  
use Apache::RequestRec ();  
use Apache::RequestIO ();  
use Apache::RequestUtil ();  
  
use Apache::Server ();  
use Apache::ServerUtil ();  
use Apache::Connection ();  
use Apache::Log ();
```

```
use APR::Table ();

use ModPerl::Registry ();

use Apache::Const -compile => ':common';
use APR::Const -compile => ':common';

1;
```

10.2.3 Perl's Command Line Switches

- `PerlSwitches` passes any Perl switches
- e.g., enable warnings and taint checking:

```
PerlSwitches -wT
```

- adjust `@INC` values:

```
PerlSwitches -I/home/stas/modperl
```

10.2.4 mod_perl 2.0 Core Handlers

10.2.5 "modperl"

SetHandler modperl

- calls the Perl*Handler's callback func.
- sets MOD_PERL, GATEWAY_INTERFACE, PATH and TZ env vars.
- No tied IO handles:

```
$r->read($line, $len, $offset);  
$r->print("Dahuuuuut!");
```

10.2.6 *"perl-script"*

`SetHandler perl-script`

- As in `mod_perl 1.0`
- Unless set to `-GlobalRequest` assumes

`PerlOptions +GlobalRequest`

- Unless set to `-SetupEnv` assumes

`PerlOptions +SetupEnv`

- Tied STDIN and STDOUT

```
my $line = <STDIN>;  
print "Dahuuuuut!";
```

- on each request restores %ENV, @INC, \$/, STDOUT's \$| and END blocks

10.2.6.1 A Simple Response Handler Example

- Printout environment variables ala `perl-script` core handler

```
PerlModule Apache::PrintEnv1
<Location /print_env1>
    SetHandler perl-script
    PerlResponseHandler Apache::PrintEnv1
</Location>
```

```
package Apache::PrintEnv1;
use strict;
use Apache::RequestRec (); # for $r->content_type
use Apache::Const -compile => 'OK';

sub handler {
    my $r = shift;

    $r->content_type('text/plain');
    for (sort keys %ENV){
        print "$_ => $ENV{$_}\n";
    }

    return Apache::OK;
}
1;
```

- Printout environment variables ala `modperl` core handler

```
PerlModule Apache::PrintEnv2
```

```
<Location /print_env2>
```

```
    SetHandler modperl
```

```
    PerlResponseHandler Apache::PrintEnv2
```

```
</Location>
```

```
package Apache::PrintEnv2;
use strict;
use Apache::RequestRec (); # for $r->content_type
use Apache::RequestIO (); # for $r->print
use Apache::Const -compile => 'OK';

sub handler {
    my $r = shift;

    $r->content_type('text/plain');
    $r->subprocess_env;
    for (sort keys %ENV){
        $r->print("$_ => $ENV{$_}\n");
    }
    return Apache::OK;
}
1;
```

- Instead of calling:

```
$r->subprocess_env;
```

- could configure the location with:

```
PerlOptions +SetupEnv
```

10.2.7 *PerlOptions Directive*

- Disable `mod_perl` for a given `VirtualHost`:

```
<VirtualHost ...>  
    PerlOptions -Enable  
</VirtualHost>
```

- Give the `VirtualHost` its own interpreter pool.

```
<VirtualHost ...>  
    PerlOptions +Clone  
    PerlInterpStart 2  
    PerlInterpMax 2  
</VirtualHost>
```

- Run different versions of the same module:

```
<VirtualHost ...>  
    ServerName dev1  
    PerlOptions +Parent  
    PerlSwitches -Mbllib=/home/dev1/lib/perl  
</VirtualHost>
```

```
<VirtualHost ...>  
    ServerName dev2  
    PerlOptions +Parent  
    PerlSwitches -Mbllib=/home/dev2/lib/perl  
</VirtualHost>
```

- disallow certain handlers/options

```
<VirtualHost ...>  
    PerlOptions -Authen -Authz -Access -Sections  
</VirtualHost>
```

- Or maybe everything but the response handler:

```
<VirtualHost ...>  
    PerlOptions None +Response  
</VirtualHost>
```

- Resolve Perl*Handlers at startup time:

PerlOptions +Autoload

PerlResponseHandler Apache::Magick

- Disable the global request_rec (Apache->request)

<Location ...>

SetHandler perl-script

PerlOptions -GlobalRequest

...

</Location>

- s/PerlSendHeader On/PerlOptions +ParseHeaders/
- Merge Handlers

```
PerlFixupHandler Apache::FixupA
```

```
<Location /inside>
```

```
    PerlOptions +MergeHandlers
```

```
    PerlFixupHandler Apache::FixupB
```

```
</Location>
```

- s/PerlSetupEnv Off/PerlOptions -SetupEnv/

10.2.8 Threads Mode Specific Directives

- PerlInterpStart
- PerlInterpMax
- PerlInterpMinSpare
- PerlInterpMaxSpare
- PerlInterpMaxRequests

- `PerlInterpScope`
- Use for the lifetime of the request (default):

`PerlInterpScope request`

- Use a separate interpreter in subrequests:

`PerlInterpScope subrequest`

- Use a separate interpreter for each handler:

`PerlInterpScope handler`

10.2.9 Retrieving Server Startup Options

```
% httpd -DONE_PROCESS
```

- Retrieve:

```
if (Apache::exists_config_define("ONE_PROCESS")) {  
    print "Running in a single mode";  
}
```

10.3 New Apache Phases

- `PerlOpenLogsHandler`
- `PerlPostConfigHandler`
- `PerlPreConnectionHandler`
- `PerlProcessConnectionHandler`
- `PerlResponseHandler`
- `PerlInputFilterHandler`
- `PerlOutputFilterHandler`