

Tutorial Handouts: 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 05:58:09 2002 GMT

1 Getting Your Feet Wet with mod_perl

1.1 Description

This chapter gives you the bare minimum information to get you started with mod_perl 2.0. For most people it's sufficient to get going.

1.2 Prerequisites

Before building mod_perl 2.0 you need to have its prerequisites installed. If you don't have them, download and install them first, using the information in the following sections. Otherwise proceed directly to the mod_perl building instructions.

The mod_perl 2.0 prerequisites are:

- **Apache**

Apache 2.0 is required. mod_perl 2.0 **does not** work with Apache 1.3.

- **Perl**

- **prefork MPM**

Requires at least Perl version 5.6.0. But we strongly suggest to use at least version 5.6.1, since 5.6.0 is quite buggy.

You don't need to have threads-support enabled in Perl. If you do have it, it **must** be *ithreads* and not *5005threads*! If you have:

```
% perl5.8.0 -V:use5005threads
use5005threads='define' ;
```

you must rebuild Perl without threads enabled or with `-Dusethreads`. Remember that threads support slows things down, so don't enable it unless you really need it.

- **threaded MPMs**

Require at least Perl version 5.8.0 with *ithreads* support built-in. That means that it should report:

```
% perl5.8.0 -V:useithreads -V:usemultiplicity
useithreads='define' ;
usemultiplicity='define' ;
```

If that's not what you see rebuild Perl with `-Dusethreads`.

1.2.1 Downloading Stable Release Sources

If you are going to install mod_perl on a production site, you want to use the officially released stable components. Since the latest stable versions change all the time you should check for the latest stable version at the listed below URLs:

- **Perl**

Download from: <http://cpan.org/src/README.html>

This direct link which symlinks to the latest release should work too:
<http://cpan.org/src/stable.tar.gz>.

For the purpose of examples in this chapter we will use the package named *perl-5.8.x.tar.gz*, where *x* should be replaced with the real version number.

- **Apache**

Download from: <http://www.apache.org/dist/httpd/>

For the purpose of examples in this chapter we will use the package named *httpd-2.x.xx.tar.gz*, where *x.xx* should be replaced with the real version number.

1.2.2 Getting Bleeding Edge CVS Sources

If you really know what you are doing you can use the cvs versions of the components. Chances are that you don't want to them on a production site. You have been warned!

- **Perl**

```
# (--delete to ensure a clean state)
% rsync -acvz --delete --force \
  rsync://ftp.linux.activestate.com/perl-current/ perl-current
```

If you are re-building Perl after rsync-ing, make sure to cleanup first:

```
% make distclean
```

before running `./Configure`.

You'll also want to install (at least) LWP if you want to fully test `mod_perl`. You can install LWP with CPAN `.pm` shell:

```
% perl -MCPAN -e 'install("LWP")'
```

- **Apache**

To download the cvs version of `httpd-2.0` and bring it to the same state of the distribution package, execute the following commands:

```
% cvs -d :pserver:anoncvs@cvs.apache.org:/home/cvspublic login
```

The password is "anoncvs".

```
% cvs -d :pserver:anoncvs@cvs.apache.org:/home/cvspublic co httpd-2.0
% cd httpd-2.0/src/lib
% cvs -d :pserver:anoncvs@cvs.apache.org:/home/cvspublic co apr
% cvs -d :pserver:anoncvs@cvs.apache.org:/home/cvspublic co apr-util
% cd ..
% ./buildconf
```

Once extracted, whenever you want to sync with the latest httpd-2.0 version and rebuild, run:

```
% cd httpd-2.0
% cvs up -dP
% make distclean && ./buildconf
```

1.2.3 Configuring and Installing Prerequisites

If you don't have the prerequisites installed yet, install them now.

- **Perl**

```
% cd perl-5.8.x
% ./Configure -des -Dusethreads
% make && make test && make install
```

If you want to debug mod_perl segmentation faults, add the following *./Configure* options:

```
-Doptimize='-g' -Dusedevel
```

- **Apache**

```
% cd httpd-2.x.xx
% ./configure --prefix=$HOME/httpd/prefork --with-mpm=prefork
% make && make install
```

1.3 Installation

First of all check that you have the prerequisites installed.

In this chapter we assume that httpd was installed under *\$HOME/httpd/prefork*.

Next, download the mod_perl 2.0 source from: <http://perl.apache.org/download/>.

Now, configure mod_perl:

```
% 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
```

where *MP_AP_PREFIX* is an Apache installation prefix, under which the *include/* directory with Apache C header files can be found.

Finally, build, test and install `mod_perl`:

```
% make && make test && make install
```

Become *root* before doing `make install` if installing system-wide.

If something goes wrong or you need to enable optional features please refer to <http://perl.apache.org/docs/2.0/user/install/install.html>.

1.4 Configuration

Enable `mod_perl` built as DSO, by adding to *httpd.conf*:

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

Next, tell Perl where to find `mod_perl2` libraries:

```
PerlModule Apache2
```

There are many other configuration options which you can find at <http://perl.apache.org/docs/2.0/user/config/config.html>.

If you want to run `mod_perl 1.0` code on `mod_perl 2.0` server enable the compatibility layer:

```
PerlModule Apache::compat
```

For more information refer to <http://perl.apache.org/docs/2.0/user/compat/compat.html>.

1.5 Server Launch and Shutdown

Apache is normally launched with `apachectl`:

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

and shut down with:

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

Check `$HOME/httpd/prefork/logs/error_log` to see that the server has started and it's a right one. It should say something similar to:

```
[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
```

1.6 Registry Scripts

To enable registry scripts add to *httpd.conf*:

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

and now assuming that we have the following script:

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

saved in */home/httpd/httpd-2.0/perl/rock.pl*. Make the script executable and readable by everybody:

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

Of course the path to the script should be readable by the server too. In the real world you probably want to have a tighter permissions, but for the purpose of testing, that things are working, this is just fine.

Now restart the server and issue a request to *http://localhost/perl/rock.pl* and you should get the response:

```
mod_perl 2.0 rocks!
```

If that didn't work check the *error_log* file.

1.7 Handler Modules

Finally check that you can run mod_perl handlers. Let's write a response handler similar to the registry script from the previous section:

```
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;
```

Save the code in the file *MyApache/Rocks.pm*, somewhere where `mod_perl` can find it. For example let's put it under */home/httpd/httpd-2.0/perl/MyApache/Rocks.pm*, and we tell `mod_perl` that */home/httpd/httpd-2.0/perl/* is in `@INC`, via a startup file which includes just:

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

and loaded from *httpd.conf*:

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

Now we can configure our module in *httpd.conf*:

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

Now restart the server and issue a request to *http://localhost/rocks* and you should get the response:

```
mod_perl 2.0 rocks!
```

If that didn't work check the *error_log* file.

2 Introducing mod_perl Handlers

2.1 Description

This chapter provides an introduction into mod_perl handlers.

2.2 Handler Anatomy

Apache distinguishes between numerous phases for which it provides hooks (because the C functions are called *ap_hook_<phase_name>*) where modules can plug various callbacks to extend and alter the default behavior of the webserver. mod_perl provides a Perl interface for most of the available hooks, so mod_perl modules writers can change the Apache behavior in Perl. These callbacks are usually referred to as *handlers* and therefore the configuration directives for the mod_perl handlers look like: `PerlFooHandler`, where `Foo` is one of the handler names. For example `PerlResponseHandler` configures the response callback.

A typical handler is simply a perl package with a *handler* subroutine. For example:

```
file:MyApache/CurrentTime.pm
-----
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;
```

This handler simply returns the current date and time as a response.

Since this is a response handler, we configure it as a such in *httpd.conf*:

```
PerlResponseHandler MyApache::CurrentTime
```

Since the response handler should be configured for a specific location, let's write a complete configuration section:

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

Now when a request is issued to *http://localhost/time* this response handler is executed and a response that includes the current time is returned to the client.

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 allows multiple modules to filter both the request and the response. Now one module can pipe its output as an input to another module as if another module was receiving the data directly from the TCP stream. The same mechanism works with the generated response.


With I/O filtering in place, simple filters, like data compression and decompression, can be easily implemented and complex filters, like SSL, are now possible without needing to modify the the server code which was the case with Apache 1.3.

In order to make the filtering mechanism efficient and avoid unnecessary copying, the *Bucket Brigades* technology was introduced.

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. The goal is to minimize the data copying where possible. Buckets come in different types, such as files, data blocks, end of stream indicators, pools, etc. To manipulate a bucket one doesn't need to know its internal representation.

The stream of data is represented by bucket brigades. When a filter is called it gets passed the brigade that was the output of the previous filter. This brigade is then manipulated by the filter (e.g., by modifying some buckets) and passed to the next filter in the stack.

The following figure depicts an imaginary bucket brigade:



The figure tries to show that after the presented bucket brigade has passed through several filters some buckets were removed, some modified and some added. Of course the handler that gets the brigade cannot tell the history of the brigade, it can only see the existing buckets in the brigade.

Bucket brigades are discussed in detail in the protocol handlers and I/O filtering chapters.

2.5 Single Phase's Multiple Handlers Behavior

For each phase there can be more than one handler assigned (also known as *hooks*, because the C functions are called *ap_hook_<phase_name>*). 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
PerlProcessConnectionHandler	RUN_FIRST
PerlPostReadRequestHandler	RUN_ALL
PerlTransHandler	RUN_FIRST
PerlInitHandler	RUN_ALL
PerlHeaderParserHandler	RUN_ALL
PerlAccessHandler	RUN_ALL
PerlAuthenHandler	RUN_FIRST
PerlAuthzHandler	RUN_FIRST
PerlTypeHandler	RUN_FIRST
PerlFixupHandler	RUN_ALL
PerlResponseHandler	RUN_FIRST
PerlLogHandler	RUN_ALL
PerlCleanupHandler	RUN_ALL
PerlInputFilterHandler	VOID
PerlOutputFilterHandler	VOID

And here is the description of the possible types:

- **VOID**

Handlers of the type VOID will be *all* executed in the order they have been registered disregarding their return values. Though in mod_perl they are expected to return `Apache::OK`.

- **RUN_FIRST**

Handlers of the type RUN_FIRST will be 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.

- **RUN_ALL**

Handlers of the type RUN_ALL will be executed in the order they have been registered until the first handler that returns something other than `Apache::OK` or `Apache::DECLINED`.

For C API declarations see *include/ap_config.h*, which includes other types which aren't exposed by the `mod_perl` handlers.

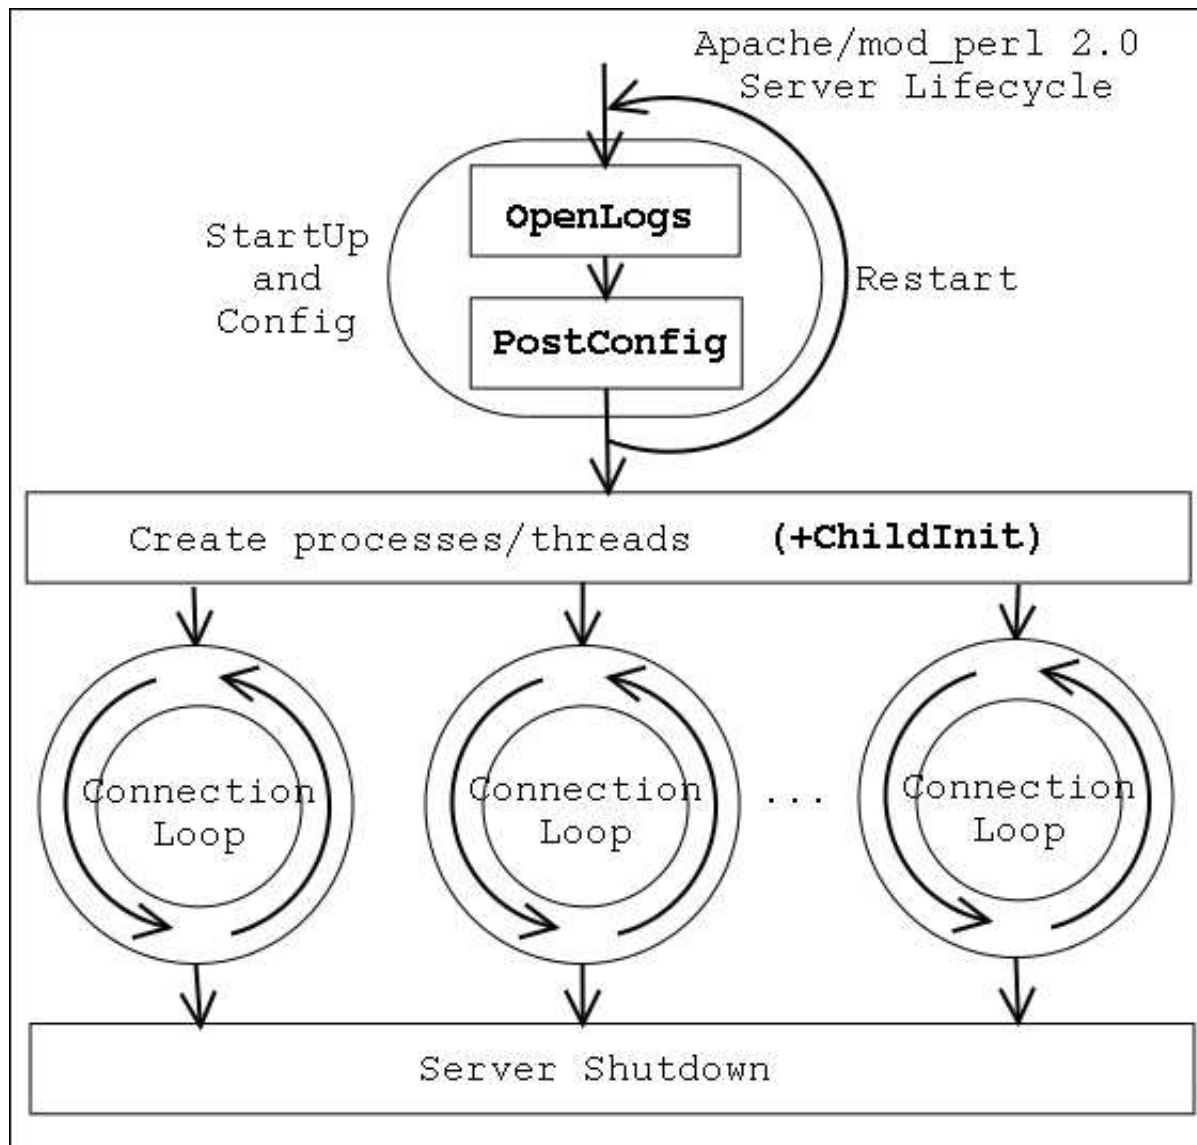
3 Server Life Cycle Handlers

3.1 Description

This chapter discusses server life cycle and the `mod_perl` handlers participating in it.

3.2 Server Life Cycle

The following diagram depicts the Apache 2.0 server life cycle and highlights which handlers are available to `mod_perl 2.0`:



Apache 2.0 starts by parsing the configuration file. After the configuration file is parsed, the `PerlOpenLogsHandler` handlers are executed if any. After that it's a turn of `PerlPostConfigHandler` handlers to be run. When the *post_config* phase is finished the server immediately restarts, to make sure

that it can survive graceful restarts after starting to serve the clients.

When the restart is completed, Apache 2.0 spawns the workers that will do the actual work. Depending on the used MPM, these can be threads, processes and a mixture of both. For example the *worker* MPM spawns a number of processes, each running a number of threads. When each child process is started PerlChildInit handlers are executed. Notice that they are run for each starting process, not a thread.

From that moment on each working thread processes connections until it's killed by the server or the server is shutdown.

3.2.1 Startup Phases Demonstration Module

Let's look at the following example that demonstrates all the startup phases:

```
file:MyApache/StartupLog.pm
-----
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

```

When we perform a server startup followed by a shutdown, the *logs/startup_log* is created if it didn't exist already (it shares the same directory with *error_log* and other standard log files), and each stage appends to it its log information. So when we perform:

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

the following is getting logged to *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

```

First of all, we can clearly see that Apache always restart itself after the first *post_config* phase is over. The logs show that the *post_config* phase is preceded by the *open_logs* phase. Only after Apache has restarted itself and has completed the *open_logs* and *post_config* phase again the *child_init* phase is run for each child process. In our example we have had the setting *StartServers=4*, therefore you can see four child processes were started.

Finally you can see that on server shutdown the *END {}* block has been executed by the parent server only.

Apache also specifies 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.

Now let's discuss each of the mentioned startup handlers and their implementation in the `MyApache::StartupLog` module in detail.

3.2.2 *PerlOpenLogsHandler*

The *open_logs* phase happens just before the *post_config* phase.

Handlers registered by `PerlOpenLogsHandler` are usually used for opening module-specific log files.

At this stage the `STDERR` stream is not yet redirected to *error_log*, and therefore any messages to that stream will be printed to the console the server is starting from (if such exists).

This phase is of type `RUN_ALL`.

The handler's configuration scope is `SRV`.

As we have seen in the `MyApache::StartupLog::open_logs` handler, the *open_logs* phase handlers accept four arguments: the configuration pool, the logging streams pool, the temporary pool and the server object:

```
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;
}
```

In our example the handler uses the function `Apache::server_root_relative()` to set the full path to the log file, which is then opened for appending and set to unbuffered mode. Finally it logs the fact that it's running in the parent process.

As you've seen in the example this handler is configured by adding to *httpd.conf*:

```
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 (which happens at the *child_init* phase).

This phase can be 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.

META: once mod_perl will have the API for that.

This phase is of type RUN_ALL.

The handler's configuration scope is SRV.

In our `MyApache::StartupLog` example we used the `post_config()` handler:

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

As you can see, its arguments are identical to the `open_logs` phase's handler. In this example handler we don't do much but logging that the configuration was completed and returning right away.

As you've seen in the example this handler is configured by adding to `httpd.conf`:

```
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!) will run 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. For example `Apache::DBI` pre-opens database connections during this phase and `Apache::Resource` sets the process' resources limits.

This phase is of type VOID.

The handler's configuration scope is SRV.

In our `MyApache::StartupLog` example we used the `child_init()` handler:

```
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. The example handler logs the pid of the child process it's run in and returns.

As you've seen in the example this handler is configured by adding to `httpd.conf`:

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

4 Protocol Handlers

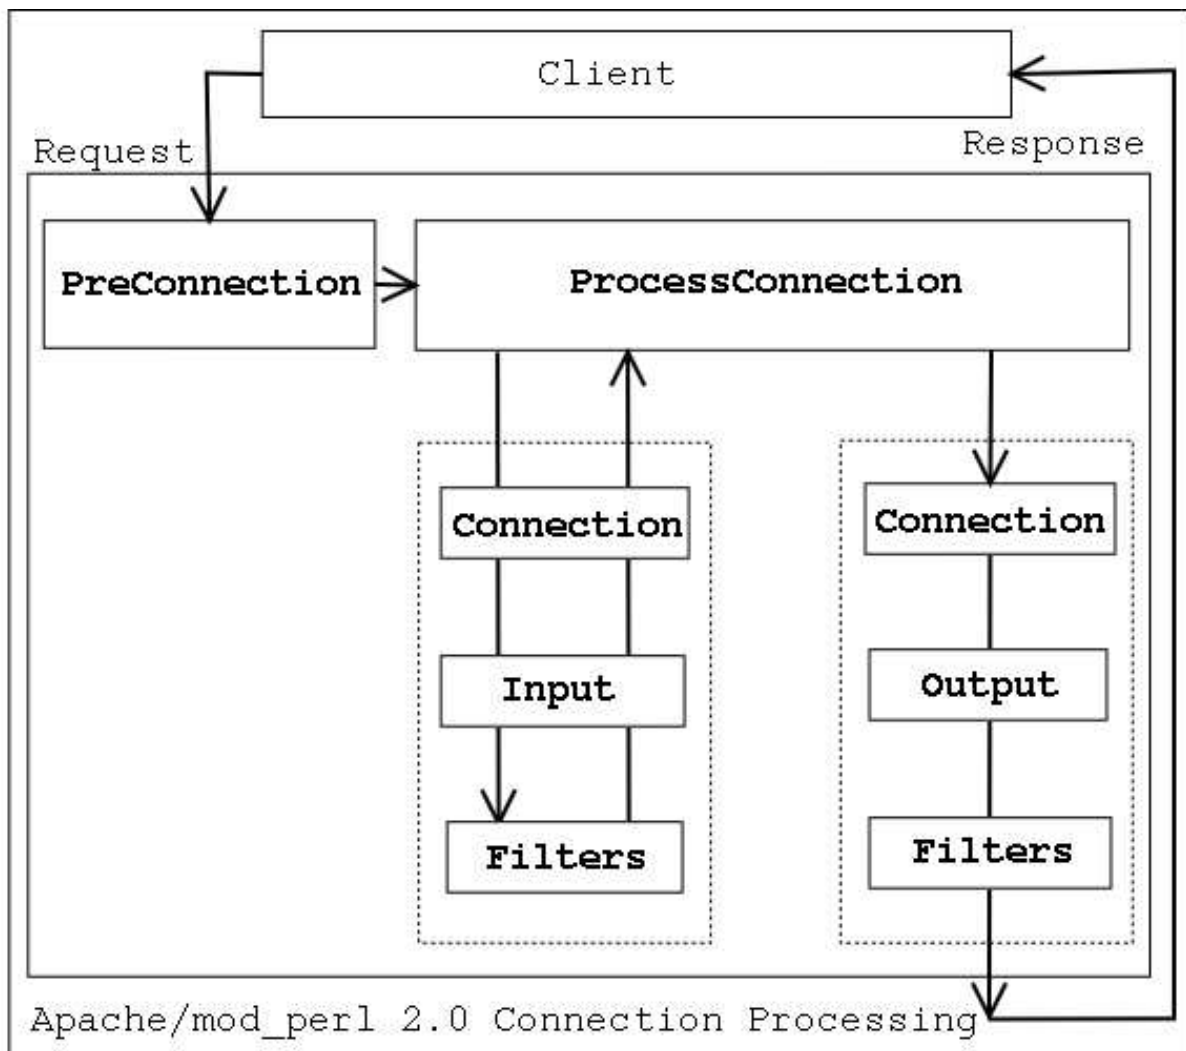
4.1 Description

This chapter explains how to implement Protocol (Connection) Handlers in mod_perl.

4.2 Connection Cycle Phases

As we saw earlier, each child server (be it a thread or a process) 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 several images are requested for the same webpage.

The following diagram depicts the connection life cycle and highlights which handlers are available to mod_perl 2.0:



When a connection is issued by a client, it's first run through `PerlPreConnectionHandler` and then passed to the `PerlProcessConnectionHandler`, which generates the response. When `PerlProcessConnectionHandler` is reading data from the client, it can be filtered by connection input filters. The generated response can be also filtered through connection output filters. Filters are usually used for modifying the data flowing through them, but can be used for other purposes as well (e.g., logging interesting information).

Now let's discuss each of the `PerlPreConnectionHandler` and `PerlProcessConnectionHandler` handlers in detail.

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.

In `mod_perl` 1.0 during code development `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 phase, so if we want to make sure that all modified Perl modules are reloaded for any protocols and its phases, it's the best to set the scope of the Perl interpreter to the lifetime of the connection via:

```
PerlInterpScope connection
```

and invoke the `Apache::Reload` handler during the *pre_connection* phase. However this development-time advantage can become a disadvantage in production--for example if a connection, handled by HTTP protocol, is configured as `KeepAlive` and 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.

Example:

A *pre_connection* handler accepts connection record and socket objects as its arguments:

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

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.

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

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

Now let's look at the following two examples of connection handlers. The first using the connection socket to read and write the data and the second using bucket brigades to accomplish the same and allow for connection filters to do their work.

4.2.2.1 Socket-based Protocol Module

To demonstrate the workings of a protocol module, we'll take a look at the `MyApache::EchoSocket` module, which simply echoes the data read back to the client. In this module we will use the implementation that works directly with the connection socket and therefore bypasses connection filters if any.

A protocol handler is configured using the `PerlProcessConnectionHandler` directive and we will use the `Listen` and `<VirtualHost>` directives to bind to the non-standard port **8010**:

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

`MyApache::EchoSocket` is then enabled when starting Apache:

```
panic% httpd
```

And we give it a whirl:

```
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.
```

Here is the code:

```
file:MyApache/EchoSocket.pm
-----
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;
```

The example handler starts with the standard *package* declaration and of course, `use strict;`. As with all Perl*Handlers, the subroutine name defaults to *handler*. However, in the case of a protocol handler, the first argument is not a *request_rec*, but a *conn_rec* blessed into the `Apache::Connection` class. We have direct access to the client socket via `Apache::Connection`'s *client_socket* method. This returns an object blessed into the `APR::Socket` class.

Inside the read/send loop, the handler attempts to read `BUFF_LEN` bytes from the client socket into the `$buff` buffer. The `$rlen` parameter will be set to the number of bytes actually read. The `APR::Socket::recv()` method returns an APR status value, but we need only to check the read length to break out of the loop if it is less than or equal to 0 bytes. The handler also breaks the loop after processing an input including nothing but new lines characters, which is how we abort the connection in the interactive mode.

If the handler receives some data, it sends it unmodified back to the client with the `APR::Socket::send()` method. When the loop is finished the handler returns `Apache::OK`, telling Apache to terminate the connection. As mentioned earlier since this handler is working directly with the connection socket, no filters can be applied.

4.2.2.2 Bucket Brigades-based Protocol Module

Now let's look at 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.

The following configuration defines a virtual host listening on port 8011 and which enables the `MyApache::EchoBB` connection handler, which will run its output through `MyApache::EchoBB::lowercase_filter` filter:

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

As before we start the httpd server:

```
panic% httpd
```

And try the new connection handler in action:

```
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

Connection closed by foreign host.
```

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

And here is the implementation of the connection and the filter handlers.

```

file:MyApache/EchoBB.pm
-----
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;

```

For the purpose of explaining how this connection handler works, we are going to simplify the handler. The whole handler can be represented by the following pseudo-code:

```

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);
}

```

The handler receives the incoming data via bucket bridges, one at a time in a loop. It then process each bridge, by retrieving the buckets contained in it, reading the data in, then creating new buckets using the received data, and attaching them to the outgoing brigade. When all the buckets from the incoming bucket brigade were transformed and attached to the outgoing bucket brigade, a flush bucket is created and added as the last bucket, so when the outgoing bucket brigade is passed out to the outgoing connection filters, it won't be buffered but sent to the client right away.

If you look at the complete handler, the loop is terminated when one of the following conditions occurs: an error happens, the end of stream bucket has been seen (no more input at the connection) or when the received data contains nothing but new line characters which we used to to tell the server to terminate the connection.

Notice that 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. To accomplish that modification simply add a code that transforms the data which has been read from the bucket before it's inserted to the outgoing brigade.

We will skip the filter discussion here, since we are going to talk in depth about filters in the dedicated to filters sections. But all you need to know at this stage is that the data sent from the connection handler is filtered by the outgoing filter and which transforms it to be all lowercase.

5 Input and Output Filters

5.1 Description

This chapter discusses `mod_perl`'s input and output filter handlers.

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`

The following sections include several examples that use the `PerlInputFilterHandler` handler.

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`

The following sections include several examples that use the `PerlOutputFilterHandler` handler.

5.2.3 *Connection vs. HTTP Request Filters*

Currently the `mod_perl` filters allow connection and request level filtering. Apache supports several other types, which `mod_perl` 2.0 will probably support in the future. `mod_perl` filter handlers specify the type of the filter using the method attributes.

Request filter handlers are declared using the `FilterRequestHandler` attribute. Consider the following 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 ();
```

The 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>

```

Now we have the request input and output filters configured.

The connection filter handler uses the `FilterConnectionHandler` attribute. Here is a similar example for the connection input and output filters.

```

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>

```

This accomplishes the configuration of the connection input and output filters.

Notice that 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 bucket brigades manipulation interface and a simpler, stream-oriented interface (META: as of this writing the latter is available only for the output filtering). The examples in the following sections will help you to understand the difference between the two interfaces.

5.3 All-in-One Filter

Before we delve into the details of how to write filters that do something with the data, let's first write a simple filter that does nothing but snooping on the data that goes through it. We are going to develop the `MyApache::FilterSnoop` handler which can snoop on request and connection filters, in input and output modes.

But first let's develop a simple response handler that simply dumps the request's *args* and *content* as strings:

```

file:MyApache/Dump.pm
-----
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:

```

file:MyApache/FilterSnoop.pm
-----
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) = @_; # filter args

    # $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;

```

This package provides two filter handlers, one for connection and another for request filtering:

```

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

```

Both handlers forward their arguments to the `snoop()` function that does the real job. We needed to add these two subroutines in order to assign the two different attributes. Plus the functions pass the filter type to `snoop()` as the first argument, which gets shifted off `@_` and the rest of the `@_` are the arguments that were originally passed to the filter handler.

It's easy to know whether a filter handler is running in the input or the output mode. The arguments `$filter` and `$bb` are always passed, whereas the arguments `$mode`, `$block`, and `$readbytes` are passed only to input filter handlers.

If we are in the input mode, we retrieve the bucket brigade and immediately link it to `$bb` which makes the brigade available to the next filter. When this filter handler returns, the next filter on the stack will get the brigade. If we forget to perform this linking our filter will become a black hole in which data simply disappears. Next we call `bb_sniff()` which returns the type and the content of the buckets in the brigade.

If we are in the output mode, `$bb` already points to the current bucket brigade. Therefore we can read the contents of the brigade right away. After that we pass the brigade to the next filter.

Finally we dump to `STDERR` the information about the type of the current mode, and the content of the bucket bridge.

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>
```

Notice that we use a virtual host because we want to install connection filters.

If we issue the following request:

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

We get the same response, when using `MyApache::FilterSnoop`, because our snooping filter didn't change anything. Though there was a lot of output printed to `error_log`. We present it all here, since it helps a lot to understand how filters work.

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
[]
```

as we saw earlier on the diagram, the connection input filter is run before the request input filter. Since our connection input filter was passing the data through unmodified and no other custom connection input filter was configured, the request input filter sees the same data. The last bucket in the brigade received by the request input filter is of type *EOS*, meaning that all the input data from the current request has been received.

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 time the brigade consists of a single bucket of type *TRANSIENT* which is allocated from the stack memory, which will eventually be converted to the *HEAP* type, before the body of the response is sent to the client.

When the content handler returns Apache sends the HTTP headers through connection output filters (notice that the request output filters don't see it):

```
>>> 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.

Also it's important to notice that the 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 Connection and HTTP Request input filters:

5.4.1 Connection Input Filters

Let's say that we want to test how our handlers behave when they are requested as HEAD requests, rather than GET. We can alter the request headers at the incoming connection level transparently to all handlers. So here is the input filter handler that does that by directly manipulating the bucket brigades:

```
file:MyApache/InputFilterGET2HEAD.pm
-----
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;

```

The filter handler is called for each bucket brigade, which in turn includes buckets with data. The gist of any filter handler is to retrieve the bucket brigade sent from the previous filter, prepare a new empty brigade, and move buckets from the former brigade to the latter optionally modifying the buckets on the way, which may include removing or adding new buckets. Of course if the filter doesn't want to modify any of the buckets it may decide to pass through the original brigade without doing any work.

In our example the handler first removes the bucket at the top of the brigade and looks at its type. If it sees an end of stream, that removed bucket is linked to the tail of the bucket brigade that will go to the next filter and it doesn't attempt to read any more buckets. If this event doesn't happen the handler reads the data from that bucket and if it finds that the data is of interest to us, it modifies the data, creates a new bucket using the modified data and links it to the tail of the outgoing brigade, while discarding the original bucket. In our case the interesting data is a such that matches the regular expression `/^GET/`. If the data is not interesting to the handler, it simply links the unmodified bucket to the outgoing brigade.

The handler looks for data like:

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

and turns it into:

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

For example, consider the following response handler:

```
file:MyApache/RequestType.pm
-----
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 and 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 really non-different from connection filters, other than that they are working on request and response bodies and have an access to a request object. The filter implementation is pretty much identical. Let's look at the request input filter that lowercases the request's body `MyApache::InputRequestFilterLC`:

```
file:MyApache/InputRequestFilterLC.pm
-----
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;

```

Now if we use the `MyApache::Dump` response handler, we have developed before in this chapter, which dumps the query string and the content body as a response, and configure the server as follows:

```

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

```

When issuing a POST request:

```
% echo "mOd_pErL 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. You can see that the query string wasn't changed.

5.5 Output Filters

mod_perl supports Connection and HTTP Request output filters:

5.5.1 Connection Output Filters

Connection filters filter **all** the data that is going through the server. Therefore if the connection is of HTTP request type, connection output filters see the headers and the body of the response, whereas request output filters see only the response body.

Since connection and request filters are very similar we will concentrate on HTTP request filters.

5.5.2 HTTP Request Output Filters

As mentioned earlier output filters can be written using the bucket brigades manipulation or the simplified stream-oriented interface.

First let's develop a response handler that send two lines of output: numerals 1234567890 and the English alphabet in a single string:

```
file:MyApache/SendAlphaNum.pm
-----
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 first filter implementation is using the stream-oriented filtering API:

```
file:MyApache/FilterReverse1.pm
-----
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;
```

Next, we 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
```

The `Apache::Filter` module loads the `read()` and `print()` methods which encapsulate the stream-oriented filtering interface.

The reversing filter is quite simple: in the loop it reads the data in the `readline()` mode in chunks up to the buffer length (1024 in our example), and then prints each line reversed while preserving the new line control characters at the end of each line. Behind the scenes `$filter->read()` retrieves the incoming brigade and gets the data from it, whereas `$filter->print()` appends to the new brigade which is then sent to the next filter in the stack. `read()` breaks the while loop, when the brigade is emptied or the end of stream is received.

In order not to distract the reader from the purpose of the example the used code is oversimplified and won't handle correctly input lines which are longer than 1024 characters and possibly using a different line termination pattern. So here is an example of a more complete handler, which does takes care of these issues:

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

    my $left_over = '';
    while ($filter->read(my $buffer, BUFF_LEN)) {
        $buffer = $left_over . $buffer;
        $left_over = '';
        while ($buffer =~ /([\r\n]*)([\r\n]*)/g) {
            $left_over = $1, last unless $2;
            $filter->print(scalar(reverse $1), $2);
        }
    }
    $filter->print(scalar reverse $left_over) if length $left_over;

    Apache::OK;
}
```

In this handler the lines longer than the buffer's length are buffered up in `$left_over` and processed only when the whole line is read in, or if there is no more input the buffered up text is flushed before the end of the handler.

5.5.2.2 Bucket Brigade-based Output Filters

The second filter implementation is using the bucket brigades API to accomplish exactly the same task as the first filter.

```
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.

The bucket brigades output filter version is just a bit more complicated than the stream-oriented one. The handler receives the incoming bucket brigade `$bb` as its second argument. Since when the handler is completed it must pass a brigade to the next filter in the stack, we create a new bucket brigade into which we are going to put the modified buckets and which eventually we pass to the next filter.

The core of the handler is in removing buckets from the head of the bucket brigade `$bb` while there are some, reading the data from the buckets, reversing and putting it into a newly created bucket which is inserted to the end of the new bucket brigade. If we see a bucket which designates the end of stream, we insert that bucket to the tail of the new bucket brigade and break the loop. Finally we pass the created brigade with modified data to the next filter and return.

6 HTTP Handlers

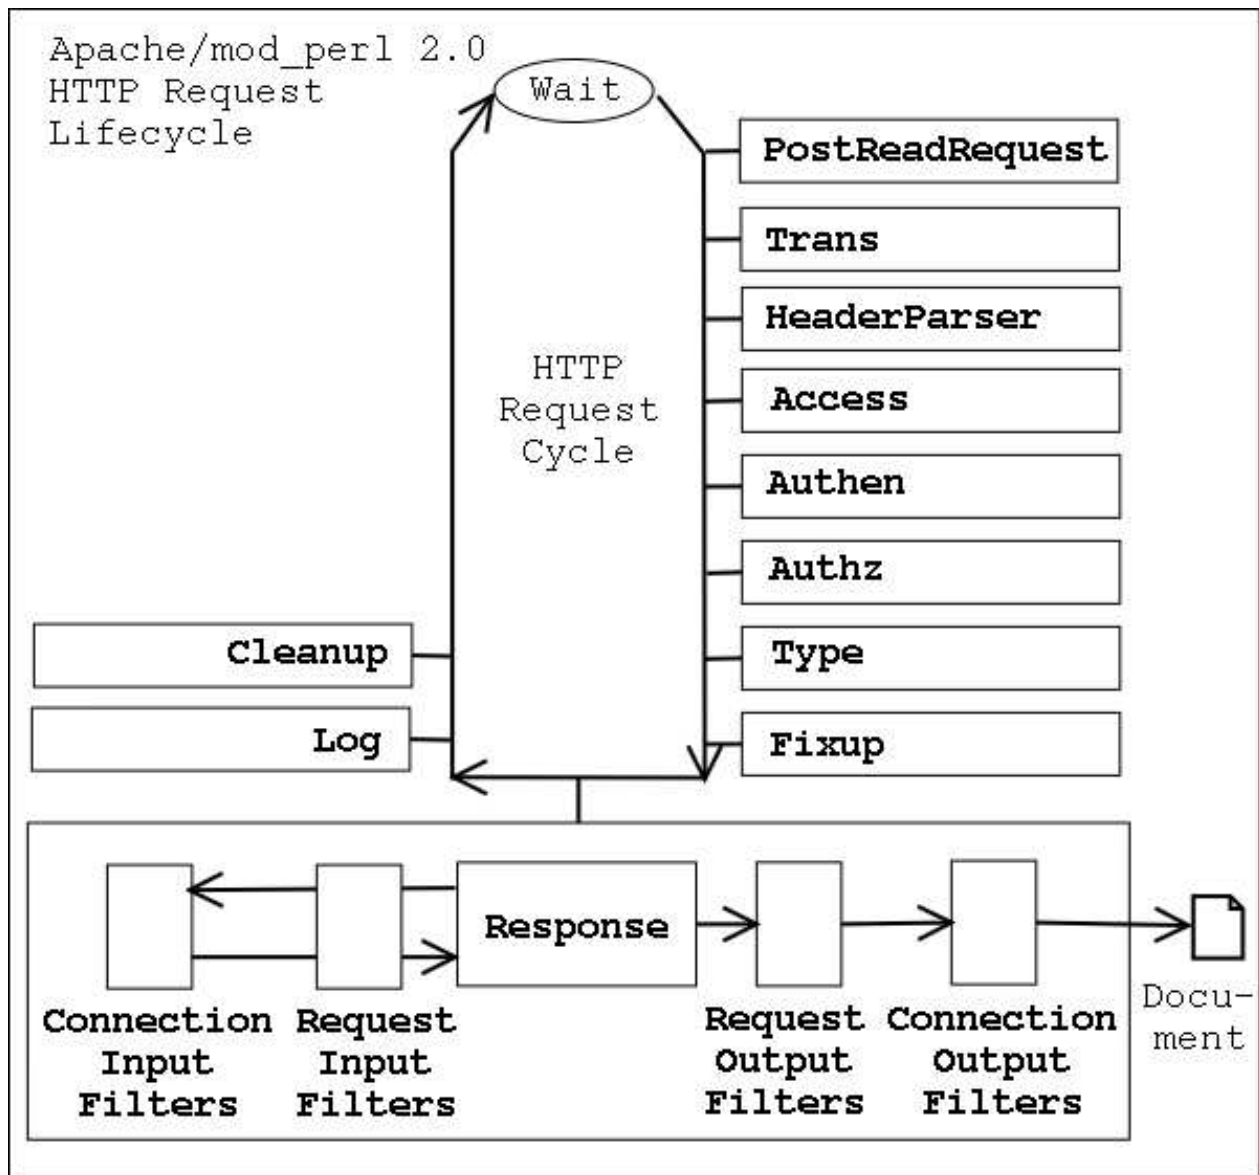
6.1 Description

This chapter explains how to implement the HTTP protocol handlers in `mod_perl`.

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*).

The following diagram depicts the HTTP request life cycle and highlights which handlers are available to `mod_perl` 2.0:



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**

It's possible that the cycle will not be completed if any of the phases terminates it, usually when an error happens.

Notice that when the response handler is reading the input data it can be filtered through request input filters, which are preceded by connection input filters if any. Similarly the generated response is first run through request output filters and eventually through connection output filters before it's sent to the client. We will talk about filters in detail later in this chapter.

Now let's discuss each of the mentioned handlers in detail.

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`, because at this phase the request has not yet been associated with a particular filename or directory.

Now, let's look at an example. Consider the following registry script:

```
touch.pl
-----
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);
```

This registry script is supposed to print when the last time *httpd.conf* has been modified, compared to the start of the request process time. If you run this script several times you might be surprised that it reports the same value all the time. Unless the request happens to be served by a recently started child process which will then report a different value. But most of the time the value won't be reported correctly.

This happens because the `-M` operator reports the difference between file's modification time and the value of a special Perl variable `$_T`. When we run scripts from the command line, this variable is always set to the time when the script gets invoked. Under `mod_perl` this variable is getting preset once when the child process starts and doesn't change since then, so all requests see the same time, when operators like `-M`, `-C` and `-A` are used.

Armed with this knowledge, in order to make our code behave similarly to the command line programs we need to reset `$_T` to the request's start time, before `-M` is used. We can change the script itself, but what if we need to do the same change for several other scripts and handlers? A simple `PerlPostReadRequestHandler` handler, which will be executed as the very first thing of each requests, comes handy here:

```
file:MyApache/TimeReset.pm
-----
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 to make things more efficient we use `$r->request_time` since the request object `$r` 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. A `PerlTransHandler` handler can alter the default translation mechanism or completely override it.

In addition to doing the translation, this stage can be 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`, because at this phase the request has not yet been associated with a particular filename or directory.

There are many useful things that can be performed at this stage. Let's look at the example handler that rewrites request URIs, similar to what `mod_rewrite` does. For example, if your web-site was originally made of static pages, and now you have moved to a dynamic page generation chances are that you don't want to change the old URIs, because you don't want to break links for those who link to your site. If the URI:

```
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:

```
file:MyApache/RewriteURI.pm
-----
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;
```

The handler matches the URI and assigns a new URI via `$r->uri()` and the query string via `$r->args()`. It then returns `Apache::DECLINED`, so the next translation handler will get invoked, if more rewrites and translations are needed.

Of course if you need to do a more complicated rewriting, this handler can be easily adjusted to do so.

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. For example this phase can be used to block evil clients targeting certain resources, while little resources were wasted so far.

This phase is of type `RUN_ALL`.

The handler's configuration scope is `DIR`.

This phase is very similar to `PerlPostReadRequestHandler`, with the only difference that it's run after the request has been mapped to the resource. Both phases are useful for doing something once per request, as early as possible. And usually you can take any `PerlPostReadRequestHandler` and turn it into `PerlHeaderParserHandler` by simply changing the directive name in *httpd.conf* and moving it inside the container where it should be executed. Moreover, because of this similarity `mod_perl` provides a special directive `PerlInitHandler` which if found outside resource containers behaves as `PerlPostReadRequestHandler`, otherwise as `PerlHeaderParserHandler`.

You already know that Apache handles the `HEAD`, `GET`, `POST` and several other HTTP methods. But did you know that you can invent your own HTTP method as long as there is a client that supports it. If you think of emails, they are very similar to HTTP messages: they have a set of headers and a body, sometimes a multi-part body. Therefore we can develop a handler that extends HTTP by adding a support for the `EMAIL` method. We can enable this protocol extension and push the real content handler during the `PerlHeaderParserHandler` phase:

```
<Location /email>
    PerlHeaderParserHandler MyApache::SendEmail
</Location>
```

and here is the `MyApache::SendEmail` handler:

```
file:MyApache/SendEmail.pm
-----
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);
    #warn $msg->as_string;
    $msg->send;
}

1;

```

Let's get the less interesting code out of the way. The function `content()` grabs the request body. The function `send_email()` sends the email over SMTP. You should adjust the constant `SMTP_HOSTNAME` to point to your outgoing SMTP server. You can replace this function with your own if you prefer to use a different method to send email.

Now to the more interesting functions. The function `handler()` returns immediately and passes the control to the next handler if the request method is not equal to `EMAIL` (set in the `METHOD` constant):

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

Next it tells Apache that this new method is a valid one and that the `perl-script` handler will do the processing. Finally it pushes the function `send_email_handler()` to the `PerlResponseHandler` list of handlers:

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

The function terminates the `header_parser` phase by:

```
return Apache::OK;
```

All other phases run as usual, so you can reuse any HTTP protocol hooks, such as authentication and fixup phases.

When the response phase starts `send_email_handler()` is invoked, assuming that no other response handlers were inserted before it. The response handler consists of three parts. Retrieve the email headers `To`, `From` and `Subject`, and the body of the message:

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

Then send the email:

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

Finally return to the client a simple response acknowledging that email has been sent and finish the response phase by returning `Apache::OK`:

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

Of course you will want to add extra validations if you want to use this code in production. This is just a proof of concept implementation.

As already mentioned 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";

```

most of the code is just a custom data. The code that does something consists of four lines at the very end. Create `HTTP::Headers` and `HTTP::Request` object. Issue the request and get the response. Finally print the response's content if it was successful or just *"failed"* if not.

Now save the client code in the file *send_http_email.pl*, adjust the *To* field, make the file executable and execute it, after you have restarted the server. You should receive an email shortly to the address set in the *To* field.

6.2.4 PerlInitHandler

When configured inside any container directive, except `<VirtualHost>`, this handler is an alias for `PerlHeaderParserHandler` described later. Otherwise it acts as an alias for `PerlPostRead-RequestHandler` 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.

The following example handler denies requests made from IPs on the blacklist.

```

file:MyApache/BlockByIP.pm
-----
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;

```

The handler retrieves the connection's IP address, looks it up in the hash of blacklisted IPs and forbids the access if found. If the IP is not blacklisted, the handler returns control to the next access checker handler, which may still block the access based on a different rule.

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 handler authenticates users by asking for a username and a password and lets them in only if the length of a string made from the supplied username and password and a single space equals to the secret length, specified by the constant `SECRET_LENGTH`.

```
file:MyApache/SecretLengthAuth.pm
-----
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;

```

First the handler retrieves the status of the authentication and the password in plain text. The status will be set to `Apache::OK` only when the user has supplied the username and the password credentials. If the status is different, we just let Apache handle this situation for us, which will usually challenge the client so it'll supply the credentials.

Once we know that we have the username and the password supplied by the client, we can proceed with the authentication. Our authentication algorithm is unusual. Instead of validating the username/password pair against a password file, we simply check that the string built from these two items plus a single space is `SECRET_LENGTH` long (14 in our example). So for example the pair *mod_perl/rules* authenticates correctly, whereas *secret/password* does not, because the latter pair will make a string of 15 characters. Of course this is not a strong authentication scheme and you shouldn't use it for serious things, but it's fun to play with. Most authentication validations simply verify the username/password against a database of valid pairs, usually this requires the password to be encrypted first, since storing passwords in clear is a bad idea.

Finally if our authentication fails the handler calls `note_basic_auth_failure()` and returns `Apache::AUTH_REQUIRED`, which sets the proper HTTP response headers that tell the client that its user that the authentication has failed and the credentials should be supplied again.

It's not enough to enable this handler for the authentication to work. You have to tell Apache what authentication scheme to use (`Basic` or `Digest`), which is specified by the `AuthType` directive, and you should also supply the `AuthName` -- the authentication realm, which is really just a string that the client usually uses as a title in the pop-up box, where the username and the password are inserted. Finally the `Require` directive is needed to specify which usernames are allowed to authenticate. If you set it to `valid-user` any username will do.

Here is the whole configuration section that requires users to authenticate before they are allowed to run the registry scripts from */perl/*:

```

<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`.

Here is the `MyApache::SecretResourceAuthz` handler which grants access to certain resources only to certain users who have already properly authenticated:

```
file:MyApache/SecretResourceAuthz.pm
-----
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;

```

This authorization handler is very similar to the authentication handler from the previous section. Here we rely on the previous phase to get users authenticated, and now as we have the username we can make decisions whether to let the user access the resource it has asked for or not. In our example we have a simple hash which maps which users are allowed to access what resources. So for example anything under */company/admin/* can be accessed only by the user *stas*, */company/report/* can be accessed by users *stas* and *boss*, whereas any other resources under */company/* can be accessed by everybody who has reached so far. If for some reason we don't get the username, or the user is not authorized to access the resource, the handler does the same thing as it does when the authentication fails, i.e. calls:

```

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

```

The configuration is similar to the one in the previous section, 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.

For example `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`.

The most important thing to remember when overriding the default *type_checker* handler, which is usually the `mod_mime` handler, is that you have to set the handler that will take care of the response phase and the response callback function or the code won't work. `mod_mime` does that based on `SetHandler` and `AddHandler` directives, and file extensions. So if you want the content handler to be run by `mod_perl`,

set either:

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

or:

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

depending on which type of response handler is wanted.

Writing a `PerlTypeHandler` handler which sets the content-type value and returns `Apache::DECLINED` so that the default handler will do the rest of the work, is not a good idea, because `mod_mime` will probably override this and other settings.

Therefore 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`.

The following fixup handler example 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.

```
file:MyApache/FileExtDispatch.pm
-----
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;

```

In the example we have used the following mapping.

```

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

```

So that *.cgi* requests will be handled by the *perl-script* handler and the *cgi_handler()* callback, *.pl* requests by *modperl* and *pl_handler()*, *.tt* (template toolkit) by *perl-script* and the *tt_handler()*, finally *.txt* request by the *default-handler* handler, which requires no callback.

Moreover the handler assumes that if the request's URI has no file extension or it does, but it's not in its mapping, the *default-handler* will be used, as if the *txt* extension was used.

After doing the mapping, the handler assigns the handler:

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

and the callback if needed:

```

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

```

In this simple example the callback functions don't do much but calling the same content handler which simply prints the name of the extension if handled by mod_perl, otherwise Apache will serve the other files using the default handler. In real world you will use callbacks to real content handlers that do real things.

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. For example:

```

<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.

Most of the Apache:: modules on CPAN are dealing with this phase. In fact most of the developers spend the majority of their time working on handlers that generate response content.

Let's write a simple response handler, that just generates some content. This time let's do something more interesting than printing "Hello world". Let's write a handler that prints itself:

```

file:MyApache/Deparse.pm
-----
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/deparse* we get the following response:

```

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

```

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

6.2.11 PerlLogHandler

The *log_transaction* phase happens no matter how the previous phases have ended up. If one of the earlier phases has aborted a request, e.g., failed authentication or 404 (file not found) errors, the rest of the phases up to and including the response phases are skipped. But this phase is always executed.

By this phase all the information about the request and the response is known, therefore the logging handlers usually record this information in various ways (e.g., logging to a flat file or a database).

This phase is of type `RUN_ALL`.

The handler's configuration scope is `DIR`.

Imagine a situation where you have to log requests into individual files, one per user. Assuming that all requests start with */users/username/*, so it's easy to categorize requests by the second URI path component. Here is the log handler that does that:

```

file:MyApache/LogPerUser.pm
-----
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;

```

First the handler tries to figure out what username the request is issued for, if it fails to match the URI, it simply returns `Apache::DECLINED`, letting other log handlers to do the logging. Though it could return `Apache::OK` since all other log handlers will be run anyway.

Next it builds the log entry, similar to the default `access_log` entry. It's comprised of remote IP, the current time, the uri, the return status and how many bytes were sent to the client as a response body.

Finally the handler appends this entry to the log file for the user the request was issued for. Usually it's safe to append short strings to the file without being afraid of messing up the file, when two files attempt to write at the same time, but just to be on the safe side the handler exclusively locks the file before performing the writing.

To configure the handler simply enable the module with the `PerlLogHandler` directive, inside the wanted section, which was `/users/` in our example:

```
<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 Description

This chapter explains how to migrate from mod_perl 1.0 to mod_perl 2.0.

7.2 Migrating from mod_perl 1.0 to mod_perl 2.0

The following sections discuss what should be done in order to migrate services from mod_perl 1.0 to 2.0 and if possible making the new services based on mod_perl 2.0 back compatible with mod_perl 1.0.

Several configuration directives were renamed or removed. Several APIs have changed, renamed, removed, or moved to new packages. Certain functions while staying exactly the same as in mod_perl 1.0, now reside in different packages. Before using them you need to find out and load those new packages containing them.

Since as of this writing mod_perl 2.0 wasn't released yet, it's possible that certain things have changed after this tutorial has been published. If something doesn't work as explained here, please refer to the documents in the mod_perl distribution or the online version at <http://perl.apache.org/docs/2.0/> for the updated documentation.

7.3 The Shortest Migration Path

mod_perl 2.0 provides two backwards-compatibility layers: one for the configuration files and the other for the code. If you are concerned to preserve the backwards compatibility with mod_perl 1.0, or simply want to try your services under mod_perl 2.0, simply enable the code compatibility layer by adding:

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

at the top of your startup file. The configuration backwards-compatibility is enabled by default.

7.4 Migrating Configuration Files

To migrate the configuration files to the mod_perl 2.0 syntax, you may need to do certain adjustments. Several configuration directives are deprecated in 2.0, but still available for backwards compatibility with mod_perl 1.0. If you don't need the backwards compatibility consider using the directives that have replaced them.

7.4.1 *PerlHandler*

`PerlHandler` was replaced with `PerlResponseHandler`.

7.4.2 *PerlSendHeader*

`PerlSendHeader` was replaced with `PerlOptions +/-ParseHeaders` directive.

```
PerlSendHeader On => PerlOptions +ParseHeaders
PerlSendHeader Off => PerlOptions -ParseHeaders
```

7.4.3 *PerlSetupEnv*

`PerlSetupEnv` was replaced with `PerlOptions +/-SetupEnv` directive.

```
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.

The default is *Off*. You cannot turn it *Off* once it's turned *On*.

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 and doesn't exist in `mod_perl 2.0`. A full tear-down and startup of interpreters is done on restart.

If you need to use the same `httpd.conf` for 1.0 and 2.0, use:

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

7.5 Code Porting

`mod_perl 2.0` is trying hard to be back compatible with `mod_perl 1.0`. However some things (mostly APIs) have been changed. In order to gain a complete compatibility with 1.0 while running under 2.0, you should load the compatibility module as early as possible:

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

at the server startup. And unless there are forgotten things or bugs, your code should work without any changes under 2.0 series.

However, unless you want to keep the 1.0 compatibility, you should try to remove the compatibility layer and adjust your code to work under 2.0 without it. You want to do it mainly for the performance improvement. The online mod_perl documentation includes a document (<http://perl.apache.org/docs/2.0/user/compat/compat.html>) that explains what APIs have changed and what new APIs should be used instead.

If you have mod_perl 1.0 and 2.0 installed on the same system and the two use the same perl libraries directory (e.g. `/usr/lib/perl5`), to use mod_perl 2.0 make sure to load first the `Apache2` module which will perform the necessary adjustments to `@INC`.

```
use Apache2; # if you have 1.0 and 2.0 installed
use Apache::compat;
```

So if before loading `Apache2.pm` the `@INC` array consisted of:

```
/home/stas/perl/ithread/lib/5.8.0/i686-linux-thread-multi
/home/stas/perl/ithread/lib/5.8.0
/home/stas/perl/ithread/lib/site_perl/5.8.0/i686-linux-thread-multi
/home/stas/perl/ithread/lib/site_perl/5.8.0
/home/stas/perl/ithread/lib/site_perl
.
```

It will now look as:

```
/home/stas/perl/ithread/lib/site_perl/5.8.0/i686-linux-thread-multi/Apache2
/home/stas/perl/ithread/lib/5.8.0/i686-linux-thread-multi
/home/stas/perl/ithread/lib/5.8.0
/home/stas/perl/ithread/lib/site_perl/5.8.0/i686-linux-thread-multi
/home/stas/perl/ithread/lib/site_perl/5.8.0
/home/stas/perl/ithread/lib/site_perl
.
```

Notice that a new directory was prepended to the search path, so if for example the code attempts to load `Apache::RequestRec` and there are two versions of this module under `/home/stas/perl/ithread/lib/site_perl/`:

```
5.8.0/i686-linux-thread-multi/Apache/RequestRec.pm
5.8.0/i686-linux-thread-multi/Apache2/Apache/RequestRec.pm
```

The mod_perl 2.0 version will be loaded first, because the directory `5.8.0/i686-linux-thread-multi/Apache2` is coming before the directory `5.8.0/i686-linux-thread-multi` in `@INC`.

7.6 ModPerl::Registry Family

In `mod_perl 2.0`, `Apache::Registry` and friends (`Apache::PerlRun`, `Apache::RegistryNG`, etc) have migrated into the `ModPerl::` namespace. The new family is based on the idea of `Apache::RegistryNG` from `mod_perl 1.0`, where you can customize pretty much all the functionality by providing your own hooks. The functionality of the modules `Apache::Registry`, `Apache::RegistryBB` and `Apache::PerlRun` hasn't changed from the user's perspective. All these modules are derived from the `ModPerl::RegistryCooker` class. So if you want to change the functionality of any of the existing sub-classes, or want to "cook" your own registry module it can be done easily. Refer to the `ModPerl::RegistryCooker` manpage for more information.

Here is a typical registry section configuration in `mod_perl 2.0`:

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

As we have explained earlier, the `ParseHeaders` option is needed if the headers are being sent via `print()` (i.e. without using `mod_perl` API) and comes as a replacement for the `PerlSendHeader` option in `mod_perl 1.0`.

Here is a simple registry script that prints the environment variables.

```
file:print_env.pl
-----
print "Content-type: text/plain\n\n";
for (sort keys %ENV){
    print "$_ => $ENV{$_}\n";
}
```

Save the file in `/home/httpd/perl/print_env.pl` and make it executable:

```
panic% chmod 0700 /home/stas/modperl/mod_perl_rules1.pl
```

Now issue a request to `http://localhost/perl/print_env.pl` and you should see all the environment variables printed out.

The only change for registry scripts from `mod_perl 1.0` is that Perl doesn't `chdir()`'s into the script's directory before executing it. This is because `chdir()` is not a thread-safe function, and as you've learned by now, `mod_perl 2.0` should run properly in the threaded environment. To accommodate for this change, the directory of the script is being pushed as the first element in `@INC` for the duration of the script's execution, so relative to the script's directory `require()` calls will succeed. This however doesn't solve the problem for other operations like `file open()` calls, when a relative to the script's directory path is used. In these cases the code needs to be changed to figure out the full path to the file at run time.

7.7 Method Handlers

In mod_perl 1.0 the method handlers could be specified by using the ($\$\$$) prototype:

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

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

Starting from Perl version 5.6, you can use subroutine attributes, and that's what mod_perl 2.0 does instead of conventional prototypes:

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

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

See the *attributes* manpage.

mod_perl 2.0 doesn't support the ($\$\$$) prototypes, mainly because several callbacks in 2.0 have more arguments than $\$r$, so the ($\$\$$) prototype doesn't make sense anymore. Therefore if you want your code to work with both mod_perl generations, you should use the subroutine attributes.

7.8 Apache::StatINC Replacement

Apache::StatINC has been replaced by Apache::Reload, which works for both mod_perl generations. To migrate to Apache::Reload simply replace:

```
PerlInitHandler Apache::StatINC
```

with:

```
PerlInitHandler Apache::Reload
```

However Apache::Reload provides an extra functionality, covered in the module's manpage.

8 That's all folks!

8.1 Thanks



8.2 References


All the information can be found at:

<http://perl.apache.org/docs/>

Further Questions?

Ask at modperl@perl.apache.org

8.3 A Shameless Plug



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

9.1 Description

This chapter presents the new features of Apache 2.0, Perl 5.6.0 - 5.8.0 and their influence on `mod_perl 2.0`.

9.2 Why `mod_perl`, the Next Generation

Here and in the rest of this document we refer to `mod_perl 1.x` series as `mod_perl 1.0` and, `2.0.x` as `mod_perl 2.0` to keep things simple. Similarly we call Apache 1.3.x series as Apache 1.3 and `2.0.x` as Apache 2.0

Since Doug MacEachern has introduced `mod_perl 1.0` in 1996, he had to adjust source code to the many changes Apache and Perl went through, while staying compatible with the older versions, leading to a very complex source code, with hundreds of `#ifdefs` and workarounds for various incompatibilities in older Perl and Apache versions. When Apache 2.0 development was underway, the new threads design was introduced, which couldn't be supported by the existing Perl version, since it required thread-safe Perl interpreters.

Think of it as a conspiracy or just a lucky coincidence, on March 10, 2002, the first Apache 2.0 alpha version was released. 13 days later, on March 23, 2002, Perl 5.6.0 has been released. And guess what, Perl 5.6.0 was the first Perl version to support the internal thread-safeness across multiple interpreters.

Since Perl 5.6.0 and Apache 2.0 were the very minimum requirements there was no need to support older version and it was a great idea to start `mod_perl 2.0` code base from scratch, incorporating the lessons learned during the 5 years of `mod_perl`'s existence.

The new version includes a mechanism for an automatic building of the Perl interface to Apache API, which allowed us to easily adjust `mod_perl 2.0` to ever changing Apache 2.0 API, during its development period. Another important feature is the `Apache::Test` framework, which was originally developed for `mod_perl 2.0`, but then was adopted by Apache 2.0 developers to test the core server features and third party modules. Moreover the tests written using the `Apache::Test` framework could be run with Apache 1.0 and 2.0, assuming that both supported the same features.

There are multiple other interesting changes that have already happened to `mod_perl` in version 2.0 and more will be developed in the future. Some of these will be covered in this document and some you will discover on your own while reading `mod_perl` documentation.

9.3 What's new in Apache 2.0

Apache 2.0 has introduces numerous new features and enhancements. Here are the most important new features:

- ***Apache Portable Runtime (APR)***

The APR presents a standard API for server applications, covering file I/O, logging, shared memory, threads, managing child processes and many other functionalities needed for developing the Apache

core and third party modules in a portable and effective way. One of the important effects is that it significantly simplifies the code that uses the APR making it much easier to review and understand the Apache code, increasing the number of revealed bugs and contributed patches.

The APR uses the concept of memory pools, which significantly simplifies the memory management code and reduces the possibility of having memory leaks, which always haunt C programmers.

- **I/O Filtering**

Apache 2.0 allows multiple modules to filter both the request and the response. Now one module can pipe its output as an input to another module as if another module was receiving the data directly from the TCP stream. The same mechanism works with the generated response.

With I/O filtering in place, things like SSL, data (de-)compression and other manipulations are done very easily.

The I/O filtering is based on the concept of bucket brigades and implemented in the APR.

- **Multi Processing Model modules (MPMs).**

In the previous Apache generation the same code base was trying to handle a management of incoming requests for different platforms, which lead to scalability problems on certain platforms, mainly on those which are different from Unix. This also lead to an undesired complexity of the code.

Apache 2.0 introduces the concept of Multi Processing Model modules, whose main responsibility is to map the incoming requests to either threads, processes or a threads/processes hybrid. Now it's possible to write different processing modules specific to various platforms. For example the Apache 2.0 on Windows is much more efficient now, since it uses *mpm_winnt* which deploys the native Windows features.

Here is a partial list of major MPMs available as of this writing.

- **prefork**

The *prefork* MPM emulates Apache 1.3's preforking model, where each request is handled by a different forked child process.

- **worker**

The *worker* MPM implements a hybrid multi-process multi-threaded approach based on the *pthreads* standard. It uses one acceptor thread, multiple worker threads.

- **mpmt_os2, netware, winnt and beos**

These MPMs also implement the hybrid multi-process/multi-threaded model, with each based on native OS thread implementations.

On platforms that support more than one MPM, it's possible to switch the used MPMs as the need change. For example on Unix it's possible to start with a preforked module. Then when the demand is growing and the code matures, it's possible to migrate to a more efficient threaded MPM, assuming that the code base is capable of running in the threaded environment.

- **New Hook Scheme**

In Apache 2.0 it's possible to dynamically register functions for each Apache hook, and allows more than one function to be registered per hook. Moreover when adding new functions, it's possible to specify where the new function should be added, e.g. a function can be pushed between two already registered functions or in front of them.

- **Protocol Modules**

The previous Apache generation could speak only the HTTP protocol. Apache 2.0 has introduced a "server framework" architecture making it possible to plug in handlers for protocols other than HTTP. The protocol module design also abstracts the transport layer so protocols such as SSL can be hooked into the server without requiring modifications to the Apache source code. This allows Apache to be extended much further than in the past, making it possible to add support for protocols such as FTP, SMTP, RPC flavors and the like. The main advantage being that protocol plugins can take advantage of Apache's portability, process/thread management, configuration mechanism and plugin API.

- **Parsed Configuration Tree**

Apache 2.0 makes the parsed configuration tree available at run time, so modules needing to read the configuration data (e.g., `mod_info`) don't have to re-parse the configuration file, but can re-use the parsed tree.

All these new features boost the Apache performance, scalability and flexibility. The APR helps the overall performance by doing lots of platform specific optimizations in the APR internals, and giving the developer the API which was already greatly optimized.

Apache 2.0 now includes special modules that can boost performance. For example the `mod_mmap_static` module loads webpages into the virtual memory and serves them directly avoiding the overhead of `open()` and `read()` system calls to pull them in from the filesystem.

The I/O layering is helping performance too, since now modules don't need to waste memory and CPU cycles to manually store the data in shared memory or *pages* in order to pass the data to another module, e.g., in order to provide response's *gzip* compression.

And of course a not least important impact of these features is the simplification and added flexibility for the core and third party Apache module developers.

9.4 What's new in Perl 5.6.0 - 5.8.0

As we have mentioned earlier Perl 5.6.0 is the minimum requirement for `mod_perl 2.0`. Though as we will see later certain new features work only with Perl 5.8.0 and higher.

These are the important changes in the recent Perl versions that had an impact on `mod_perl`. For a complete list of changes see the corresponding to the used version *perldelta* manpage.

The 5.6 Perl generation has introduced the following features:

- The beginnings of support for running multiple interpreters concurrently in different threads. In conjunction with the `perl_clone()` API call, which can be used to selectively duplicate the state of any given interpreter, it is possible to compile a piece of code once in an interpreter, clone that interpreter one or more times, and run all the resulting interpreters in distinct threads. See the *perlembed* and *perl561delta* manpages.
- The core support for declaring subroutine attributes, which is used by `mod_perl 2.0`'s *method handlers*. See the *attributes* manpage.
- The *warnings* pragma, which allows to force the code to be super clean, via the setting:

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

which will abort any code that generates warnings. This pragma also allows a fine control over what warnings should be reported. See the *perllexwarn* manpage.

- Certain `CORE::` functions now can be overridden via `CORE::GLOBAL::` namespace. For example `mod_perl` now can override `CORE::exit()` via `CORE::GLOBAL::exit`. See the *perlsub* manpage.
- The `XSLoader` extension as a simpler alternative to `DynaLoader`. See the *XSLoader* manpage.
- The large file support. If you have filesystems that support "large files" (files larger than 2 gigabytes), you may now also be able to create and access them from Perl. See the *perl561delta* manpage.
- Multiple performance enhancements were made. See the *perl561delta* manpage.
- Numerous memory leaks were fixed. See the *perl561delta* manpage.
- Improved security features: more potentially unsafe operations taint their results for improved security. See the *perlsec* and *perl561delta* manpages.
- Available on new platforms: GNU/Hurd, Rhapsody/Darwin, EPOC.

Overall multiple bugs and problems very fixed in the Perl 5.6.1, so if you plan on running the 5.6 generation, you should run at least 5.6.1. It is possible that when this tutorial is printed 5.6.2 will be out.

The Perl 5.8.0 has introduced the following features:

- The introduced in 5.6.0 experimental PerlIO layer has been stabilized and become the default IO layer in 5.8.0. Now the IO stream can be filtered through multiple layers. See the *perlpio* and *perliol* manpages.

For example this allows mod_perl to inter-operate with the APR IO layer and even use the APR IO layer in Perl code. See the *APR::PerlIO* manpage.

Another example of using the new feature is the extension of the `open()` functionality to create anonymous temporary files via:

```
open my $fh, "+>", undef or die $!;
```

That is a literal `undef()`, not an undefined value. See the `open()` entry in the *perlfunc* manpage.

- More overridable via `CORE:::GLOBAL:::` keywords. See the *perlsub* manpage.
- The signal handling in Perl has been notoriously unsafe because signals have been able to arrive at inopportune moments leaving Perl in inconsistent state. Now Perl delays signal handling until it is safe.
- `File::Temp` was added to allow a creation of temporary files and directories in an easy, portable, and secure way. See the *File::Temp* manpage.
- A new command-line option, `-t` is available. It is the little brother of `-T`: instead of dying on taint violations, lexical warnings are given. This is only meant as a temporary debugging aid while securing the code of old legacy applications. **This is not a substitute for `-T`**. See the *perlrun* manpage.

A new special variable `$_{^TAINT}` was introduced. It indicates whether taint mode is enabled. See the *perlvar* manpage.

- Threads implementation is much improved since 5.6.
- A much better support for Unicode.
- Numerous bugs and memory leaks fixed. For example now you can localize the tied `Apache::DBI` filehandles without leaking memory.
- Available on new platforms: AtheOS, Mac OS Classic, Mac OS X, MinGW, NCR MP-RAS, NonStop-UX, NetWare and UTS. The following platforms are again supported: BeOS, DYNIX/ptx, POSIX-BC, VM/ESA, z/OS (OS/390).

9.5 What's new in mod_perl 2.0

The new features introduced by Apache 2.0 and Perl 5.6 and 5.8 generations provide the base of the new mod_perl 2.0 features. In addition mod_perl 2.0 re-implements itself from scratch providing such new features as new build and testing framework. Let's look at the major changes since mod_perl 1.0.

9.5.1 *Threads Support*

In order to adapt to the Apache 2.0 threads architecture (for threaded MPMs), `mod_perl 2.0` needs to use thread-safe Perl interpreters, also known as "ithreads" (Interpreter Threads). This mechanism can be enabled at compile time and ensures that each Perl interpreter uses its private `PerlInterpreter` structure for storing its symbol tables, stacks and other Perl runtime mechanisms. When this separation is engaged any number of threads in the same process can safely perform concurrent callbacks into Perl. This of course requires each thread to have its own `PerlInterpreter` object, or at least that each instance is only accessed by one thread at any given time.

The first `mod_perl` generation has only a single `PerlInterpreter`, which is constructed by the parent process, then inherited across the forks to child processes. `mod_perl 2.0` has a configurable number of `PerlInterpreters` and two classes of interpreters, *parent* and *clone*. A *parent* is like that in `mod_perl 1.0`, where the main interpreter created at startup time compiles any pre-loaded Perl code. A *clone* is created from the parent using the Perl API `perl_clone()` function. At request time, *parent* interpreters are only used for making more *clones*, as the *clones* are the interpreters which actually handle requests. Care is taken by Perl to copy only mutable data, which means that no runtime locking is required and read-only data such as the syntax tree is shared from the *parent*, which should reduce the overall `mod_perl` memory footprint.

Rather than create a `PerlInterpreter` per-thread by default, `mod_perl` creates a pool of interpreters. The pool mechanism helps cut down memory usage a great deal. As already mentioned, the syntax tree is shared between all cloned interpreters. If your server is serving more than `mod_perl` requests, having a smaller number of `PerlInterpreters` than the number of threads will clearly cut down on memory usage. Finally and perhaps the biggest win is memory re-use: as calls are made into Perl subroutines, memory allocations are made for variables when they are used for the first time. Subsequent use of variables may allocate more memory, e.g. if a scalar variable needs to hold a longer string than it did before, or an array has new elements added. As an optimization, Perl hangs onto these allocations, even though their values "go out of scope". `mod_perl 2.0` has a much better control over which `PerlInterpreters` are used for incoming requests. The interpreters are stored in two linked lists, one for available interpreters and another for busy ones. When needed to handle a request, one interpreter is taken from the head of the available list and put back into the head of the same list when done. This means if for example you have 10 interpreters configured to be cloned at startup time, but no more than 5 are ever used concurrently, those 5 continue to reuse Perl's allocations, while the other 5 remain much smaller, but ready to go if the need arises.

The interpreters pool mechanism has been abstracted into an API known as "tipool", *Thread Item Pool*. This pool can be used to manage any data structure, in which you wish to have a smaller number than the number of configured threads. For example a replacement for `Apache::DBI` based on the *tipool* will allow to reuse database connections between multiple threads of the same process.

9.5.2 *Thread-environment Issues*

The only thing you have to worry about your code is that it's thread-safe and that you don't use functions that affect all threads.

Perl 5.8.0 itself is a thread-safe. That means that operations like `push()`, `map()`, `chomp()`, `=`, `/`, `+=`, etc. are thread-safe. Operations that involve system calls, may or may not be thread-safe. It all depends on whether the underlying C libraries used by the perl functions are thread-safe.

For example the function `localtime()` is not thread-safe when the implementation of `asctime(3)` is not thread-safe. Other usually problematic functions include `readdir()`, `srand()`, etc.

Another important issue that shouldn't be missed is what some people refer to as *thread-locality*. Certain functions executed in a single thread affect the whole process and therefore all other threads running inside that process. For example if you `chdir()` in one thread, all other thread now see the current working directory of that thread that `chdir()`'ed to that directory. Other functions with similar effects include `umask()`, `chroot()`, etc. Currently there is no cure for this problem. You have to find these functions in your code and replace them with different workarounds.

9.5.3 *Perl interface to the APR and Apache APIs*

As we have mentioned earlier, Apache 2.0 uses two APIs:

- the Apache Portable APR (APR) API, which implements a portable and efficient API to handle generically work with files, threads, processes, shared memory, etc.
- the Apache API, which handles issues specific to the web server.

`mod_perl 2.0` provides its own very flexible special purpose XS code generator, which is capable of doing things none of the existing generators can handle. It's possible that in the future this generator will be generalized and used for other projects of a high complexity.

This generator creates the Perl glue code for the public APR and Apache API, almost without a need for any extra code, but a few thin wrappers to make the API more Perl-ish.

In particular, since APR can be used outside of Apache, the Perl `APR::` modules can be used outside of Apache as well.

9.5.4 *Other New Features*

In addition to the already mentioned new features, the following are of a major importance:

- Apache 2.0 protocol modules are supported. Later we will see an example of a protocol module running on top of `mod_perl 2.0`.
- `mod_perl 2.0` provides a very simply to use interface to the Apache filtering API. We will present a filter module example later on.
- A feature-full and flexible `Apache::Test` framework was developed especially for `mod_perl` testing. While used to test the core `mod_perl` features, it is used by third-party module writers to easily test their modules. Moreover `Apache::Test` was adopted by Apache and currently used to test both Apache 1.3, 2.0 and other ASF projects. Anything that runs top of Apache can be tested with `Apache::Test`, be the target written in Perl, C, PHP, etc.

- The support of the new MPMs model makes mod_perl 2.0 can scale better on wider range of platforms. For example if you've happened to try mod_perl 1.0 on Win32 you probably know that the requests had to be serialized, i.e. only a single request could be processed at a time, rendering the Win32 platform unusable with mod_perl as a heavy production service. Thanks to the new Apache MPM design, now mod_perl 2.0 can be used efficiently on Win32 platforms using its native *win32* MPM.

9.5.5 Optimizations

The rewrite of mod_perl gives us the chances to build a smarter, stronger and faster implementation based on lessons learned over the 4.5 years since mod_perl was introduced. There are optimizations which can be made in the mod_perl source code, some which can be made in the Perl space by optimizing its syntax tree and some a combination of both. In this section we'll take a brief look at some of the optimizations that are being considered.

The details of these optimizations from the most part are hidden from mod_perl users, the exception being that some will only be turned on with configuration directives. A few of which include:

- "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 Description

This chapter provides an in-depth mod_perl 2.0 configuration details.

10.2 mod_perl configuration directives

Similar to mod_perl 1.0, in order to use mod_perl 2.0 a few configuration settings should be added to *httpd.conf*. They are quite similar to 1.0 settings but some directives were renamed and new directives were added.

10.3 Enabling mod_perl

To enable mod_perl built as DSO add to *httpd.conf*:

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

This setting specifies the location of the mod_perl module relative to the `ServerRoot` setting, therefore you should put it somewhere after `ServerRoot` is specified.

If mod_perl has been statically linked it's automatically enabled.

Win32 users need to make sure that the path to the Perl binary (e.g., *C:\Perl\bin*) is in the `PATH` environment variable.

10.4 Accessing the mod_perl 2.0 Modules

In order to prevent from inadvertently loading mod_perl 1.0 modules mod_perl 2.0 Perl modules are installed into dedicated directories under *Apache2/*. The `Apache2` module prepends the locations of the mod_perl 2.0 libraries to `@INC`, which are the same as the core `@INC`, but with *Apache2/* appended. This module has to be loaded just after mod_perl has been enabled. This can be accomplished with:

```
use Apache2 ();
```

in the startup file. Only if you don't use a startup file you can add:

```
PerlModule Apache2
```

to *httpd.conf*, due to the order the `PerlRequire` and `PerlModule` directives are processed.

10.5 Startup File

Next usually a startup file with Perl code is loaded:

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

It's used to adjust Perl modules search paths in @INC, pre-load commonly used modules, pre-compile constants, etc. Here is a typical *startup.pl* for mod_perl 2.0:

```
file:startup.pl
-----
use Apache2 ();

use lib qw(/home/httpd/perl);

# enable if the mod_perl 1.0 compatibility is needed
# use Apache::compat ();

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;
```

In this file the Apache2 modules is loaded, so the 2.0 modules will be found. Afterwards @INC is adjusted to include non-standard directories with Perl modules:

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

If you need to use the backwards compatibility layer load:

```
use Apache::compat ();
```

Next we preload the commonly used mod_perl 2.0 modules and precompile common constants.

Finally as usual the *startup.pl* file must be terminated with `1 ;`.

10.6 Perl's Command Line Switches

Now you can pass any Perl's command line switches in *httpd.conf* using the `PerlSwitches` directive. For example to enable warnings and Taint checking add:

```
PerlSwitches -wT
```

As an alternative to using `use lib` in *startup.pl* to adjust `@INC`, now you can use the command line switch `-I` to do that:

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

You could also use `-Mlib=/home/stas/modperl` which is the exact equivalent as `use lib`, but it's broken on certain platforms/version (e.g. Darwin/5.6.0). `use lib` is removing duplicated entries, whereas `-I` does not.

10.7 mod_perl 2.0 Handlers

mod_perl 2.0 provides two types of handlers: `modperl` and `perl-script`.

10.7.1 modperl

Configured as:

```
SetHandler modperl
```

The bare `mod_perl` handler type, which just calls the `Perl*Handler`'s callback function. If you don't need the features provided by the *perl-script* handler, with the `modperl` handler, you can gain even more performance. (This handler isn't available in `mod_perl 1.0`.)

Unless the `Perl*Handler` callback running under the `modperl` handler is configured with:

```
PerlOptions +SetupEnv
```

or calls:

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

in a void context (which has the same effect as `PerlOptions +SetupEnv` for the handler that called it), only the following environment variables are accessible via `%ENV`:

- `MOD_PERL` and `GATEWAY_INTERFACE` (always)
- `PATH` and `TZ` (if you had them defined in the shell or *httpd.conf*)

Therefore if you don't want to add the overhead of populating `%ENV`, when you simply want to pass some configuration variables from *httpd.conf*, consider using `PerlSetVar` and `PerlAddVar` instead of `PerlSetEnv` and `PerlPassEnv`. In your code you can retrieve the values using the `dir_config()` method. For example if you set in *httpd.conf*:

```
<Location /print_env2>
  SetHandler modperl
  PerlResponseHandler Apache::VarTest
  PerlSetVar VarTest VarTestValue
</Location>
```

this value can be retrieved inside `Apache::VarTest::handler()` with:

```
$r->dir_config('VarTest');
```

Alternatively use the Apache core directives `SetEnv` and `PassEnv`, which always populate `r->subprocess_env`, but this doesn't happen until the Apache *fixups* phase, which could be too late for your needs.

10.7.2 *perl-script*

Configured as:

```
SetHandler perl-script
```

Most `mod_perl` handlers use the *perl-script* handler. Among other things it does:

- `PerlOptions +GlobalRequest` is in effect unless:

```
PerlOptions -GlobalRequest
```

is specified.

- `PerlOptions +SetupEnv` is in effect unless:

```
PerlOptions -SetupEnv
```

is specified.

- `STDIN` and `STDOUT` get tied to the request object `$r`, which makes possible to read from `STDIN` and print directly to `STDOUT` via `CORE::print()`, instead of implicit calls like `$r->print()`.
- Several special global Perl variables are saved before the handler is called and restored afterwards (similar to `mod_perl 1.0`). This includes: `%ENV`, `@INC`, `$/`, `STDOUT`'s `$|` and `END` blocks array (`PL_endav`).

10.8 PerlOptions Directive

The directive `PerlOptions` provides fine-grained configuration for what were compile-time only options in the first `mod_perl` generation. It also provides control over what class of `PerlInterpreter` is used for a `<VirtualHost>` or location configured with `<Location>`, `<Directory>`, etc.

Options are enabled by prepending `+` and disabled with `-`. The options include:

10.8.1 *Enable*

On by default, can be used to disable `mod_perl` for a given `VirtualHost`. For example:

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

10.8.2 Clone

Share the parent Perl interpreter, but give the `VirtualHost` its own interpreter pool. For example should you wish to fine tune interpreter pools for a given virtual host:

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

This might be worthwhile in the case where certain hosts have their own sets of large-ish modules, used only in each host. By tuning each host to have its own pool, that host will continue to reuse the Perl allocations in their specific modules.

When cloning a Perl interpreter, to inherit base Perl interpreter's `PerlSwitches` use:

```
<VirtualHost ...>
  ...
  PerlSwitches +inherit
</VirtualHost>
```

10.8.3 Parent

Create a new parent Perl interpreter for the given `VirtualHost` and give it its own interpreter pool (implies the `Clone` option).

A common problem with `mod_perl 1.0` was the shared namespace between all code within the process. Consider two developers using the same server and each wants to run a different version of a module with the same name. This example will create two *parent* Perl interpreters, one for each `<VirtualHost>`, each with its own namespace and pointing to a different paths in `@INC`:

META: is `-Mlib` portable? (problems with `-Mlib` on Darwin/5.6.0?)

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

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

Or even for a given location, for something like "dirty" cgi scripts:

```
<Location /cgi-bin>
  PerlOptions +Parent
  PerlInterpMaxRequests 1
  PerlInterpStart 1
  PerlInterpMax 1
  PerlResponseHandler ModPerl::Registry
</Location>
```

will use a fresh interpreter with its own namespace to handle each request.

10.8.4 Perl*Handler

Disable Perl*Handlers, all compiled-in handlers are enabled by default. The option name is derived from the Perl*Handler name, by stripping the Perl and Handler parts of the word. So PerlLogHandler becomes Log which can be used to disable PerlLogHandler:

```
PerlOptions -Log
```

Suppose one of the hosts does not want to allow users to configure PerlAuthenHandler, PerlAuthzHandler, PerlAccessHandler and <Perl> sections:

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

Or maybe everything but the response handler:

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

10.8.5 AutoLoad

Resolve Perl*Handlers at startup time, which includes loading the modules from disk if not already loaded.

In mod_perl 1.0, configured Perl*Handlers which are not a fully qualified subroutine names are resolved at request time, loading the handler module from disk if needed. In mod_perl 2.0, configured Perl*Handlers are resolved at startup time. By default, modules are not auto-loaded during startup-time resolution. It is possible to enable this feature with:

```
PerlOptions +Autoload
```

Consider this configuration:

```
PerlResponseHandler Apache::Magick
```

In this case, `Apache::Magick` is the package name, and the subroutine name will default to *handler*. If the `Apache::Magick` module is not already loaded, `PerlOptions +AutoLoad` will attempt to pull it in at startup time. With this option enabled you don't have to explicitly load the handler modules. For example you don't need to add:

```
PerlModule Apache::Magick
```

in our example.

10.8.6 GlobalRequest

Setup the global `request_rec` for use with `Apache->request`. This setting is needed for example if you use `CGI.pm` to process the incoming request.

This setting is enabled by default for sections configured as:

```
<Location ...>
    SetHandler perl-script
    ...
</Location>
```

And can be disabled with:

```
<Location ...>
    SetHandler perl-script
    PerlOptions -GlobalRequest
    ...
</Location>
```

10.8.7 ParseHeaders

Scan output for HTTP headers, same functionality as `mod_perl 1.0's PerlSendHeader`, but more robust. This option is usually needs to be enabled for registry scripts which send the HTTP header with:

```
print "Content-type: text/html\n\n";
```

10.8.8 MergeHandlers

Turn on merging of `Perl*Handler` arrays. For example with a setting:

```
PerlFixupHandler Apache::FixupA

<Location /inside>
    PerlFixupHandler Apache::FixupB
</Location>
```

a request for */inside* only runs `Apache::FixupB` (`mod_perl 1.0` behavior). But with this configuration:

```
PerlFixupHandler Apache::FixupA

<Location /inside>
    PerlOptions +MergeHandlers
    PerlFixupHandler Apache::FixupB
</Location>
```

a request for */inside* will run both `Apache::FixupA` and `Apache::FixupB` handlers.

10.8.9 SetupEnv

Set up environment variables for each request ala `mod_cgi`.

When this option is enabled, `mod_perl` fiddles with the environment to make it appear as if the code is called under the `mod_cgi` handler. For example, the `$ENV{QUERY_STRING}` environment variable is initialized with the contents of `Apache::args()`, and the value returned by `Apache::server_hostname()` is put into `$ENV{SERVER_NAME}`.

But `%ENV` population is expensive. Those who have moved to the Perl Apache API no longer need this extra `%ENV` population, and can gain by disabling it. A code using the `CGI.pm` module require `PerlOptions +SetupEnv` because that module relies on a properly populated CGI environment table.

This option is enabled by default for sections configured as:

```
<Location ...>
    SetHandler perl-script
    ...
</Location>
```

Since this option adds an overhead to each request, if you don't need this functionality you can turn it off for a certain section:

```
<Location ...>
    SetHandler perl-script
    PerlOptions -SetupEnv
    ...
</Location>
```

or globally:

```
PerlOptions -SetupEnv
<Location ...>
    ...
</Location>
```

and then it'll affect the whole server. It can still be enabled for sections that need this functionality.

When this option is disabled you can still read environment variables set by you. For example when you use the following configuration:

```

PerlOptions -SetupEnv
PerlModule Modperl::Registry
<Location /perl>
  PerlSetEnv TEST hi
  SetHandler perl-script
  PerlHandler ModPerl::Registry
  Options +ExecCGI
</Location>

```

and you issue a request for this script:

```

setupenvoff.pl
-----
use Data::Dumper;
my $r = Apache->request();
$r->send_http_header('text/plain');
print Dumper(\%ENV);

```

you should see something like this:

```

$VAR1 = {
  'GATEWAY_INTERFACE' => 'CGI-Perl/1.1',
  'MOD_PERL' => 'mod_perl/2.0.1',
  'PATH' => 'bin:/usr/bin',
  'TEST' => 'hi'
};

```

Notice that we have got the value of the environment variable *TEST*.

10.9 Handlers Directives

See *The handlers chapters*.

10.10 Threads Mode Specific Directives

These directives are enabled only in a threaded mod_perl+Apache combo:

10.10.1 *PerlInterpStart*

The number of interpreters to clone at startup time.

10.10.2 *PerlInterpMax*

If all running interpreters are in use, mod_perl will clone new interpreters to handle the request, up until this number of interpreters is reached. when *PerlInterpMax* is reached, mod_perl will block (via *COND_WAIT()*) until one becomes available (signaled via *COND_SIGNAL()*).

10.10.3 PerlInterpMinSpare

The minimum number of available interpreters this parameter will clone interpreters up to `PerlInterpMax`, before a request comes in.

10.10.4 PerlInterpMaxSpare

`mod_perl` will throttle down the number of interpreters to this number as those in use become available.

10.10.5 PerlInterpMaxRequests

The maximum number of requests an interpreter should serve, the interpreter is destroyed when the number is reached and replaced with a fresh clone.

10.10.6 PerlInterpScope

As mentioned, when a request in a threaded mpm is handled by `mod_perl`, an interpreter must be pulled from the interpreter pool. The interpreter is then only available to the thread that selected it, until it is released back into the interpreter pool. By default, an interpreter will be held for the lifetime of the request, equivalent to this configuration:

```
PerlInterpScope request
```

For example, if a `PerlAccessHandler` is configured, an interpreter will be selected before it is run and not released until after the logging phase.

Interpreters will be shared across sub-requests by default, however, it is possible to configure the interpreter scope to be per-sub-request on a per-directory basis:

```
PerlInterpScope subrequest
```

With this configuration, an autoindex generated page, for example, would select an interpreter for each item in the listing that is configured with a `Perl*Handler`.

It is also possible to configure the scope to be per-handler:

```
PerlInterpScope handler
```

With this configuration, an interpreter will be selected before `PerlAccessHandlers` are run, and put back immediately afterwards, before Apache moves onto the authentication phase. If a `PerlFixupHandler` is configured further down the chain, another interpreter will be selected and again put back afterwards, before `PerlResponseHandler` is run.

For protocol handlers, the interpreter is held for the lifetime of the connection. However, a C protocol module might hook into `mod_perl` (e.g. `mod_ftp`) and provide a `request_rec` record. In this case, the default scope is that of the request. Should a `mod_perl` handler want to maintain state for the lifetime of an ftp connection, it is possible to do so on a per-virtualhost basis:

PerlInterpScope connection

10.11 mod_perl Directives Argument Types and Allowed Location

The following table shows where in the configuration files mod_perl configuration directives are allowed to appear, what kind and how many arguments they expect:

General directives:

Directive	Arguments	Scope
PerlSwitches	ITERATE	SRV
PerlRequire	ITERATE	SRV
PerlModule	ITERATE	SRV
PerlOptions	ITERATE	DIR
PerlSetVar	TAKE2	DIR
PerlAddVar	ITERATE2	DIR
PerlSetEnv	TAKE2	DIR
PerlPassEnv	TAKE1	SRV
<Perl> Sections	RAW_ARGS	SRV
PerlTrace	TAKE1	SRV

Handler assignment directives:

Directive	Arguments	Scope
PerlOpenLogsHandler	ITERATE	SRV
PerlPostConfigHandler	ITERATE	SRV
PerlChildInitHandler	ITERATE	SRV
PerlPreConnectionHandler	ITERATE	SRV
PerlProcessConnectionHandler	ITERATE	SRV
PerlPostReadRequestHandler	ITERATE	SRV
PerlTransHandler	ITERATE	SRV
PerlInitHandler	ITERATE	DIR
PerlHeaderParserHandler	ITERATE	DIR
PerlAccessHandler	ITERATE	DIR
PerlAuthenHandler	ITERATE	DIR
PerlAuthzHandler	ITERATE	DIR
PerlTypeHandler	ITERATE	DIR
PerlFixupHandler	ITERATE	DIR
PerlResponseHandler	ITERATE	DIR
PerlLogHandler	ITERATE	DIR
PerlCleanupHandler	ITERATE	DIR
PerlInputFilterHandler	ITERATE	DIR
PerlOutputFilterHandler	ITERATE	DIR

Perl Interpreter management directives:

Directive	Arguments	Scope
PerlInterpStart	TAKE1	SRV
PerlInterpMax	TAKE1	SRV
PerlInterpMinSpare	TAKE1	SRV
PerlInterpMaxSpare	TAKE1	SRV
PerlInterpMaxRequests	TAKE1	SRV
PerlInterpScope	TAKE1	DIR

mod_perl 1.0 back-compatibility directives:

Directive	Arguments	Scope
PerlHandler	ITERATE	DIR
PerlSendHeader	FLAG	DIR
PerlSetupEnv	FLAG	DIR
PerlTaintCheck	FLAG	SRV
PerlWarn	FLAG	SRV

The *Arguments* column represents the type of arguments directives accepts, where:

- **ITERATE**

Expects a list of arguments.

- **ITERATE2**

Expects one argument, followed by at least one or more arguments.

- **TAKE1**

Expects one argument only.

- **TAKE2**

Expects two arguments only.

- **FLAG**

One of On or Off (case insensitive).

- **RAW_ARGS**

The function parses the command line by itself.

The *Scope* column shows the location the directives are allowed to appear in:

- **SRV**

Global configuration and <VirtualHost> (mnemonic: *SeRVer*). These directives are defined as `RSRC_CONF` in the source code.

- **DIR**

<Directory>, <Location>, <Files> and all their regular expression variants (mnemonic: *DIRectory*). These directives can also appear in *.htaccess* files. These directives are defined as `OR_ALL` in the source code.

These directives can also appear in the global server configuration and <VirtualHost>.

Apache specifies other allowed location types which are currently not used by the core `mod_perl` directives and their definition can be found in *include/httpd_config.h* (hint: search for `RSRC_CONF`).

Also see *Single Phase's Multiple Handlers Behavior*.

10.12 Server Startup Options Retrieval

Inside *httpd.conf* one can do conditional configuration based on the define options passed at the server startup. For example:

```
<IfDefine PERLDB>
  <Perl>
    use Apache::DB ();
    Apache::DB->init;
  </Perl>

  <Location />
    PerlFixupHandler Apache::DB
  </Location>
</IfDefine>
```

So only when the server is started as:

```
% httpd C<-DPERLDB> ...
```

The configuration inside `IfDefine` will have an effect. If you want to have some configuration section to have an effect if a certain define wasn't defined use `!`, for example here is the opposite of the previous example:

```
<IfDefine !PERLDB>
  # ...
</IfDefine>
```

If you need to access any of the startup defines in the Perl code you use `Apache::exists_config_define`. For example in a startup file you can say:

```
use Apache::ServerUtil ();
if (Apache::exists_config_define("PERLDB")) {
  require Apache::DB;
  Apache::DB->init;
}
```

For example to check whether the server has been started in a single mode use:

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

10.12.1 MODPERL2 Define Option

When running under mod_perl 2.0 a special configuration define MODPERL2 is enabled internally, as if the server had been started with `-DMODPERL2`. For example this can be used to write a configuration file which needs to do something different whether it's running under mod_perl 1.0 or 2.0:

```
<IfDefine MODPERL2>
    # 2.0 configuration
</IfDefine>
<IfDefine !MODPERL2>
    # else
</IfDefine>
```

From within Perl code this can be tested with `Apache::exists_config_define()`, for example:

```
if (Apache::exists_config_define("MODPERL2")) {
    # some 2.0 specific code
}
```

10.13 Perl Interface to the Apache Configuration Tree

META:

Apache::Directive

t/response/TestApache/confree.pm

10.14 Debug Directives

10.14.1 PerlTrace [level]

set the trace level. This directive is enabled when mod_perl is compiled with the `MP_TRACE` option. level is either:

```
all
```

which sets maximum logging and debugging levels;

a combination of one or more option letters (or option numerical equivalents) from the following list:

```
d ( 1) directive processing
f ( 2) filters
g ( 4) Perl runtime interaction
h ( 8) handlers
i (16) interpreter pool management
m (32) memory allocations
s (64) perl sections
t (128) benchmark-ish timings
```

When `level` is not specified, the tracing level will be set to the value of the `MOD_PERL_TRACE` environment variable.

Table of Contents:

Tutorial Handouts: mod_perl 2.0 By Example	1
Getting Your Feet Wet with mod_perl	2
1 Getting Your Feet Wet with mod_perl	2
1.1 Description	3
1.2 Prerequisites	3
1.2.1 Downloading Stable Release Sources	3
1.2.2 Getting Bleeding Edge CVS Sources	4
1.2.3 Configuring and Installing Prerequisites	5
1.3 Installation	5
1.4 Configuration	6
1.5 Server Launch and Shutdown	6
1.6 Registry Scripts	7
1.7 Handler Modules	7
Introducing mod_perl Handlers	9
2 Introducing mod_perl Handlers	9
2.1 Description	10
2.2 Handler Anatomy	10
2.3 mod_perl Handlers Categories	11
2.4 Bucket Brigades	11
2.5 Single Phase's Multiple Handlers Behavior	13
Server Life Cycle Handlers	15
3 Server Life Cycle Handlers	15
3.1 Description	16
3.2 Server Life Cycle	16
3.2.1 Startup Phases Demonstration Module	17
3.2.2 PerlOpenLogsHandler	19
3.2.3 PerlPostConfigHandler	19
3.2.4 PerlChildInitHandler	20
Protocol Handlers	21
4 Protocol Handlers	21
4.1 Description	22
4.2 Connection Cycle Phases	22
4.2.1 PerlPreConnectionHandler	23
4.2.2 PerlProcessConnectionHandler	24
4.2.2.1 Socket-based Protocol Module	24
4.2.2.2 Bucket Brigades-based Protocol Module	26
Input and Output Filters	30
5 Input and Output Filters	30
5.1 Description	31
5.2 I/O Filtering	31
5.2.1 PerlInputFilterHandler	31
5.2.2 PerlOutputFilterHandler	31
5.2.3 Connection vs. HTTP Request Filters	31
5.3 All-in-One Filter	33

5.4	Input Filters	39
5.4.1	Connection Input Filters	39
5.4.2	HTTP Request Input Filters	42
5.5	Output Filters	44
5.5.1	Connection Output Filters	44
5.5.2	HTTP Request Output Filters	44
5.5.2.1	Stream-oriented Output Filter	45
5.5.2.2	Bucket Brigade-based Output Filters	46
	HTTP Handlers	49
6	HTTP Handlers	49
6.1	Description	50
6.2	HTTP Request Cycle Phases	50
6.2.1	PerlPostReadRequestHandler	52
6.2.2	PerlTransHandler	53
6.2.3	PerlHeaderParserHandler	55
6.2.4	PerlInitHandler	58
6.2.5	PerlAccessHandler	59
6.2.6	PerlAuthenHandler	60
6.2.7	PerlAuthzHandler	62
6.2.8	PerlTypeHandler	63
6.2.9	PerlFixupHandler	64
6.2.10	PerlResponseHandler	66
6.2.11	PerlLogHandler	67
6.2.12	PerlCleanupHandler	69
	mod_perl 1.0 to mod_perl 2.0 Migration	70
7	mod_perl 1.0 to mod_perl 2.0 Migration	70
7.1	Description	71
7.2	Migrating from mod_perl 1.0 to mod_perl 2.0	71
7.3	The Shortest Migration Path	71
7.4	Migrating Configuration Files	71
7.4.1	PerlHandler	71
7.4.2	PerlSendHeader	72
7.4.3	PerlSetupEnv	72
7.4.4	PerlTaintCheck	72
7.4.5	PerlWarn	72
7.4.6	PerlFreshRestart	72
7.5	Code Porting	72
7.6	ModPerl::Registry Family	74
7.7	Method Handlers	75
7.8	Apache::StatINC Replacement	75
	That's all folks!	76
8	That's all folks!	76
8.1	Thanks	77
8.2	References	77
8.3	A Shameless Plug	77

Bonus Chapter: What's new under the sun?	79
9 Bonus Chapter: What's new under the sun?	79
9.1 Description	80
9.2 Why mod_perl, the Next Generation	80
9.3 What's new in Apache 2.0	80
9.4 What's new in Perl 5.6.0 - 5.8.0	83
9.5 What's new in mod_perl 2.0	84
9.5.1 Threads Support	85
9.5.2 Thread-environment Issues	85
9.5.3 Perl interface to the APR and Apache APIs	86
9.5.4 Other New Features	86
9.5.5 Optimizations	87
Bonus Chapter: mod_perl 2.0 Server Configuration	88
10 Bonus Chapter: mod_perl 2.0 Server Configuration	88
10.1 Description	89
10.2 mod_perl configuration directives	89
10.3 Enabling mod_perl	89
10.4 Accessing the mod_perl 2.0 Modules	89
10.5 Startup File	89
10.6 Perl's Command Line Switches	90
10.7 mod_perl 2.0 Handlers	91
10.7.1 modperl	91
10.7.2 perl-script	92
10.8 PerlOptions Directive	92
10.8.1 Enable	92
10.8.2 Clone	93
10.8.3 Parent	93
10.8.4 Perl*Handler	94
10.8.5 AutoLoad	94
10.8.6 GlobalRequest	95
10.8.7 ParseHeaders	95
10.8.8 MergeHandlers	95
10.8.9 SetupEnv	96
10.9 Handlers Directives	97
10.10 Threads Mode Specific Directives	97
10.10.1 PerlInterpStart	97
10.10.2 PerlInterpMax	97
10.10.3 PerlInterpMinSpare	98
10.10.4 PerlInterpMaxSpare	98
10.10.5 PerlInterpMaxRequests	98
10.10.6 PerlInterpScope	98
10.11 mod_perl Directives Argument Types and Allowed Location	99
10.12 Server Startup Options Retrieval	101
10.12.1 MODPERL2 Define Option	102
10.13 Perl Interface to the Apache Configuration Tree	102
10.14 Debug Directives	102
10.14.1 PerlTrace [level]	102