

1 Debugging mod_perl C Internals

1.1 Description

This document explains how to debug C code under mod_perl, including mod_perl core itself.

For certain debugging purposes you may find useful to read first the following notes on mod_perl internals: Apache 2.0 Integration and mod_perl-specific functionality flow.

1.2 Debug notes

META: needs more organization (if you grok any of the following, patches are welcome)

META: there is a new compile-time option in perl-5.8.8+: -DDEBUG_LEAKING_SCALARS, which prints out the addresses of leaked SVs and new_SV() can be used to discover where those SVs were allocated. (see perlhack.pod for more info)

META: httpd has quite a lot of useful debug info: <http://httpd.apache.org/dev/debugging.html> (need to add this link to mp1 docs as well)

META: profiling: need a new entry of profiling. + running mod_perl under gprof: Defining GPROF when compiling uses the moncontrol() function to disable gprof profiling in the parent, and enable it only for request processing in children (or in one_process mode).

META: Jeff Trawick wrote a few useful debug modules, for httpd-2.1: mod_backtrace (similar to bt in gdb, but doesn't require the core file) and mod_whatkilledus (gives the info about the request that caused the segfault). http://httpd.apache.org/~trawick/exception_hook.html

1.2.1 Entering Single Server Mode

Most of the time, when debugging Apache or mod_perl, one needs to start Apache in a single server mode and not allow it to detach itself from the initial process. This is accomplished with:

```
% httpd -DONE_PROCESS -DNO_DETACH
```

1.2.2 Setting gdb Breakpoints with mod_perl Built as DSO

If mod_perl is built as a DSO module, you cannot set the breakpoint in the mod_perl source files when the *httpd* program gets loaded into the debugger. The reason is simple: At this moment *httpd* has no idea about mod_perl module yet. After the configuration file is processed and the mod_perl DSO module is loaded then the breakpoints in the source of mod_perl itself can be set.

The trick is to break at *apr_dso_load*, let it load *libmodperl.so*, then you can set breakpoints anywhere in the modperl code:

```
% gdb httpd
(gdb) b apr_dso_load
(gdb) run -DONE_PROCESS
[New Thread 1024 (LWP 1600)]
```

```
[Switching to Thread 1024 (LWP 1600)]

Breakpoint 1, apr_dso_load (res_handle=0xbffffb48c, path=0x811adcc
"/home/stas/apache.org/modperl-perlmodule/src/modules/perl/libmodperl.so",
pool=0x80e1a3c) at dso.c:138
141     void *os_handle = dlopen(path, RTLD_NOW | RTLD_GLOBAL);
(gdb) finish
...
Value returned is $1 = 0
(gdb) b modperl_hook_init
(gdb) continue
```

This example shows how to set a breakpoint at *modperl_hook_init*.

To automate things you can put those in the *.gdb-jump-to-init* file:

```
b apr_dso_load
run -DONE_PROCESS -d `pwd`/t -f `pwd`/t/conf/httpd.conf
finish
b modperl_hook_init
continue
```

and then start the debugger with:

```
% gdb /home/stas/httpd-2.0/bin/httpd -command \
`pwd`/t/.gdb-jump-to-init
```

1.2.3 Starting the Server Fast under gdb

When the server is started under gdb, it first loads the symbol tables of the dynamic libraries that it sees going to be used. Some versions of gdb may take ages to complete this task, which makes the debugging very irritating if you have to restart the server all the time and it doesn't happen immediately.

The trick is to set the `auto-solib-add` flag to 0:

```
set auto-solib-add 0
```

as early as possible in *~/gdbinit* file.

With this setting in effect, you can load only the needed dynamic libraries with *sharedlibrary* gdb command. Remember that in order to set a breakpoint and step through the code inside a certain dynamic library you have to load it first. For example consider this gdb commands file:

```
.gdb-commands
-----
file ~/httpd/prefork/bin/httpd
handle SIGPIPE pass
handle SIGPIPE nostop
set auto-solib-add 0
b ap_run_pre_config
run -d `pwd`/t -f `pwd`/t/conf/httpd.conf \
-DONE_PROCESS -DAPACHE2 -DPERL_USEITHREADS
sharedlibrary mod_perl
```

1.2.3 Starting the Server Fast under gdb

```
b modperl_hook_init
# start: modperl_hook_init
continue
# restart: ap_run_pre_config
continue
# restart: modperl_hook_init
continue
b apr_poll
continue

# load APR/PerlIO/PerlIO.so
sharedlibrary PerlIO
b PerlIOAPR_open
```

which can be used as:

```
% gdb -command=.gdb-commands
```

This script stops in *modperl_hook_init()*, so you can step through the *mod_perl* startup. We had to use the *ap_run_pre_config* so we can load the *libmodperl.so* library as explained earlier. Since *httpd* restarts on the start, we have to *continue* until we hit *modperl_hook_init* second time, where we can set the breakpoint at *apr_poll*, the very point where *httpd* polls for new request and run again *continue* so it'll stop at *apr_poll*. This particular script passes over *modperl_hook_init()*, since we run the *continue* command a few times to reach the *apr_poll* breakpoint. See the *Precooked gdb Startup Scripts* section for standalone script examples.

When *gdb* stops at the function *apr_poll* it's a time to start the client, that will issue a request that will exercise the server execution path we want to debug. For example to debug the implementation of *APR::Pool* we may run:

```
% t/TEST -run apr/pool
```

which will trigger the run of a handler in *t/response/TestAPR/pool.pm* which in turn tests the *APR::Pool* code.

But before that if we want to debug the server response we need to set breakpoints in the libraries we want to debug. For example if we want to debug the function *PerlIOAPR_open* which resides in *APR/PerlIO/PerlIO.so* we first load it and then we can set a breakpoint in it. Notice that *gdb* may not be able to load a library if it wasn't referenced by any of the code. In this case we have to load this library at the server startup. In our example we load:

```
PerlModule APR::PerlIO
```

in *httpd.conf*. To check which libraries' symbol tables can be loaded in *gdb*, run (when the server has been started):

```
gdb> info sharedlibrary
```

which also shows which libraries are loaded already.

Also notice that you don't have to type the full path of the library when trying to load them, even a partial name will suffice. In our commands file example we have used `sharedlibrary mod_perl` instead of saying `sharedlibrary mod_perl.so`.

If you want to set breakpoints and step through the code in the Perl and APR core libraries you should load their appropriate libraries:

```
gdb> sharedlibrary libperl
gdb> sharedlibrary libapr
gdb> sharedlibrary libaprutil
```

Setting *auto-solib-add* to 0 makes the debugging process unusual, since originally gdb was loading the dynamic libraries automatically, whereas now it doesn't. This is the price one has to pay to get the debugger starting the program very fast. Hopefully the future versions of gdb will improve.

Just remember that if you try to *step-in* and debugger doesn't do anything, that means that the library the function is located in wasn't loaded. The solution is to create a commands file as explained in the beginning of this section and craft the startup script the way you need to avoid extra typing and mistakes when repeating the same debugging process again and again.

Under threaded mpms (e.g. worker), it's possible that you won't be able to debug unless you tell gdb to load the symbols from the threads library. So for example if on your OS that library is called *libpthread.so* make sure to run:

```
sharedlibrary libpthread
```

somewhere after the program has started. See the Precooked gdb Startup Scripts section for examples.

Another important thing is that whenever you want to be able to see the source code for the code you are stepping through, the library or the executable you are in must have the debug symbols present. That means that the code has to be compiled with `-g` option for the gcc compiler. For example if I want to set a breakpoint in `/lib/libc.so`, I can do that by loading:

```
gdb> sharedlibrary /lib/libc.so
```

But most likely that this library has the debug symbols stripped off, so while gdb will be able to break at the breakpoint set inside this library, you won't be able to step through the code. In order to do so, recompile the library to add the debug symbols.

If debug code in response handler you usually start the client after the server was started, when doing this a lot you may find it annoying to need to wait before the client can be started. Therefore you can use a few tricks to do it in one command. If the server starts fast you can use `sleep()`:

```
% ddd -command=.debug-modperl-init & ; \
sleep 2 ; t/TEST -verbose -run apr/pool
```

or the Apache::Test framework's `-ping=block` option:

```
% ddd -command=.debug-modperl-init & ; \
t/TEST -verbose -run -ping=block apr/pool
```

which will block till the server starts responding, and only then will try to run the test.

1.2.4 Precooked gdb Startup Scripts

Here are a few startup scripts you can use with gdb to accomplish one of the common debugging tasks. To execute the startup script, simply run:

```
% gdb -command=.debug-script-filename
```

They can be run under gdb and any of the gdb front-ends. For example to run the scripts under ddd substitute gdb with ddd:

```
% ddd -command=.debug-script-filename
```

● Debugging mod_perl Initialization

The *code/.debug-modperl-init*:

```
# This gdb startup script breaks at the modperl_hook_init() function,
# which is useful for debug things at the modperl init phase.
#
# Invoke as:
# gdb -command=.debug-modperl-init
#
# see ADJUST notes for things that may need to be adjusted

# ADJUST: the path to the httpd executable if needed
file ~/httpd/worker/bin/httpd
handle SIGPIPE nostop
handle SIGPIPE pass
set auto-solib-add 0

define myrun
    tbreak main
    break ap_run_pre_config
    # ADJUST: the httpd.conf file's path if needed
    # ADJUST: add -DPERL_USEITHREADS to debug threaded mpms
    run -d `pwd`/t -f `pwd`/t/conf/httpd.conf -DONE_PROCESS -DAPACHE2
    continue
end

define modperl_init
    sharedlibrary mod_perl
    b modperl_hook_init
    continue
end

define sharedap
    # ADJUST: uncomment next line to debug threaded mpms
    #sharedlibrary libpthread
    sharedlibrary apr
```

```

        sharedlibrary aprutil
        #sharedlibrary mod_ssl.so
    continue
end

define sharedperl
    sharedlibrary libperl
end

# start the server and run till modperl_hook_init on start
myrun
modperl_init

# ADJUST: uncomment to reach modperl_hook_init on restart
#continue
#continue

# ADJUST: uncomment if you need to step through the code in apr libs
#sharedap

# ADJUST: uncomment if you need to step through the code in perlib
#sharedperl

startup script breaks at the modperl_hook_init() function, which is useful for debugging code
at the modperl's initialization phase.

```

- **Debugging mod_perl's Hooks Registration With httpd**

Similar to the previous startup script, the *code/.debug-modperl-register*:

```

# This gdb startup script allows to break at the very first invocation
# of mod_perl initialization, just after it was loaded. When the
# perl_module is loaded, and its pointer struct is added via
# ap_add_module(), the first hook that will be called is
# modperl_register_hooks().
#
# Invoke as:
# gdb -command=.debug-modperl-register
#
# see ADJUST notes for things that may need to be adjusted

define sharedap
    sharedlibrary apr
    sharedlibrary aprutil
    #sharedlibrary mod_ssl.so
end

define sharedperl
    sharedlibrary libperl
end

### Run ###

# ADJUST: the path to the httpd executable if needed
file ~/httpd/prefork/bin/httpd
handle SIGPIPE nostop

```

1.2.4 Precooked gdb Startup Scripts

```
handle SIGPIPE pass
set auto-solib-add 0

tbreak main

# assuming that mod_dso is compiled in
b load_module

# ADJUST: the httpd.conf file's path if needed
# ADJUST: add -DPERL_USEITHREADS to debug threaded mpms
run -d `pwd`/t -f `pwd`/t/conf/httpd.conf \
-DONE_PROCESS -DNO_DETACH -DAPACHE2

# skip over 'tbreak main'
continue

# In order to set the breakpoint in mod_perl.so, we need to get to
# the point where it's loaded.
#
# With static mod_perl, the bp can be set right away
#

# With DSO mod_perl, mod_dso's load_module() loads the mod_perl.so
# object and it immediately calls ap_add_module(), which calls
# modperl_register_hooks(). So if we want to bp at the latter, we need
# to stop at load_module(), set the 'bp modperl_register_hooks' and
# then continue.

# Assuming that 'LoadModule perl_module' is the first LoadModule
# directive in httpd.conf, you need just one 'continue' after
# 'ap_add_module'. If it's not the first one, you need to add as many
# 'continue' commands as the number of 'LoadModule foo' before
# perl_module, but before setting the 'ap_add_module' bp.
#
# If mod_perl is compiled statically, everything is already preloaded,
# so you can set modperl_* the breakpoints right away

b ap_add_module
continue

sharedlibrary mod_perl
b modperl_register_hooks
continue

#b modperl_hook_init
#b modperl_config_srv_create
#b modperl_startup
#b modperl_init_vhost
#b modperl_dir_config
#b modperl_cmd_load_module
#modperl_config_apply_PerlModule

# ADJUST: uncomment next line to debug threaded mpms
#sharedlibrary libpthread

# ADJUST: uncomment if you need to step through the code in apr libs
```

```
#sharedap

# ADJUST: uncomment if you need to step through the code in perlib
#sharedperl
```

startup script breaks at the `modperl_register_hooks()`, which is the very first hook called in the `mod_perl` land. Therefore use this one if you need to start debugging at an even earlier entry point into `mod_perl`.

Refer to the notes inside the script to adjust it for a specific *httpd.conf* file.

● Debugging mod_perl XS Extensions

The *code/.debug-modperl-xs*:

```
# This gdb startup script breaks at the mpxs_Apache__Filter_print()
# function from the XS code, as an example how you can debug the code
# in XS extensions.
#
# Invoke as:
# gdb -command=.debug-modperl-xs
# and then run:
# t/TEST -v -run -ping=block filter/api
#
# see ADJUST notes for things that may need to be adjusted

# ADJUST: the path to the httpd executable if needed
file /home/stas/httpd/worker/bin/httpd
handle SIGPIPE nostop
handle SIGPIPE pass
set auto-solib-add 0

define myrun
    tbreak main
    break ap_run_pre_config
    # ADJUST: the httpd.conf file's path if needed
    # ADJUST: add -DPERL_USEITHREADS to debug threaded mpms
    run -d `pwd`/t -f `pwd`/t/conf/httpd.conf \
    -DONE_PROCESS -DNO_DETACH -DAPACHE2
    continue
end

define sharedap
    # ADJUST: uncomment next line to debug threaded mpms
    #sharedlibrary libpthread
    sharedlibrary apr
    sharedlibrary aprutil
    #sharedlibrary mod_ssl.so
    continue
end

define sharedperl
```

```

        sharedlibrary libperl
    end

    define gopoll
        b apr_poll
        continue
        continue
    end

    define mybp
        # load Apache/Filter.so
        sharedlibrary Filter
        b mpxs_Apache__Filter_print
        # no longer needed and they just make debugging harder under threads
        disable 2
        disable 3
        continue
    end

    myrun
    gopoll
    mybp

    # ADJUST: uncomment if you need to step through the code in apr libs
    #sharedap

    # ADJUST: uncomment if you need to step through the code in perlib
    #sharedperl

```

startup script breaks at the `mpxs_Apache2__Filter_print()` function implemented in `xs/Apache2/Filter/Apache2__Filter.h`. This is an example of debugging code in XS Extensions. For this particular example the complete test case is:

```

% ddd -command=.debug-modperl-xs & \
  t/TEST -v -run -ping=block filter/api

```

When `filter/api` test is running it calls `mpxs_Apache2__Filter_print()` which is when the breakpoint is reached.

- **Debugging code in shared objects created by `Inline.pm`**

This is not strictly related to `mod_perl`, but sometimes when trying to reproduce a problem (e.g. for a p5p bug-report) outside `mod_perl`, the code has to be written in C. And in certain cases, `Inline` can be just the right tool to do it quickly. However if you want to interactively debug the library that it creates, it might get tricky. So similar to the previous sections, here is a `gdb code/.debug-inline`:

```

# save this file as .debug and execute this as:
# gdb -command=.debug
# or if you prefer gui
# ddd -command=.debug
#
# NOTE: Adjust the path to the perl executable
# also this perl should be built with debug enabled
file /usr/bin/perl

```

```

# If you need to debug with gdb a live script and not a library, you
# are going to have a hard time to set any breakpoint in the C code.
# the workaround is force Inline to compile and load .so, by putting
# all the code in the BEGIN {} block and call Inline->init from there.
#
# you also need to prevent from Inline deleting autogenerated .xs so
# you can step through the C source code, and of course you need to
# add '-g' so .so won't be stripped of debug info
#
# here is a sample perl script that can be used with this gdb script
#
# test.pl
# #-----#
# use strict;
# use warnings;
#
# BEGIN {
#     use Inline Config =>
#         #FORCE_BUILD => 1,
#         CLEAN_AFTER_BUILD => 0;
#
#     use Inline C => Config =>
#         OPTIMIZE => '-g';
#
#     use Inline C => <init;
#
# }
#
# my_bp();

tb main
# NOTE: adjust the name of the script that you run
run test.pl

# when Perl_runops_debug breakpoint is hit Inline will already load
# the autogenerated .so, so we can set the bp in it (that's only if
# you have run 'Inline->init' inside the BEGIN {} block

b S_run_body
continue
b Perl_runops_debug
continue

# here you set your breakpoints
b my_bp
continue

```

startup script that will save you a lot of time. All the details and a sample perl script are inside the gdb script.

1.3 Analyzing Dumped Core Files

When your application dies with the *Segmentation fault* error (which is generated by the SIGSEGV signal) and optionally generates a *core* file you can use `gdb` or a similar debugger to find out what caused the *Segmentation fault* (or a *segfault* as we often call it).

1.3.1 Getting Ready to Debug

In order to debug the *core* file you may need to recompile Perl and `mod_perl` with debugging symbols. Usually you have to recompile only `mod_perl`, but if the *core* dump happens in the *libperl.so* library and you want to see the whole backtrace, you need to recompile Perl as well. It may also occur inside `httpd` or 3rd party module, in which case you will need to recompile those. The following notes should help to accomplish the right thing:

- **mod_perl**

rebuild `mod_perl` with `MP_DEBUG=1`.

```
% perl Makefile.PL MP_DEBUG=1 ...
% make && make test && make install
```

Building `mod_perl` with `PERL_DEBUG=1` will:

1. add `-g` to `EXTRA_CFLAGS`
2. turn on `MP_TRACE` (tracing)
3. Set `PERL_DESTRUCT_LEVEL=2`
4. Link against *libperl.so* if `$Config{archlibexp}/CORE/libperl${Config{lib_ext}}` exists.

- **httpd**

If the *segfault* happens inside `ap_` or `apr_` calls, rebuild `httpd` with `--enable-maintainer-mode`:

```
% CFLAGS="-g" ./configure --enable-maintainer-mode ...
% make && make install
```

- **perl**

If the *segfault* happens inside `Perl_` calls, rebuild `perl` with `-Doptimize='-g'`:

```
% ./Configure -Doptimize='-g' ...
% make && make test && make install
```

Remember to recompile `mod_perl` if you've recompiled `perl`.

- **3rd party perl modules**

if the trace happens in one of the 3rd party perl modules, make sure to rebuild them, now that you've perl re-built with debugging flags. They will automatically pick the right compile flags from perl.

Now the software is ready for a proper debug.

1.3.2 *Causing a SegFault*

Most likely you already have the segfault situation, but sometimes you want to create one. For example sometimes you need to make sure that your system is configured to dump core files.

For that purpose you can use `Debug::DumpCore` available from CPAN:
<http://search.cpan.org/dist/Debug-FaultAutoBT/>

```
% perl -MDebug::DumpCore -eDebug::DumpCore::segv
Segmentation fault (core dumped)
```

Notice that you could use Perl's `CORE::dump` to achieve the same goal:

```
% perl -le 'dump'
Abort (core dumped)
```

but the generated in that case backtrace is not very useful for learning purposes. If all you want to test is whether your system is configured to dump core files then Perl's `CORE::dump` will do just fine.

1.3.3 *Getting the core File Dumped*

Now let's get the *core* file dumped from within the `mod_perl` server. Sometimes the program aborts abnormally via the `SIGSEGV` signal (*Segmentation Fault*), but no *core* file is dumped. And without the *core* file it's hard to find the cause of the problem, unless you run the program inside `gdb` or another debugger in first place. In order to get the *core* file, the application has to:

1. have the effective UID the same as real UID (the same goes for GID). Which is the case of `mod_perl` unless you modify these settings in the program.
2. be running from a directory which at the moment of the *Segmentation fault* is writable by the process that received this signal. Notice that the program might change its current directory during its run, so it's possible that the *core* file will need to be dumped in a different directory from the one the program was originally started from.

Under `Apache ServerRoot` is used as the default directory. Since that directory is usually not writable by the user running Apache, it's possible to use the directive `CoreDumpDirectory` (available since Apache 2.0.45) to tell Apache to dump the core file elsewhere.

3. be started from a shell process with sufficient resource allocations for the *core* file to be dumped. You can override the default setting from within a shell script if the process is not started manually. In addition you can use `BSD::Resource` to manipulate the setting from within the code as well.

You can use `ulimit` for `bash` and `limit` for `csch` to check and adjust the resource allocation. For example inside `bash`, you may set the core file size to unlimited:

```
panic% ulimit -c unlimited
```

or for `csch`:

```
panic% limit coredumpsize unlimited
```

For example you can set an upper limit on the *core* file size to 8MB with:

```
panic% ulimit -c 8388608
```

So if the core file is bigger than 8MB it will be not created.

4. Of course you have to make sure that you have enough disk space to create a big core file (`mod_perl` *core* files tend to be of a few MB in size).

Note that when you are running the program under a debugger like `gdb`, which traps the `SIGSEGV` signal, the *core* file will not be dumped. Instead it allows you to examine the program stack and other things without having the *core* file.

So let's write a simple script that uses `Debug::DumpCore`:

```
core_dump.pl
-----
use strict;
use warnings FATAL => 'all';

use Apache2::RequestRec ();
use Apache2::RequestIO ();
use Debug::DumpCore ();
use Cwd;

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

my $dir = getcwd();
$r->print("The core should be found at $dir/core.$$\\n");
$r->rflush;

Debug::DumpCore::segv();
```

In this script we load the `Apache2::RequestRec`, `Apache2::RequestIO`, `Debug::DumpCore` and `Cwd` modules, then we acquire the Apache request object and set the HTTP response header. Now we come to the real part -- we get the current working directory, print out the location of the *core* file that we are about to dump and finally we call `Debug::DumpCore::segv()` which dumps the *core* file.

Before we run the script we make sure that the shell sets the *core* file size to be unlimited, start the server in single server mode as a non-root user and generate a request to the script:

```
panic% cd /home/httpd/bin
panic% limit coredumpsize unlimited
panic% ./httpd -DONE_PROCESS -DNO_DETACH
# issue a request here
Segmentation fault (core dumped)
```

Our browser prints out:

```
The core should be found at /home/httpd/bin/core.12345
```

And indeed the core file appears where we were told it will be:

```
panic% ls -l /home/httpd/bin/core.12345
-rw----- 1 stas stas 13758464 Nov 23 18:33 /home/httpd/bin/core.12345
```

As you can see it's about 14MB *core* file. Notice that mod_perl was started as user *stas*, which had write permission for directory */home/httpd/bin*.

Notice that on certain platforms you get no PID digits appended to the core file name, so sometimes, it'll be just *core*.

1.3.4 Analyzing the core File

First we start gdb:

```
panic% gdb /home/httpd/bin/httpd /home/httpd/bin/core.12345
```

with the location of the mod_perl executable and the core file as the arguments.

To see the backtrace you run the *where* or the *bt* command:

```
(gdb) bt
#0 0x407ab26c in crash_now_for_real (
    suicide_message=0x407ad300 "Cannot stand this life anymore")
    at DumpCore.xs:10
#1 0x407ab293 in crash_now (
    suicide_message=0x407ad300 "Cannot stand this life anymore",
    attempt_num=42) at DumpCore.xs:17
#2 0x407ab39b in XS_Debug__DumpCore_segV (my_perl=0x86a9298, cv=0x8d36750)
    at DumpCore.xs:26
#3 0x40540649 in Perl_pp_entersub () from ../libperl.so
...
#7 0x404530cc in modperl_callback () from ../mod_perl.so
```

Well, you can see the last commands, but our perl and mod_perl are probably without the debug symbols. This is not the kind of trace you should send as a part of your bug report, because a lot of important information that should aid resolve the reported problem is missing.

Therefore the next step is to recompile Perl and mod_perl (and may be Apache) with debug symbols as explained earlier in this chapter.

Now when we repeat the process of starting the server, issuing a request and getting the core file, after which we run `gdb` again against the executable and the dumped `core.6789` file.

```
panic% gdb /home/httpd/bin/httpd /home/httpd/bin/core.6789
```

Now we can see the whole backtrace:

```
(gdb) bt
#0 0x407ab26c in crash_now_for_real (
    suicide_message=0x407ad300 "Cannot stand this life anymore")
    at DumpCore.xs:10
#1 0x407ab293 in crash_now (
    suicide_message=0x407ad300 "Cannot stand this life anymore",
    attempt_num=42) at DumpCore.xs:17
#2 0x407ab39b in XS_Debug__DumpCore_segv (my_perl=0x86a9298, cv=0x8d36750)
    at DumpCore.xs:26
#3 0x40540649 in Perl_pp_entersub (my_perl=0x86a9298) at pp_hot.c:2890
#4 0x4051ca4d in Perl_runops_debug (my_perl=0x86a9298) at dump.c:1449
#5 0x404c1ea3 in S_call_body (my_perl=0x86a9298, myop=0xbfffed90, is_eval=0)
    at perl.c:2298
#6 0x404c19cf in Perl_call_sv (my_perl=0x86a9298, sv=0x8cd0914, flags=4)
    at perl.c:2216
#7 0x404530cc in modperl_callback (my_perl=0x86a9298, handler=0x81ba6d8,
    p=0x8d16828, r=0x8d16860, s=0x813d238, args=0x8d018d8)
    at modperl_callback.c:102
#8 0x404539ce in modperl_callback_run_handlers (idx=6, type=4, r=0x8d16860,
    c=0x0, s=0x813d238, pconf=0x0, plog=0x0, ptemp=0x0,
    run_mode=MP_HOOK_RUN_FIRST) at modperl_callback.c:263
#9 0x40453c2d in modperl_callback_per_dir (idx=6, r=0x8d16860,
    run_mode=MP_HOOK_RUN_FIRST) at modperl_callback.c:351
#10 0x4044c728 in modperl_response_handler_run (r=0x8d16860, finish=0)
    at mod_perl.c:911
#11 0x4044cadb in modperl_response_handler_cgi (r=0x8d16860) at mod_perl.c:1006
#12 0x080db2bc in ap_run_handler (r=0x8d16860) at config.c:151
#13 0x080dba19 in ap_invoke_handler (r=0x8d16860) at config.c:363
#14 0x080a9953 in ap_process_request (r=0x8d16860) at http_request.c:246
#15 0x080a3ef8 in ap_process_http_connection (c=0x8d10920) at http_core.c:250
#16 0x080e7efc in ap_run_process_connection (c=0x8d10920) at connection.c:42
#17 0x080e82f8 in ap_process_connection (c=0x8d10920, csd=0x8d10848)
    at connection.c:175
#18 0x080d9b6d in child_main (child_num_arg=0) at prefork.c:609
#19 0x080d9c44 in make_child (s=0x813d238, slot=0) at prefork.c:649
#20 0x080d9d6a in startup_children (number_to_start=2) at prefork.c:721
#21 0x080da177 in ap_mpm_run (_pconf=0x81360a8, plog=0x817e1c8, s=0x813d238)
    at prefork.c:940
#22 0x080e0de8 in main (argc=11, argv=0xbffff284) at main.c:619
```

That's the perfect back trace to send as a part of the bug report.

Reading the trace from bottom to top, we can see that it starts with Apache calls, followed by `mod_perl` calls which end up in `modperl_callback()` which calls the Perl program via `Perl_call_sv`.

Notice that in our example we knew what script has caused the Segmentation fault. In a real world the chances are that you will find the *core* file without any clue to which of handler or script has triggered it. The special *curinfo* gdb macro comes to help:

For perl enabled with threads that's:

```
define curinfo
  printf "%d:%s\n", my_perl->Tcurcop->cop_line, \
    my_perl->Tcurcop->cop_file
end
```

For a non-threaded version that's:

```
define curinfo
  printf "%d:%s\n", PL_curcop->cop_line, \
    ((XPV*)(*(XPVGV*)PL_curcop->cop_filegv->sv_any)\
->xgv_gp->gp_sv->sv_any)->xpv_pv
end
```

Simply past the correct version at the gdb prompt (in this example the perl is threaded):

```
(gdb) define curinfo
Type commands for definition of "curinfo".
End with a line saying just "end".
> printf "%d:%s\n", my_perl->Tcurcop->cop_line, \
    my_perl->Tcurcop->cop_file
>end
```

and now we can call it:

```
(gdb) curinfo
No symbol "my_perl" in current context.
```

Oops, the function where the segfault has happened doesn't have the perl context, so we need to look at the backtrace and find the first function which accepts the *my_perl* argument (this is because we use a threaded perl). In this example this is the second frame:

```
#2 0x407ab39b in XS_Debug__DumpCore_segfv (my_perl=0x86a9298, cv=0x8d36750)
    at DumpCore.xs:26
```

therefore we need to go two frames up:

```
(gdb) up 2
#2 0x407ab39b in XS_Debug__DumpCore_segfv (my_perl=0x86a9298, cv=0x8d36750)
    at DumpCore.xs:26
26      in DumpCore.xs
```

and now we call *curinfo* again:

```
(gdb) curinfo
14:/home/httpd/cgi-bin/core_dump.pl
```

Et voilà, we can see that the segfault was triggered on line 14 of *core_dump.pl*, which has the line:

```
Debug::DumpCore::segv();
```

And we are done.

These are the bits of information that are important to extract and include in your bug report in order for us to be able to reproduce and resolve a problem. In this example it was the full backtrace, the filename and line where the faulty function was called (the faulty function is `Debug::DumpCore::segv()`) and the actual line where the Segmentation fault occurred (`crash_now_for_real` at `DumpCore.xs:10`).

1.3.5 Analyzing the core File Automatically

If the core file(s) are found in the `mod_perl` source directory, when running *t/REPORT* the core file backtraces will be automatically extracted and added to the report if the perl module `Devel::GDB` is installed.

See the function `dump_core_file()` in *Apache-Test/lib/Apache/TestReport.pm* if you want to see how it is invoked or refer to the `Devel::GDB` manpage.

1.3.6 Obtaining core Files under Solaris

There are two ways to get core files under Solaris. The first is by configuring the system to allow core dumps, the second is by stopping the process when it receives the SIGSEGV signal and "manually" obtaining the core file.

1.3.6.1 Configuring Solaris to Allow core Dumps

By default, Solaris 8 won't allow a setuid process to write a core file to the file system. Since apache starts as root and spawns children as 'nobody', core dumps won't produce core files unless you modify the system settings.

To see the current settings, run the `coreadm` command with no parameters and you'll see:

```
% coreadm
  global core file pattern:
    init core file pattern: core
      global core dumps: disabled
    per-process core dumps: enabled
  global setid core dumps: disabled
per-process setid core dumps: disabled
  global core dump logging: disabled
```

These settings are stored in the `/etc/coreadm.conf` file, but you should set them with the `coreadm` utility. As super-user, you can run `coreadm` with `-g` to set the pattern and path for core files (you can use a few variables here) and `-e` to enable some of the disabled items. After setting a new pattern, enabling global, global-setid, and log, and rebooting the system (reboot is required), the new settings look like:

```
% coreadm
  global core file pattern: /usr/local/apache/cores/core.%f.%p
  init core file pattern: core
    global core dumps: enabled
    per-process core dumps: enabled
  global setid core dumps: enabled
per-process setid core dumps: disabled
  global core dump logging: enabled
```

Now you'll start to see core files in the designated cores directory and they will look like *core.httpd.2222* where httpd is the name of the executable and the 2222 is the process id. The new core files will be read/write for root only to maintain some security, and you should probably do this on development systems only.

1.3.6.2 Manually Obtaining core Dumps

On Solaris the following method can be used to generate a core file.

1. Use truss(1) as *root* to stop a process on a segfault:

```
panic% truss -f -l -t \!all -s \!SIGALRM -S SIGSEGV -p <pid>
```

or, to monitor all httpd processes (from bash):

```
panic% for pid in `ps -eaf -o pid,comm | fgrep httpd | cut -d'/' -f1`;
do truss -f -l -t \!all -s \!SIGALRM -S SIGSEGV -p $pid 2>&1 &
done
```

The used truss(1) options are:

- `-f` - follow forks.
- `-l` - (that's an el) includes the thread-id and the pid (the pid is what we want).
- `-t` - specifies the syscalls to trace,
- `!all` - turns off the tracing of syscalls specified by `-t`
- `-s` - specifies signals to trace and the `!SIGALRM` turns off the numerous alarms Apache creates.
- `-S` - specifies signals that stop the process.
- `-p` - is used to specify the pid.

Instead of attaching to the process, you can start it under truss(1):

```
panic% truss -f -l -t \!all -s \!SIGALRM -S SIGSEGV \
  /usr/local/bin/httpd -f httpd.conf 2>&1 &
```

2. Watch the *error_log* file for reaped processes, as when they get SIGSEGV signals. When the process is reaped it's stopped but not killed.

3. Use `gcore(1)` to get a *core* of stopped process or attach to it with `gdb(1)`. For example if the process id is 662:

```
panic% gcore 662
gcore: core.662 dumped
```

Now you can load this *core* file in `gdb(1)`.

4. `kill -9` the stopped process. Kill the `truss(1)` processes as well, if you don't need to trap other segfaults.

Obviously, this isn't great to be doing on a production system since `truss(1)` stops the process after it dumps core and prevents Apache from reaping it. So, you could hit the clients/threads limit if you segfault a lot.

1.4 Debugging Threaded MPMs

1.4.1 Useful Information from *gdb* Manual

Debugging programs with multiple threads: http://sources.redhat.com/gdb/current/onlinedocs/gdb_5.html#SEC25

Stopping and starting multi-thread programs: http://sources.redhat.com/gdb/current/onlinedocs/gdb_6.html#SEC40

1.4.2 *libpthread*

when using:

```
set auto-solib-add 0
```

make sure to:

```
sharedlibrary libpthread
```

(or whatever the shared library is used on your OS) without which you may have problems to debug the threaded mpm `mod_perl`.

1.5 Defining and Using Custom *gdb* Macros

GDB provides two ways to store sequences of commands for execution as a unit: user-defined commands and command files. See: http://sources.redhat.com/gdb/current/onlinedocs/gdb_21.html

Apache 2.0 source comes with a nice pack of macros and can be found in *httpd-2.0/gdbinit*. To use it issue:

```
gdb> source /wherever/httpd-2.0/.gdbinit
```

Now if for example you want to dump the contents of the bucket brigade, you can do:

```
gdb> dump_brigade my_brigade
```

where `my_brigade` is the pointer to the bucket brigade that you want to debug.

mod_perl 1.0 has a similar file (*modperl/.gdbinit*) mainly including handy macros for dumping Perl data-structures, however it works only with non-threaded Perls. But otherwise it's useful in debugging mod_perl 2.0 as well.

1.6 Expanding C Macros

Perl, mod_perl and httpd C code makes an extensive use of C macros, which sometimes use many other macros in their definitions, so it becomes quite a task to figure out how to figure out what a certain macro expands to, especially when the macro expands to different values in different environments. Luckily there are ways to automate the expansion process.

1.6.1 Expanding C Macros with make

The mod_perl *Makefile*'s include a rule for macro expansions which you can find by looking for the `c.i.` rule. To expand all macros in a certain C file, you should run `make filename.i`, which will create *filename.i* with all macros expanded in it. For example to create *apr_perlio.i* with all macros used in *apr_perlio.c*:

```
% cd modperl-2.0/xs/APR/PerlIO
% make apr_perlio.i
```

the *apr_perlio.i* file now lists all the macros:

```
% less apr_perlio.i
# 1 "apr_perlio.c"
# 1 "<built-in>"
#define __VERSION__ "3.1.1 (Mandrake Linux 8.3 3.1.1-0.4mdk)"
...
```

1.6.2 Expanding C Macros with gdb

With gcc-3.1 or higher and gdb-5.2-dev or higher you can expand macros in gdb, when you step through the code. e.g.:

```
(gdb) macro expand pTHX_
expands to: PerlInterpreter *my_perl __attribute__((unused)),
(gdb) macro expand PL_dirty
expands to: (*Perl_Tdirty_ptr(my_perl))
```

For each library that you want to use this feature with you have to compile it with:

```
CFLAGS="-gdwarf-2 -g3"
```

or whatever is appropriate for your system, refer to the gcc manpage for more info.

To compile perl with this debug feature, pass `-Doptimize='-gdwarf-2 -g3'` to `./Configure`. For Apache run:

```
CFLAGS="-gdwarf-2 -g3" ./configure [...]
```

for `mod_perl` you don't have to do anything, as it'll pick the `$Config{optimize}` Perl flags automatically, if Perl is compiled with `-DDEBUGGING` (which is implied on most systems, if you use `-Doptimize='-g'` or similar.)

Notice that this will make your libraries **huge!** e.g. on Linux 2.4 Perl 5.8.0's normal `libperl.so` is about 0.8MB on linux, compiled with `-Doptimize='-g'` about 2.7MB and with `-Doptimize='-gdwarf-2 -g3'` 12.5MB. `httpd` is also becomes about 10 times bigger with this feature enabled. `mod_perl.so` instead of 0.2k becomes 11MB. You get the idea. Of course since you may want this only during the development/debugging, that shouldn't be a problem.

The complete details are at: http://sources.redhat.com/gdb/current/onlinedocs/gdb_10.html#SEC69

1.7 Maintainers

Maintainer is the person(s) you should contact with updates, corrections and patches.

Stas Bekman [<http://stason.org/>]

1.8 Authors

- Stas Bekman [<http://stason.org/>]

Only the major authors are listed above. For contributors see the Changes file.

Table of Contents:

| | | |
|---------|--|----|
| 1 | Debugging mod_perl C Internals | 1 |
| 1.1 | Description | 2 |
| 1.2 | Debug notes | 2 |
| 1.2.1 | Entering Single Server Mode | 2 |
| 1.2.2 | Setting gdb Breakpoints with mod_perl Built as DSO | 2 |
| 1.2.3 | Starting the Server Fast under gdb | 3 |
| 1.2.4 | Precooked gdb Startup Scripts | 6 |
| 1.3 | Analyzing Dumped Core Files | 12 |
| 1.3.1 | Getting Ready to Debug | 12 |
| 1.3.2 | Causing a SegFault | 13 |
| 1.3.3 | Getting the core File Dumped | 13 |
| 1.3.4 | Analyzing the core File | 15 |
| 1.3.5 | Analyzing the core File Automatically | 18 |
| 1.3.6 | Obtaining core Files under Solaris | 18 |
| 1.3.6.1 | Configuring Solaris to Allow core Dumps | 18 |
| 1.3.6.2 | Manually Obtaining core Dumps | 19 |
| 1.4 | Debugging Threaded MPMs | 20 |
| 1.4.1 | Useful Information from gdb Manual | 20 |
| 1.4.2 | libpthread | 20 |
| 1.5 | Defining and Using Custom gdb Macros | 20 |
| 1.6 | Expanding C Macros | 21 |
| 1.6.1 | Expanding C Macros with make | 21 |
| 1.6.2 | Expanding C Macros with gdb | 21 |
| 1.7 | Maintainers | 22 |
| 1.8 | Authors | 22 |