

mod_security

<http://www.modsecurity.org>

Copyright (c) 2002,2003 Ivan Ristic

Reference Manual v1.5.1

10 July 2003

Table of Contents

Reference Manual.....	1
Introduction.....	4
Licencing.....	4
Contact.....	4
Installation.....	4
CVS Access.....	5
Download.....	5
Nightly tarball.....	5
Releases.....	5
Installing from source.....	5
DSO.....	5
Static installation with Apache 1.x.....	6
Static installation with Apache 2.x.....	6
Installing from binary.....	7
Apache 1.x.....	7
Apache 2.x.....	7
Configuration.....	7
Turning filtering on and off.....	7
POST scanning.....	7
Default action.....	7
Filter inheritance.....	8
URL Encoding validation.....	8
Byte range check.....	8
Allowing others to see mod_security.....	8
Debugging.....	9
Request filtering.....	10
Simple filtering.....	10
Path normalisation.....	10
Regular expressions.....	10
Inverted expressions.....	11
Advanced filtering.....	11
Argument filtering exceptions.....	13
Actions.....	13
Specifying actions.....	13
Built-in actions.....	14
pass.....	14
deny.....	14
status.....	14
redirect.....	14
exec.....	14
log.....	14
nolog.....	14
Request headers added by mod_security.....	15
Making mod_security talk to your firewall.....	15
Common problems.....	15
Directory traversal.....	16
Cross site scripting attacks.....	16
SQL/database attacks.....	16
Operating system command execution.....	17

Buffer overflow attacks.....	17
Audit logging.....	17
Testing utility.....	18
Technology specific.....	19
PHP.....	19
Preventing register_global problems.....	19
Additional Examples.....	20
File upload.....	20
Securing FormMail.....	20

Introduction

ModSecurity is an open source intrusion detection and prevention engine for web applications. It operates embedded into the web server, acting as a powerful umbrella - shielding applications from attacks.

ModSecurity integrates with the web server, increasing your power to deal with web attacks. Some of its features worth mentioning are:

- **Request filtering**; incoming requests are analysed as they come in, and before they get handled by the web server or other modules.
- **Anti-evasion techniques**; paths and parameters are normalised before analysis takes place in order to fight evasion techniques.
- **Understanding of the HTTP protocol**; since the engine understands HTTP, it performs very specific and fine granulated filtering.
- **POST payload analysis**; the engine will intercept the contents transmitted using the POST method, too.
- **Audit logging**; full details of every request (including POST) can be logged for later analysis.
- **HTTPS filtering**; since the engine is embedded in the web server, it gets access to request data after decryption takes place.

Licensing

Mod_security is available under two licences. Users can choose to use the software under the terms of the GNU General Public License (<http://www.gnu.org/licenses/gpl.html>), as an Open Source / Free Software product. Alternatively, a commercial licence is available for distribution with closed-source products. For more information on commercial licencing please contact us.

Contact

ModSecurity is being developed by Ivan Ristic. Comments and questions are welcome. Please send your emails to ivanr@webkreator.com.

Installation

Before you begin with installation you will need to choose your preferred installation method. First you need to choose whether to install the latest mod_security version directly from CVS (latest features, but possibly unstable) or use the latest stable release (stable but with less features).

If you choose a stable release, it might be possible to install mod_security from binary, or compile it from the source code.

Finally, when compiling from source you need to choose a statically compiled module that

integrates with the web server or a shared library module that is loaded into the web server at run time.

The next few pages will give you more information on benefits of choosing one way over another.

CVS Access

If you want to access the up to date version of the module you need to get it through CVS. The list of changes made from the last release is normally available on the web site. CVS for mod_security is being hosted by SourceForge, <http://www.sf.net>. You can access it directly or view it through web using this address: <http://cvs.sourceforge.net/cgi-bin/viewcvs.cgi/mod-security/>

To download the source code to your computer you need to execute the following two commands:

```
cv$ -d:pserver:anonymous@cvs.sourceforge.net:/cvsroot/mod-security login
cv$ -z3 -d:pserver:anonymous@cvs.sourceforge.net:/cvsroot/mod-security co
mod_security
```

The first line will log you in as an anonymous user, and the second will download all files available in the mod_security repository.

Download

Nightly tarball

If you don't like CVS but you still want the latest version you can download the latest nightly tarball from the following address:

```
http://cvs.sourceforge.net/cvstarballs/mod-security-cvsroot.tar.gz
```

New features are added to mod_security one by one, with regression tests being run after each change. This should ensure that the version available from CVS is always usable.

Releases

You can download stable mod_security releases from the web site, go to <http://www.modsecurity.org/download/>. You may even be able to find binary releases for some platforms.

Installing from source

When installing from source you have two choices: to install the module into the web server itself, or to compile mod_security into a dynamic shared object (DSO).

DSO

Installing as DSO is easier, and the procedure is the same for both Apache branches:

1. Unpack mod_security source code
2. Compile the module into the web server:

```
/apachehome/bin/apxs -cia mod_security.c
```

3. Stop and start the server

```
/apachehome/bin/apachectl stop  
/apachehome/bin/apachectl start
```

Static installation with Apache 1.x

When a module is compiled statically, it gets embedded into the body of the web server. This method results in a slightly faster executable but the the compilation method is a bit more complicated, too:

1. Copy the file `mod_security.c` to `/src/modules/extra`
2. Configure Apache distribution with two additional configuration options:

```
--activate-module=src/modules/extra/mod_security  
--enable-module=security
```

3. Compile and install as usual

Static installation with Apache 2.x

At the moment it seems that there is no definitive documentation on how to compile a module into the body of an Apache 2.x web server. The procedure we use is as follows (valid for Apache 2.0.45, the last version at the time of writing):

1. Configure Apache instructing it to include `mod_security`:

```
./configure --prefix=(my prefix) --with-module=mappers:security
```

2. Copy the file `mod_security.c` to `modules/mappers/`
3. Edit `modules/mappers/modules.mk` and add the following to it:

```
mod_security.la: mod_security.lo  
    $(MOD_LINK) mod_security.lo
```

^ this is a TAB character

Also, add "`mod_security.la`" to the "`static =`" line:

```
static = mod_negotiation.la mod_security.la mod_dir.la ...
```

4. You can proceed to compile the server:

```
make
```

You can verify that the module was included in the binary by executing:

```
./httpd -l
```

The output should show `mod_security.c` on the list. If everything appears in order, continue with the installation:

```
make install
/apachehome/apachectl stop
/apachehome/apachectl start
```

Installing from binary

In some circumstances, you will want to install the module as a binary.

Apache 1.x

1. Copy the `mod_security.so` (on Unix) or the `mod_security.dll` (on Windows) file to `libexec/` and
2. Add the following line to the `httpd.conf` file

```
LoadModule security_module libexec/mod_security.so
```

Apache 2.x

1. Copy the `mod_security.so` (on Unix) or the `mod_security.dll` (on Windows) file to `modules/` and
2. Add the following line to the `httpd.conf` file

```
LoadModule security_module modules/mod_security.so
```

Configuration

Turning filtering on and off

Filtering engine is turned off by default. To use it, you need to turn it on:

```
SecFilterEngine On
```

POST scanning

POST payload scanning is turned off by default. To use it, you need to turn it on:

```
SecFilterScanPOST On
```

Default action

Whenever a filter is matched against a request, an action (or a series of actions) is taken. Individual filters can each have their own actions but in practice you will want to define a set of actions for all filters. You can do this with the configuration directive `SecFilterDefaultAction`. For example:

```
SecFilterDefaultAction "deny,log,status:404"
```

This directive accepts only one parameter, a series of actions separated with commas. Actions you specify here will be applied on every filter match, where individual actions were not defined.

Filter inheritance

Filters defined in parent folders are usually inherited by subfolders. This behavior is acceptable (and required) in most cases, but not all. Sometimes you need to relax checks in some part of the site. By using the following directive:

```
SecFilterInheritance Off
```

you can instruct `mod_security` to disregard parent filters so that you can start from the scratch.

URL Encoding validation

Special characters need to be encoded before they can be transmitted in the URL. Any character can be replaced using the three character combination `%XY`, where `XY` is an hexadecimal character code (see <http://www.rfc-editor.org/rfc/rfc1738.txt> for more details). Hexadecimal numbers only allow letters A to F, but attackers sometimes use other letters in order to trick the decoding algorithm. `Mod_security` checks all supplied encodings in order to verify they are valid.

You can turn URL encoding validation with the following line:

```
SecFilterCheckURLEncoding Off
```

Note: This directive will not check encoding in a POST payload when "multipart/form-data" encoding (file upload) is used.

Byte range check

You can force requests to consist only of bytes from a certain byte range. This can be useful to avoid stack overflow attacks (since they usually contain "random" binary content).

To only allow bytes from 32 to 126 (inclusive), use the following directive:

```
SecFilterForceByteRange 32 126
```

Default range values are 0 and 255, i.e. all byte values are allowed.

Note: This directive will not check byte range in a POST payload when "multipart/form-data" encoding (file upload) is used.

Allowing others to see `mod_security`

Normally, attackers won't be able to tell whether your web server is running `mod_security` or not. You can give yourself away by sending specific messages, or by using unusual HTTP codes (e.g. 406). If you want to stay hidden your best bet is to use HTTP 500, which stands for "Internal Server Error". Attackers that encounter such a response might think that your application has crashed.

There is another school of thought on this matter, which says that you should not hide the fact that you are running `mod_security`. The theory says that if they see it they will know you pay attention and that breaking into will be very difficult. And that they will go away looking for a weaker target. Or maybe they will become more determined and challenged.

By default Apache will return the information on itself with every request it serves. `Mod_security` keeps quiet by default, but you can allow others to see it by adding the following line to your configuration:

```
SecServerResponseToken On
```

The result will be similar to this:

```
[ivanr@wvx conf]$ telnet 0 8080
Trying 0.0.0.0...
Connected to 0.
Escape character is '^]'.
GET / HTTP/1.0

HTTP/1.1 406 Not Acceptable
Date: Mon, 19 May 2003 18:13:51 GMT
Server: Apache/2.0.45 (Unix) mod_security/1.5
Content-Length: 351
Connection: close
Content-Type: text/html; charset=iso-8859-1

<!DOCTYPE HTML PUBLIC "-//IETF//DTD HTML 2.0//EN">
<html><head>
<title>406 Not Acceptable</title>
</head><body>
<h1>Not Acceptable</h1>
<p>An appropriate representation of the requested resource / could not be found
on this server.</p>
<hr />
<address>Apache/2.0.45 (Unix) mod_security/1.4.2 Server at wxw.dyndns.org Port
80</address>
</body></html>

Connection closed by foreign host.
```

You should note that Apache itself supports two runtime directives, *ServerTokens* and *ServerSignature*. Using these directives you can completely hide the information on your server, no matter what you told `mod_security` to do.

Debugging

Use *SecFilterDebugLog* to choose a file where debug output will be written. If the parameter does not start with the forward slash, Apache home path will be prepended to it.

```
SecFilterDebugLog logs/modsec_debug_log
```

You can control how detailed the debug log is with *SecFilterDebugLevel*:

```
SecFilterDebugLevel 4
```

- 0 - none
- 1 - significant events (these will also be reported in the `error_log`)

- 2 - info messages
- 3 - more detailed info messages

Request filtering

When `mod_security` is On, every incoming request is intercepted and analysed prior to fulfillment. The analysis first performs a series of built-in checks in order to validate the request. These checks can be controlled using configuration directives, as you have already seen. The other part of the analysis consists from a series of user-defined filters that are matched against the request. Whenever there is a positive match, certain actions are taken.

Simple filtering

The most simplest form of filtering `mod_security` offers is, well, simple. It looks like this:

```
SecFilter KEYWORD
```

For each simple filter like this, `mod_security` will look for the keyword in the request. The search is pretty broad; it will be applied to the first line of the request (the one that looks like this "GET / index.php?parameter=value HTTP/1.0"). In case of POST requests, the body will be searched, too (provided you turn the `SecFilterScanPOST` on).

Path normalisation

Filter are not applied to raw request data, but on a normalised copy. We do this because attackers can (and do) apply various evasion techniques to avoid detection. For example, you might want to setup a filter that detects shell command execution:

```
SecFilter /bin/sh
```

But the attacker may use `/bin/./sh` and avoid the filter.

These transformations are applied in the current version of `mod_security`:

- Reduce `./` to `/`
- Verify URL encoding (optional)
- Only allow a byte range (optional)
- Decodes URL encoding
- Reduce `//` to `/`

Regular expressions

Remember the simple filtering facility, the one that looked like this:

```
SecFilter KEYWORD [ACTIONS]
```

There is a twist to it. When specifying the keyword you are not limited to a single word. Instead

you can put any regular expression that you want. To make most out of this (now) powerful tool you need to understand regular expressions well. I recommend that you start with one of the following resources:

- Perl Regular Expressions, <http://www.perldoc.com/perl5.6/pod/perlre.html>
- Mastering Regular Expressions, <http://www.oreilly.com/catalog/regex/>
- Google search on regular expressions, <http://www.google.com/search?q=regular%20expressions>

Inverted expressions

If exclamation mark is the first character of a regular expression, the filter will treat that regular expression as inverted. For example, the following:

```
SecFilter "!php"
```

will reject all requests that do not contain the word "php".

Advanced filtering

While *SecFilter* allows you to start quickly, you will soon discover that the search it performs is too broad, and doesn't work very well. Another directive,

```
SecFilterSelective LOCATION KEYWORD [ACTIONS]
```

allows you to choose exactly where you want the search to be performed. The KEYWORD and the ACTIONS bits are the same as in *SecFilter*. The LOCATION bit requires further explanation.

The LOCATION parameter consist of a series of location identifiers separated with a pipe. For example:

```
SecFilterSelective "REMOTE_ADDR|REMOTE_HOST" KEYWORD
```

will search only the IP address of the client and the host name. The list of possible location includes all CGI variables, and some more. Here is the full list:

- REMOTE_ADDR
- REMOTE_HOST
- REMOTE_USER
- REMOTE_IDENT
- REQUEST_METHOD
- SCRIPT_FILENAME
- PATH_INFO
- QUERY_STRING

- AUTH_TYPE
- DOCUMENT_ROOT
- SERVER_ADMIN
- SERVER_NAME
- SERVER_ADDR
- SERVER_PORT
- SERVER_PROTOCOL
- SERVER_SOFTWARE
- TIME_YEAR
- TIME_MON
- TIME_DAY
- TIME_HOUR
- TIME_MIN
- TIME_SEC
- TIME_WDAY
- TIME
- API_VERSION
- THE_REQUEST
- REQUEST_URI
- REQUEST_FILENAME
- IS_SUBREQ

There are some special locations:

- POST_PAYLOAD – filter the body of the POST request
- ARGS - filter arguments, the same as "QUERY_STRING|POST_PAYLOAD"
- ARGS_NAMES – variable/parameter names only
- ARGS_VALUES – variable/parameter values only

And even more special:

- HTTP_header – search request header "header"
- ENV_variable – search environment variable "variable"
- ARG_variable – search request variable/parameter "variable"

Argument filtering exceptions

The ARG_variable location names support inverted usage when used together with the ARG location. For example:

```
SecFilterSelective "ARGS|!ARG_param" KEYWORD
```

will search all arguments except the one named "param".

Actions

There are several types of actions:

- A primary action will make a decision whether to continue with the request or not. There can exist only one primary action. If you put several primary actions in the parameter, the last action to be seen will be executed. Primary actions are *deny*, *pass*, and *redirect*.
- Secondary actions will be performed on a filter match independently on the decision made by primary actions. There can be any number of secondary actions. One secondary action is *exec*.
- Flow actions can change the flow of rules, causing mod_security to jump to another rule, or to skip one or several rules. No flow actions exist at the moment.
- Parameters are not really actions, but a method of attaching parameters to filters. Some of this parameters can be used by real actions. For example, *status*, supplies the response code to the primary action *deny*.

Specifying actions

There are three places where you can put actions. One is the SecFilterDefaultAction directive, where you define actions you want executed on every filter match:

```
SecFilterDefaultAction "deny,log,status:500"
```

This example defines three examples separated using commas. First two actions consist of a single word. Third action requires a parameter. Use double colon to separate the parameter from the action name.

You can also specify per-filter actions. Both filtering directives accept a set of actions as an optional parameter. Examples of why you would want this can be found further in the text.

Built-in actions

pass

Allow request to continue on filter match. This action is useful when you want to log a filter match but otherwise do not want to take action.

```
SecFilter KEYWORD "pass,log"
```

deny

Interrupt request processing when a filter is matched. Unless *status* action is used, `mod_security` will immediately return a HTTP 500 error code. If a request is denied a header "mod_security-action" will be added to the list of request headers. It will contain the status code sent by `mod_security`.

status

Use the supplied HTTP status code when request is denied:

```
SecFilter KEYWORD "status:404"
```

will return a "Page not found" if triggered. Apache *ErrorDocument* directive works well here. If you have previously defined a custom error page for a given status then it will be displayed to the user.

redirect

Redirect the user to the given URL. For example:

```
SecFilter KEYWORD "redirect:http://www.modsecurity.org"
```

This configuration directive will always override HTTP status code, or the deny keyword.

exec

Execute a binary on filter match. Full path to the binary is required:

```
SecFilter KEYWORD "exec:/home/ivanr/report-attack.pl"
```

This directive does not effect the primary action (allow or deny). This action will always call script with no parameters, but providing all information in the environment. All the usual CGI environment variables will be there.

You can have one binary executed per filter match. Execution will add a header "mod_security-executed" to the list of request headers.

log

Log filter match to Apache error log.

nolog

Do not log filter match to Apache error log.

Request headers added by mod_security

Wherever possible, mod_security will add information to request headers, allowing your scripts to find and use them. Obviously, you will have to configure mod_security not to reject requests in order for your scripts to be executed at all.

This is the list of headers added:

- **mod_security-executed**; with the path to the binary executed
- **mod_security-action**; with the status code returned
- **mod_security-message**; the message about the problem detected, the same as the message added to the error log

Making mod_security talk to your firewall

In some cases, after detecting a particularly dangerous attack or a series of attacks you will want to prevent further attacks coming from the same source. You can do this by modifying the firewall to reject all traffic coming from a particular IP address.

This method can be very dangerous since it can result in a denial of service (DOS) attack. For example, an attacker can use a proxy to launch attacks. Rejecting all access from a proxy server can be very dangerous since all legitimate users will be rejected, too.

Since most proxies send information describing the original client (some information on this is available here <http://www.webkreator.com/cms/view.php/1685.html>, under the "Stop hijacking" header), we can try to be smart and find the real IP address. While this can work, consider the following scenario:

- The attacker is accessing the application directly but is pretending to be a proxy server, citing a random (or valid) IP address as the real source IP address. If we start rejecting requests based on that deduced information, the attacker will simply change the IP address and continue. As a result we might have banned legitimate users while the attacker is still free searching for application holes.

Therefore this method can be useful only if you do not allow access to the application through proxies, or allow access only through proxies that are well known and, more importantly, trusted.

If you still want to ban requests based on IP address (in spite of all our warnings), you will need to write a small script that will be executed on a filter match. The script should extract the IP address of the attacker from environment variables, and then make a call to iptables or ipchains to ban the IP address. We will include a sample script doing this with a future version of mod_security.

Other features

Chroot support

Standard approach

Starting with v1.5.1 `mod_security` includes support for Apache chrooting. Chrooting is a process of confining an application into a special part of the file system, sometimes called the "jail". Once the `chroot` call is performed, the application can no longer access what lies outside the jail. Only the root user can escape the jail, and a vital part of the chrooting process is not allowing anything root related (root processes or `suid` root binaries) inside the jail. The idea is that if an attacker manages to break in through the web server he won't have much to do because he, too, will be in jail, with no means to escape.

Applications do not have to support chrooting. Any application can be chrooted using the `chroot` binary. The following line:

```
chroot /chroot/apache /usr/local/web/bin/apachectl start
```

will start Apache but only after replacing the file system with what lies beneath `/chroot/apache`.

Unfortunately, things are not as simple as this. The problem is that applications typically require shared libraries, various files, and other binaries to function properly. So, to make them function you must make copies of required files and make them available inside the jail. This is no easy task (take a look at <http://penguin.epfl.ch/chroot.html> for detailed instructions on how to chroot an Apache web server).

The `mod_security` way

While I was chrooting an Apache the other day I realised that I was bored with the process and I started looking for ways to simplify it. As a result, I built the chrooting functionality into the `mod_security` module itself, making the whole process less complicated. With `mod_security` under your belt, you only need to add one line to the configuration file:

```
SecChrootDir /chroot/apache
```

and your web server will be chrooted successfully.

Apart from simplicity, `mod_security` chrooting brings another significant advantage. Unlike external chrooting (mentioned previously) `mod_security` chrooting *requires no additional files to exist in jail*. The `chroot` call is made after web server initialisation but before forking. Because of this, all shared libraries are already loaded, all web server modules are initialised, and log files are opened. You only need your data in jail.

There are some cases, however, when you will need additional files in jail, and that is if you intend to execute CGI scripts or system binaries. They may have their own file requirements. If you fall within this category then you need to proceed with the external `chroot` procedure as normal but you still won't have to think of the Apache itself.

Required module ordering for chroot support

As mentioned above, the chroot call must be performed at a specific moment in Apache initialisation, only after all other modules are initialised. This means that `mod_security` must be first on the list of modules. To ensure that, you will probably need to make some changes to module ordering, using the following configuration directives:

```
ClearModuleList
AddModule mod_security.c
AddModule ...
AddModule ...
AddModule ...
```

The first directive clears the list. You must put `mod_security` next, followed by all other modules you intend to use (except `http_core.c`, which is always automatically added and you do not have to worry about it). You can find out the list of built-in modules by executing the `httpd` binary with the `-l` switch:

```
./httpd -l
```

Solving common security problems

As an example of `mod_security` capabilities we will demonstrate how you can use it to detect and prevent the most common security problems. We won't go into detail here about problems themselves but a very good description is available in the Open Web Application Security Project's guide, available at <http://www.owasp.org>.

Directory traversal

If your scripts are dealing with the file system then you need to pay attention to certain meta characters and constructs. For example, a character combination `"../"` in a path is a request to go up one directory level.

In normal operation there is no need for this character combination to occur in requests and you can forbid them with the following filter:

```
SecFilter "\.\/"
```

Cross site scripting attacks

Cross site scripting attacks (XSS) occur when an attacker injects HTML or/and Javascript code into your Web pages and then that code gets executed by other users. This is usually done by adding HTML to places where you would not expect them. A successful XSS attack can result in the attacker obtaining the cookie of your session and gaining full access to the application!

Proper defense against this attack is parameter filtering (and thus removing the offending HTML/Javascript) but often you must protect existing applications without changing them. This can be done with one of the following filters:

```
SecFilter "<( |\n)*script"
SecFilter "<(.\n)+>"
```

The first filter will protect only against Javascript injection with the "<script>" tag. The second filter is more general, and disallows any HTML code in parameters.

You need to be careful when applying filters like this since many application want HTML in parameters (e.g. CMS applications, forums, etc). You can this with selective filtering. For example, you can have the second filter from above as a general site-wide rule, but later relax rules for a particular script with the following code:

```
<Location /cms/article-update.php>
  SecFilterInheritance Off
  # other filters here ...
  SecFilterSelective "ARGS|!ARG_body" "<(.\|\\n)+>"
</Location>
```

This code fragment will only accept HTML in a named parameter "body". In reality you will probably add a few more named parameters to the list.

SQL/database attacks

Most Web applications nowadays rely heavily on databases for data manipulation. Unless great care is taken to perform database access safely, an attacker can inject arbitrary SQL commands directly into the database. This can result in the attacker reading sensitive data, changing it, or even deleting it from the database altogether.

Filters like:

```
SecFilter "delete[[:space:]]+from"
SecFilter "insert[[:space:]]+into"
SecFilter "select.+from"
```

can protect you from most SQL-related attacks. These are only examples, you need to craft your filters carefully depending on the actual database engine you use.

Operating system command execution

Web applications are sometimes written to execute operating system commands to perform operations. A persistent attacker may find a hole in the concept, allowing him to execute arbitrary commands on the system.

A filter like this:

```
SecFilter "bin/"
```

will detect attempts to execute binaries residing in various folders on a Unix-related operating system.

Buffer overflow attacks

Buffer overflow is a technique of overflowing the execution stack of a program and adding assembly instructions in an attempt to get them executed. In some circumstances it may be possible to prevent these types of attack by using the line similar to:

as it will only accept requests that consists of bytes from this range. Whether you use this type of protection or not depends on your application and the used character encoding.

Audit logging

Standard Apache logging will not help much if you need to trace back steps of a particular user or an attacker. The problem is that only a very small subset of each request is written to a log file. This problem can be remedied with the audit logging feature of `mod_security`. These two directives:

```
SecAuditEngine On
SecAuditLog logs/audit_log
```

will let `mod_security` know that you want a full audit log stored into the log file `audit_log`. Here is an example of how a request is logged:

```
=====
Request: 192.168.0.2 - - [[18/May/2003:11:20:43 +0100]] "GET /cgi-
bin/printenv?p1=666 HTTP/1.0" 406 822
Handler: cgi-script
-----
GET /cgi-bin/printenv?p1=666 HTTP/1.0
Host: wkx.dyndns.org:8080
User-Agent: mod_security regression test utility
Connection: Close
mod_security-message: Access denied with code 406. Pattern match "666" at
ARGS_SELECTIVE.
mod_security-action: 406

HTTP/1.0 406 Not Acceptable
=====
```

You can see that on the first line you get what you normally get from Apache. The second line contains the name of the handler that was supposed to handle the request. Full request (with additional `mod_security` headers) is given after the separator, and the response headers (in this case there is only one line) is given after one empty line.

When POST filtering is on, the POST payload will always be included in the audit log. Actual response will never be included (at least not in this version).

At this time, the audit logging part of the module will log Apache 1.x error messages, on the line below the "Handler:" line. The line will always begin with "Error:". This functionality will be added to the Apache 2.x version of the module if possible.

Choosing what to log

Starting with v1.5.1 the `SecAuditEngine` parameter accepts one of four values:

- On – log all requests

- Off – do not log requests at all
- RelevantOnly – only log relevant requests. Relevant requests are those requests that caused a filter match.
- DynamicOrRelevant – log dynamically generated or relevant requests. A request is considered dynamic if its handler is not null.

Getting mod_security to log dynamic requests could require a little bit of work. In theory, a response to a request is generated by a handler, and if there is a handler attached to a request it can be considered to be of a dynamic nature. In practice, however, Apache can be configured to server dynamic pages without a handler (it will choose the module based on the mime type). This will happen, for example, if you configure PHP as instructed in the main distribution:

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

While this works, it won't serve our purpose here. However, if you replace the above line with the following:

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

PHP will work just as well and audit logger will be able to do its job.

Testing utility

A small testing utility was developed as part of the mod_security effort. It provides a simple and easy way to send crafted HTTP requests to a server, and to determine whether the attack was successfully detected or not.

Calling the utility without parameters will result in its usage printed:

```
[ivanr@wkk tests]$ ./run-test.pl
Usage: ./run-test.pl host[:port] testfile1, testfile2, ...
```

First parameter is the hostname of the server, with port being optional. All other parameters are filenames of files containing crafted HTTP requests.

To make your life a little bit easier, the utility will generate certain request headers:

- Host: (hostname)
- User-Agent: mod_security regression testing utility
- Connection: Close

You can include them in the request if you want, the utility will not add them if they are already there.

Here is how an HTTP request looks like:

```
# 01 Simple keyword filter
#
# mod_security is configured not to allow
```

```
# the "/cgi-bin/keyword" pattern
#
GET /cgi-bin/keyword HTTP/1.0
```

This request consists only of the first line, with no additional headers. You can create as complicated requests as you wish. Here is one example of a POST method usage:

```
# 10 Keyword in POST
#
POST /cgi-bin/printenv HTTP/1.0
Content-Type: application/x-www-form-urlencoded
Content-Length: 5

p=333
```

Lines that are at the beginning of the file and begin with # will be treated as comments. The first line is special, and it should contain the name of the test.

The utility expects status 200 as a result and will treat such responses as successes. If you want some other response you need to tell it by writing the expected response code on the first line (anywhere on the line). Like this:

```
# 14 Redirect action (requires 302)
#
GET /cgi-bin/test.cgi?p=xxx HTTP/1.0
```

The brackets and the "requires" keyword are not required but are recommended for better readability.

Implementation notes

Apache 2

Apache 2 is very different from Apache 1. At this time, Apache 1 version of mod_security still gets more attention and more testing (this is partly due to the fact that Apache 2 is not very well documented).

Apache 2 introduced a notion of input and output filters, module parts that can observe and, optionally, change content as it comes in and out of the server. Since Apache 1 has no such API there are some differences in behaviour between versions of mod_security for Apache 1 and Apache 2.

Apache 1 versions perform all checks in one go. By the time checking starts the module acquires all data it requires to perform this task completely.

With Apache 2, checks are performed in two phases. In phase one all checks except checks against the POST payload are performed. Phase two includes the POST payload checks. Phase one is performed before the server starts executing the request. Phase two is performed in parallel with request execution. If mod_security determines that a POST payload matches a filter it will rudely interrupt request execution. While this is not as good as I would want it to be, in practice it fulfills the task i.e. it prevents the malicious content from reaching the application code.

Note: The two-phase check will produce unexpected behaviour at times. For example, if you have a *CGI script* that does not read the POST payload it will completely execute even if you have a filter against POST_PAYLOAD and there is a filter match. Future versions of mod_security will correct this problem if possible.

Technology specific

PHP

Preventing register_globals problems

Nowadays it is widely accepted that using the register_globals feature of PHP leads to security problems, but it wasn't always like this (if you don't know what this feature is then you are probably not using it; but, hey, read on the discussion is informative). In fact, the register_globals feature was turned on by default until version 4.2.0. As a result of that, many applications that exist depend on this feature (for more details have a look at http://www.php.net/register_globals).

If you can choose, it is better to refactor and rewrite the code to not use this feature. But if you cannot afford to do that for some reason or another, you can use mod_security to protect an application from a known vulnerability. Problematic bits of code usually look like this:

```
<?php
// this is the beginning of the page
if ($authorised) {
    // do something protected
}
// the rest of the page here
?>
```

And the attacker would take advantage of this simply by adding an additional parameter to the URL. For example, <http://www.modsecurity.org/examples/test.php?authorised=1>.

Rejecting all requests that explicitly supply the parameter in question will be sufficient to protect the application from all attackers:

```
<Location "/vulnerable-application/">
    SecFilter ARG_authorized "!$^"
</Location>
```

The filter above rejects all requests where the variable "authorized" is not empty. You can also see that we've added the <Location>...</Location> directives to limit filter only to those parts of the web server that really need it.

Additional Examples

File upload

Forbid file upload for the application as a whole, but allow it in a subfolder:

```
# Reject requests with header "Content-Type" set to "multipart/form-data"
SecFilterSelective "HTTP_CONTENT_TYPE" multipart/form-data
```

```
# Only for the script that performs upload
<Location "/upload.php">
    # Do not inherit filters from the parent folder
    SecFilterInheritance Off
</Location>
```

Securing FormMail

Earlier versions of FormMail could be abused to send email to any recipient (I've been told that there is a new version that can be secured properly).

```
# Only for the FormMail script
<Location /cgi-bin/FormMail>
    # Reject request where the value of parameter "recipient"
    # does not end with "@webkreator.com"
    SecFilterSelective "ARG_recipient" "!@webkreator.com$">
</Location>
```