



ModSecurity® Reference Manual

Version 2.5.9 (March 5, 2009)

Copyright © 2004-2009 Breach Security, Inc. (<http://www.breach.com>)

Table of Contents

- Introduction8
 - HTTP Traffic Logging8
 - Real-Time Monitoring and Attack Detection8
 - Attack Prevention and Just-in-time Patching8
 - Flexible Rule Engine9
 - Embedded-mode Deployment9
 - Network-based Deployment9
 - Portability9
 - Licensing10
- ModSecurity Core Rules™11
 - Overview11
 - Core Rules Content11
- Installation12
- Configuration Directives14
 - SecAction14
 - SecArgumentSeparator14
 - SecAuditEngine15
 - SecAuditLog15
 - SecAuditLog216
 - SecAuditLogParts16
 - SecAuditLogRelevantStatus17

SecAuditLogStorageDir	18
SecAuditLogType	18
SecCacheTransformations (Deprecated/Experimental)	19
SecChrootDir	19
SecComponentSignature	20
SecContentInjection	20
SecCookieFormat	20
SecDataDir	21
SecDebugLog	21
SecDebugLogLevel	21
SecDefaultAction	22
SecGeoLookupDb	22
SecGuardianLog	23
SecMarker	23
SecPdfProtect	24
SecPdfProtectMethod	24
SecPdfProtectSecret	25
SecPdfProtectTimeout	25
SecPdfProtectTokenName	25
SecRequestBodyAccess	26
SecRequestBodyLimit	26
SecRequestBodyNoFilesLimit	26
SecRequestBodyInMemoryLimit	27
SecResponseBodyLimit	27
SecResponseBodyLimitAction	27
SecResponseBodyMimeType	28
SecResponseBodyMimeTypesClear	28
SecResponseBodyAccess	29
SecRule	29
SecRuleInheritance	32
SecRuleEngine	33
SecRuleRemoveById	33
SecRuleRemoveByMsg	34
SecRuleScript (Experimental)	34
SecRuleUpdateActionById	36
SecServerSignature	36
SecTmpDir	36
SecUploadDir	37
SecUploadFileMode	37
SecUploadKeepFiles	37
SecWebAppId	38
Processing Phases	39

Phase Request Headers	41
Phase Request Body	41
Phase Response Headers	41
Phase Response Body	42
Phase Logging	42
Variables	43
ARGS	43
ARGS_COMBINED_SIZE	44
ARGS_NAMES	44
ARGS_GET	44
ARGS_GET_NAMES	44
ARGS_POST	44
ARGS_POST_NAMES	44
AUTH_TYPE	44
ENV	45
FILES	45
FILES_COMBINED_SIZE	45
FILES_NAMES	45
FILES_SIZES	45
FILES_TMPNAMES	45
GEO	46
HIGHEST_SEVERITY	46
MATCHED_VAR	46
MATCHED_VAR_NAME	47
MODSEC_BUILD	47
MULTIPART_CRLF_LF_LINES	47
MULTIPART_STRICT_ERROR	47
MULTIPART_UNMATCHED_BOUNDARY	48
PATH_INFO	48
QUERY_STRING	48
REMOTE_ADDR	49
REMOTE_HOST	49
REMOTE_PORT	49
REMOTE_USER	49
REQBODY_PROCESSOR	49
REQBODY_PROCESSOR_ERROR	50
REQBODY_PROCESSOR_ERROR_MSG	50
REQUEST_BASENAME	50
REQUEST_BODY	50
REQUEST_COOKIES	51
REQUEST_COOKIES_NAMES	51
REQUEST_FILENAME	51

REQUEST_HEADERS	51
REQUEST_HEADERS_NAMES	52
REQUEST_LINE	52
REQUEST_METHOD	52
REQUEST_PROTOCOL	52
REQUEST_URI	52
REQUEST_URI_RAW	53
RESPONSE_BODY	53
RESPONSE_CONTENT_LENGTH	53
RESPONSE_CONTENT_TYPE	53
RESPONSE_HEADERS	53
RESPONSE_HEADERS_NAMES	54
RESPONSE_PROTOCOL	54
RESPONSE_STATUS	54
RULE	54
SCRIPT_BASENAME	55
SCRIPT_FILENAME	55
SCRIPT_GID	55
SCRIPT_GROUPNAME	55
SCRIPT_MODE	55
SCRIPT_UID	56
SCRIPT_USERNAME	56
SERVER_ADDR	56
SERVER_NAME	56
SERVER_PORT	56
SESSION	56
SESSIONID	57
TIME	57
TIME_DAY	57
TIME_EPOCH	57
TIME_HOUR	57
TIME_MIN	58
TIME_MON	58
TIME_SEC	58
TIME_WDAY	58
TIME_YEAR	58
TX	58
USERID	59
WEBAPPID	59
WEBSERVER_ERROR_LOG	59
XML	59
Transformation functions	60

base64Decode	61
base64Encode	61
compressWhitespace	61
cssDecode	61
escapeSeqDecode	62
hexDecode	62
hexEncode	62
htmlEntityDecode	62
jsDecode	62
length	62
lowercase	62
md5	63
none	63
normalisePath	63
normalisePathWin	63
parityEven7bit	63
parityOdd7bit	63
parityZero7bit	63
removeNulls	63
removeWhitespace	63
replaceComments	63
replaceNulls	64
urlDecode	64
urlDecodeUni	64
urlEncode	64
sha1	64
trimLeft	64
trimRight	64
trim	64
Actions	65
allow	65
append	66
auditlog	66
block	67
capture	67
chain	68
ctl	68
deny	69
deprecatevar	69
drop	70
exec	70
expirevar	71

id	71
initcol	72
log	72
logdata	72
msg	73
multiMatch	73
noauditlog	73
nolog	74
pass	74
pause	74
phase	75
prepend	75
proxy	75
redirect	76
rev	76
sanitiseArg	76
sanitiseMatched	76
sanitiseRequestHeader	77
sanitiseResponseHeader	77
severity	77
setuid	78
setsid	78
setenv	78
setvar	79
skip	79
skipAfter	80
status	80
t	81
tag	81
xmlns	81
Operators	82
beginsWith	82
contains	82
endsWith	82
eq	82
ge	83
geoLookup	83
gt	83
inspectFile	83
le	84
lt	84
pm	84

pmFromFile	84
rbl	85
rx	85
streq	85
validateByteRange	86
validateDTD	86
validateSchema	86
validateUrlEncoding	87
validateUtf8Encoding	87
verifyCC	88
within	88
Macro Expansion	89
Persistent Storage	90
Miscellaneous Topics	91
Impedance Mismatch	91

Introduction

ModSecurity is a web application firewall (WAF). With over 70% of attacks now carried out over the web application level, organisations need all the help they can get in making their systems secure. WAFs are deployed to establish an increased external security layer to detect and/or prevent attacks before they reach web applications. ModSecurity provides protection from a range of attacks against web applications and allows for HTTP traffic monitoring and real-time analysis with little or no changes to existing infrastructure.

HTTP Traffic Logging

Web servers are typically well-equipped to log traffic in a form useful for marketing analyses, but fall short logging traffic to web applications. In particular, most are not capable of logging the request bodies. Your adversaries know this, and that is why most attacks are now carried out via POST requests, rendering your systems blind. ModSecurity makes full HTTP transaction logging possible, allowing complete requests and responses to be logged. Its logging facilities also allow fine-grained decisions to be made about exactly what is logged and when, ensuring only the relevant data is recorded. As some of the request and/or response may contain sensitive data in certain fields, ModSecurity can be configured to mask these fields before they are written to the audit log.

Real-Time Monitoring and Attack Detection

In addition to providing logging facilities, ModSecurity can monitor the HTTP traffic in real time in order to detect attacks. In this case, ModSecurity operates as a web intrusion detection tool, allowing you to react to suspicious events that take place at your web systems.

Attack Prevention and Just-in-time Patching

ModSecurity can also act immediately to prevent attacks from reaching your web applications. There are three commonly used approaches:

1. Negative security model. A negative security model monitors requests for anomalies, unusual behaviour, and common web application attacks. It keeps anomaly scores for each request, IP addresses, application sessions, and user accounts. Requests with high anomaly scores are either logged or rejected altogether.
2. Positive security model. When a positive security model is deployed, only requests that are known to be valid are accepted, with everything else rejected. This model requires knowledge of the web applications you are protecting. Therefore a positive security model works best with applications that are heavily used but rarely updated so that maintenance of the model is minimized.
3. Known weaknesses and vulnerabilities. Its rule language makes ModSecurity an ideal external patching tool. External patching (sometimes referred to as Virtual Patching) is about reducing the window of opportunity. Time needed to patch application vulnerabilities often runs to

weeks in many organisations. With ModSecurity, applications can be patched from the outside, without touching the application source code (and even without any access to it), making your systems secure until a proper patch is applied to the application.

Flexible Rule Engine

A flexible rule engine sits in the heart of ModSecurity. It implements the ModSecurity Rule Language, which is a specialised programming language designed to work with HTTP transaction data. The ModSecurity Rule Language is designed to be easy to use, yet flexible: common operations are simple while complex operations are possible. Certified ModSecurity Rules, included with ModSecurity, contain a comprehensive set of rules that implement general-purpose hardening, protocol validation and detection of common web application security issues. Heavily commented, these rules can be used as a learning tool.

Embedded-mode Deployment

ModSecurity is an embeddable web application firewall, which means it can be deployed as part of your existing web server infrastructure provided your web servers are Apache-based. This deployment method has certain advantages:

1. No changes to existing network. It only takes a few minutes to add ModSecurity to your existing web servers. And because it was designed to be completely passive by default, you are free to deploy it incrementally and only use the features you need. It is equally easy to remove or deactivate it if required.
2. No single point of failure. Unlike with network-based deployments, you will not be introducing a new point of failure to your system.
3. Implicit load balancing and scaling. Because it works embedded in web servers, ModSecurity will automatically take advantage of the additional load balancing and scalability features. You will not need to think of load balancing and scaling unless your existing system needs them.
4. Minimal overhead. Because it works from inside the web server process there is no overhead for network communication and minimal overhead in parsing and data exchange.
5. No problem with encrypted or compressed content. Many IDS systems have difficulties analysing SSL traffic. This is not a problem for ModSecurity because it is positioned to work when the traffic is decrypted and decompressed.

Network-based Deployment

ModSecurity works equally well when deployed as part of an Apache-based reverse proxy server, and many of our customers choose to do so. In this scenario, one installation of ModSecurity can protect any number of web servers (even the non-Apache ones).

Portability

ModSecurity is known to work well on a wide range of operating systems. Our customers are successfully running it on Linux, Windows, Solaris, FreeBSD, OpenBSD, NetBSD, AIX, Mac OS X, and HP-UX.

Licensing

ModSecurity is available under two licenses. Users can choose to use the software under the terms of the GNU General Public License version 2 (licence text is included with the distribution), as an Open Source / Free Software product. A range of commercial licenses is also available, together with a range of commercial support contracts. For more information on commercial licensing please contact Breach Security.

Note

ModSecurity, mod_security, ModSecurity Pro, and ModSecurity Core Rules are trademarks or registered trademarks of Breach Security, Inc.

ModSecurity Core Rules™

Overview

ModSecurity is a web application firewall engine that provides very little protection on its own. In order to become useful, ModSecurity must be configured with rules. In order to enable users to take full advantage of ModSecurity out of the box, Breach Security, Inc. is providing a free certified rule set for ModSecurity 2.x. Unlike intrusion detection and prevention systems, which rely on signatures specific to known vulnerabilities, the Core Rules provide generic protection from unknown vulnerabilities often found in web applications, which are in most cases custom coded. The Core Rules are heavily commented to allow it to be used as a step-by-step deployment guide for ModSecurity. The latest Core Rules can be found at the ModSecurity website - <http://www.modsecurity.org/projects/rules/>.

Core Rules Content

In order to provide generic web applications protection, the Core Rules use the following techniques:

- HTTP protection - detecting violations of the HTTP protocol and a locally defined usage policy.
- Common Web Attacks Protection - detecting common web application security attack.
- Automation detection - Detecting bots, crawlers, scanners and other surface malicious activity.
- Trojan Protection - Detecting access to Trojans horses.
- Error Hiding - Disguising error messages sent by the server.

Installation

ModSecurity installation consists of the following steps:

1. ModSecurity 2.x works with Apache 2.0.x or better.
2. Make sure you have `mod_unique_id` installed.
`mod_unique_id` is packaged with Apache httpd.
3. Install the latest version of libxml2, if it isn't already installed on the server.
<http://xmlsoft.org/downloads.html>
4. Optionally install the latest version of Lua in the 5.1.x branch, if it isn't already installed on the server and you will be using the new Lua engine.
<http://www.lua.org/download.html>

Note that ModSecurity requires the dynamic libraries. These are not built by default in the source distribution, so the binary distribution is recommended.

5. Stop Apache httpd
6. Unpack the ModSecurity archive
7. Building differs for UNIX (or UNIX-like) operating systems and Windows.
 - UNIX
 - a. Run the configure script to generate a Makefile. Typically no options are needed.

```
./configure
```

Options are available for more customization (use `./configure --help` for a full list), but typically you will only need to specify the location of the `apxs` command installed by Apache httpd with the `--with-apxs` option.

```
./configure --with-apxs=/path/to/httpd-2.x.y/bin/apxs
```

Note

There are certain configure options that are meant for debugging an other development use. If enabled, these options can substantially impact performance. These options include all `-debug-*` options as well as the `--enable-performance-measurements` options.

-
- b. Compile with: `make`
 - c. Optionally test with: `make test`

Note

This step is still a bit experimental. If you have problems, please send the full output and error from the build to the support list. Most common issues are related to not finding the required headers and/or libraries.

-
- d. Optionally build the ModSecurity Log Collector with: `make mlogc`
 - e. Optionally install `mlogc`: Review the `INSTALL` file included in the

apache2/mlogc-src directory in the distribution.

- f. Install the ModSecurity module with: `make install`
- Windows (MS VC++ 8)
 - a. Edit `Makefile.win` to configure the Apache base and library paths.
 - b. Compile with: `nmake -f Makefile.win`
 - c. Install the ModSecurity module with: `nmake -f Makefile.win install`
 - d. Copy the `libxml2.dll` and `lua5.1.dll` to the Apache `bin` directory. Alternatively you can follow the step below for using `LoadFile` to load these libraries.
8. Edit the main Apache `httpd` config file (usually `httpd.conf`)
On UNIX (and Windows if you did not copy the DLLs as stated above) you must load `libxml2` and `lua5.1` before ModSecurity with something like this:

```
LoadFile /usr/lib/libxml2.so
LoadFile /usr/lib/liblua5.1.so
```
- Load the ModSecurity module with:

```
LoadModule security2_module modules/mod_security2.so
```
9. Configure ModSecurity
10. Start Apache `httpd`
11. You should now have ModSecurity 2.x up and running.

Note

If you have compiled Apache yourself you might experience problems compiling ModSecurity against PCRE. This is because Apache bundles PCRE but this library is also typically provided by the operating system. I would expect most (all) vendor-packaged Apache distributions to be configured to use an external PCRE library (so this should not be a problem).

You want to avoid Apache using the bundled PCRE library and ModSecurity linking against the one provided by the operating system. The easiest way to do this is to compile Apache against the PCRE library provided by the operating system (or you can compile it against the latest PCRE version you downloaded from the main PCRE distribution site). You can do this at configure time using the `--with-pcre` switch. If you are not in a position to recompile Apache, then, to compile ModSecurity successfully, you'd still need to have access to the bundled PCRE headers (they are available only in the Apache source code) and change the include path for ModSecurity (as you did in step 7 above) to point to them (via the `--with-pcre` ModSecurity configure option).

Do note that if your Apache is using an external PCRE library you can compile ModSecurity with `WITH_PCRE_STUDY` defined, which would possibly give you a slight performance edge in regular expression processing.

Configuration Directives

The following section outlines all of the ModSecurity directives. Most of the ModSecurity directives can be used inside the various Apache Scope Directives such as `VirtualHost`, `Location`, `LocationMatch`, `Directory`, etc... There are others, however, that can only be used once in the main configuration file. This information is specified in the Scope sections below. The first version to use a given directive is given in the Version sections below.

These rules, along with the Core rules files, should be contained in files outside of the `httpd.conf` file and called up with Apache "Include" directives. This allows for easier updating/migration of the rules. If you create your own custom rules that you would like to use with the Core rules, you should create a file called `-modsecurity_crs_15_customrules.conf` and place it in the same directory as the Core rules files. By using this file name, your custom rules will be called up after the standard ModSecurity Core rules configuration file but before the other Core rules. This allows your rules to be evaluated first which can be useful if you need to implement specific "allow" rules or to correct any false positives in the Core rules as they are applied to your site.

Note

It is highly encouraged that you do not edit the Core rules files themselves but rather place all changes (such as `SecRuleRemoveByID`, etc...) in your custom rules file. This will allow for easier upgrading as newer Core rules are released by Breach Security on the ModSecurity website.

SecAction

Description: Unconditionally processes the action list it receives as the first and only parameter. It accepts one parameter, the syntax of which is identical to the third parameter of `SecRule`.

Syntax: `SecAction action1,action2,action3`

Example Usage: `SecAction nolog,initcol:RESOURCE=%{REQUEST_FILENAME}`

Processing Phase: Any

Scope: Any

Version: 2.0.0

Dependencies/Notes: None

`SecAction` is best used when you unconditionally execute an action. This is explicit triggering whereas the normal Actions are conditional based on data inspection of the request/response. This is a useful directive when you want to run certain actions such as `initcol` to initialize collections.

SecArgumentSeparator

Description: Specifies which character to use as separator for application/x-www-form-urlencoded content. Defaults to `&`. Applications are sometimes (very rarely) written

to use a semicolon (;).

Syntax: SecArgumentSeparator character

Example Usage: SecArgumentSeparator ;

Processing Phase: Any

Scope: Main

Version: 2.0.0

Dependencies/Notes: None

This directive is needed if a backend web application is using a non-standard argument separator. If this directive is not set properly for each web application, then ModSecurity will not be able to parse the arguments appropriately and the effectiveness of the rule matching will be significantly decreased.

SecAuditEngine

Description: Configures the audit logging engine.

Syntax: SecAuditEngine On|Off|RelevantOnly

Example Usage: SecAuditEngine On

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: Can be set/changed with the "ctl" action for the current transaction.

Example: The following example shows the various audit directives used together.

```
SecAuditEngine RelevantOnly
SecAuditLog logs/audit/audit.log
SecAuditLogParts ABCFHZ
SecAuditLogType concurrent
SecAuditLogStorageDir logs/audit
SecAuditLogRelevantStatus ^(?:5|4\d{^4})
```

Possible values are:

- On - log all transactions by default.
- Off - do not log transactions by default.
- RelevantOnly - by default only log transactions that have triggered a warning or an error, or have a status code that is considered to be relevant (see SecAuditLogRelevantStatus).

SecAuditLog

Description: Defines the path to the main audit log file.

Syntax: SecAuditLog /path/to/auditlog

Example Usage: SecAuditLog /usr/local/apache/logs/audit.log

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: This file is open on startup when the server typically still runs as *root*. You should not allow non-root users to have write privileges for this file or for the directory it is stored in..

This file will be used to store the audit log entries if serial audit logging format is used. If concurrent audit logging format is used this file will be used as an index, and contain a record of all audit log files created. If you are planning to use Concurrent audit logging and sending your audit log data off to a remote Console host or commercial ModSecurity Management Appliance, then you will need to configure and use the ModSecurity Log Collector (mlogc) and use the following format for the audit log:

```
SecAuditLog "|/path/to/mlogc /path/to/mlogc.conf"
```

SecAuditLog2

Description: Defines the path to the secondary audit log index file when concurrent logging is enabled. See SecAuditLog2 for more details.

Syntax: SecAuditLog2 /path/to/auditlog2

Example Usage: SecAuditLog2 /usr/local/apache/logs/audit2.log

Processing Phase: N/A

Scope: Any

Version: 2.1.2

Dependencies/Notes: A main audit log must be defined via SecAuditLog before this directive may be used. Additionally, this log is only used for replicating the main audit log index file when concurrent audit logging is used. It will *not* be used for non-concurrent audit logging.

SecAuditLogParts

Description: Defines which part of each transaction are going to be recorded in audit log. Each part is assigned a single letter. If a letter appears in the list then the equivalent part of each transactions will be recorded. See below for the list of all parts.

Syntax: SecAuditLogParts PARTS

Example Usage: SecAuditLogParts ABCFHZ

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: At this time ModSecurity does not log response bodies of stock Apache responses (e.g. 404), or the Server and Date response headers.

Default: ABCFHZ.

Note

Please refer to the ModSecurity Data Formats document for a detailed description of every available part.

Available audit log parts:

- A - audit log header (mandatory)
- B - request headers
- C - request body (present only if the request body exists and ModSecurity is configured to intercept it)
- D - RESERVED for intermediary response headers, not implemented yet.
- E - intermediary response body (present only if ModSecurity is configured to intercept response bodies, and if the audit log engine is configured to record it). Intermediary response body is the same as the actual response body unless ModSecurity intercepts the intermediary response body, in which case the actual response body will contain the error message (either the Apache default error message, or the ErrorDocument page).
- F - final response headers (excluding the Date and Server headers, which are always added by Apache in the late stage of content delivery).
- G - RESERVED for the actual response body, not implemented yet.
- H - audit log trailer
- I - This part is a replacement for part C. It will log the same data as C in all cases except when `multipart/form-data` encoding is used. In this case it will log a fake `application/x-www-form-urlencoded` body that contains the information about parameters but not about the files. This is handy if you don't want to have (often large) files stored in your audit logs.
- J - RESERVED. This part, when implemented, will contain information about the files uploaded using `multipart/form-data` encoding.
- K - This part contains a full list of every rule that matched (one per line) in the order they were matched. The rules are fully qualified and will thus show inherited actions and default operators. Supported as of v2.5.0
- Z - final boundary, signifies the end of the entry (mandatory)

SecAuditLogRelevantStatus

Description: Configures which response status code is to be considered relevant for the purpose of audit logging.

Syntax: `SecAuditLogRelevantStatus REGEX`

Example Usage: `SecAuditLogRelevantStatus ^(?:5|4\d{^4})`

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: Must have the `SecAuditEngine` set to `RelevantOnly`. The parameter is a regular expression.

The main purpose of this directive is to allow you to configure audit logging for only transactions that generate the specified HTTP Response Status Code. This directive is often used to decrease the total size of the audit log file. Keep in mind that if this parameter is used, then successful attacks that result in a 200 OK status code will not be logged.

SecAuditLogStorageDir

Description: Configures the storage directory where concurrent audit log entries are to be stored.

Syntax: `SecAuditLogStorageDir /path/to/storage/dir`

Example Usage: `SecAuditLogStorageDir /usr/local/apache/logs/audit`

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: `SecAuditLogType` must be set to `Concurrent`. The directory must already be created before starting Apache and it must be writable by the web server user as new files are generated at runtime.

As with all logging mechanisms, ensure that you specify a file system location that has adequate disk space and is not on the root partition.

SecAuditLogType

Description: Configures the type of audit logging mechanism to be used.

Syntax: `SecAuditLogType Serial|Concurrent`

Example Usage: `SecAuditLogType Serial`

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: Must specify `SecAuditLogStorageDir` if you use concurrent logging.

Possible values are:

1. `Serial` - all audit log entries will be stored in the main audit logging file. This is more convenient for casual use but it is slower as only one audit log entry can be written to the file at any one file.
2. `Concurrent` - audit log entries will be stored in separate files, one for each transaction. Concurrent logging is the mode to use if you are going to send the audit log data off to a remote ModSecurity Console host.

SecCacheTransformations (Deprecated/Experimental)

Description: Controls caching of transformations. Caching is off by default starting with 2.5.6, when it was deprecated and downgraded back to experimental.

Syntax: SecCacheTransformations On|Off [options]

Example Usage: SecCacheTransformations On "minlen:64,maxlen:0"

Processing Phase: N/A

Scope: Any

Version: 2.5.0

Dependencies/Notes: N/A

First parameter:

- On - cache transformations (per transaction, per phase) allowing identical transformations to be performed only once. (default)
- Off - do not cache any transformations, forcing all transformations to be performed for each rule executed.

The following options are allowed (comma separated):

- incremental:on|off - enabling this option will cache every transformation instead of just the final transformation. (default: off)
- maxitems:N - do not allow more than N transformations to be cached. The cache will then be disabled. A zero value is interpreted as "unlimited". This option may be useful to limit caching for a form with a large number of ARGS. (default: 512)
- minlen:N - do not cache the transformation if the value's length is less than N bytes. (default: 32)
- maxlen:N - do not cache the transformation if the value's length is more than N bytes. A zero value is interpreted as "unlimited". (default: 1024)

SecChrootDir

Description: Configures the directory path that will be used to jail the web server process.

Syntax: SecChrootDir /path/to/chroot/dir

Example Usage: SecChrootDir /chroot

Processing Phase: N/A

Scope: Main

Version: 2.0.0

Dependencies/Notes: This feature is not available on Windows builds. The internal chroot functionality provided by ModSecurity works great for simple setups. One example of a simple setup is Apache serving static files only, or running scripts using modules.builds. Some problems you might encounter

with more complex setups:

1. DNS lookups do not work (this is because this feature requires a shared library that is loaded on demand, after chroot takes place).
2. You cannot send email from PHP because it uses sendmail and sendmail is outside the jail.
3. In some cases Apache graceful (reload) no longer works.

You should be aware that the internal chroot feature might not be 100% reliable. Due to the large number of default and third-party modules available for the Apache web server, it is not possible to verify the internal chroot works reliably with all of them. A module, working from within Apache, can do things that make it easy to break out of the jail. In particular, if you are using any of the modules that fork in the module initialisation phase (e.g. `mod_fastcgi`, `mod_fcgid`, `mod_cgid`), you are advised to examine each Apache process and observe its current working directory, process root, and the list of open files. Consider what your options are and make your own decision.

SecComponentSignature

Description: Appends component signature to the ModSecurity signature.

Syntax: `SecComponentSignature "COMPONENT_NAME/X.Y.Z (COMMENT)"`

Example usage: `SecComponentSignature "Core Rules/1.2.3"`

Processing Phase: N/A

Scope: Main

Version: 2.5.0

Dependencies/Notes: This directive should be used to make the presence of significant ModSecurity components known. The entire signature will be recorded in transaction audit log. It should be used by ModSecurity module and rule set writers to make debugging easier.

SecContentInjection

Description: Enables content injection using actions append and prepend.

Syntax: `SecContentInjection (On|Off)`

Example Usage: `SecContentInjection On`

Processing Phase: N/A

Scope: Any

Version: 2.5.0

Dependencies/Notes: N/A

SecCookieFormat

Description: Selects the cookie format that will be used in the current configuration context.

Syntax: `SecCookieFormat 0|1`

Example Usage: `SecCookieFormat 0`

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: None

Possible values are:

- 0 - use version 0 (Netscape) cookies. This is what most applications use. It is the default value.
- 1 - use version 1 cookies.

SecDataDir

Description: Path where persistent data (e.g. IP address data, session data, etc) is to be stored.

Syntax: SecDataDir /path/to/dir

Example Usage: SecDataDir /usr/local/apache/logs/data

Processing Phase: N/A

Scope: Main

Dependencies/Notes: This directive is needed when initcol, setsid an setuid are used. Must be writable by the web server user.

SecDebugLog

Description: Path to the ModSecurity debug log file.

Syntax: SecDebugLog /path/to/modsec-debug.log

Example Usage: SecDebugLog /usr/local/apache/logs/modsec-debug.log

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: None

SecDebugLogLevel

Description: Configures the verbosity of the debug log data.

Syntax: SecDebugLogLevel 0|1|2|3|4|5|6|7|8|9

Example Usage: SecDebugLogLevel 4

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: Levels 1 - 3 are always sent to the Apache error log. Therefore you can always use level 0 as the default logging level in production. Level 5 is useful when debugging. It is not advisable to use higher logging levels in production as excessive logging can slow down server significantly.

Possible values are:

- 0 - no logging.
- 1 - errors (intercepted requests) only.
- 2 - warnings.
- 3 - notices.
- 4 - details of how transactions are handled.
- 5 - as above, but including information about each piece of information handled.
- 9 - log everything, including very detailed debugging information.

SecDefaultAction

Description: Defines the default action to take on a rule match.

Syntax: SecDefaultAction action1,action2,action3

Example Usage: SecDefaultAction log,auditlog,deny,status:403,phase:2

Processing Phase: Any

Scope: Any

Version: 2.0.0

Dependencies/Notes: Rules following a SecDefaultAction directive will inherit this setting unless a specific action is specified for an individual rule or until another SecDefaultAction is specified. Take special note that in the logging disruptive actions are not allowed, but this can inadvertently be inherited using a disruptive action in SecDefaultAction.

The default value is minimal (differing from previous versions):

```
SecDefaultAction phase:2,log,auditlog,pass
```

Note

SecDefaultAction must specify a disruptive action and a processing phase and cannot contain metadata actions.

Warning

SecDefaultAction is *not* inherited across configuration contexts. (For an example of why this may be a problem for you, read the following ModSecurity Blog entry <http://blog.modsecurity.org/2008/07/modsecurity-tri.html>).

SecGeoLookupDb

Description: Defines the path to the geographical database file.

Syntax: SecGeoLookupDb /path/to/db

Example Usage: SecGeoLookupDb /usr/local/geo/data/GeoLiteCity.dat

Processing Phase: N/A

Scope: Any

Version: 2.5.0

Dependencies/Notes: Check out maxmind.com for free database files.

SecGuardianLog

Description: Configuration directive to use the httpd-guardian script to monitor for Denial of Service (DoS) attacks.

Syntax: SecGuardianLog |/path/to/httpd-guardian

Example Usage: SecGuardianLog |/usr/local/apache/bin/httpd-guardian

Processing Phase: N/A

Scope: Main

Version: 2.0.0

Dependencies/Notes: By default httpd-guardian will defend against clients that send more than 120 requests in a minute, or more than 360 requests in five minutes.

Since 1.9, ModSecurity supports a new directive, SecGuardianLog, that is designed to send all access data to another program using the piped logging feature. Since Apache is typically deployed in a multi-process fashion, making information sharing difficult, the idea is to deploy a single external process to observe all requests in a stateful manner, providing additional protection.

Development of a state of the art external protection tool will be a focus of subsequent ModSecurity releases. However, a fully functional tool is already available as part of the Apache httpd tools project [<http://www.apachesecurity.net/tools/>]. The tool is called httpd-guardian and can be used to defend against Denial of Service attacks. It uses the blacklist tool (from the same project) to interact with an iptables-based (Linux) or pf-based (*BSD) firewall, dynamically blacklisting the offending IP addresses. It can also interact with SnortSam (<http://www.snortsam.net>). Assuming httpd-guardian is already configured (look into the source code for the detailed instructions) you only need to add one line to your Apache configuration to deploy it:

```
SecGuardianLog |/path/to/httpd-guardian
```

SecMarker

Description: Adds a fixed rule marker in the ruleset to be used as a target in a skipAfter action. A SecMarker directive essentially creates a rule that does nothing and whose only purpose it to carry the given ID.

Syntax: SecMarker ID

Example Usage: SecMarker 9999

Processing Phase: Any

Scope: Any

Version: 2.5.0

Dependencies/Notes: None

```
SecRule REQUEST_URI "^/$" \
    "chain,t:none,t:urlDecode,t:lowercase,t:normalisePath,skipAfter:99"
SecRule REMOTE_ADDR "^127\.0\.0\.1$" "chain"
SecRule REQUEST_HEADERS:User-Agent \
    "^Apache \(\(internal dummy connection\)\)$" "t:none"
SecRule &REQUEST_HEADERS:Host "@eq 0" \
    "deny,log,status:400,id:08,severity:4,msg:'Missing a Host Header'"
SecRule &REQUEST_HEADERS:Accept "@eq 0" \
    "log,deny,log,status:400,id:15,msg:'Request Missing an Accept Header'"

SecMarker 99
```

SecPdfProtect

Description: Enables the PDF XSS protection functionality. Once enabled access to PDF files is tracked. Direct access attempts are redirected to links that contain one-time tokens. Requests with valid tokens are allowed through unmodified. Requests with invalid tokens are also allowed through but with forced download of the PDF files. This implementation uses response headers to detect PDF files and thus can be used with dynamically generated PDF files that do not have the .pdf extension in the request URI.

Syntax: SecPdfProtect On|Off

Example Usage: SecPdfProtect On

Processing Phase: N/A

Scope: Any

Version: 2.5.0

Dependencies/Notes: None

SecPdfProtectMethod

Description: Configure desired protection method to be used when requests for PDF files are detected. Possible values are TokenRedirection and ForcedDownload. The token redirection approach will attempt to redirect with tokens where possible. This allows PDF files to continue to be opened inline but only works for GET requests. Forced download always causes PDF files to be delivered as opaque binaries and attachments. The latter will always be used for non-GET requests. Forced download is considered to be more secure but may cause usability problems for users ("This PDF won't open anymore!").

Syntax: SecPdfProtectMethod method

Example Usage: SecPdfProtectMethod TokenRedirection

Processing Phase: N/A

Scope: Any

Version: 2.5.0

Dependencies/Notes: None

Default: TokenRedirection

SecPdfProtectSecret

Description: Defines the secret that will be used to construct one-time tokens. You should use a reasonably long value for the secret (e.g. 16 characters is good). Once selected the secret should not be changed as it will break the tokens that were sent prior to change. But it's not a big deal even if you change it. It will just force download of PDF files with tokens that were issued in the last few seconds.

Syntax: SecPdfProtectSecret secret

Example Usage: SecPdfProtectSecret MyRandomSecretString

Processing Phase: N/A

Scope: Any

Version: 2.5.0

Dependencies/Notes: None

SecPdfProtectTimeout

Description: Defines the token timeout. After token expires it can no longer be used to allow access to PDF file. Request will be allowed through but the PDF will be delivered as attachment.

Syntax: SecPdfProtectTimeout timeout

Example Usage: SecPdfProtectTimeout 10

Processing Phase: N/A

Scope: Any

Version: 2.5.0

Dependencies/Notes: None

Default: 10

SecPdfProtectTokenName

Description: Defines the name of the token. The only reason you would want to change the name of the token is if you wanted to hide the fact you are running ModSecurity. It's a good reason but it won't really help as the adversary can look into the algorithm used for PDF protection and figure it out anyway. It does raise the bar slightly so go ahead if you want to.

Syntax: SecPdfProtectTokenName name

Example Usage: SecPdfProtectTokenName PDFTOKEN

Processing Phase: N/A

Scope: Any

Version: 2.5.0

Dependencies/Notes: None

Default: PDFTOKEN

SecRequestBodyAccess

Description: Configures whether request bodies will be buffered and processed by ModSecurity by default.

Syntax: SecRequestBodyAccess On|Off

Example Usage: SecRequestBodyAccess On

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: This directive is required if you plan to inspect POST_PAYLOAD. This directive must be used along with the "phase:2" processing phase action and REQUEST_BODY variable/location. If any of these 3 parts are not configured, you will not be able to inspect the request bodies.

Possible values are:

- On - access request bodies.
- Off - do not attempt to access request bodies.

SecRequestBodyLimit

Description: Configures the maximum request body size ModSecurity will accept for buffering.

Syntax: SecRequestBodyLimit NUMBER_IN_BYTES

Example Usage: SecRequestBodyLimit 134217728

Scope: Any

Version: 2.0.0

Dependencies/Notes: 131072 KB (134217728 bytes) is the default setting. Anything over this limit will be rejected with status code 413 Request Entity Too Large. There is a hard limit of 1 GB.

SecRequestBodyNoFilesLimit

Description: Configures the maximum request body size ModSecurity will accept for buffering, excluding the size of files being transported in the request. This directive comes handy to further reduce susceptibility to DoS attacks when someone is sending request bodies of very large sizes. Web applications that require file uploads must configure SecRequestBodyLimit to a high value. Since large files are streamed to disk file uploads will not increase memory consumption. However, it's still possible for someone to take advantage of a large request body limit and send non-upload requests with large body sizes. This directive eliminates that loophole.

Syntax: SecRequestBodyNoFilesLimit NUMBER_IN_BYTES

Example Usage: SecRequestBodyLimit 131072

Scope: Any

Version: 2.5.0

Dependencies/Notes: 1 MB (1048576 bytes) is the default setting. This value is very conservative. For most applications you should be able to reduce it down to 128 KB or lower. Anything over the limit will be rejected with status code 413 `Request Entity Too Large`. There is a hard limit of 1 GB.

SecRequestBodyInMemoryLimit

Description: Configures the maximum request body size ModSecurity will store in memory.

Syntax: `SecRequestBodyInMemoryLimit NUMBER_IN_BYTES`

Example Usage: `SecRequestBodyInMemoryLimit 131072`

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: None

By default the limit is 128 KB:

```
# Store up to 128 KB in memory
SecRequestBodyInMemoryLimit 131072
```

SecResponseBodyLimit

Description: Configures the maximum response body size that will be accepted for buffering.

Syntax: `SecResponseBodyLimit NUMBER_IN_BYTES`

Example Usage: `SecResponseBodyLimit 524228`

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: Anything over this limit will be rejected with status code 500 Internal Server Error. This setting will not affect the responses with MIME types that are not marked for buffering. There is a hard limit of 1 GB.

By default this limit is configured to 512 KB:

```
# Buffer response bodies of up to 512 KB in length
SecResponseBodyLimit 524288
```

SecResponseBodyLimitAction

Description: Controls what happens once a response body limit, configured with `SecResponseBodyLimit`, is encountered. By default ModSecurity will reject a response body that is longer than spe-

cified. Some web sites, however, will produce very long responses making it difficult to come up with a reasonable limit. Such sites would have to raise the limit significantly to function properly defying the purpose of having the limit in the first place (to control memory consumption). With the ability to choose what happens once a limit is reached site administrators can choose to inspect only the first part of the response, the part that can fit into the desired limit, and let the rest through. Some could argue that allowing parts of responses to go uninspected is a weakness. This is true in theory but only applies to cases where the attacker controls the output (e.g. can make it arbitrary long). In such cases, however, it is not possible to prevent leakage anyway. The attacker could compress, obfuscate, or even encrypt data before it is sent back, and therefore bypass any monitoring device.

Syntax: `SecResponseBodyLimitAction Reject|ProcessPartial`

Example Usage: `SecResponseBodyLimitAction ProcessPartial`

Processing Phase: N/A

Scope: Any

Version: 2.5.0

Dependencies/Notes: None

SecResponseBodyMimeType

Description: Configures which MIME types are to be considered for response body buffering.

Syntax: `SecResponseBodyMimeType mime/type`

Example Usage: `SecResponseBodyMimeType text/plain text/html`

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: Multiple `SecResponseBodyMimeType` directives can be used to add MIME types.

The default value is `text/plain text/html`:

```
SecResponseBodyMimeType text/plain text/html
```

SecResponseBodyMimeTypeClear

Description: Clears the list of MIME types considered for response body buffering, allowing you to start populating the list from scratch.

Syntax: `SecResponseBodyMimeTypeClear`

Example Usage: `SecResponseBodyMimeTypeClear`

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: None

SecResponseBodyAccess

Description: Configures whether response bodies are to be buffer and analysed or not.

Syntax: SecResponseBodyAccess On|Off

Example Usage: SecResponseBodyAccess On

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: This directive is required if you plan to inspect HTML responses. This directive must be used along with the "phase:4" processing phase action and RESPONSE_BODY variable/location. If any of these 3 parts are not configured, you will not be able to inspect the response bodies.

Possible values are:

- On - access response bodies (but only if the MIME type matches, see above).
- Off - do not attempt to access response bodies.

SecRule

Description: SecRule is the main ModSecurity directive. It is used to analyse data and perform actions based on the results.

Syntax: SecRule VARIABLES OPERATOR [ACTIONS]

Example Usage: SecRule REQUEST_URI "attack" \
"phase:1,t:none,t:urlDecode,t:lowercase,t:normalisePath"

Processing Phase: Any

Scope: Any

Version: 2.0.0

Dependencies/Notes: None

In general, the format of this rule is as follows:

```
SecRule VARIABLES OPERATOR [ACTIONS]
```

The second part, OPERATOR, specifies how they are going to be checked. The third (optional) part, ACTIONS, specifies what to do whenever the operator used performs a successful match against a variable.

Variables in rules

The first part, VARIABLES, specifies which variables are to be checked. For example, the following rule will reject a transaction that has the word *dirty* in the URI:

```
SecRule ARGS dirty
```

Each rule can specify one or more variables:

```
SecRule ARGS|REQUEST_HEADERS:User-Agent dirty
```

There is a third format supported by the selection operator - XPath expression. XPath expressions can only be used against the special variable `XML`, which is available only if the request body was processed as XML.

```
SecRule XML:/xPath/Expression dirty
```

Note

Not all collections support all selection operator format types. You should refer to the documentation of each collection to determine what is and isn't supported.

Collections

A variable can contain one or many pieces of data, depending on the nature of the variable and the way it is used. We've seen examples of both approaches in the previous section. When a variable can contain more than one value we refer to it as a *collection*.

Collections are always expanded before a rule is run. For example, the following rule:

```
SecRule ARGS dirty
```

will be expanded to:

```
SecRule ARGS:p dirty  
SecRule ARGS:q dirty
```

in a request that has only two parameters, named `p` and `q`.

Collections come in several flavours:

Read-only	Created at runtime using transaction data. For example: <code>ARGS</code> (contains a list of all request parameter values) and <code>REQUEST_HEADERS</code> (contains a list of all request header values).
Transient Read/Write	The <code>TX</code> collection is created (empty) for every transaction. Rules can read from it and write to it (using the <code>setvar</code> action, for example), but the information stored in this collection will not survive the end of transaction.
Persistent Read/Write	There are several collections that can be written to, but which are persisted to the storage backend. These collections are used to track clients across transactions. Examples of collections that fall into this type are <code>IP</code> , <code>SESSION</code> and <code>USER</code> .

Operators in rules

In the simplest possible case you will use a regular expression pattern as the second rule parameter. This is what we've done in the examples above. If you do this ModSecurity assumes you want to use the `rx` (regular expression) operator. You can also explicitly specify the operator you want to use by using `@`, followed by the name of an operator, at the beginning of the second `SecRule` parameter:

```
SecRule ARGS "@rx dirty"
```

Note how we had to use double quotes to delimit the second rule parameter. This is because the second parameter now has whitespace in it. Any number of whitespace characters can follow the name of the operator. If there are any non-whitespace characters there, they will all be treated as a special parameter to the operator. In the case of the regular expression operator the special parameter is the pattern that will be used for comparison.

The `@` can be the second character if you are using negation to negate the result returned by the operator:

```
SecRule &ARGS "!@rx ^0$"
```

Operator negation

Operator results can be negated by using an exclamation mark at the beginning of the second parameter. The following rule matches if the word `dirty` does *not* appear in the `User-Agent` request header:

```
SecRule REQUEST_HEADERS:User-Agent !dirty
```

You can use the exclamation mark in combination with any parameter. If you do, the exclamation mark needs to go first, followed by the explicit operator reference. The following rule has the same effect as the previous example:

```
SecRule REQUEST_HEADERS:User-Agent "@!rx dirty"
```

If you need to use negation in a rule that is going to be applied to several variables then it may not be immediately clear what will happen. Consider the following example:

```
SecRule ARGS:p|ARGS:q !dirty
```

The above rule is identical to:

```
SecRule ARGS:p !dirty  
SecRule ARGS:q !dirty
```

Warning

Negation is applied to operations against individual operations, not against the entire variable list.

Actions in rules

The third parameter, `ACTIONS`, can be omitted only because there is a helper feature that specifies the default action list. If the parameter isn't omitted the actions specified in the parameter will be merged with the default action list to create the actual list of actions that will be processed on a rule match.

SecRuleInheritance

Description: Configures whether the current context will inherit rules from the parent context (configuration options are inherited in most cases - you should look up the documentation for every directive to determine if it is inherited or not).

Syntax: `SecRuleInheritance On|Off`

Example Usage: `SecRuleInheritance Off`

Processing Phase: Any

Scope: Any

Version: 2.0.0

Dependencies/Notes: Resource-specific contexts (e.g. `Location`, `Directory`, etc) cannot override *phase1* rules configured in the main server or in the virtual server. This is because phase 1 is run early in the request processing process, before Apache maps request to resource. Virtual host context can override phase 1 rules configured in the main server.

Example: The following example shows where ModSecurity may be enabled in the main Apache configuration scope, however you might want to configure your `VirtualHosts` differently. In the first example, the first `VirtualHost` is not inheriting the ModSecurity main config directives and in the second one it is.

```
SecRuleEngine On
SecDefaultAction log,pass,phase:2
...

<VirtualHost *:80>
ServerName appl.com
ServerAlias www.appl.com
SecRuleInheritance Off
SecDefaultAction log,deny,phase:1,redirect:http://www.site2.com
...
</VirtualHost>

<VirtualHost *:80>
ServerName app2.com
ServerAlias www.app2.com
SecRuleInheritance On SecRule ARGS "attack"
...
</VirtualHost>
```

Possible values are:

- On - inherit rules from the parent context.
- Off - do not inherit rules from the parent context.

Note

Configuration contexts are an Apache concept. Directives <Directory>, <Files>, <Location> and <VirtualHost> are all used to create configuration contexts. For more information please go to the Apache documentation section Configuration Sections [<http://httpd.apache.org/docs/2.0/sections.html>].

SecRuleEngine

Description: Configures the rules engine.

Syntax: SecRuleEngine On|Off|DetectionOnly

Example Usage: SecRuleEngine On

Processing Phase: Any

Scope: Any

Version: 2.0.0

Dependencies/Notes: This directive can also be controlled by the ctl action (ctl:ruleEngine=off) for per rule processing.

Possible values are:

- On - process rules.
- Off - do not process rules.
- DetectionOnly - process rules but never intercept transactions, even when rules are configured to do so.

SecRuleRemoveById

Description: Removes matching rules from the parent contexts.

Syntax: SecRuleUpdateActionById RULEID ACTIONLIST

Example Usage: SecRuleRemoveById 1 2 "9000-9010"

Processing Phase: Any

Scope: Any

Version: 2.0.0

Dependencies/Notes: This directive supports multiple parameters, where each parameter can either be a rule ID, or a range. Parameters that contain spaces must be delimited using double quotes.

```
SecRuleRemoveById 1 2 5 10-20 "400-556" 673
```

SecRuleRemoveByMsg

Description: Removes matching rules from the parent contexts.

Syntax: SecRuleRemoveByMsg REGEX

Example Usage: SecRuleRemoveByMsg "FAIL"

Processing Phase: Any

Scope: Any

Version: 2.0.0

Dependencies/Notes: This directive supports multiple parameters. Each parameter is a regular expression that will be applied to the message (specified using the msg action).

SecRuleScript (Experimental)

Description: This directive creates a special rule that executes a Lua script to decide whether to match or not. The main difference from SecRule is that there are no targets nor operators. The script can fetch any variable from the ModSecurity context and use any (Lua) operator to test them. The second optional parameter is the list of actions whose meaning is identical to that of SecRule.

Syntax: SecRuleScript /path/to/script.lua [ACTIONS]

Example Usage: SecRuleScript "/path/to/file.lua" "block"

Processing Phase: Any

Scope: Any

Version: 2.5.0

Dependencies/Notes: None

Note

All Lua scripts are compiled at configuration time and cached in memory. To reload scripts you must reload the entire ModSecurity configuration by restarting Apache.

Example script:

```
-- Your script must define the main entry
-- point, as below.
function main()
  -- Log something at level 1. Normally you shouldn't be
  -- logging anything, especially not at level 1, but this is
  -- just to show you can. Useful for debugging.
  m.log(1, "Hello world!");

  -- Retrieve one variable.
  local var1 = m.getvar("REMOTE_ADDR");

  -- Retrieve one variable, applying one transformation function.
  -- The second parameter is a string.
```

```
local var2 = m.getvar("ARGS", "lowercase");

-- Retrieve one variable, applying several transformation functions.
-- The second parameter is now a list. You should note that m.getvar()
-- requires the use of comma to separate collection names from
-- variable names. This is because only one variable is returned.
local var3 = m.getvar("ARGS.p", { "lowercase", "compressWhitespace" } );

-- If you want this rule to match return a string
-- containing the error message. The message must contain the name
-- of the variable where the problem is located.
-- return "Variable ARGS:p looks suspicious!"

-- Otherwise, simply return nil.
return nil;
end
```

In this first example we were only retrieving one variable at the time. In this case the name of the variable is known to you. In many cases, however, you will want to examine variables whose names you won't know in advance, for example script parameters.

Example showing use of `m.getvars()` to retrieve many variables at once:

```
function main()
  -- Retrieve script parameters.
  local d = m.getvars("ARGS", { "lowercase", "htmlEntityDecode" } );

  -- Loop through the paramters.
  for i = 1, #d do
    -- Examine parameter value.
    if (string.find(d[i].value, "<script")) then
      -- Always specify the name of the variable where the
      -- problem is located in the error message.
      return ("Suspected XSS in variable " .. d[i].name .. ".");
    end
  end

  -- Nothing wrong found.
  return nil;
end
```

Note

Go to <http://www.lua.org/> to find more about the Lua programming language. The reference manual too is available online, at <http://www.lua.org/manual/5.1/>.

Note

Lua support is marked as *experimental* as the way the programming interface may continue to evolve while we are working for the best implementation style. Any user input into the programming interface is appreciated.

SecRuleUpdateActionById

Description: Updates the action list of the specified rule.

Syntax: SecRuleRemoveById RULEID ACTIONLIST

Example Usage: SecRuleUpdateActionById 12345 deny,status:403

Processing Phase: Any

Scope: Any

Version: 2.5.0

Dependencies/Notes: This directive merges the specified action list with the rule's action list. There are two limitations. The rule ID cannot be changed, nor can the phase. Further note that actions that may be specified multiple times are appended to the original.

```
SecAction \  
  "t:lowercase,phase:2,id:12345,pass,msg:'The Message',log,auditlog"  
SecRuleUpdateActionById 12345 "t:compressWhitespace,deny,status:403,msg:'A new message'
```

The example above will cause the rule to be executed as if it was specified as follows:

```
SecAction \  
  "t:lowercase,phase:2,id:12345,log,auditlog,t:compressWhitespace,deny,status:403,msg:'A new message'"
```

SecServerSignature

Description: Instructs ModSecurity to change the data presented in the "Server:" response header token.

Syntax: SecServerSignature "WEB SERVER SOFTWARE"

Example Usage: SecServerSignature "Netscape-Enterprise/6.0"

Processing Phase: N/A

Scope: Main

Version: 2.0.0

Dependencies/Notes: In order for this directive to work, you must set the Apache ServerTokens directive to Full. ModSecurity will overwrite the server signature data held in this memory space with the data set in this directive. If ServerTokens is not set to Full, then the memory space is most likely not large enough to hold the new data we are looking to insert.

SecTmpDir

Description: Configures the directory where temporary files will be created.

Syntax: SecTmpDir /path/to/dir

Example Usage: SecTmpDir /tmp

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: Needs to be writable by the Apache user process. This is the directory location where Apache will swap data to disk if it runs out of memory (more data than what was specified in the SecRequestBodyInMemoryLimit directive) during inspection.

SecUploadDir

Description: Configures the directory where intercepted files will be stored.

Syntax: SecUploadDir /path/to/dir

Example Usage: SecUploadDir /tmp

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: This directory must be on the same filesystem as the temporary directory defined with SecTmpDir. This directive is used with SecUploadKeepFiles.

SecUploadFileMode

Description: Configures the mode (permissions) of any uploaded files using an octal number (as used in chmod).

Syntax: SecUploadFileMode octal_mode| "default"

Example Usage: SecUploadFileMode 0640

Processing Phase: N/A

Scope: Any

Version: 2.1.6

Dependencies/Notes: This feature is not available on operating systems not supporting octal file modes. The default mode (0600) only grants read/write access to the account writing the file. If access from another account is needed (using clamd is a good example), then this directive may be required. However, use this directive with caution to avoid exposing potentially sensitive data to unauthorized users. Using the value "default" will revert back to the default setting.

SecUploadKeepFiles

Description: Configures whether or not the intercepted files will be kept after transaction is processed.

Syntax: SecUploadKeepFiles On|Off|RelevantOnly

Example Usage: SecUploadKeepFiles On

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: This directive requires the storage directory to be defined (using `SecUploadDir`).

Possible values are:

- On - Keep uploaded files.
- Off - Do not keep uploaded files.
- RelevantOnly - This will keep only those files that belong to requests that are deemed relevant.

SecWebAppId

Description: Creates a partition on the server that belongs to one web application.

Syntax: `SecWebAppId "NAME"`

Example Usage: `SecWebAppId "WebApp1"`

Processing Phase: N/A

Scope: Any

Version: 2.0.0

Dependencies/Notes: Partitions are used to avoid collisions between session IDs and user IDs. This directive must be used if there are multiple applications deployed on the same server. If it isn't used, a collision between session IDs might occur. The default value is `default`. Example:

```
<VirtualHost *:80>
ServerName app1.com
ServerAlias www.app1.com
SecWebAppId "App1"
SecRule REQUEST_COOKIES:PHPSESSID !^$ chain,nolog,pass
SecAction setsid:%{REQUEST_COOKIES.PHPSESSID}
...
</VirtualHost>

<VirtualHost *:80>
ServerName app2.com
ServerAlias www.app2.com
SecWebAppId "App2"
SecRule REQUEST_COOKIES:PHPSESSID !^$ chain,nolog,pass
SecAction setsid:%{REQUEST_COOKIES.PHPSESSID}
...
</VirtualHost>
```

In the two examples configurations shown, `SecWebAppId` is being used in conjunction with the Apache `VirtualHost` directives. What this achieves is to create more unique collection names when being hosted on one server. Normally, when `setsid` is used, ModSecurity will create a collection with the name "SES-

SION" and it will hold the value specified. With using SecWebAppId as shown in the examples, however, the name of the collection would become "App1_SESSION" and "App2_SESSION".

SecWebAppId is relevant in two cases:

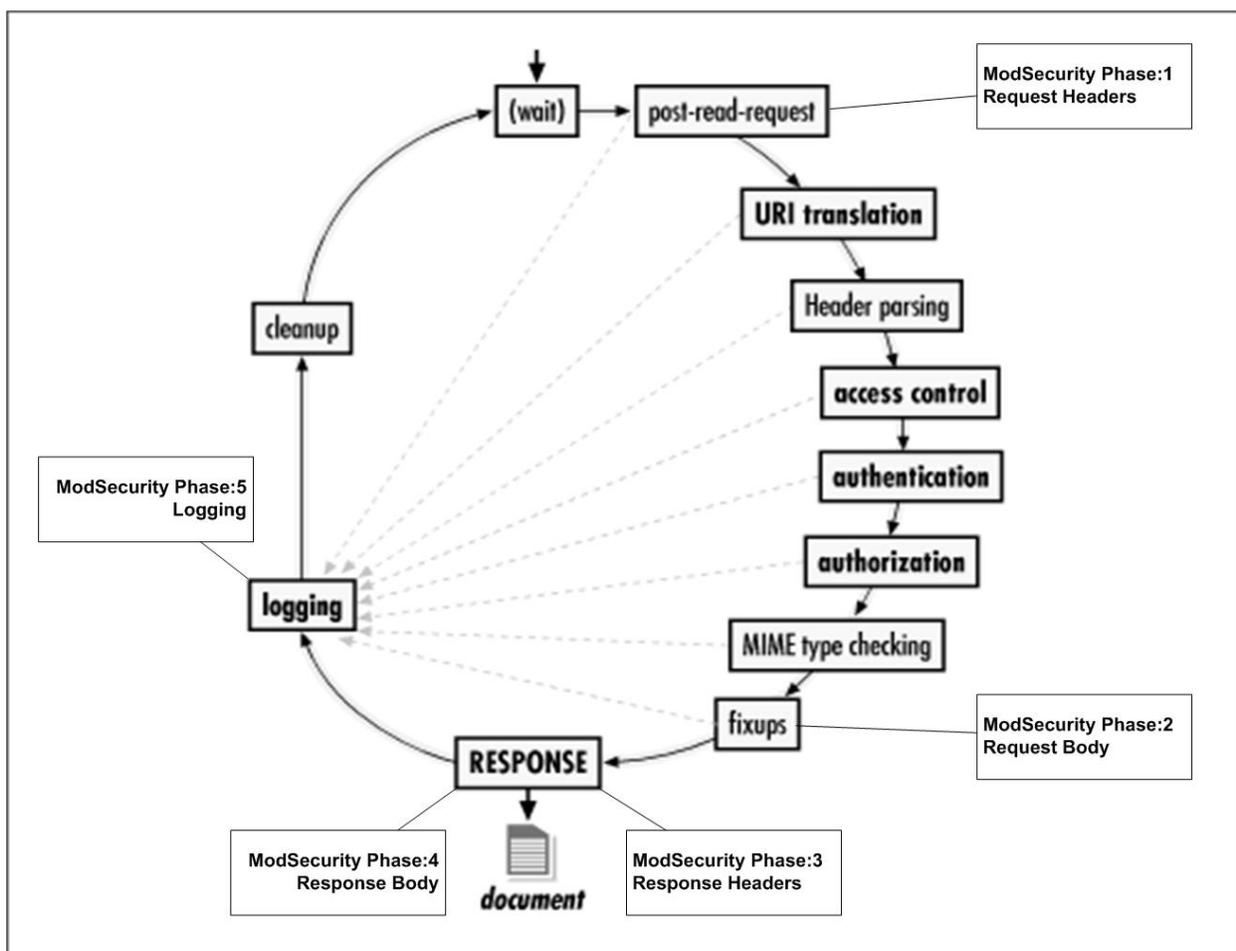
1. You are logging transactions/alerts to the ModSecurity Console and you want to use the web application ID to search only the transactions belonging to that application.
2. You are using the data persistence facility (collections SESSION and USER) and you need to avoid collisions between sessions and users belonging to different applications.

Processing Phases

ModSecurity 2.x allows rules to be placed in one of the following five phases:

1. Request headers (REQUEST_HEADERS)
2. Request body (REQUEST_BODY)
3. Response headers (RESPONSE_HEADERS)
4. Response body (RESPONSE_BODY)
5. Logging (LOGGING)

Below is a diagram of the standard Apache Request Cycle. In the diagram, the 5 ModSecurity processing phases are shown.



In order to select the phase a rule executes during, use the phase action either directly in the rule or in using the `SecDefaultAction` directive:

```
SecDefaultAction "log,pass,phase:2"
SecRule REQUEST_HEADERS:Host "!^$" "deny,phase:1"
```

Note

Keep in mind that rules are executed according to phases, so even if two rules are adjacent in a configuration file, but are set to execute in different phases, they would not happen one after the other. The order of rules in the configuration file is important only within the rules of each phase. This is especially important when using the `skip` and `skipAfter` actions.

Note

The `LOGGING` phase is special. It is executed at the end of each transaction no matter what happened in the previous phases. This means it will be processed even if the request was intercepted or the `allow` action was used to pass the transaction through.

Phase Request Headers

Rules in this phase are processed immediately after Apache completes reading the request headers (post-read-request phase). At this point the request body has not been read yet, meaning not all request arguments are available. Rules should be placed in this phase if you need to have them run early (before Apache does something with the request), to do something before the request body has been read, determine whether or not the request body should be buffered, or decide how you want the request body to be processed (e.g. whether to parse it as XML or not).

Note

Rules in this phase can not leverage Apache scope directives (`Directory`, `Location`, `LocationMatch`, etc...) as the post-read-request hook does not have this information yet. The exception here is the `VirtualHost` directive. If you want to use ModSecurity rules inside Apache locations, then they should run in Phase 2. Refer to the Apache Request Cycle/ModSecurity Processing Phases diagram.

Phase Request Body

This is the general-purpose input analysis phase. Most of the application-oriented rules should go here. In this phase you are guaranteed to have received the request arguments (provided the request body has been read). ModSecurity supports three encoding types for the request body phase:

- `application/x-www-form-urlencoded` - used to transfer form data
- `multipart/form-data` - used for file transfers
- `text/xml` - used for passing XML data

Other encodings are not used by most web applications.

Phase Response Headers

This phase takes place just before response headers are sent back to the client. Run here if you want to observe the response before that happens, and if you want to use the response headers to determine if you want to buffer the response body. Note that some response status codes (such as 404) are handled earlier

in the request cycle by Apache and may not be able to be triggered as expected. Additionally, there are some response headers that are added by Apache at a later hook (such as Date, Server and Connection) that we would not be able to trigger on or sanitize. This should work appropriately in a proxy setup or within phase:5 (logging).

Phase Response Body

This is the general-purpose output analysis phase. At this point you can run rules against the response body (provided it was buffered, of course). This is the phase where you would want to inspect the outbound HTML for information disclosure, error messages or failed authentication text.

Phase Logging

This phase is run just before logging takes place. The rules placed into this phase can only affect how the logging is performed. This phase can be used to inspect the error messages logged by Apache. You cannot deny/block connections in this phase as it is too late. This phase also allows for inspection of other response headers that weren't available during phase:3 or phase:4. Note that you must be careful not to inherit a disruptive action into a rule in this phase as this is a configuration error in ModSecurity 2.5.0 and later versions.

Variables

The following variables are supported in ModSecurity 2.x:

ARGS

ARGS is a collection and can be used on its own (means all arguments including the POST Payload), with a static parameter (matches arguments with that name), or with a regular expression (matches all arguments with name that matches the regular expression). To look at only the query string or body arguments, see the `ARGS_GET` and `ARGS_POST` collections.

Some variables are actually collections, which are expanded into more variables at runtime. The following example will examine all request arguments:

```
SecRule ARGS dirty
```

Sometimes, however, you will want to look only at parts of a collection. This can be achieved with the help of the *selection operator*(colon). The following example will only look at the arguments named `p` (do note that, in general, requests can contain multiple arguments with the same name):

```
SecRule ARGS:p dirty
```

It is also possible to specify exclusions. The following will examine all request arguments for the word *dirty*, except the ones named `z` (again, there can be zero or more arguments named `z`):

```
SecRule ARGS|!ARGS:z dirty
```

There is a special operator that allows you to count how many variables there are in a collection. The following rule will trigger if there is more than zero arguments in the request (ignore the second parameter for the time being):

```
SecRule &ARGS !^0$
```

And sometimes you need to look at an array of parameters, each with a slightly different name. In this case you can specify a regular expression in the selection operator itself. The following rule will look into all arguments whose names begin with `id_`:

```
SecRule ARGS:/^id_/ dirty
```

Note

Using `ARGS:p` will not result in any invocations against the operator if argument `p` does not exist.

In ModSecurity 1.X, the `ARGS` variable stood for `QUERY_STRING + POST_PAYLOAD`, whereas now it expands to individual variables.

ARGS_COMBINED_SIZE

This variable allows you to set more targeted evaluations on the total size of the Arguments as compared with normal Apache LimitRequest directives. For example, you could create a rule to ensure that the total size of the argument data is below a certain threshold (to help prevent buffer overflow issues). Example: Block request if the size of the arguments is above 25 characters.

```
SecRule REQUEST_FILENAME "^/cgi-bin/login\.php" \
    "chain,log,deny,phase:2,t:none,t:lowercase,t:normalisePath"
SecRule ARGS_COMBINED_SIZE "@gt 25"
```

ARGS_NAMES

Is a collection of the argument names. You can search for specific argument names that you want to block. In a positive policy scenario, you can also whitelist (using an inverted rule with the ! character) only authorized argument names. Example: This example rule will only allow 2 argument names - p and a. If any other argument names are injected, it will be blocked.

```
SecRule REQUEST_FILENAME "/index.php" \
    "chain,log,deny,status:403,phase:2,t:none,t:lowercase,t:normalisePath"
SecRule ARGS_NAMES "!^(p|a)$" "t:none,t:lowercase"
```

ARGS_GET

ARGS_GET is similar to ARGS, but only contains arguments from the query string.

ARGS_GET_NAMES

ARGS_GET_NAMES is similar to ARGS_NAMES, but only contains argument names from the query string.

ARGS_POST

ARGS_POST is similar to ARGS, but only contains arguments from the POST body.

ARGS_POST_NAMES

ARGS_POST_NAMES is similar to ARGS_NAMES, but only contains argument names from the POST body.

AUTH_TYPE

This variable holds the authentication method used to validate a user. Example:

```
SecRule AUTH_TYPE "basic" log,deny,status:403,phase:1,t:lowercase
```

Note

This data will not be available in a proxy-mode deployment as the authentication is not local. In a proxy-mode deployment, you would need to inspect the `REQUEST_HEADERS:Authorization` header.

ENV

Collection, requires a single parameter (after colon). The `ENV` variable is set with `setenv` and does not give access to the CGI environment variables. Example:

```
SecRule REQUEST_FILENAME "printenv" pass,setenv:tag=suspicious
SecRule ENV:tag "suspicious"
```

FILES

Collection. Contains a collection of original file names (as they were called on the remote user's file system). Note: only available if files were extracted from the request body. Example:

```
SecRule FILES "\.conf$" log,deny,status:403,phase:2
```

FILES_COMBINED_SIZE

Single value. Total size of the uploaded files. Note: only available if files were extracted from the request body. Example:

```
SecRule FILES_COMBINED_SIZE "@gt 1000" log,deny,status:403,phase:2
```

FILES_NAMES

Collection w/o parameter. Contains a list of form fields that were used for file upload. Note: only available if files were extracted from the request body. Example:

```
SecRule FILES_NAMES "^upfile$" log,deny,status:403,phase:2
```

FILES_SIZES

Collection. Contains a list of file sizes. Useful for implementing a size limitation on individual uploaded files. Note: only available if files were extracted from the request body. Example:

```
SecRule FILES_SIZES "@gt 100" log,deny,status:403,phase:2
```

FILES_TMPNAMES

Collection. Contains a collection of temporary files' names on the disk. Useful when used together with

@inspectFile. Note: only available if files were extracted from the request body. Example:

```
SecRule FILES_TMPNAMES "@inspectFile /path/to/inspect_script.pl"
```

GEO

GEO is a collection populated by the @geoLookups operator. It can be used to match geographical fields looked up by an IP address or hostname.

Available since 2.2.0.

Fields:

- *COUNTRY_CODE*: Two character country code. EX: US, UK, etc.
- *COUNTRY_CODE3*: Up to three character country code.
- *COUNTRY_NAME*: The full country name.
- *COUNTRY_CONTINENT*: The two character continent that the country is located. EX: EU
- *REGION*: The two character region. For US, this is state. For Canada, providence, etc.
- *CITY*: The city name.
- *POSTAL_CODE*: The postal code.
- *LATITUDE*: The latitude.
- *LONGITUDE*: The longitude.
- *DMA_CODE*: The metropolitan area code. (US only)
- *AREA_CODE*: The phone system area code. (US only)

Example:

```
SecRule REMOTE_ADDR "@geoLookup" "chain,drop,msg:'Non-UK IP address'"  
SecRule GEO:COUNTRY_CODE "!@streq UK"
```

HIGHEST_SEVERITY

This variable holds the highest severity of any rules that have matched so far. Severities are numeric values and thus can be used with comparison operators such as @lt, etc.

Note

Higher severities have a lower numeric value.

A value of 255 indicates no severity has been set.

```
SecRule HIGHEST_SEVERITY "@le 2" "phase:2,deny,status:500,msg:'severity %{HIGHEST_SEVERITY}'"
```

MATCHED_VAR

This variable holds the value of the variable that was matched against. It is similar to the TX:0, except it can be used for all operators and does not require that the `capture` action be specified.

```
SecRule ARGS pattern chain,deny
...
SecRule MATCHED_VAR "further scrutiny"
```

MATCHED_VAR_NAME

This variable holds the full name of the variable that was matched against.

```
SecRule ARGS pattern setvar:tx.mymatch=%{MATCHED_VAR_NAME}
...
SecRule TX:MYMATCH "@eq ARGS:param" deny
```

MODSEC_BUILD

This variable holds the ModSecurity build number. This variable is intended to be used to check the build number prior to using a feature that is available only in a certain build. Example:

```
SecRule MODSEC_BUILD "!@ge 02050102" skipAfter:12345
SecRule ARGS "@pm some key words" id:12345,deny,status:500
```

MULTIPART_CRLF_LF_LINES

This flag variable will be set to 1 whenever a multi-part request uses mixed line terminators. The multipart/form-data RFC requires CRLF sequence to be used to terminate lines. Since some client implementations use only LF to terminate lines you might want to allow them to proceed under certain circumstances (if you want to do this you will need to stop using MULTIPART_STRICT_ERROR and check each multi-part flag variable individually, avoiding MULTIPART_LF_LINE). However, mixing CRLF and LF line terminators is dangerous as it can allow for evasion. Therefore, in such cases, you will have to add a check for MULTIPART_CRLF_LF_LINES.

MULTIPART_STRICT_ERROR

MULTIPART_STRICT_ERROR will be set to 1 when any of the following variables is also set to 1: REQBODY_PROCESSOR_ERROR, MULTIPART_BOUNDARY_QUOTED, MULTIPART_BOUNDARY_WHITESPACE, MULTIPART_DATA_BEFORE, MULTIPART_DATA_AFTER, MULTIPART_HEADER_FOLDING, MULTIPART_LF_LINE, MULTIPART_SEMICOLON_MISSING. Each of these variables covers one unusual (although sometimes legal) aspect of the request body in multipart/form-data format. Your policies should *always* contain a rule to check either this variable (easier) or one or more individual variables (if you know exactly what you want to accomplish). Depending on the rate of false positives and your default policy you should decide whether to block or

just warn when the rule is triggered.

The best way to use this variable is as in the example below:

```
SecRule MULTIPART_STRICT_ERROR "!@eq 0" \
"phase:2,t:none,log,deny,msg:'Multipart request body \
failed strict validation: \
PE %{REQBODY_PROCESSOR_ERROR}, \
BQ %{MULTIPART_BOUNDARY_QUOTED}, \
BW %{MULTIPART_BOUNDARY_WHITESPACE}, \
DB %{MULTIPART_DATA_BEFORE}, \
DA %{MULTIPART_DATA_AFTER}, \
HF %{MULTIPART_HEADER_FOLDING}, \
LF %{MULTIPART_LF_LINE}, \
SM %{MULTIPART_SEMICOLON_MISSING}' "
```

The multipart/form-data parser was upgraded in ModSecurity v2.1.3 to actively look for signs of evasion. Many variables (as listed above) were added to expose various facts discovered during the parsing process. The MULTIPART_STRICT_ERROR variable is handy to check on all abnormalities at once. The individual variables allow detection to be fine-tuned according to your circumstances in order to reduce the number of false positives. Detailed analysis of various evasion techniques covered will be released as a separated document at a later date.

MULTIPART_UNMATCHED_BOUNDARY

Set to 1 when, during the parsing phase of a multipart/request-body, ModSecurity encounters what feels like a boundary but it is not. Such an event may occur when evasion of ModSecurity is attempted.

The best way to use this variable is as in the example below:

```
SecRule MULTIPART_UNMATCHED_BOUNDARY "!@eq 0" \
"phase:2,t:none,log,deny,msg:'Multipart parser detected a possible unmatched boundary.' "
```

Change the rule from blocking to logging-only if many false positives are encountered.

PATH_INFO

Besides passing query information to a script/handler, you can also pass additional data, known as extra path information, as part of the URL. Example:

```
SecRule PATH_INFO "^/(bin|etc|sbin|opt|usr) "
```

QUERY_STRING

This variable holds form data passed to the script/handler by appending data after a question mark. Warning: Not URL-decoded. Example:

```
SecRule QUERY_STRING "attack"
```

REMOTE_ADDR

This variable holds the IP address of the remote client. Example:

```
SecRule REMOTE_ADDR "^192\.168\.1\.101$"
```

REMOTE_HOST

If HostnameLookUps are set to On, then this variable will hold the DNS resolved remote host name. If it is set to Off, then it will hold the remote IP address. Possible uses for this variable would be to deny known bad client hosts or network blocks, or conversely, to allow in authorized hosts. Example:

```
SecRule REMOTE_HOST "\.evil\.network\.org$"
```

REMOTE_PORT

This variable holds information on the source port that the client used when initiating the connection to our web server. Example: in this example, we are evaluating to see if the REMOTE_PORT is less than 1024, which would indicate that the user is a privileged user (root).

```
SecRule REMOTE_PORT "@lt 1024" phase:1,log,pass,setenv:remote_port=privileged
```

REMOTE_USER

This variable holds the username of the authenticated user. If there are no password (basic|digest) access controls in place, then this variable will be empty. Example:

```
SecRule REMOTE_USER "admin"
```

Note

This data will not be available in a proxy-mode deployment as the authentication is not local.

REQBODY_PROCESSOR

Built-in processors are URLENCODED, MULTIPART, and XML. Example:

```
SecRule REQBODY_PROCESSOR "^XML$" chain
SecRule XML "@validatedDTD /opt/apache-frontent/conf/xml.dtd"
```

REQBODY_PROCESSOR_ERROR

Possible values are 0 (no error) or 1 (error). This variable will be set by request body processors (typically the `multipart/request-data` parser or the XML parser) when they fail to properly parse a request payload.

Example:

```
SecRule REQBODY_PROCESSOR_ERROR "@eq 1" deny,phase:2
```

Note

Your policies *must* have a rule to check `REQBODY_PROCESSOR_ERROR` at the beginning of phase 2. Failure to do so will leave the door open for impedance mismatch attacks. It is possible, for example, that a payload that cannot be parsed by ModSecurity can be successfully parsed by more tolerant parser operating in the application. If your policy dictates blocking then you should reject the request if error is detected. When operating in detection-only mode your rule should alert with high severity when request body processing fails.

REQBODY_PROCESSOR_ERROR_MSG

Empty, or contains the error message from the processor. Example:

```
SecRule REQBODY_PROCESSOR_ERROR_MSG "failed to parse" t:lowercase
```

REQUEST_BASENAME

This variable holds just the filename part of `REQUEST_FILENAME` (e.g. `index.php`).

Example:

```
SecRule REQUEST_BASENAME "^login\\.php$" phase:2,t:none,t:lowercase
```

Note

Please note that anti-evasion transformations are not applied to this variable by default. `REQUEST_BASENAME` will recognise both `/` and `\` as path separators.

REQUEST_BODY

This variable holds the data in the request body (including `POST_PAYLOAD` data). `REQUEST_BODY` should be used if the original order of the arguments is important (`ARGS` should be used in all other cases). Example:

```
SecRule REQUEST_BODY "^username=\w{25,}&password=\w{25,}&Submit\=login$" phase:2,t:none,t:lowercase
```

Note

This variable is only available if the URLENCODED request body processor parsed a request body. This will occur by default when an `application/x-www-form-urlencoded` is detected, or the URLENCODED request body parser is forced. As of 2.5.7 it is possible to force the presence of the REQUEST_BODY variable, but only when there is no request body processor defined, using the `ctl:forceRequestBodyVariable` option in the REQUEST_HEADERS phase.

REQUEST_COOKIES

This variable is a collection of all of the cookie data. Example: the following example is using the Ampersand special operator to count how many variables are in the collection. In this rule, it would trigger if the request does not include any Cookie headers.

```
SecRule &REQUEST_COOKIES "@eq 0"
```

REQUEST_COOKIES_NAMES

This variable is a collection of the cookie names in the request headers. Example: the following rule will trigger if the JSESSIONID cookie is not present.

```
SecRule &REQUEST_COOKIES_NAMES:JSESSIONID "@eq 0"
```

REQUEST_FILENAME

This variable holds the relative REQUEST_URI minus the QUERY_STRING part (e.g. `/index.php`). Example:

```
SecRule REQUEST_FILENAME "^/cgi-bin/login\.php$" phase:2,t:none,t:normalisePath
```

Note

Please note that anti-evasion transformations are not used on REQUEST_FILENAME by default.

REQUEST_HEADERS

This variable can be used as either a collection of all of the request headers or can be used to specify individual headers (by using `REQUEST_HEADERS:Header-Name`). Example: the first example uses REQUEST_HEADERS as a collection and is applying the `validateUrlEncoding` operator against all headers.

```
SecRule REQUEST_HEADERS "@validateUrlEncoding"
```

Example: the second example is targeting only the Host header.

```
SecRule REQUEST_HEADERS:Host "^[\\d\\.]+$" \
    "deny,log,status:400,msg:'Host header is a numeric IP address'"
```

REQUEST_HEADERS_NAMES

This variable is a collection of the names of all of the request headers. Example:

```
SecRule REQUEST_HEADERS_NAMES "^x-forwarded-for" \
    "log,deny,status:403,t:lowercase,msg:'Proxy Server Used'"
```

REQUEST_LINE

This variable holds the complete request line sent to the server (including the REQUEST_METHOD and HTTP version data). Example: this example rule will trigger if the request method is something other than GET, HEAD, POST or if the HTTP is something other than HTTP/0.9, 1.0 or 1.1.

```
SecRule REQUEST_LINE "!^(?:(:?(:pos|ge)t|head))|http/(0\\.9|1\\.0|1\\.1)$" t:none,t:lowercase
```

REQUEST_METHOD

This variable holds the request method used by the client.

The following example will trigger if the request method is either CONNECT or TRACE.

```
SecRule REQUEST_METHOD "^(?:connect|trace)$" t:none,t:lowercase
```

REQUEST_PROTOCOL

This variable holds the request protocol version information. Example:

```
SecRule REQUEST_PROTOCOL "!^http/(0\\.9|1\\.0|1\\.1)$" t:none,t:lowercase
```

REQUEST_URI

This variable holds the full URL including the QUERY_STRING data (e.g. /index.php?p=X), however it will never contain a domain name, even if it was provided on the request line. It also does not include either the REQUEST_METHOD or the HTTP version info.

Example:

```
SecRule REQUEST_URI "attack" phase:1,t:none,t:urlDecode,t:lowercase,t:normalisePath
```

Note

Please note that anti-evasion transformations are not used on REQUEST_URI by default.

REQUEST_URI_RAW

Same as REQUEST_URI but will contain the domain name if it was provided on the request line (e.g. http://www.example.com/index.php?p=X).

Example:

```
SecRule REQUEST_URI_RAW "http:/" phase:1,t:none,t:urlDecode,t:lowercase,t:normalisePath
```

Note

Please note that anti-evasion transformations are not used on REQUEST_URI_RAW by default.

RESPONSE_BODY

This variable holds the data for the response payload.

Example:

```
SecRule RESPONSE_BODY "ODBC Error Code"
```

RESPONSE_CONTENT_LENGTH

Response body length in bytes. Can be available starting with phase 3 but it does not have to be (as the length of response body is not always known in advance.) If the size is not known this variable will contain a zero. If RESPONSE_CONTENT_LENGTH contains a zero in phase 5 that means the actual size of the response body was 0.

The value of this variable can change between phases if the body is modified. For example, in embedded mode `mod_deflate` can compress the response body between phases 4 and 5.

RESPONSE_CONTENT_TYPE

Response content type. Only available starting with phase 3.

RESPONSE_HEADERS

This variable is similar to the REQUEST_HEADERS variable and can be used in the same manner. Example:

```
SecRule RESPONSE_HEADERS:X-Cache "MISS"
```

Note

This variable may not have access to some headers when running in embedded-mode. Headers such as Server, Date, Connection and Content-Type are added during a later Apache hook just prior to sending the data to the client. This data should be available, however, either during ModSecurity phase:5 (logging) or when running in proxy-mode.

RESPONSE_HEADERS_NAMES

This variable is a collection of the response header names. Example:

```
SecRule RESPONSE_HEADERS_NAMES "Set-Cookie"
```

Note

Same limitations as RESPONSE_HEADERS with regards to access to some headers in embedded-mode.

RESPONSE_PROTOCOL

This variable holds the HTTP response protocol information. Example:

```
SecRule RESPONSE_PROTOCOL "^HTTP\/0\.9"
```

RESPONSE_STATUS

This variable holds the HTTP response status code as generated by Apache. Example:

```
SecRule RESPONSE_STATUS "^[45]"
```

Note

This directive may not work as expected in embedded-mode as Apache handles many of the stock response codes (404, 401, etc...) earlier in Phase 2. This variable should work as expected in a proxy-mode deployment.

RULE

This variable provides access to the id, rev, severity, logdata, and msg fields of the rule that triggered the action. Only available for expansion in action strings (e.g.setvar:tx.varname=%{rule.id}). Example:

```
SecRule &REQUEST_HEADERS:Host "@eq 0" "log,deny,setvar:tx.varname=%{rule.id}"
```

SCRIPT_BASENAME

This variable holds just the local filename part of SCRIPT_FILENAME. Example:

```
SecRule SCRIPT_BASENAME "^login\.php$"
```

Note

This variable is not available in proxy mode.

SCRIPT_FILENAME

This variable holds the full path on the server to the requested script. (e.g. SCRIPT_NAME plus the server path). Example:

```
SecRule SCRIPT_FILENAME "^/usr/local/apache/cgi-bin/login\.php$"
```

Note

This variable is not available in proxy mode.

SCRIPT_GID

This variable holds the group id (numerical value) of the group owner of the script. Example:

```
SecRule SCRIPT_GID "!\^46$"
```

Note

This variable is not available in proxy mode.

SCRIPT_GROUPNAME

This variable holds the group name of the group owner of the script. Example:

```
SecRule SCRIPT_GROUPNAME "!\^apache$"
```

Note

This variable is not available in proxy mode.

SCRIPT_MODE

This variable holds the script's permissions mode data (numerical - 1=execute, 2=write, 4=read and 7=read/write/execute). Example: will trigger if the script has the WRITE permissions set.

```
SecRule SCRIPT_MODE "^(2|3|6|7)$"
```

Note

This variable is not available in proxy mode.

SCRIPT_UID

This variable holds the user id (numerical value) of the owner of the script. Example: the example rule below will trigger if the UID is not 46 (the Apache user).

```
SecRule SCRIPT_UID "!^46$"
```

Note

This variable is not available in proxy mode.

SCRIPT_USERNAME

This variable holds the username of the owner of the script. Example:

```
SecRule SCRIPT_USERNAME "!^apache$"
```

Note

This variable is not available in proxy mode.

SERVER_ADDR

This variable contains the IP address of the server. Example:

```
SecRule SERVER_ADDR "^192\.168\.1\.100$"
```

SERVER_NAME

This variable contains the server's hostname or IP address. Example:

```
SecRule SERVER_NAME "hostname\.com$"
```

Note

This data is taken from the Host header submitted in the client request.

SERVER_PORT

This variable contains the local port that the web server is listening on. Example:

```
SecRule SERVER_PORT "^80$"
```

SESSION

This variable is a collection, available only after `setuid` is executed. Example: the following example

shows how to initialize a SESSION collection with `setuid`, how to use `setvar` to increase the `session.score` values, how to set the `session.blocked` variable and finally how to deny the connection based on the `session:blocked` value.

```
SecRule REQUEST_COOKIES:PHPSESSID !^$ chain,nolog,pass
SecAction setuid:%{REQUEST_COOKIES.PHPSESSID}
SecRule REQUEST_URI "^/cgi-bin/finger$" \
    "phase:2,t:none,t:lowercase,t:normalisePath,pass,log,setvar:session.score=+10"
SecRule SESSION:SCORE "@gt 50" "pass,log,setvar:session.blocked=1"
SecRule SESSION:BLOCKED "@eq 1" "log,deny,status:403"
```

SESSIONID

This variable is the value set with `setuid`. Example:

```
SecRule SESSIONID !^$ chain,nolog,pass
SecRule REQUEST_COOKIES:PHPSESSID !^$
SecAction setuid:%{REQUEST_COOKIES.PHPSESSID}
```

TIME

This variable holds a formatted string representing the time (hour:minute:second). Example:

```
SecRule TIME "^\(([1](8|9))|([2](0|1|2|3))):\d{2}:\d{2}$"
```

TIME_DAY

This variable holds the current date (1-31). Example: this rule would trigger anytime between the 10th and 20th days of the month.

```
SecRule TIME_DAY "^\(([1](0|1|2|3|4|5|6|7|8|9))|20)$"
```

TIME_EPOCH

This variable holds the time in seconds since 1970. Example:

```
SecRule TIME_EPOCH "@gt 1000"
```

TIME_HOUR

This variable holds the current hour (0-23). Example: this rule would trigger during "off hours".

```
SecRule TIME_HOUR "^(0|1|2|3|4|5|6|[1](8|9)|[2](0|1|2|3))$"
```

TIME_MIN

This variable holds the current minute (0-59). Example: this rule would trigger during the last half hour of every hour.

```
SecRule TIME_MIN "^(3|4|5)"
```

TIME_MON

This variable holds the current month (0-11). Example: this rule would match if the month was either November (10) or December (11).

```
SecRule TIME_MON "^1"
```

TIME_SEC

This variable holds the current second count (0-59). Example:

```
SecRule TIME_SEC "@gt 30"
```

TIME_WDAY

This variable holds the current weekday (0-6). Example: this rule would trigger only on week-ends (Saturday and Sunday).

```
SecRule TIME_WDAY "^(0|6)$"
```

TIME_YEAR

This variable holds the current four-digit year data. Example:

```
SecRule TIME_YEAR "^2006$"
```

TX

Transaction Collection. This is used to store pieces of data, create a transaction anomaly score, and so on. Transaction variables are set for 1 request/response cycle. The scoring and evaluation will not last past the current request/response process. Example: In this example, we are using setvar to increase the tx.score value by 5 points. We then have a follow-up run that will evaluate the transactional score this request and then it will decided whether or not to allow/deny the request through.

The following is a list of reserved names in the TX collection:

- TX:0 - The matching value when using the @rx or @pm operator with the capture action.

- TX:1-TX:9 - The captured subexpression value when using the @rx operator with capturing parens and the capture action.

```
SecRule WEBSERVER_ERROR_LOG "does not exist" "phase:5,pass,setvar:tx.score=+5"
SecRule TX:SCORE "@gt 20" deny,log
```

USERID

This variable is the value set with `setuid`. Example:

```
SecAction setuid:%{REMOTE_USER},nolog
SecRule USERID "Admin"
```

WEBAPPID

This variable is the value set with `SecWebAppId`. Example:

```
SecWebAppId "WebApp1"
SecRule WEBAPPID "WebApp1" "chain,log,deny,status:403"
SecRule REQUEST_HEADERS:Transfer-Encoding "!^$"
```

WEBSERVER_ERROR_LOG

Contains zero or more error messages produced by the web server. Access to this variable is in phase:5 (logging). Example:

```
SecRule WEBSERVER_ERROR_LOG "File does not exist" "phase:5,setvar:tx.score=+5"
```

XML

Can be used standalone (as a target for `validateDTD` and `validateSchema`) or with an XPath expression parameter (which makes it a valid target for any function that accepts plain text). Example using XPath:

```
SecDefaultAction log,deny,status:403,phase:2
SecRule REQUEST_HEADERS:Content-Type ^text/xml$ \
  phase:1,t:lowercase,nolog,pass,ctl:requestBodyProcessor=XML
SecRule REQBODY_PROCESSOR "!^XML$" skipAfter:12345
SecRule XML:/employees/employee/name/text() Fred
SecRule XML:/xq:employees/employee/name/text() Fred \
  id:12345,xmlns:xq=http://www.example.com/employees
```

The first XPath expression does not use namespaces. It would match against payload such as this one:

```
<employees>
  <employee>
    <name>Fred Jones</name>
    <address location="home">
      <street>900 Aurora Ave.</street>
      <city>Seattle</city>
      <state>WA</state>
      <zip>98115</zip>
    </address>
    <address location="work">
      <street>2011 152nd Avenue NE</street>
      <city>Redmond</city>
      <state>WA</state>
      <zip>98052</zip>
    </address>
    <phone location="work">(425)555-5665</phone>
    <phone location="home">(206)555-5555</phone>
    <phone location="mobile">(206)555-4321</phone>
  </employee>
</employees>
```

The second XPath expression does use namespaces. It would match the following payload:

```
<xq:employees xmlns:xq="http://www.example.com/employees">
  <employee>
    <name>Fred Jones</name>
    <address location="home">
      <street>900 Aurora Ave.</street>
      <city>Seattle</city>
      <state>WA</state>
      <zip>98115</zip>
    </address>
    <address location="work">
      <street>2011 152nd Avenue NE</street>
      <city>Redmond</city>
      <state>WA</state>
      <zip>98052</zip>
    </address>
    <phone location="work">(425)555-5665</phone>
    <phone location="home">(206)555-5555</phone>
    <phone location="mobile">(206)555-4321</phone>
  </employee>
</xq:employees>
```

Note the different namespace used in the second example.

To learn more about XPath we suggest the following resources:

1. XPath Standard [<http://www.w3.org/TR/xpath>]
2. XPath Tutorial [<http://www.zvon.org/xxl/XPathTutorial/General/examples.html>]

Transformation functions

When ModSecurity receives request or response information, it makes a copy of this data and places it into memory. It is on this data in memory that transformation functions are applied. The raw request/response data is never altered. Transformation functions are used to transform a variable before testing it in a rule.

Note

There are no default transformation functions as there were in previous versions of ModSecurity.

The following rule will ensure that an attacker does not use mixed case in order to evade the ModSecurity rule:

```
SecRule ARG:p "xp_cmdshell" "t:lowercase"
```

multiple transformation actions can be used in the same rule, for example the following rule also ensures that an attacker does not use URL encoding (%xx encoding) for evasion. Note the order of the transformation functions, which ensures that a URL encoded letter is first decoded and then translated to lower case.

```
SecRule ARG:p "xp_cmdshell" "t:urlDecode,t:lowercase"
```

One can use the SecDefaultAction command to ensure the translation occurs for every rule until the next. Note that transformation actions are additive, so if a rule explicitly list actions, the translation actions set by SecDefaultAction are still performed.

```
SecDefaultAction t:urlDecode,t:lowercase
```

The following transformation functions are supported:

base64Decode

This function decodes a base64-encoded string.

base64Encode

This function encodes input string using base64 encoding.

compressWhitespace

It converts whitespace characters (32, \f, \t, \n, \r, \v, 160) to spaces (ASCII 32) and then compresses multiple consecutive space characters into one.

cssDecode

Decodes CSS-encoded characters, as specified at <http://www.w3.org/TR/REC-CSS2/syndata.html>. This

function uses only up to two bytes in the decoding process, meaning it is useful to uncover ASCII characters (that wouldn't normally be encoded) encoded using CSS encoding, or to counter evasion which is a combination of a backslash and non-hexadecimal characters (e.g. `ja\vascript` is equivalent to `javascript`).

escapeSeqDecode

This function decode ANSI C escape sequences: `\a`, `\b`, `\f`, `\n`, `\r`, `\t`, `\v`, `\\`, `\?`, `\'`, `\"`, `\xHH` (hexadecimal), `\0000` (octal). Invalid encodings are left in the output.

hexDecode

This function decodes a hex-encoded string.

hexEncode

This function encodes input as hex-encoded string.

htmlEntityDecode

This function decodes HTML entities present in input. The following variants are supported:

- `&#xHH` and `&#xHH;` (where H is any hexadecimal number)
- `&#DDD` and `&#DDD;` (where D is any decimal number)
- `"` and `"`;
- ` sp` and ` sp;`;
- `<` and `<`;
- `>` and `>`;

This function will convert any entity into a single byte only, possibly resulting in a loss of information. It is thus useful to uncover bytes that would otherwise not need to be encoded, but it cannot do anything with the characters from the range above 255.

jsDecode

Decodes JavaScript escape sequences. If a `\uHHHH` code is in the range of `FF01-FF5E` (the full width ASCII codes), then the higher byte is used to detect and adjust the lower byte. Otherwise, only the lower byte will be used and the higher byte zeroed.

length

This function converts the input to its numeric length (count of bytes).

lowercase

This function converts all characters to lowercase using the current C locale.

md5

This function calculates an MD5 hash from input. Note that the computed hash is in a raw binary form and may need encoded into text to be usable (for example: `t:md5`, `t:hexEncode`).

none

Not an actual transformation function, but an instruction to ModSecurity to remove all transformation functions associated with the current rule.

normalisePath

This function will remove multiple slashes, self-references and directory back-references (except when they are at the beginning of the input).

normalisePathWin

Same as `normalisePath`, but will first convert backslash characters to forward slashes.

parityEven7bit

This function calculates even parity of 7-bit data replacing the 8th bit of each target byte with the calculated parity bit.

parityOdd7bit

This function calculates odd parity of 7-bit data replacing the 8th bit of each target byte with the calculated parity bit.

parityZero7bit

This function calculates zero parity of 7-bit data replacing the 8th bit of each target byte with a zero parity bit which allows inspection of even/odd parity 7bit data as ASCII7 data.

removeNulls

This function removes NULL bytes from input.

removeWhitespace

This function removes all whitespace characters from input.

replaceComments

This function replaces each occurrence of a C-style comments (`/* ... */`) with a single space (multiple consecutive occurrences of a space will not be compressed). Unterminated comments will too

be replaced with a space (ASCII 32). However, a standalone termination of a comment (`*/`) will not be acted upon.

replaceAll

This function is enabled by default. It replaces NULL bytes in input with spaces (ASCII 32).

urlDecode

This function decodes an URL-encoded input string. Invalid encodings (i.e. the ones that use non-hexadecimal characters, or the ones that are at the end of string and have one or two characters missing) will not be converted. If you want to detect invalid encodings use the `@validateUrlEncoding` operator. The transformation function should not be used against variables that have already been URL-decoded unless it is your intention to perform URL decoding twice!

urlDecodeUni

In addition to decoding `%xx` like `urlDecode`, `urlDecodeUni` also decodes `%uXXXX` encoding. If the code is in the range of `FF01-FF5E` (the full width ASCII codes), then the higher byte is used to detect and adjust the lower byte. Otherwise, only the lower byte will be used and the higher byte zeroed.

urlEncode

This function encodes input using URL encoding.

sha1

This function calculates a SHA1 hash from input. Note that the computed hash is in a raw binary form and may need encoded to be usable (for example: `t:sha1`, `t:hexEncode`).

trimLeft

This function removes whitespace from the left side of input.

trimRight

This function removes whitespace from the right side of input.

trim

This function removes whitespace from both the left and right sides of input.

Actions

Each action belongs to one of five groups:

Disruptive actions	Cause ModSecurity to do something. In many cases something means block transaction, but not in all. For example, the <code>allow</code> action is classified as a disruptive action, but it does the opposite of blocking. There can only be one disruptive action per rule (if there are multiple disruptive actions present, or inherited, only the last one will take effect), or rule chain (in a chain, a disruptive action can only appear in the first rule).
Non-disruptive actions	Do something, but that something does not and cannot affect the rule processing flow. Setting a variable, or changing its value is an example of a non-disruptive action. Non-disruptive action can appear in any rule, including each rule belonging to a chain.
Flow actions	These actions affect the rule flow (for example <code>skip</code> or <code>skipAfter</code>).
Meta-data actions	Meta-data actions are used to provide more information about rules. Examples include <code>id</code> , <code>rev</code> , <code>severity</code> and <code>msg</code> .
Data actions	Not really actions, these are mere containers that hold data used by other actions. For example, the <code>status</code> action holds the status that will be used for blocking (if it takes place).

`allow`

Description: Stops rule processing on a successful match and allows the transaction to proceed.

Action Group: Disruptive

Example:

```
SecRule REMOTE_ADDR "^192\.168\.1\.100$" nolog,phase:1,allow
```

Prior to ModSecurity 2.5 the `allow` action would only affect the current phase. An `allow` in phase 1 would skip processing the remaining rules in phase 1 but the rules from phase 2 would execute. Starting with v2.5.0 `allow` was enhanced to allow for fine-grained control of what is done. The following rules now apply:

1. If used on its own, like in the example above, `allow` will affect the entire transaction, stopping processing of the current phase but also skipping over all other phases apart from the logging phase. (The logging phase is special; it is designed to always execute.)
2. If used with parameter "phase", `allow` will cause the engine to stop processing the current phase. Other phases will continue as normal.
3. If used with parameter "request", `allow` will cause the engine to stop processing the current

phase. The next phase to be processed will be phase `RESPONSE_HEADERS`.

Examples:

```
# Do not process request but process response.
SecAction phase:1,allow:request
```

```
# Do not process transaction (request and response).
SecAction phase:1,allow
```

If you want to allow a response through, put a rule in phase `RESPONSE_HEADERS` and simply use `allow` on its own:

```
# Allow response through.
SecAction phase:3,allow
```

append

Description: Appends text given as parameter to the end of response body. For this action to work content injection must be enabled by setting `SecContentInjection` to `On`. Also make sure you check the content type of the response before you make changes to it (e.g. you don't want to inject stuff into images).

Action Group: Non-disruptive

Processing Phases: 3 and 4.

Example:

```
SecRule RESPONSE_CONTENT_TYPE "^text/html" "nolog,pass,append:'<hr>Footer' "
```

Note

While macro expansion is allowed in the additional content, you are strongly cautioned against inserting user defined data fields.

auditlog

Description: Marks the transaction for logging in the audit log.

Action Group: Non-disruptive

Example:

```
SecRule REMOTE_ADDR "^192\.168\.1\.100$" auditlog,phase:1,allow
```

Note

The `auditlog` action is now explicit if `log` is already specified.

block

Description: Performs the default disruptive action.

Action Group: Disruptive

It is intended to be used by ruleset writers to signify that the rule was intended to block and leaves the "how" up to the administrator. This action is currently a placeholder which will just be replaced by the action from the last `SecDefaultAction` in the same context. Using the `block` action with the `SecRuleUpdateActionById` directive allows a rule to be reverted back to the previous `SecDefaultAction` disruptive action.

In future versions of ModSecurity, more control and functionality will be added to define "how" to block.

Examples:

In the following example, the second rule will "deny" because of the `SecDefaultAction` disruptive action. The intent being that the administrator could easily change this to another disruptive action without editing the actual rules.

```
### Administrator defines "how" to block (deny,status:403)...
SecDefaultAction phase:2,deny,status:403,log,auditlog

### Included from a rulest...
# Intent is to warn for this User Agent
SecRule REQUEST_HEADERS:User-Agent "perl" "phase:2,pass,msg:'Perl based user agent identified'"
# Intent is to block for this User Agent, "how" described in SecDefaultAction
SecRule REQUEST_HEADERS:User-Agent "nikto" "phase:2,block,msg:'Nikto Scanners Identified'"
```

In the following example, The rule is reverted back to the `pass` action defined in the `SecDefaultAction` directive by using the `SecRuleUpdateActionById` directive in conjunction with the `block` action. This allows an administrator to override an action in a 3rd party rule without modifying the rule itself.

```
### Administrator defines "how" to block (deny,status:403)...
SecDefaultAction phase:2,pass,log,auditlog

### Included from a rulest...
SecRule REQUEST_HEADERS:User-Agent "nikto" "id:1,phase:2,deny,msg:'Nikto Scanners Identified'"

### Added by the administrator
SecRuleUpdateActionById 1 "block"
```

capture

Description: When used together with the regular expression operator, capture action will create copies of regular expression captures and place them into the transaction variable collection. Up to ten captures will be copied on a successful pattern match, each with a name consisting of a digit from 0 to 9.

Action Group: Non-disruptive

Example:

```
SecRule REQUEST_BODY "^username=(\w{25,})" phase:2,capture,t:none,chain
SecRule TX:1 "(?:(?:a(dmin|nonymous)))"
```

Note

The 0 data captures the entire REGEX match and 1 captures the data in the first parens, etc...

chain

Description: Chains the rule where the action is placed with the rule that immediately follows it. The result is called a *rule chain*. Chained rules allow for more complex rule matches where you want to use a number of different VARIABLES to create a better rule and to help prevent false positives.

Action Group: Flow

Example:

```
# Refuse to accept POST requests that do
# not specify request body length. Do note that
# this rule should be preceded by a rule that verifies
# only valid request methods (e.g. GET, HEAD and POST) are used.
SecRule REQUEST_METHOD ^POST$ chain,t:none
SecRule REQUEST_HEADER:Content-Length ^$ t:none
```

Note

In programming language concepts, think of chained rules somewhat similar to AND conditional statements. The actions specified in the first portion of the chained rule will only be triggered if all of the variable checks return positive hits. If one aspect of the chained rule is negative, then the entire rule chain is negative. Also note that disruptive actions, execution phases, metadata actions (id, rev, msg), skip and skipAfter actions can only be specified on by the chain starter rule.

ctl

Description: The ctl action allows configuration options to be updated for the transaction.

Action Group: Non-disruptive

Example:

```
# Parse requests with Content-Type "text/xml" as XML
SecRule REQUEST_CONTENT_TYPE ^text/xml nolog,pass,ctl:requestBodyProcessor=XML
```

Note

The following configuration options are supported:

1. auditEngine
2. auditLogParts

3. `debugLogLevel`
4. `ruleRemoveById` (single rule ID, or a single rule ID range accepted as parameter)
5. `requestBodyAccess`
6. `forceRequestBodyVariable`
7. `requestBodyLimit`
8. `requestBodyProcessor`
9. `responseBodyAccess`
10. `responseBodyLimit`
11. `ruleEngine`

With the exception of `requestBodyProcessor` and `forceRequestBodyVariable`, each configuration option corresponds to one configuration directive and the usage is identical.

The `requestBodyProcessor` option allows you to configure the request body processor. By default ModSecurity will use the `URLENCODED` and `MULTIPART` processors to process an `application/x-www-form-urlencoded` and a `multipart/form-data` bodies, respectively. A third processor, `XML`, is also supported, but it is never used implicitly. Instead you must tell ModSecurity to use it by placing a few rules in the `REQUEST_HEADERS` processing phase. After the request body was processed as `XML` you will be able to use the `XML`-related features to inspect it.

Request body processors will not interrupt a transaction if an error occurs during parsing. Instead they will set variables `REQBODY_PROCESSOR_ERROR` and `REQBODY_PROCESSOR_ERROR_MSG`. These variables should be inspected in the `REQUEST_BODY` phase and an appropriate action taken.

The `forceRequestBodyVariable` option allows you to configure the `REQUEST_BODY` variable to be set when there is no request body processor configured. This allows for inspection of request bodies of unknown types.

deny

Description: Stops rule processing and intercepts transaction.

Action Group: Disruptive

Example:

```
SecRule REQUEST_HEADERS:User-Agent "nikto" "log,deny,msg:'Nikto Scanners Identified'"
```

deprecatevar

Description: Decrement counter based on its age.

Action Group: Non-Disruptive

Example: The following example will decrement the counter by 60 every 300 seconds.

```
SecAction deprecatevar:session.score=60/300
```

Note

Counter values are always positive, meaning the value will never go below zero.

drop

Description: Immediately initiate a "connection close" action to tear down the TCP connection by sending a FIN packet.

Action Group: Disruptive

Example: The following example initiates an IP collection for tracking Basic Authentication attempts. If the client goes over the threshold of more than 25 attempts in 2 minutes, it will DROP subsequent connections.

```
SecAction initcol:ip=%{REMOTE_ADDR},nolog
SecRule ARGS:login "!^$" \
    nolog,phase:1,setvar:ip.auth_attempt=+1,deprecatevar:ip.auth_attempt=20/120
SecRule IP:AUTH_ATTEMPT "@gt 25" \
    log,drop,phase:1,msg:'Possible Brute Force Attack'
```

Note

This action is currently not available on Windows based builds. This action is extremely useful when responding to both Brute Force and Denial of Service attacks in that, in both cases, you want to minimize both the network bandwidth and the data returned to the client. This action causes error message to appear in the log "(9)Bad file descriptor: core_output_filter: writing data to the network"

exec

Description: Executes an external script/binary supplied as parameter. As of v2.5.0, if the parameter supplied to exec is a Lua script (detected by the .lua extension) the script will be processed *internally*. This means you will get direct access to the internal request context from the script. Please read the SecRuleScript documentation for more details on how to write Lua scripts.

Action Group: Non-disruptive

Example:

```
# The following is going to execute /usr/local/apache/bin/test.sh
# as a shell script on rule match.
SecRule REQUEST_URI "^/cgi-bin/script\.pl" \
    "phase:2,t:none,t:lowercase,t:normalisePath,log,exec:/usr/local/apache/bin/test.sh"

# The following is going to process /usr/local/apache/conf/exec.lua
# internally as a Lua script on rule match.
SecRule ARGS:p attack log,exec:/usr/local/apache/conf/exec.lua
```

Note

The `exec` action is executed independently from any disruptive actions. External scripts will always be called with no parameters. Some transaction information will be placed in environment variables. All the usual CGI environment variables will be there. You should be aware that forking a threaded process results in all threads being replicated in the new process. Forking can therefore incur larger overhead in multi-threaded operation. The script you execute must write something (anything) to `stdout`. If it doesn't ModSecurity will assume execution didn't work.

expirevar

Description: Configures a collection variable to expire after the given time in seconds.

Action Group: Non-disruptive

Example:

```
SecRule REQUEST_COOKIES:JSESSIONID "!^$" nolog,phase:1,pass,chain
SecAction setsid:%{REQUEST_COOKIES:JSESSIONID}
SecRule REQUEST_URI "^/cgi-bin/script\.pl" \
    "phase:2,t:none,t:lowercase,t:normalisePath,log,allow,\
setvar:session.suspicious=1,expirevar:session.suspicious=3600,phase:1"
```

Note

You should use `expirevar` actions at the same time that you use `setvar` actions in order to keep the indented expiration time. If they are used on their own (perhaps in a `SecAction` directive) the expire time could get re-set. When variables are removed from collections, and there are no other changes, collections are not written to disk at the end of request. This is because the variables can always be expired again when the collection is read again on a subsequent request.

id

Description: Assigns a unique ID to the rule or chain.

Action Group: Meta-data

Example:

```
SecRule &REQUEST_HEADERS:Host "@eq 0" \
    "log,id:60008,severity:2,msg:'Request Missing a Host Header'"
```

Note

These are the reserved ranges:

- 1-99,999; reserved for local (internal) use. Use as you see fit but do not use this range for rules that are distributed to others.
- 100,000-199,999; reserved for internal use of the engine, to assign to rules that do not have explicit IDs.
- 200,000-299,999; reserved for rules published at modsecurity.org.

- 300,000-399,999; reserved for rules published at gotroot.com.
- 400,000-419,999; unused (available for reservation).
- 420,000-429,999; reserved for ScallyWhack [<http://projects.otaku42.de/wiki/ScallyWhack>].
- 430,000-899,999; unused (available for reservation).
- 900,000-999,999; reserved for the Core Rules [<http://www.modsecurity.org/projects/rules/>] project.
- 1,000,000 and above; unused (available for reservation).

initcol

Description: Initialises a named persistent collection, either by loading data from storage or by creating a new collection in memory.

Action Group: Non-disruptive

Example: The following example initiates IP address tracking.

```
SecAction initcol:ip=%{REMOTE_ADDR},nolog
```

Note

Collections are loaded into memory when the initcol action is encountered. The collection in storage will be persisted (and the appropriate counters increased) *only* if it was changed during transaction processing.

See the "Persistent Storage" section for further details.

log

Description: Indicates that a successful match of the rule needs to be logged.

Action Group: Non-disruptive

Example:

```
SecAction initcol:ip=%{REMOTE_ADDR},log
```

Note

This action will log matches to the Apache error log file and the ModSecurity audit log.

logdata

Description: Allows a data fragment to be logged as part of the alert message.

Action Group: Non-disruptive

Example:

```
SecRule &ARGS:p "@eq 0" "log,logdata:'%{TX.0}' "
```

Note

The logdata information appears in the error and/or audit log files and is not sent back to the client in response headers. Macro expansion is preformed so you may use variable names such as %{TX.0}, etc. The information is properly escaped for use with logging binary data.

msg

Description: Assigns a custom message to the rule or chain.

Action Group: Meta-data

Example:

```
SecRule &REQUEST_HEADERS:Host "@eq 0" \  
    "log,id:60008,severity:2,msg:'Request Missing a Host Header'"
```

Note

The msg information appears in the error and/or audit log files and is not sent back to the client in response headers.

multiMatch

Description: If enabled ModSecurity will perform multiple operator invocations for every target, before and after every anti-evasion transformation is performed.

Action Group: Non-disruptive

Example:

```
SecDefaultAction log,deny,phase:1,t:removeNulls,t:lowercase  
SecRule ARGS "attack" multiMatch
```

Note

Normally, variables are evaluated once, only after all transformation functions have completed. With multiMatch, variables are checked against the operator before and after every transformation function that changes the input.

noauditlog

Description: Indicates that a successful match of the rule should not be used as criteria whether the transaction should be logged to the audit log.

Action Group: Non-disruptive

Example:

```
SecRule REQUEST_HEADERS:User-Agent "Test" allow,noauditlog
```

Note

If the SecAuditEngine is set to On, all of the transactions will be logged. If it is set to RelevantOnly, then

you can control it with the `noauditlog` action. Even if the `noauditlog` action is applied to a specific rule and a rule either before or after triggered an audit event, then the transaction will be logged to the audit log. The correct way to disable audit logging for the entire transaction is to use `ctl:auditEngine=Off`

nolog

Description: Prevents rule matches from appearing in both the error and audit logs.

Action Group: Non-disruptive

Example:

```
SecRule REQUEST_HEADERS:User-Agent "Test" allow,nolog
```

Note

The `nolog` action also implies `noauditlog`.

pass

Description: Continues processing with the next rule in spite of a successful match.

Action Group: Disruptive

Example1:

```
SecRule REQUEST_HEADERS:User-Agent "Test" log,pass
```

When using `pass` with `SecRule` with multiple targets, *all* targets will be processed and *all* non-disruptive actions will trigger for *every* match found. In the second example the `TX:test` target would be incremented by 1 for each matching argument.

Example2:

```
SecRule ARGS "test" log,pass,setvar:TX.test=+1
```

Note

The transaction will not be interrupted but a log will be generated for each matching target (unless logging has been suppressed).

pause

Description: Pauses transaction processing for the specified number of milliseconds.

Action Group: Non-disruptive

Example:

```
SecRule REQUEST_HEADERS:User-Agent "Test" log,deny,status:403,pause:5000
```

Note

This feature can be of limited benefit for slowing down Brute Force Scanners, however use with care. If you are under a Denial of Service type of attack, the pause feature may make matters worse as this feature will cause child processes to sit idle until the pause is completed.

phase

Description: Places the rule (or the rule chain) into one of five available processing phases.

Action Group: Meta-data

Example:

```
SecDefaultAction log,deny,phase:1,t:removeNulls,t:lowercase
SecRule REQUEST_HEADERS:User-Agent "Test" log,deny,status:403
```

Note

Keep in mind that if you specify the incorrect phase, the target variable that you specify may be empty. This could lead to a false negative situation where your variable and operator (RegEx) may be correct, but it misses malicious data because you specified the wrong phase.

prepend

Description: Prepends text given as parameter to the response body. For this action to work content injection must be enabled by setting `SecContentInjection` to `On`. Also make sure you check the content type of the response before you make changes to it (e.g. you don't want to inject stuff into images).

Action Group: Non-disruptive

Processing Phases: 3 and 4.

Example:

```
SecRule RESPONSE_CONTENT_TYPE ^text/html "phase:3,nolog,pass,prepend:'Header<br>' "
```

Note

While macro expansion is allowed in the additional content, you are strongly cautioned against inserting user defined data fields.

proxy

Description: Intercepts transaction by forwarding request to another web server using the proxy backend.

Action Group: Disruptive

Example:

```
SecRule REQUEST_HEADERS:User-Agent "Test" log,proxy:http://www.honeypothost.com/
```

Note

For this action to work, `mod_proxy` must also be installed. This action is useful if you would like to proxy matching requests onto a honeypot webserver.

redirect

Description: Intercepts transaction by issuing a redirect to the given location.

Action Group: Disruptive

Example:

```
SecRule REQUEST_HEADERS:User-Agent "Test" \
    log,redirect:http://www.hostname.com/failed.html
```

Note

If the `status` action is present and its value is acceptable (301, 302, 303, or 307) it will be used for the redirection. Otherwise status code 302 will be used.

rev

Description: Specifies rule revision.

Action Group: Meta-data

Example:

```
SecRule REQUEST_METHOD "^PUT$" "id:340002,rev:1,severity:2,msg:'Restricted HTTP function'"
```

Note

This action is used in combination with the `id` action to allow the same rule ID to be used after changes take place but to still provide some indication the rule changed.

sanitiseArg

Description: Sanitises (replaces each byte with an asterisk) a named request argument prior to audit logging.

Action Group: Non-disruptive

Example:

```
SecAction nolog,phase:2,sanitiseArg:password
```

Note

The sanitize actions do not sanitize any data within the actual raw requests but only on the copy of data within memory that is set to log to the audit log. It will not sanitize the data in the `modsec_debug.log` file (if the log level is set high enough to capture this data).

sanitiseMatched

Description: Sanitises the variable (request argument, request header, or response header) that caused a rule match.

Action Group: Non-disruptive

Example: This action can be used to sanitise arbitrary transaction elements when they match a condition. For example, the example below will sanitise any argument that contains the word *password* in the name.

```
SecRule ARGS_NAMES password nolog,pass,sanitiseMatched
```

Note

Same note as sanitiseArg.

sanitiseRequestHeader

Description: Sanitises a named request header.

Action Group: Non-disruptive

Example: This will sanitise the data in the Authorization header.

```
SecAction log,phase:1,sanitiseRequestHeader:Authorization
```

Note

Same note as sanitiseArg.

sanitiseResponseHeader

Description: Sanitises a named response header.

Action Group: Non-disruptive

Example: This will sanitise the Set-Cookie data sent to the client.

```
SecAction log,phase:3,sanitiseResponseHeader:Set-Cookie
```

Note

Same note as sanitiseArg.

severity

Description: Assigns severity to the rule it is placed with.

Action Group: Meta-data

Example:

```
SecRule REQUEST_METHOD "^PUT$" "id:340002,rev:1,severity:CRITICAL,msg:'Restricted HTTP func"
```

Note

Severity values in ModSecurity follow those of syslog, as below:

- 0 - EMERGENCY
- 1 - ALERT
- 2 - CRITICAL
- 3 - ERROR
- 4 - WARNING
- 5 - NOTICE
- 6 - INFO
- 7 - DEBUG

It is possible to specify severity levels using either the numerical values or the text values. You should always specify severity levels using the text values. The use of the numerical values is deprecated (as of v2.5.0) and may be removed in one of the subsequent major updates.

setuid

Description: Special-purpose action that initialises the USER collection.

Action Group: Non-disruptive

Example:

```
SecAction setuid:%{REMOTE_USER},nolog
```

Note

After initialisation takes place the variable USERID will be available for use in the subsequent rules.

setsid

Description: Special-purpose action that initialises the SESSION collection.

Action Group: Non-disruptive

Example:

```
# Initialise session variables using the session cookie value
SecRule REQUEST_COOKIES:PHPSESSID !^$ chain,nolog,pass
SecAction setsid:%{REQUEST_COOKIES.PHPSESSID}
```

Note

On first invocation of this action the collection will be empty (not taking the predefined variables into account - see `initcol` for more information). On subsequent invocations the contents of the collection (session, in this case) will be retrieved from storage. After initialisation takes place the variable SESSIONID will be available for use in the subsequent rules. This action understands each application maintains its own set of sessions. It will utilise the current web application ID to create a session namespace.

setenv

Description: Creates, removes, or updates an environment variable.

Action Group: Non-disruptive

Examples:

To create a new variable (if you omit the value 1 will be used):

```
setenv:name=value
```

To remove a variable:

```
setenv:!name
```

Note

This action can be used to establish communication with other Apache modules.

setvar

Description: Creates, removes, or updates a variable in the specified collection.

Action Group: Non-disruptive

Examples:

To create a new variable:

```
setvar:tx.score=10
```

To remove a variable prefix the name with exclamation mark:

```
setvar:!tx.score
```

To increase or decrease variable value use + and - characters in front of a numerical value:

```
setvar:tx.score+=5
```

skip

Description: Skips one or more rules (or chains) on successful match.

Action Group: Flow

Example:

```
SecRule REQUEST_URI "^/$" \  
"phase:2,chain,t:none,skip:2"  
SecRule REMOTE_ADDR "^127\.0\.0\.1$" "chain"  
SecRule REQUEST_HEADERS:User-Agent "^Apache \((internal dummy connection)\)$" "t:none"  
SecRule &REQUEST_HEADERS:Host "@eq 0" \  
    "deny,log,status:400,id:960008,severity:4,msg:'Request Missing a Host Header'"  
SecRule &REQUEST_HEADERS:Accept "@eq 0" \  

```

```
"log,deny,log,status:400,id:960015,msg:'Request Missing an Accept Header'"
```

Note

Skip only applies to the current processing phase and not necessarily the order in which the rules appear in the configuration file. If you group rules by processing phases, then skip should work as expected. This action can not be used to skip rules within one chain. Accepts a single parameter denoting the number of rules (or chains) to skip.

skipAfter

Description: Skips rules (or chains) on successful match resuming rule execution after the specified rule ID or marker (see `SecMarker`) is found.

Action Group: Flow

Example:

```
SecRule REQUEST_URI "^/$" "chain,t:none,skipAfter:960015"
SecRule REMOTE_ADDR "^127\.0\.0\.1$" "chain"
SecRule REQUEST_HEADERS:User-Agent "^Apache \((internal dummy connection)\)$" "t:none"
SecRule &REQUEST_HEADERS:Host "@eq 0" \
    "deny,log,status:400,id:960008,severity:4,msg:'Request Missing a Host Header'"
SecRule &REQUEST_HEADERS:Accept "@eq 0" \
    "log,deny,log,status:400,id:960015,msg:'Request Missing an Accept Header'"
```

Note

skipAfter only applies to the current processing phase and not necessarily the order in which the rules appear in the configuration file. If you group rules by processing phases, then skip should work as expected. This action can not be used to skip rules within one chain. Accepts a single parameter denoting the last rule ID to skip.

status

Description: Specifies the response status code to use with actions `deny` and `redirect`.

Action Group: Data

Example:

```
SecDefaultAction log,deny,status:403,phase:1
```

Note

Status actions defined in Apache scope locations (such as `Directory`, `Location`, etc..) may be superseded by `phase:1` action settings. The Apache `ErrorDocument` directive will be triggered if present in the configuration. Therefore if you have previously defined a custom error page for a given status then it will be executed and its output presented to the user.

t

Description: This action can be used which transformation function should be used against the specified variables before they (or the results, rather) are run against the operator specified in the rule.

Action Group: Non-disruptive

Example:

```
SecDefaultAction log,deny,phase:1,t:removeNulls,t:lowercase
SecRule REQUEST_COOKIES:SESSIONID "47414e81cbbef3cf8366e84eeacba091" \
    log,deny,status:403,t:md5,t:hexEncode
```

Note

Any transformation functions that you specify in a SecRule will be in addition to previous ones specified in SecDefaultAction. Use of "t:none" will remove all transformation functions for the specified rule.

tag

Description: Assigns custom text to a rule or chain.

Action Group: Meta-data

Example:

```
SecRule REQUEST_FILENAME "\b(?:n(?:map|et|c)|w(?:guest|sh)|cmd(?:32)?|telnet|rcmd|ftp)\.exe" \
    "t:none,t:lowercase,deny,msg:'System Command Access',id:'950002',\
    tag:'WEB_ATTACK/FILE_INJECTION',tag:'OWASP/A2',severity:'2'"
```

Note

The tag information appears in the error and/or audit log files. Its intent is to be used to automate classification of rules and the alerts generated by rules. Multiple tags can be used per rule/chain.

xmlns

Description: This action should be used together with an XPath expression to register a namespace.

Action Group: Data

Example:

```
SecRule REQUEST_HEADERS:Content-Type "text/xml" \
    "phase:1,pass,ctl:requestBodyProcessor=XML,ctl:requestBodyAccess=On, \
    xmlns:xsd="http://www.w3.org/2001/XMLSchema"
SecRule XML:/soap:Envelope/soap:Body/ql:getInput/id() "123" phase:2,deny
```

Operators

A number of operators can be used in rules, as documented below. The operator syntax uses the @ symbol followed by the specific operator name.

beginsWith

Description: This operator is a string comparison and returns true if the parameter value is found at the beginning of the input. Macro expansion is performed so you may use variable names such as %{TX.1}, etc.

Example:

```
SecRule REQUEST_LINE "!@beginsWith GET" t:none,deny,status:403
SecRule REQUEST_ADDR "^(.*)\\.\\d+$" deny,status:403,capture,chain
SecRule ARGS:gw "!@beginsWith %{TX.1}"
```

contains

Description: This operator is a string comparison and returns true if the parameter value is found anywhere in the input. Macro expansion is performed so you may use variable names such as %{TX.1}, etc.

Example:

```
SecRule REQUEST_LINE "!@contains .php" t:none,deny,status:403
SecRule REQUEST_ADDR "^(.*)$" deny,status:403,capture,chain
SecRule ARGS:ip "!@contains %{TX.1}"
```

endsWith

Description: This operator is a string comparison and returns true if the parameter value is found at the end of the input. Macro expansion is performed so you may use variable names such as %{TX.1}, etc.

Example:

```
SecRule REQUEST_LINE "!@endsWith HTTP/1.1" t:none,deny,status:403
SecRule ARGS:route "!@endsWith %{REQUEST_ADDR}" t:none,deny,status:403
```

eq

Description: This operator is a numerical comparison and stands for "equal to."

Example:

```
SecRule &REQUEST_HEADERS_NAMES "@eq 15"
```

ge

Description: This operator is a numerical comparison and stands for "greater than or equal to."

Example:

```
SecRule &REQUEST_HEADERS_NAMES "@ge 15"
```

geoLookup

Description: This operator looks up various data fields from an IP address or hostname. The results will be captured in the GEO collection.

You must provide a database via `SecGeoLookupDb` before this operator can be used.

See the GEO variable for an example and more information on various fields available.

gt

Description: This operator is a numerical comparison and stands for "greater than."

Example:

```
SecRule &REQUEST_HEADERS_NAMES "@gt 15"
```

inspectFile

Description: Executes the external script/binary given as parameter to the operator against every file extracted from the request. As of v2.5.0, if the supplied filename is not absolute it is treated as relative to the directory in which the configuration file resides. Also as of v2.5.0, if the filename is determined to be a Lua script (based on its extension) the script will be processed by the internal engine. As such it will have full access to the ModSecurity context.

Example of using an external binary/script:

```
# Execute external script to validate uploaded files.
SecRule FILES_TMPNAMES "@inspectFile /opt/apache/bin/inspect_script.pl"
```

Example of using Lua script:

```
SecRule FILES_TMPNAMES "@inspectFile inspect.lua"
```

Script `inspect.lua`:

```
function main(filename)
    -- Do something to the file to verify it. In this example, we
    -- read up to 10 characters from the beginning of the file.
    local f = io.open(filename, "rb");
    local d = f:read(10);
```

```
f:close();

-- Return null if there is no reason to believe there is anything
-- wrong with the file (no match). Returning any text will be taken
-- to mean a match should be triggered.
return null;
end
```

le

Description: This operator is a numerical comparison and stands for "less than or equal to."

Example:

```
SecRule &REQUEST_HEADERS_NAMES "@le 15"
```

lt

Description: This operator is a numerical comparison and stands for "less than."

Example:

```
SecRule &REQUEST_HEADERS_NAMES "@lt 15"
```

pm

Description: Phrase Match operator. This operator uses a set based matching engine (Aho-Corasick) for faster matches of keyword lists. It will match any one of its arguments anywhere in the target value.

Example:

```
SecRule REQUEST_HEADERS:User-Agent "@pm WebZIP WebCopier Webster WebStripper SiteSnagger Pr
```

The above would deny access with 403 if any of the words matched within the User-Agent HTTP header value.

pmFromFile

Description: Phrase Match operator. This operator uses a set based matching engine (Aho-Corasick) for faster matches of keyword lists. This operator is the same as @pm except that it takes a list of files as arguments. It will match any one of the phrases listed in the file(s) anywhere in the target value.

Notes:

1. The contents of the files should be one phrase per line. End of line markers will be stripped from the phrases, however, whitespace will not be trimmed from phrases in the file. Empty lines and comment lines (beginning with a '#') are ignored.

2. To allow easier inclusion of phrase files with rulesets, relative paths may be used to the phrase files. In this case, the path of the file containing the rule is prepended to the phrase file path.

Example:

```
SecRule REQUEST_HEADERS:User-Agent "@pm /path/to/blacklist1 blacklist2" "deny,status:403"
```

The above would deny access with 403 if any of the patterns in the two files matched within the User-Agent HTTP header value. The `blacklist2` file would need to be placed in the same path as the file containing the rule.

rbl

Description: Look up the parameter in the RBL given as parameter. Parameter can be an IPv4 address, or a hostname.

Example:

```
SecRule REMOTE_ADDR "@rbl sc.surbl.org"
```

rx

Description: Regular expression operator. This is the default operator, so if the "@" operator is not defined, it is assumed to be rx.

Example:

```
SecRule REQUEST_HEADERS:User-Agent "@rx nikto"
```

Note

Regular expressions are handled by the PCRE library (<http://www.pcre.org>). ModSecurity compiles its regular expressions with the following settings:

1. The entire input is treated as a single line, even when there are newline characters present.
2. All matches are case-sensitive. If you do not care about case sensitivity you either need to implement the `lowercase` transformation function, or use the per-pattern (`?i`) modifier, as allowed by PCRE.
3. The `PCRE_DOTALL` and `PCRE_DOLLAR_ENDONLY` flags are set during compilation, meaning a single dot will match any character, including the newlines and a `$` end anchor will not match a trailing newline character.

streq

Description: This operator is a string comparison and returns true if the parameter value matches the input exactly. Macro expansion is performed so you may use variable names such as `%{TX.1}`, etc.

Example:

```
SecRule ARGS:foo "!@streq bar" t:none,deny,status:403
SecRule REQUEST_ADDR "^(.*)$" deny,status:403,capture,chain
SecRule REQUEST_HEADERS:Ip-Address "!@streq %{TX.1}"
```

validateByteRange

Description: Validates the byte range used in the variable falls into the specified range.

Example:

```
SecRule ARG:text "@validateByteRange 10, 13, 32-126"
```

Note

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). Default range values are 0 and 255, i.e. all byte values are allowed. This directive does not check byte range in a POST payload when multipart/form-data encoding (file upload) is used. Doing so would prevent binary files from being uploaded. However, after the parameters are extracted from such request they are checked for a valid range.

validateByteRange is similar to the ModSecurity 1.X SecFilterForceByteRange Directive however since it works in a rule context, it has the following differences:

- You can specify a different range for different variables.
- It has an "event" context (id, msg....)
- It is executed in the flow of rules rather than being a built in pre-check.

validateDTD

Description: Validates the DOM tree generated by the XML request body processor against the supplied DTD.

Example:

```
SecDefaultAction log,deny,status:403,phase:2
SecRule REQUEST_HEADERS:Content-Type ^text/xml$ \
  phase:1,t:lowercase,nolog,pass,ctl:requestBodyProcessor=XML
SecRule REQBODY_PROCESSOR "!^XML$" nolog,pass,skipAfter:12345
SecRule XML "@validateDTD /path/to/apache2/conf/xml.dtd" "deny,id:12345"
```

Note

This operator requires request body to be processed as XML.

validateSchema

Description: Validates the DOM tree generated by the XML request body processor against the supplied XML Schema.

Example:

```
SecDefaultAction log,deny,status:403,phase:2
SecRule REQUEST_HEADERS:Content-Type ^text/xml$ \
    phase:1,t:lowercase,nolog,pass,ctl:requestBodyProcessor=XML
SecRule REQUEST_BODY_PROCESSOR "!^XML$" nolog,pass,skipAfter:12345
SecRule XML "@validateSchema /path/to/apache2/conf/xml.xsd" "deny,id:12345"
```

Note

This operator requires request body to be processed as XML.

validateUrlEncoding

Description: Verifies the encodings used in the variable (if any) are valid.

Example:

```
SecRule ARGS "@validateUrlEncoding"
```

Note

URL encoding is an HTTP standard for encoding byte values within a URL. The byte is escaped with a % followed by two hexadecimal values (0-F). This directive does not check encoding in a POST payload when the `multipart/form-data` encoding (file upload) is used. It is not necessary to do so because URL encoding is not used for this encoding.

validateUtf8Encoding

Description: Verifies the variable is a valid UTF-8 encoded string.

Example:

```
SecRule ARGS "@validateUtf8Encoding"
```

Note

UTF-8 encoding is valid on most web servers. Integer values between 0-65535 are encoded in a UTF-8 byte sequence that is escaped by percents. The short form is two bytes in length.

check for three types of errors:

- Not enough bytes. UTF-8 supports two, three, four, five, and six byte encodings. ModSecurity will locate cases when a byte or more is missing.
- Invalid encoding. The two most significant bits in most characters are supposed to be fixed to 0x80. Attackers can use this to subvert Unicode decoders.
- Overlong characters. ASCII characters are mapped directly into the Unicode space and are thus

represented with a single byte. However, most ASCII characters can also be encoded with two, three, four, five, and six characters thus tricking the decoder into thinking that the character is something else (and, presumably, avoiding the security check).

verifyCC

Description: This operator verifies a given regular expression as a potential credit card number. It first matches with a single generic regular expression then runs the resulting match through a Luhn checksum algorithm to further verify it as a potential credit card number.

Example:

```
SecRule ARGS "@verifyCC \d{13,16}" \
    "phase:2,sanitiseMatched,log,auditlog,pass,msg:'Potential credit card number'
```

within

Description: This operator is a string comparison and returns true if the input value is found anywhere within the parameter value. Note that this is similar to @contains, except that the target and match values are reversed. Macro expansion is performed so you may use variable names such as %{TX.1}, etc.

Example:

```
SecRule REQUEST_METHOD "!@within get,post,head" t:lowercase,deny,status:403

SecAction "pass,setvar:'tx.allowed_methods=get,post,head'"
SecRule REQUEST_METHOD "!@within %{tx.allowed_methods}" t:lowercase,deny,status:403
```

Macro Expansion

Macros allow for using place holders in rules that will be expanded out to their values at runtime. Currently only variable expansion is supported, however more options may be added in future versions of ModSecurity.

Format:

```
%{VARIABLE}
%{COLLECTION.VARIABLE}
```

Macro expansion can be used in actions such as `initcol`, `setuid`, `setvar`, `setenv`, `logdata`. Operators that are evaluated at runtime support expansion and are noted above. Such operators include `@beginsWith`, `@endsWith`, `@contains`, `@within` and `@streq`. You cannot use macro expansion for operators that are "compiled" such as `@pm`, `@rx`, etc. as these operators have their values fixed at configure time for efficiency.

Some values you may want to expand include: `TX`, `REMOTE_ADDR`, `USERID`, `HIGHEST_SEVERITY`, `MATCHED_VAR`, `MATCHED_VAR_NAME`, `MULTI-PART_STRICT_ERROR`, `RULE`, `SESSION`, `USERID`, among others.

Persistent Storage

At this time it is only possible to have three collections in which data is stored persistently (i.e. data available to multiple requests). These are: `IP`, `SESSION` and `USER`.

Every collection contains several built-in variables that are available and are read-only unless otherwise specified:

1. `CREATE_TIME` - date/time of the creation of the collection.
2. `IS_NEW` - set to 1 if the collection is new (not yet persisted) otherwise set to 0.
3. `KEY` - the value of the `initcol` variable (the client's IP address in the example).
4. `LAST_UPDATE_TIME` - date/time of the last update to the collection.
5. `TIMEOUT` - date/time in seconds when the collection will be updated on disk from memory (if no other updates occur). This variable may be set if you wish to specify an explicit expiration time (default is 3600 seconds).
6. `UPDATE_COUNTER` - how many times the collection has been updated since creation.
7. `UPDATE_RATE` - is the average rate updates per minute since creation.

To create a collection to hold session variables (`SESSION`) use action `setsid`. To create a collection to hold user variables (`USER`) use action `setuid`. To create a collection to hold client address variables (`IP`) use action `initcol`.

Note

ModSecurity implements atomic updates of persistent variables only for integer variables (counters) at this time. Variables are read from storage whenever `initcol` is encountered in the rules and persisted at the end of request processing. Counters are adjusted by applying a delta generated by re-reading the persisted data just before being persisted. This keeps counter data consistent even if the counter was modified and persisted by another thread/process during the transaction.

Note

ModSecurity uses a Berkley Database (SDBM) for persistent storage. This type of database is generally limited to storing a maximum of 1008 bytes per key. This may be a limitation if you are attempting to store a considerable amount of data in variables for a single key. Some of this limitation is planned to be reduced in a future version of ModSecurity.

Miscellaneous Topics

Impedance Mismatch

Web application firewalls have a difficult job trying to make sense of data that passes by, without any knowledge of the application and its business logic. The protection they provide comes from having an independent layer of security on the outside. Because data validation is done twice, security can be increased without having to touch the application. In some cases, however, the fact that everything is done twice brings problems. Problems can arise in the areas where the communication protocols are not well specified, or where either the device or the application do things that are not in the specification. In such cases it may be possible to design payload that will be interpreted in one way by one device and in another by the other device. This problem is better known as Impedance Mismatch. It can be exploited to evade the security devices.

While we will continue to enhance ModSecurity to deal with various evasion techniques the problem can only be minimized, but never solved. With so many different application backend chances are some will always do something completely unexpected. The only solution is to be aware of the technologies in the backend when writing rules, adapting the rules to remove the mismatch. See the next section for some examples.

PHP Peculiarities for ModSecurity Users

When writing rules to protect PHP applications you need to pay attention to the following facts:

1. When "register_globals" is set to "On" request parameters are automatically converted to script variables. In some PHP versions it is even possible to override the \$GLOBALS array.
2. Whitespace at the beginning of parameter names is ignored. (This is very dangerous if you are writing rules to target specific named variables.)
3. The remaining whitespace (in parameter names) is converted to underscores. The same applies to dots and to a "[" if the variable name does not contain a matching closing bracket. (Meaning that if you want to exploit a script through a variable that contains an underscore in the name you can send a parameter with a whitespace or a dot instead.)
4. Cookies can be treated as request parameters.
5. The discussion about variable names applies equally to the cookie names.
6. The order in which parameters are taken from the request and the environment is EGPCS (environment, GET, POST, Cookies, built-in variables). This means that a POST parameter will overwrite the parameters transported on the request line (in QUERY_STRING).
7. When "magic_quotes_gpc" is set to "On" PHP will use backslash to escape the following characters: single quote, double quote, backslash, and the nul byte.
8. If "magic_quotes_sybase" is set to "On" only the single quote will be escaped using another single quote. In this case the "magic_quotes_gpc" setting becomes irrelevant. The "magic_quotes_sybase" setting completely overrides the "magic_quotes_gpc" behaviour but "ma-

`gpc_quotes_gpc` still must be set to "On" for the Sybase-specific quoting to be work.

9. PHP will also automatically create nested arrays for you. For example "`p[x][y]=1`" results in a total of three variables.