

SnortTMUsers Manual

2.4.0RC1

The Snort Project

26th August 2005

Copyright ©1998-2003 Martin Roesch

Copyright ©2001-2003 Chris Green

Copyright ©2003-2005 Sourcefire, Inc.

Contents

1	Snort Overview	3
1.1	Getting Started	3
1.2	Sniffer Mode	3
1.3	Packet Logger Mode	4
1.4	Network Intrusion Detection Mode	5
1.4.1	NIDS Mode Output Options	5
1.4.2	Understanding Standard Alert Output	6
1.4.3	High Performance Configuration	6
1.4.4	Changing Alert Order	7
1.5	Inline Mode	7
1.5.1	Snort Inline Rule Application Order	8
1.5.2	New STREAM4 Options for Use with Snort Inline	8
1.5.3	Replacing Packets with Snort Inline	8
1.5.4	Installing Snort Inline	9
1.5.5	Running Snort Inline	9
1.5.6	Using the Honeynet Snort Inline Toolkit	9
1.5.7	Troubleshooting Snort Inline	9
1.6	Miscellaneous	10
1.6.1	Running in Daemon mode	10
1.6.2	Obfuscating IP address printouts	10
1.6.3	Specifying multiple-instance identifiers	10
1.7	More Information	10
2	Configuring Snort	12
2.0.1	Includes	12
2.0.2	Variables	12
2.0.3	Config	13
2.1	Preprocessors	16
2.1.1	Frag2	16
2.1.2	Frag3	16
2.1.3	Stream4	19

2.1.4	Flow	23
2.1.5	Portscan	23
2.1.6	Flow-Portscan	23
2.1.7	sfPortscan	23
2.1.8	Telnet Decode	29
2.1.9	RPC Decode	29
2.1.10	Performance Monitor	30
2.1.11	HTTP Inspect	32
2.1.12	ASN.1 Detection	40
2.1.13	X-Link2State Mini-Preprocessor	41
2.2	Event Thresholding	41
2.3	Output Modules	41
2.3.1	alert_syslog	42
2.3.2	alert_fast	43
2.3.3	alert_full	43
2.3.4	alert_unixsock	44
2.3.5	log_tcpdump	44
2.3.6	database	44
2.3.7	csv	46
2.3.8	unified	47
2.3.9	alert_prelude	47
2.3.10	log null	48
3	Writing Snort Rules	
	How to Write Snort Rules and Keep Your Sanity	49
3.1	The Basics	49
3.2	Rules Headers	50
3.2.1	Rule Actions	50
3.2.2	Protocols	50
3.2.3	IP Addresses	51
3.2.4	Port Numbers	51
3.2.5	The Direction Operator	53
3.2.6	Activate/Dynamic Rules	53
3.3	Rule Options	54
3.4	Meta-Data Rule Options	54
3.4.1	msg	54
3.4.2	reference	54
3.4.3	sid	55
3.4.4	rev	55
3.4.5	classtype	56

3.4.6	Priority	57
3.5	Payload Detection Rule Options	57
3.5.1	content	57
3.5.2	nocase	58
3.5.3	rawbytes	59
3.5.4	depth	59
3.5.5	offset	59
3.5.6	distance	60
3.5.7	within	60
3.5.8	uricontent	60
3.5.9	isdataat	61
3.5.10	pcre	61
3.5.11	byte.test	62
3.5.12	byte.jump	63
3.5.13	ftpbounce	65
3.5.14	regex	65
3.5.15	content-list	65
3.6	Non-payload Detection Rule Options	66
3.6.1	fragoffset	66
3.6.2	ttl	66
3.6.3	tos	66
3.6.4	id	67
3.6.5	ipopts	67
3.6.6	fragbits	68
3.6.7	dsiz	68
3.6.8	flags	69
3.6.9	flow	69
3.6.10	flowbits	70
3.6.11	seq	71
3.6.12	ack	71
3.6.13	window	71
3.6.14	itype	72
3.6.15	icode	72
3.6.16	icmp_id	72
3.6.17	icmp_seq	72
3.6.18	rpc	73
3.6.19	ip_proto	73
3.6.20	sameip	73
3.7	Post-Detection Rule Options	74

3.7.1	logto	74
3.7.2	session	74
3.7.3	resp	75
3.7.4	react	75
3.7.5	tag	76
3.8	Event Thresholding	77
3.8.1	Standalone Options	77
3.8.2	Standalone Format	78
3.8.3	Rule Keyword Format	78
3.8.4	Rule Keyword Format	78
3.8.5	Examples	79
3.9	Event Suppression	81
3.9.1	Format	81
3.9.2	Examples	81
3.10	Snort Multi-Event Logging (Event Queue)	82
3.10.1	Event Queue Configuration Options	82
3.10.2	Event Queue Configuration Examples	82
3.11	Writing Good Rules	83
3.11.1	Content Matching	83
3.11.2	Catch the Vulnerability, Not the Exploit	83
3.11.3	Catch the Oddities of the Protocol in the Rule	83
3.11.4	Optimizing Rules	84
3.11.5	testing numerical values	85
4	Making Snort Faster	88
4.1	MMAPEd pcap	88
5	Snort Development	89
5.1	Submitting Patches	89
5.2	Snort dataflow	89
5.2.1	Preprocessors	89
5.2.2	Detection Plugins	90
5.2.3	Output Plugins	90
5.3	The Snort Team	90

Chapter 1

Snort Overview

This manual is based on *Writing Snort Rules* by Martin Roesch and further work from Chris Green <cmg@snort.org>. It is now maintained by Brian Caswell <bmc@snort.org>. If you have a better way to say something or find that something in the documentation is outdated, drop us a line and we will update it. If you would like to submit patches for this document, you can find the latest version of the documentation in L^AT_EX format in the Snort CVS repository at `/doc/snort_manual.tex`. Small documentation updates are the easiest way to help out the Snort Project.

1.1 Getting Started

Snort really isn't very hard to use, but there are a lot of command line options to play with, and it's not always obvious which ones go together well. This file aims to make using Snort easier for new users.

Before we proceed, there are a few basic concepts you should understand about Snort. Snort can be configured to run in three modes:

- *Sniffer mode*, which simply reads the packets off of the network and displays them for you in a continuous stream on the console (screen).
- *Packet Logger mode*, which logs the packets to disk.
- *Network Intrusion Detection System (NIDS) mode*, the most complex and configurable configuration, which allows Snort to analyze network traffic for matches against a user-defined rule set and performs several actions based upon what it sees.
- *Inline Mode*, which obtains packets from iptables instead of from libpcap and then causes iptables to drop or pass packets based on Snort rules that use inline-specific rule types.

1.2 Sniffer Mode

First, let's start with the basics. If you just want to print out the TCP/IP packet headers to the screen (i.e. sniffer mode), try this:

```
./snort -v
```

This command will run Snort and just show the IP and TCP/UDP/ICMP headers, nothing else. If you want to see the application data in transit, try the following:

```
./snort -vd
```

This instructs Snort to display the packet data as well as the headers. If you want an even more descriptive display, showing the data link layer headers, do this:

```
./snort -vde
```

(As an aside, these switches may be divided up or smashed together in any combination. The last command could also be typed out as:

```
./snort -d -v -e
```

and it would do the same thing.)

1.3 Packet Logger Mode

OK, all of these commands are pretty cool, but if you want to record the packets to the disk, you need to specify a logging directory and Snort will automatically know to go into packet logger mode:

```
./snort -dev -l ./log
```

Of course, this assumes you have a directory named `log` in the current directory. If you don't, Snort will exit with an error message. When Snort runs in this mode, it collects every packet it sees and places it in a directory hierarchy based upon the IP address of one of the hosts in the datagram.

If you just specify a plain `-l` switch, you may notice that Snort sometimes uses the address of the remote computer as the directory in which it places packets and sometimes it uses the local host address. In order to log relative to the home network, you need to tell Snort which network is the home network:

```
./snort -dev -l ./log -h 192.168.1.0/24
```

This rule tells Snort that you want to print out the data link and TCP/IP headers as well as application data into the directory `./log`, and you want to log the packets relative to the 192.168.1.0 class C network. All incoming packets will be recorded into subdirectories of the log directory, with the directory names being based on the address of the remote (non-192.168.1) host.

NOTE

Note that if both the source and destination hosts are on the home network, they are logged to a directory with a name based on the higher of the two port numbers or, in the case of a tie, the source address.

If you're on a high speed network or you want to log the packets into a more compact form for later analysis, you should consider logging in binary mode. Binary mode logs the packets in tcpdump format to a single binary file in the logging directory:

```
./snort -l ./log -b
```

Note the command line changes here. We don't need to specify a home network any longer because binary mode logs everything into a single file, which eliminates the need to tell it how to format the output directory structure. Additionally, you don't need to run in verbose mode or specify the `-d` or `-e` switches because in binary mode the entire packet is logged, not just sections of it. All you really need to do to place Snort into logger mode is to specify a logging directory at the command line using the `-l` switch—the `-b` binary logging switch merely provides a modifier that tells Snort to log the packets in something other than the default output format of plain ASCII text.

Once the packets have been logged to the binary file, you can read the packets back out of the file with any sniffer that supports the tcpdump binary format (such as tcpdump or Ethereal). Snort can also read the packets back by using the `-r` switch, which puts it into playback mode. Packets from any tcpdump formatted file can be processed through Snort in any of its run modes. For example, if you wanted to run a binary log file through Snort in sniffer mode to dump the packets to the screen, you can try something like this:

```
./snort -dv -r packet.log
```

You can manipulate the data in the file in a number of ways through Snort's packet logging and intrusion detection modes, as well as with the BPF interface that's available from the command line. For example, if you only wanted to see the ICMP packets from the log file, simply specify a BPF filter at the command line and Snort will only see the ICMP packets in the file:

```
./snort -dvr packet.log icmp
```

For more info on how to use the BPF interface, read the Snort and tcpdump man pages.

1.4 Network Intrusion Detection Mode

To enable Network Intrusion Detection (NIDS) mode so that you don't record every single packet sent down the wire, try this:

```
./snort -dev -l ./log -h 192.168.1.0/24 -c snort.conf
```

where `snort.conf` is the name of your rules file. This will apply the rules configured in the `snort.conf` file to each packet to decide if an action based upon the rule type in the file should be taken. If you don't specify an output directory for the program, it will default to `/var/log/snort`.

One thing to note about the last command line is that if Snort is going to be used in a long term way as an IDS, the `-v` switch should be left off the command line for the sake of speed. The screen is a slow place to write data to, and packets can be dropped while writing to the display.

It's also not necessary to record the data link headers for most applications, so you can usually omit the `-e` switch, too.

```
./snort -d -h 192.168.1.0/24 -l ./log -c snort.conf
```

This will configure Snort to run in its most basic NIDS form, logging packets that trigger rules specified in the `snort.conf` in plain ASCII to disk using a hierarchical directory structure (just like packet logger mode).

1.4.1 NIDS Mode Output Options

There are a number of ways to configure the output of Snort in NIDS mode. The default logging and alerting mechanisms are to log in decoded ASCII format and use full alerts. The full alert mechanism prints out the alert message in addition to the full packet headers. There are several other alert output modes available at the command line, as well as two logging facilities.

Alert modes are somewhat more complex. There are seven alert modes available at the command line: full, fast, socket, syslog, console, cmg, and none. Six of these modes are accessed with the `-A` command line switch. These options are:

Option	Description
-A fast	Fast alert mode. Writes the alert in a simple format with a timestamp, alert message, source and destination IPs/ports.
-A full	Full alert mode. This is the default alert mode and will be used automatically if you do not specify a mode.
-A unsock	Sends alerts to a UNIX socket that another program can listen on.
-A none	Turns off alerting.
-A console	Sends “fast-style” alerts to the console (screen).
-A cmg	Generates “cmg style” alerts.

Packets can be logged to their default decoded ASCII format or to a binary log file via the -b command line switch. To disable packet logging altogether, use the -N command line switch.

For output modes available through the configuration file, see Section 2.7.

NOTE

Command line logging options override any output options specified in the configuration file. This allows debugging of configuration issues quickly via the command line.

To send alerts to syslog, use the -s switch. The default facilities for the syslog alerting mechanism are LOG_AUTHPRIV and LOG_ALERT. If you want to configure other facilities for syslog output, use the output plugin directives in the rules files. See Section 2.3.1 for more details on configuring syslog output.

For example, use the following command line to log to default (decoded ASCII) facility and send alerts to syslog:

```
./snort -c snort.conf -l ./log -h 192.168.1.0/24 -s
```

As another example, use the following command line to log to the default facility in /var/log/snort and send alerts to a fast alert file:

```
./snort -c snort.conf -A fast -h 192.168.1.0/24
```

1.4.2 Understanding Standard Alert Output

When Snort generates an alert message, it will usually look like the following:

```
[**] [116:56:1] (snort_decoder): T/TCP Detected [**]
```

The first number is the Generator ID, this tells the user what component of Snort generated this alert. For a list of GIDs, please read etc/generators in the Snort source. In this case, we know that this event came from the “decode” (116) component of Snort.

The second number is the Snort ID (sometimes referred to as Signature ID). For a list of preprocessor SIDs, please see etc/gen-msg.map. Rule-based SIDs are written directly into the rules with the “sid” option. In this case, “56” represents a T/TCP event.

The third number is the revision ID. This number is primarily used when writing signatures, as each rendition of the rule should increment this number with the “rev” option.

1.4.3 High Performance Configuration

If you want Snort to go *fast* (like keep up with a 1000 Mbps connection), you need to use unified logging and a unified log reader such as *barnyard*. This allows Snort to log alerts in a binary form as fast as possible while another program performs the slow actions, such as writing to a database.

If you want a text file that's easily parsable, but still somewhat fast, try using binary logging with the “fast” output mechanism.

This will log packets in tcpdump format and produce minimal alerts. For example:

```
./snort -b -A fast -c snort.conf
```

1.4.4 Changing Alert Order

The default way in which Snort applies its rules to packets may not be appropriate for all installations. The Alert rules are applied first, then the Pass rules, and finally, Log rules are applied. This sequence is somewhat counterintuitive, but it's a more foolproof method than allowing a user to write a hundred alert rules that are then disabled by an errant pass rule. For more information on rule types, see Section 3.2.1.

If you know what you're doing, you can use the `-o` switch to change the default rule application behavior to apply Pass rules, then Alert rules, then Log rules:

```
./snort -d -h 192.168.1.0/24 -l ./log -c snort.conf -o
```

1.5 Inline Mode

Snort 2.3.0 RC1 integrated the intrusion prevention system (IPS) capability of `snort_inline` into the official Snort project. `Snort_inline` obtains packets from iptables instead of libpcap and then uses new rule types to help iptables pass or drop packets based on Snort rules.

In order for `snort_inline` to work properly, you must download and compile the iptables code to include “make install-devel” (<http://www.iptables.org>). This will install the `libipq` library that allows `snort_inline` to interface with iptables. Also, you must build and install LibNet, which is available from <http://www.packetfactory.net>.

There are three rule types you can use when running Snort with `snort_inline`:

- **drop** - The drop rule type will tell iptables to drop the packet and log it via usual Snort means.
- **reject** - The reject rule type will tell iptables to drop the packet, log it via usual Snort means, and send a TCP reset if the protocol is TCP or an icmp port unreachable if the protocol is UDP.
- **sdrop** - The sdrop rule type will tell iptables to drop the packet. Nothing is logged.

NOTE

You can also replace sections of the packet payload when using `snort_inline`. See Section 1.5.3 for more information.

When using a **reject** rule, there are two options you can use to send TCP resets:

- You can use a RAW socket (the default behavior for `snort_inline`), in which case you must have an interface that has an IP address assigned to it. If there is not an interface with an IP address assigned with access to the source of the packet, the packet will be logged and the reset packet will never make it onto the network.
- You can also now perform resets via a physical device when using iptables. We take the `indev` name from `ip_queue` and use this as the interface on which to send resets. We no longer need an IP loaded on the bridge, and can remain pretty stealthy as the `config layer2_resets` in `snort_inline.conf` takes a source MAC address which we substitute for the MAC of the bridge. For example:

```
config layer2resets
```

tells snort_inline to use layer2 resets and uses the MAC address of the bridge as the source MAC in the packet, and:

```
config layer2resets: 00:06:76:DD:5F:E3
```

will tell snort_inline to use layer2 resets and uses the source MAC of 00:06:76:DD:5F:E3 in the reset packet.

1.5.1 Snort Inline Rule Application Order

The current rule application order is:

```
->activation->dynamic->drop->sdrop->reject->alert->pass->log
```

This will ensure that a drop rule has precedence over an alert or log rule. You can use the -o flag to the rule application order to:

```
->activation->dynamic->pass->drop->sdrop->reject->alert->log
```

1.5.2 New STREAM4 Options for Use with Snort Inline

When using snort_inline, you can use two additional stream4 options:

- **inline_state** (no arguments)

This option causes Snort to drop TCP packets that are not associated with an existing TCP session, and is not a valid TCP initiator.

- **midstream_drop_alerts** (no arguments)

By default, when running in inline mode, Snort will silently drop any packets that were picked up in midstream and would have caused an alert to be generated, if not for the 'flow: established' option. This is to mitigate stick/snot type attacks when the user hasn't enabled inline_state. If you want to see the alerts that are silently dropped, enable this keyword. Note that by enabling this keyword, you have opened yourself up to stick/snot-type attacks.

For more information about Stream4, see Section 2.1.3.

1.5.3 Replacing Packets with Snort Inline

Additionally, Jed Haile's content replace code allows you to modify packets before they leave the network. For example:

```
alert tcp any any <> any 80 (msg: "tcp replace"; content:"GET"; replace:"BET");
alert udp any any <> any 53 (msg: "udp replace"; \
    content: "yahoo"; replace: "xxxxx");
```

These rules will comb tcp port 80 traffic looking for GET, and udp port 53 traffic looking for yahoo. Once they are found, they are replaced with BET and xxxxx, respectively. The only catch is that the replace must be the same length as the content.

1.5.4 Installing Snort Inline

To install Snort inline, use the following command:

```
./configure --enable-inline  
make  
make install
```

1.5.5 Running Snort Inline

First, you need to ensure that the `ip_queue` module is loaded. Then, you need to send traffic to `snort_inline` using the `QUEUE` target. For example,

```
iptables -A OUTPUT -p tcp --dport 80 -j QUEUE
```

sends all TCP traffic leaving the firewall going to port 80 to the `QUEUE` target. This is what sends the packet from kernel space to user space (`snort_inline`). A quick way to get all outbound traffic going to the `QUEUE` is to use the `rc.firewall` script created and maintained by the HoneyNet Project (<http://www.honeynet.org/papers/honeynet/tools/>) This script is well-documented and allows you to direct packets to `snort_inline` by simply changing the `QUEUE` variable to yes.

Finally, start `snort_inline`.

```
snort_inline -QDc ../etc/drop.conf -l /var/log/snort
```

You can use the following command line options:

- `-Q` - Gets packets from iptables.
- `-D` - Runs `snort_inline` in daemon mode. The process ID is stored at `/var/run/snort_inline.pid`
- `-c` - Reads the following configuration file.
- `-l` - Logs to the following directory.

Ideally, `snort_inline` will be run using only its own `drop.rules`. If you want to use Snort for just alerting, a separate process should be running with its own ruleset.

1.5.6 Using the HoneyNet Snort Inline Toolkit

The HoneyNet Snort Inline Toolkit is a statically compiled `snort_inline` binary put together by the HoneyNet Project for the Linux operating system. It comes with a set of `drop.rules`, the `snort_inline` binary, a `snort-inline` rotation shell script, and a good README. It can be found at:

<http://www.honeynet.org/papers/honeynet/tools/>

1.5.7 Troubleshooting Snort Inline

If you run `snort_inline` and see something like this:

```
Initializing Output Plugins!  
Reading from iptables  
Log directory = /var/log/snort  
Initializing Inline mode  
InlineInit: : Failed to send netlink message: Connection refused
```

More than likely, the `ip_queue` module is not loaded or `ip_queue` support is not compiled into your kernel. Either recompile your kernel to support `ip_queue`, or load the module.

The `ip_queue` module is loaded by executing:

```
insmod ip_queue
```

Also, if you want to ensure `snort_inline` is getting packets, you can start it in the following manner:

```
snort_inline -Qvc <configuration file>
```

This will display the header of every packet that `snort_inline` sees.

1.6 Miscellaneous

1.6.1 Running in Daemon mode

If you want to run Snort in daemon mode, you can add `-D` switch to any combination described in the previous sections. Please notice that if you want to be able to restart Snort by sending a `SIGHUP` signal to the daemon, you *must* specify the full path to the Snort binary when you start it, for example:

```
/usr/local/bin/snort -d -h 192.168.1.0/24 \  
-l /var/log/snortlogs -c /usr/local/etc/snort.conf -s -D
```

Relative paths are not supported due to security concerns.

1.6.2 Obfuscating IP address printouts

If you need to post packet logs to public mailing lists, you might want to use the `-O` switch. This switch obfuscates your IP addresses in packet printouts. This is handy if you don't want people on the mailing list to know the IP addresses involved. You can also combine the `-O` switch with the `-h` switch to only obfuscate the IP addresses of hosts on the home network. This is useful if you don't care who sees the address of the attacking host. For example, you could use the following command to read the packets from a log file and dump them to the screen, obfuscating only the addresses from the 192.168.1.0/24 class C network:

```
./snort -d -v -r snort.log -O -h 192.168.1.0/24
```

1.6.3 Specifying multiple-instance identifiers

In Snort v2.4, a `-G` command line option was added that specifies an instance identifier for the event logs. This option can be used when running multiple instances of snort, either on different CPUs, or on same CPU but different interface. Each snort instance will use the value specified to generate unique event IDs. Users can specify either a decimal value (`-G 1`) or hex value preceded by `0x` (`-G 0x11`).

1.7 More Information

Chapter 2 contains much information about many configuration options available in the configuration file. The Snort manual page and the output of `snort -?` contain information that can help you get Snort running in several different modes.

NOTE

In many shells, a backslash (\) is needed to escape the ?, so you may have to type `snort -\?` instead of `snort -?` for a list of Snort command line options.

The Snort web page (<http://www.snort.org>) and the Snort User's mailing list (<http://marc.theaimsgroup.com/?l=snort-users> at snort-users@lists.sourceforge.net provide informative announcements as well as a venue for community discussion and support. There's a lot to Snort, so sit back with a beverage of your choosing and read the documentation and mailing list archives.

Chapter 2

Configuring Snort

2.0.1 Includes

The `include` keyword allows other rule files to be included within the rules file indicated on the Snort command line. It works much like an `#include` from the C programming language, reading the contents of the named file and adding the contents in the place where the include statement appears in the file.

Format

```
include: <include file path/name>
```

NOTE

Note that there is no semicolon at the end of this line.

Included files will substitute any predefined variable values into their own variable references. See Section (2.0.2) for more information on defining and using variables in Snort rule files.

2.0.2 Variables

Variables may be defined in Snort. These are simple substitution variables set with the `var` keyword as shown in Figure 2.1.

Format

```
var: <name> <value>
```

```
var MY_NET [192.168.1.0/24,10.1.1.0/24]  
alert tcp any any -> $MY_NET any (flags:S; msg:"SYN packet";)
```

Figure 2.1: Example of Variable Definition and Usage

Rule variable names can be modified in several ways. You can define meta-variables using the `$` operator. These can be used with the variable modifier operators `?` and `-`, as described in the following table:

Variable Syntax	Description
<code>\$var</code>	Defines a meta-variable.
<code>\$(var)</code>	Replaces with the contents of variable <code>var</code> .
<code>\$(var:-default)</code>	Replaces the contents of the variable <code>var</code> with “default” if <code>var</code> is undefined.
<code>\$(var:?message)</code>	Replaces with the contents of variable <code>var</code> or prints out the error message and exits.

See Figure 2.2 for an example of advanced variable usage in action.

```
var MY_NET 192.168.1.0/24
log tcp any any -> $(MY_NET:?MY_NET is undefined!) 23
```

Figure 2.2: Figure Advanced Variable Usage Example

2.0.3 Config

Many configuration and command line options of Snort can be specified in the configuration file.

Format

```
config <directive> [: <value>]
```

Directives

Table 2.1: Config Directives

Command	Example	Description
<code>order</code>	<code>config order: pass alert log activation</code>	Changes the order that rules are evaluated.
<code>alertfile</code>	<code>config alertfile: alerts</code>	Sets the alerts output file.
<code>classification</code>	<code>config classification: misc-activity,Misc activity,3</code>	See Table 3.2 for a list of classifications.
<code>dump_chars_only</code>	<code>config dump_chars_only</code>	Turns on character dumps (<code>snort -C</code>).
<code>dump_payload</code>	<code>config dump_payload</code>	Dumps application layer (<code>snort -d</code>).
<code>decode_data_link</code>	<code>config decode_data_link</code>	Decodes Layer2 headers (<code>snort -e</code>).
<code>bpf_file</code>	<code>config bpf_file: filters.bpf</code>	Specifies BPF filters (<code>snort -F</code>).
<code>daemon</code>	<code>config daemon</code>	Forks as a daemon (<code>snort -D</code>).
<code>interface</code>	<code>config interface: xl0</code>	Sets the network interface (<code>snort -i</code>).
<code>alert_with_interface_name</code>	<code>config alert_with_interface_name</code>	Appends interface name to alert (<code>snort -I</code>).
<code>logdir</code>	<code>config logdir: /var/log/snort</code>	Sets the logdir (<code>snort -l</code>).
<code>umask</code>	<code>config umask: 022</code>	Sets umask when running (<code>snort -m</code>).
<code>pkt_count</code>	<code>config pkt_count: 13</code>	Exits after N packets (<code>snort -n</code>).
<code>nolog</code>	<code>config nolog</code>	Disables logging. Note: Alerts will still occur. (<code>snort -N</code>).
<code>obfuscate</code>	<code>config obfuscate</code>	Obfuscates IP Addresses (<code>snort -O</code>).
<code>no_promisc</code>	<code>config no_promisc</code>	Disables promiscuous mode (<code>snort -p</code>).
<code>quiet</code>	<code>config quiet</code>	Disables banner and status reports (<code>snort -q</code>).
<code>chroot</code>	<code>config chroot: /home/snort</code>	Chroots to specified dir (<code>snort -t</code>).

checksum_mode	config checksum_mode : all	Types of packets to calculate checksums. Values: none, noip, notcp, noicmp, noudp, ip, tcp, udp, icmp or all.
set_gid	config set_gid: 30	Changes GID to specified GID (snort -g).
set_uid	set_uid: snort_user	Sets UID to <id> (snort -u).
utc	config utc	Uses UTC instead of local time for timestamps (snort -U).
verbose	config verbose	Uses verbose logging to STDOUT (snort -v).
dump_payload_verbose	config dump_payload_verbose	Dumps raw packet starting at link layer (snort -X).
show_year	config show_year	Shows year in timestamps (snort -y).
stateful	config stateful	Sets assurance mode for stream4 (est). See the stream4 reassemble configuration 2.3.
min_ttl	config min_ttl:30	Sets a Snort-wide minimum ttl to ignore all traffic.
disable_decode_alerts	config disable_decode_alerts	Turns off the alerts generated by the decode phase of Snort.
disable_tcpopt_experimental_alerts	config disable_tcpopt_experimental_alerts	Turns off alerts generated by experimental TCP options.
disable_tcpopt_experimental_alerts	config disable_tcpopt_experimental_alerts	Turns off alerts generated by experimental TCP options.
disable_tcpopt_obsolete_alerts	config disable_tcpopt_obsolete_alerts	Turns off alerts generated by obsolete TCP options.
disable_tcpopt_ttcp_alerts	config disable_tcpopt_ttcp_alerts	Turns off alerts generated by T/TCP options.
disable_ttcp_alerts	config disable_ttcp_alerts	Turns off alerts generated by T/TCP options.
disable_tcpopt_alerts	config disable_tcpopt_alerts	Disables option length validation alerts.
disable_ipopt_alerts	config disable_ipopt_alerts	Disables IP option length validation alerts.
disable_decode_drops	config disable_decode_drops	Disables the dropping of bad packets identified by decoder (only applicable in inline mode).
disable_tcpopt_experimental_drops	config disable_tcpopt_experimental_drops	Disables the dropping of bad packets with obsolete TCP option (only applicable in inline mode).
disable_ttcp_drops	disable_ttcp_drops	Disables the dropping of bad packets with TCP echo option (only applicable in inline mode).
disable_tcpopt_drops	config disable_tcpopt_drops	Disables the dropping of bad packets with bad/truncated TCP option (only applicable in inline mode).
disable_ipopt_drops	config disable_ipopt_drops	Disables the dropping of bad packets with bad/truncated IP options (only applicable in inline mode).
flowbits_size	config flowbits_size: 128	Specifies the maximum number of flowbit tags that can be used within a ruleset.

event_queue	config event_queue: max_queue 512 log 100 order_events priority	Specifies conditions about Snort's event queue. You can use the following options: <ul style="list-style-type: none"> • max_queue <integer> (max events supported) • log <integer> (number of events to log) • order_events [priority content_length] (how to order events within the queue) See Section 3.10 for more information and examples.
layer2resets	config layer2resets: 00:06:76:DD:5F:E3	This option is only available when running in inline mode. See Section 1.5.
detection	config detection: search-method ac no_stream_inserts max_queue_events 128	Makes changes to the detection engine. The following options can be used: <ul style="list-style-type: none"> • search-method<ac mwm lowmem> • no_stream_inserts • max_queue_events<integer>
asn1	config asn1:256	Specifies the maximum number of nodes to track when doing ASN1 decoding. See Section 2.1.12 for more information and examples.
snaplen	config snaplen: 2048	Set the snaplength of packet, same effect as -P <snaplen> option.
read_bin_file	config read_bin_file: test_alert.pcap	Specifies a pcap file to use (instead of reading from network), same effect as -r <tf> option.
reference	config reference: myref http://myurl.com/?id=	Adds a new reference system to Snort.
ignore_ports	config ignore_ports: udp 1:17 53	Specifies ports to ignore (useful for ignoring noisy NFS traffic). Specify the protocol (tcp, udp, ip, or icmp), followed by a list of ports. Port ranges are supported.

2.1 Preprocessors

Preprocessors were introduced in version 1.5 of Snort. They allow the functionality of Snort to be extended by allowing users and programmers to drop modular plugins into Snort fairly easily. Preprocessor code is run before the detection engine is called, but after the packet has been decoded. The packet can be modified or analyzed in an out-of-band manner using this mechanism.

Preprocessors are loaded and configured using the `preprocessor` keyword. The format of the preprocessor directive in the Snort rules file is:

```
preprocessor <name>: <options>
```

```
preprocessor minfrag: 128
```

Figure 2.3: Preprocessor Directive Format Example

2.1.1 Frag2

NOTE

Frag2 is deprecated in Snort 2.4.0 and later in favor of Frag3. See Section 2.1.2 for more information about Frag3.

2.1.2 Frag3

The Frag3 preprocessor is a target-based IP defragmentation module for Snort. Frag3 is intended as a replacement for the frag2 defragmentation module and was designed with the following goals:

1. Faster execution than frag2 with less complex data management.
2. Target-based host modeling anti-evasion techniques.

The frag2 preprocessor used splay trees extensively for managing the data structures associated with defragmenting packets. Splay trees are excellent data structures to use when you have some assurance of locality of reference for the data that you are handling but in high speed, heavily fragmented environments the nature of the splay trees worked against the system and actually hindered performance. Frag3 uses the sfxhash data structure and linked lists for data handling internally which allows it to have much more predictable and deterministic performance in any environment which should aid us in managing heavily fragmented environments.

Target-based analysis is a relatively new concept in network-based intrusion detection. The idea of a target-based system is to model the actual targets on the network instead of merely modeling the protocols and looking for attacks within them. When IP stacks are written for different operating systems, they are usually implemented by people who read the RFCs and then their interpretation of what the RFC outlines into code. Unfortunately, there are ambiguities in the way that the RFCs define some of the edge conditions that may occur and when this happens different people implement certain aspects of their IP stacks differently. For an IDS this is a big problem.

In an environment where the attacker can determine what style of IP defragmentation is being used on a particular target, the attacker can try to fragment packets such that the target will put them back together in a specific manner while any passive systems trying to model the host traffic have to guess which way the target OS is going to handle the overlaps and retransmits. As I like to say, if the attacker has more information about the targets on a network than the IDS does, it is possible to evade the IDS. This is where

the idea for “target-based IDS” came from. For more detail on this issue and how it affects IDSes, check out the famous Ptacek & Newsham paper at <http://www.snort.org/docs/idspaper/>.

The basic idea behind target-based IDS is that we tell the IDS information about hosts on the network so that it can avoid Ptacek & Newsham style evasion attacks based on information about how an individual target IP stack operates. Vern Paxson and Umesh Shankar did a great paper on this very topic in 2003 that detailed mapping the hosts on a network and determining how their various IP stack implementations handled the types of problems seen in IP defragmentation and TCP stream reassembly. Check it out at <http://www.icir.org/vern/papers/activemap-oak03.pdf>.

We can also present the IDS with topology information to avoid TTL-based evasions and a variety of other issues, but that’s a topic for another day. Once we have this information we can start to really change the game for these complex modeling problems.

Frag3 was implemented to showcase and prototype a target-based module within Snort to test this idea.

Frag 3 Configuration

Frag3 configuration is somewhat more complex than frag2. There are at least two preprocessor directives required to activate frag3, a global configuration directive and an engine instantiation. There can be an arbitrary number of engines defined at startup with their own configuration, but only one global configuration.

Global configuration

- Preprocessor name: `frag3_global`
- Available options:
 - `max_fragments <number>` - Maximum simultaneous fragments to track, default is 8192
 - `memcap <bytes>` - Memory cap for self preservation, default is 4MB
 - `prealloc_fragments <number>` - alternate memory management mode, use preallocated fragment nodes (faster in some situations)

Engine Configuration

- Preprocessor name: `frag3_engine`
- Available options:
 - `timeout <seconds>` - Timeout for fragments, fragments in the engine for longer than this period will be automatically dropped. Default is 60 seconds.
 - `ttl_limit <hops>` - Max TTL delta acceptable for packets based on the first packet in the fragment. Default is 5.
 - `min_ttl <value>` - Minimum acceptable TTL value for a fragment packet. Default is 1.
 - `detect_anomalies` - Detect fragment anomalies
 - `bind_to <ip-list>` - IP List to bind this engine to. This engine will only run for packets with destination addresses contained within the IP List. Default value is `all`.
 - `policy <type>` - Select a target-based defragmentation mode. Available types are `first`, `last`, `bsd`, `bsd-right`, `linux`. Default type is `bsd`.

The Paxson Active Mapping paper introduced the terminology frag3 is using to describe policy types. The known mappings are as follows. Anyone who develops more mappings and would like to add to this list please feel free to send us an email!

Platform	Type
AIX 2	BSD
AIX 4.3 8.9.3	BSD
Cisco IOS	Last
FreeBSD	BSD
HP JetDirect (printer)	BSD-right
HP-UX B.10.20	BSD
HP-UX 11.00	First
IRIX 4.0.5F	BSD
IRIX 6.2	BSD
IRIX 6.3	BSD
IRIX64 6.4	BSD
Linux 2.2.10	linux
Linux 2.2.14-5.0	linux
Linux 2.2.16-3	linux
Linux 2.2.19-6.2.10smp	linux
Linux 2.4.7-10	linux
Linux 2.4.9-31SGI 1.0.2smp	linux
Linux 2.4 (RedHat 7.1-7.3)	linux
MacOS (version unknown)	First
NCD Thin Clients	BSD
OpenBSD (version unknown)	linux
OpenBSD (version unknown)	linux
OpenVMS 7.1	BSD
OS/2 (version unknown)	BSD
OSF1 V3.0	BSD
OSF1 V3.2	BSD
OSF1 V4.0,5.0,5.1	BSD
SunOS 4.1.4	BSD
SunOS 5.5.1,5.6,5.7,5.8	First
Tru64 Unix V5.0A,V5.1	BSD
Vax/VMS	BSD
Windows (95/98/NT4/W2K/XP)	First

format

```
preprocessor frag3_global
preprocessor frag3_engine
```

Figure 2.4: Example configuration (Basic)

```
preprocessor frag3_global: prealloc_nodes 8192
preprocessor frag3_engine: policy linux, bind_to 192.168.1.0/24
preprocessor frag3_engine: policy first, bind_to [10.1.47.0/24,172.16.8.0/24]
preprocessor frag3_engine: policy last, detect_anomalies
```

Figure 2.5: Example configuration (Advanced)

Note in the advanced example (Figure 2.5), there are three engines specified running with linux, first and last policies assigned. The first two engines are bound to specific IP address ranges and the last one applies to all other traffic, packets that don't fall within the address requirements of the first two engines automatically fall through to the third one.

Frag 3 Alert Output

Frag3 is capable of detecting eight different types of anomalies. Its event output is packet-based so it will work with all output modes of Snort. Read the documentation in the doc/signatures directory with filenames that begin with “123-” for information on the different event types.

2.1.3 Stream4

The Stream4 module provides TCP stream reassembly and stateful analysis capabilities to Snort. Robust stream reassembly capabilities allow Snort to ignore “stateless” attacks (which include the types of attacks that Stick and Snot produce). Stream4 also gives large scale users the ability to track many simultaneous TCP streams. Stream4 is set to handle 8192 simultaneous TCP connections in its default configuration, however, it scales to handle over 100,000 simultaneous connections.

Stream4 contains two configurable modules, the Stream4 preprocessor and the associated Stream4 reassemble plugin. The stream4.reassemble options are listed below.

NOTE

Additional options can be used if Snort is running in inline mode. See Section 1.5.2 for more information.

Stream4 Format

```
preprocessor stream4: [noinspect], [asynchronous_link], [keepstats [<machine|binary>]], \  
    [detect_scans], [log_flushed_streams], [detect_state_problems], \  
    [disable_evasion_alerts], [timeout <seconds>], [memcap <bytes>], \  
    [max_sessions <num sessions>], [cache_clean_percent <% of sessions>], \  
    [cache_clean_sessions <num of sessions>], [ttl_limit <count>], \  
    [self_preservation_threshold], [self_preservation_period], \  
    [suspend_threshold], [suspend_period], [enforce_state], \  
    [state_protection], [server_inspect_limit <bytes>]
```

Option	Description
asynchronous_link	Use state transitions based only on one-sided conversations (ack/sequence numbers).
cache_clean_percent	Purge this percent of least-recently used sessions from the session cache (cache_clean_sessions).
cache_clean_sessions	Purge this number of least-recently used sessions from the session cache.
detect_scans	Turns on alerts for portscan events.
detect_state_problems	Turns on alerts for stream events of note, such as evasive RST or SYN packet, and out of window sequence numbers.
enforce_state	Enforces statefulness so that sessions aren't picked up mid-stream.
favor_old	Favor old segments based on sequence number over a new segment.
favor_new	Favor new segments based on sequence number over a old segment.
flush_behavior <default large_window random>	Use specified flush behavior. Random flush points defined by flush_base and flush_range.
flush_base <number>	Lowest allowed random flushpoint. The default value is 512. If flush_behavior is less than 0.
flush_range <number>	Space within random flushpoints are generated. The default value is 1024. If flush_behavior is less than 0.
flush_seed <number>	Random seed for flushpoints. The default value is computed from the current time. Only used if flush_behavior is less than 0.
keepstats	Records session summary information in <logdir>/session.log. If not specified, output is human readable.
log_flushed_streams	Log the packets that are part of reassembled stream.
disable_evasion_alerts	Turns off alerts for events such as TCP overlap.
timeout <seconds>	Amount of time to keep an inactive stream in the state table. If not flushed will automatically be picked up again if more activity is seen. The default value is 30 seconds.
memcap <bytes>	Number of bytes used to store packets for reassembly.
max_sessions	Maximum number of simultaneous sessions.
noinspect	Disables stateful inspection.
ttl_limit	Sets the delta value that will set off an evasion alert.
self_preservation_threshold	Limit on number of sessions before entering self-preservation mode (no data on the default ports).
self_preservation_period	Length of time (seconds) to remain in self-preservation mode.
suspend_threshold	Limit on number of sessions before entering suspend mode (no data on the default ports).
suspend_period	Length of time (seconds) to remain in suspend mode.
server_inspect_limit	Restrict inspection of server traffic to this many bytes until an alert is seen (ie: client packet with data).
state_protection	Protect self against DoS attacks.

Stream4_reassemble Format

```
preprocessor stream4_reassemble: [clientonly], [serveronly], [both], [noalerts], \
    [favor_old], [favor_new], [flush_on_alert], \
    [overlap_limit], [ports <portlist>]
```


Option	Description
<code>clientonly</code>	Provides reassembly for the client side of a connection only.
<code>serveronly</code>	Provides reassembly for the server side of a connection only.
<code>both</code>	Reassemble for client and server sides of connection.
<code>noalerts</code>	Won't alert on events that may be insertion or evasion attacks.
<code>favor_old/favor_new</code>	Use newer or older packet when reassembling.
<code>flush_on_alert</code>	Flush a stream when an individual packet causes an alert.
<code>overlap_limit</code>	Alert when the number of overlapping data bytes reaches a threshold.
<code>ports <portlist></code>	Provides reassembly for a whitespace-separated list of ports. By default, reassembly is performed for ports 21, 23, 25, 42, 53, 80, 110, 111, 135, 136, 137, 139, 143, 445, 513, 1443, 1521, and 3306. To perform reassembly for all ports, use <code>all</code> as the port list.

Notes

Just setting the `Stream4` and `Stream4.reassemble` directives without arguments in the `snort.conf` file will set them up in their default configurations shown in Table 2.2 and Table 2.3.

`Stream4` introduces a new command line switch: `-z`. On TCP traffic, if the `-z` switch is specified, Snort will only alert on streams that have been established via a three way handshake or streams where cooperative bidirectional activity has been observed (i.e., where some traffic went one way and something other than a RST or FIN was seen going back to the originator). With `-z` turned on, Snort completely ignores TCP-based Stick/Snot attacks.

NOTE

As of Snort v2.1.1, `-z` has been superseded by the `flow: established` rule option. Please see section 3.6.9 for more information.

Table 2.2: Stream4 Defaults

Option	Default
<code>session timeout (timeout)</code>	30 seconds
<code>session memory cap (memcap)</code>	8388608 bytes
<code>stateful inspection (noinspect)</code>	active (<code>noinspect</code> disabled)
<code>stream stats (keepstats)</code>	inactive
<code>state problem alerts (detect_state_problems)</code>	inactive (<code>detect_state_problems</code> disabled)
<code>evasion alerts (disable_evasion_alerts)</code>	inactive (<code>disable_evasion_alerts</code> enabled)
<code>asynchronous link (asynchronous.link)</code>	inactive
<code>log flushed streams (log_flushed_streams)</code>	inactive
<code>max sessions (max_sessions)</code>	8192
<code>session cache purge (cache_clean_sessions)</code>	5
<code>session cache purge percent (cache_clean_percent)</code>	inactive
<code>self preservation threshold (self_preservation_threshold)</code>	50 sessions/sec
<code>self preservation period (self_preservation_period)</code>	90 seconds
<code>suspend threshold (suspend_threshold)</code>	200 sessions/sec
<code>suspend period (suspend_period)</code>	30 seconds
<code>state protection (state_protection)</code>	inactive
<code>server inspect limit (server_inspect_limit)</code>	-1 (inactive)

Table 2.3: Stream4_reassemble Defaults

Option	Default
reassemble client (clientonly)	active
reassemble server (serveronly)	inactive
reassemble both (both)	inactive
reassemble ports (ports)	21 23 25 42 53 80 110 111 135 136 137 139 143 445 513 1433 1521 3306
reassemble alerts (noalerts)	active (noalerts disabled)
favor old packet (favor_old)	active
favor new packet (favor_new)	inactive
flush on alert (flush_on_alert)	inactive
overlap limit (overlap_limit)	-1 (inactive)

2.1.4 Flow

The Flow tracking module is meant to start unifying the state keeping mechanisms of Snort into a single place. As of Snort 2.1.0, only a portscan detector is implemented, but in the long term, many of the stateful subsystems of Snort will be migrated over to becoming flow plugins. With the introduction of flow, this effectively makes the conversation preprocessor obsolete.

An IPv4 flow is unique when the IP protocol (`ip_proto`), source IP (`sip`), source port (`sport`), destination IP (`dip`), and destination port (`dport`) are the same. The `dport` and `sport` are 0 unless the protocol is TCP or UDP.

Format

```
preprocessor flow: [memcap <bytes>], [rows <count>], \  
                  [stats_interval <seconds>], [hash <1|2>]
```

Table 2.4: Flow Options

Option	Description
<code>memcap</code>	Number of bytes to allocate.
<code>rows</code>	Number of rows for the flow hash table. ^a
<code>stats_interval</code>	Interval (in seconds) to dump statistics to STDOUT. Set this to 0 to disable.
<code>hash</code>	Hashing method to use. ^b

^aThis number can be increased, at the cost of using more memory, to enhance performance. Increasing rows provide a larger hash table.

^b1 - hash by byte, 2 - hash by integer (faster, not as much of a chance to become diverse). The hash table has a pseudo-random salt picked to make algorithmic complexity attacks much more difficult.

Example Configuration

```
preprocessor flow: stats_interval 0 hash 2
```

2.1.5 Portscan

⚠NOTE

The "Portscan" preprocessor was deprecated in Snort 2.2, in favor of Flow Portscan, which was deprecated in Snort 2.3, in favor of sfPortscan.

2.1.6 Flow-Portscan

⚠NOTE

The Flow-Portscan preprocessor was deprecated in Snort 2.3, in favor of sfPortscan.

2.1.7 sfPortscan

The sfPortscan module, developed by Sourcefire, is designed to detect the first phase in a network attack: Reconnaissance. In the Reconnaissance phase, an attacker determines what types of network protocols or services a host supports. This is the traditional place where a portscan takes place. This phase assumes the

attacking host has no prior knowledge of what protocols or services are supported by the target, otherwise this phase would not be necessary.

As the attacker has no beforehand knowledge of its intended target, most queries sent by the attacker will be negative (meaning that the service ports are closed). In the nature of legitimate network communications, negative responses from hosts are rare, and rarer still are multiple negative responses within a given amount of time. Our primary objective in detecting portscans is to detect and track these negative responses.

One of the most common portscanning tools in use today is Nmap. Nmap encompasses many, if not all, of the current portscanning techniques. sfPortscan was designed to be able to detect the different types of scans Nmap can produce.

sfPortscan will currently alert for the following types of Nmap scans:

- TCP Portscan
- UDP Portscan
- IP Portscan

These alerts are for one→one portscans, which are the traditional types of scans; one host scans multiple ports on another host. Most of the port queries will be negative, since most hosts have relatively few services available.

sfPortscan also alerts for the following types of decoy portscans:

- TCP Decoy Portscan
- UDP Decoy Portscan
- IP Decoy Portscan

Decoy portscans are much like the Nmap portscans described above, only the attacker has spoofed source address inter-mixed with the real scanning address. This tactic helps hide the true identity of the attacker.

sfPortscan alerts for the following types of distributed portscans:

- TCP Distributed Portscan
- UDP Distributed Portscan
- IP Distributed Portscan

These are many→one portscans. Distributed portscans occur when multiple hosts query one host for open services. This is used to evade an IDS and obfuscate command and control hosts.

NOTE

Negative queries will be distributed among scanning hosts, so we track this type of scan through the scanned host.

sfPortscan alerts for the following types of portsweeps:

- TCP Portsweep
- UDP Portsweep
- IP Portsweep
- ICMP Portsweep

These alerts are for one→many portsweeps. One host scans a single port on multiple hosts. This usually occurs when a new exploit comes out and the attacker is looking for a specific service.

NOTE

The characteristics of a portsweep scan may not result in many negative responses. For example, if an attacker portsweeps a web farm for port 80, we will most likely not see many negative responses.

sfPortscan alerts on the following filtered portscans and portsweeps:

- TCP Filtered Portscan
- UDP Filtered Portscan
- IP Filtered Portscan
- TCP Filtered Decoy Portscan
- UDP Filtered Decoy Portscan
- IP Filtered Decoy Portscan
- TCP Filtered Portsweep
- UDP Filtered Portsweep
- IP Filtered Portsweep
- ICMP Filtered Portsweep
- TCP Filtered Distributed Portscan
- UDP Filtered Distributed Portscan
- IP Filtered Distributed Portscan

“Filtered” alerts indicate that there were no network errors (ICMP unreachables or TCP RSTs) or responses on closed ports have been suppressed. It’s also a good indicator of whether the alert is just a very active legitimate host. Active hosts, such as NATs, can trigger these alerts because they can send out many connection attempts within a very small amount of time. A filtered alert may go off before responses from the remote hosts are received.

sfPortscan only generates one alert for each host pair in question during the time window (more on windows below). On TCP scan alerts, sfPortscan will also display any open ports that were scanned. On TCP sweep alerts however, sfPortscan will only track open ports after the alert has been triggered. Open port events are not individual alerts, but tags based off the original scan alert.

sfPortscan Configuration

You may want to use the following line in your snort.conf to disable evasion alerts within stream4 because some scan packets can cause these alerts to be generated:

```
preprocessor stream4: disable_evasion_alerts
```

Use of the Flow preprocessor is required for sfPortscan. Flow gives portscan direction in the case of connectionless protocols like ICMP and UDP. You should enable the Flow preprocessor in your snort.conf by using the following:

```
preprocessor flow: stats_interval 0 hash 2
```

The parameters you can use to configure the portscan module are:

1. proto <protocol>

Available options:

- tcp
- udp
- icmp
- ip_proto
- all

2. scan_type <scan_type>

Available options:

- portscan
- portsweep
- decoy_portscan
- distributed_portscan
- all

3. sense_level <level>

Available options:

- low - “Low” alerts are only generated on error packets sent from the target host, and because of the nature of error responses, this setting should see very few false positives. However, this setting will never trigger a Filtered Scan alert because of a lack of error responses. This setting is based on a static time window of 60 seconds, after which this window is reset.
- medium - “Medium” alerts track Connection Counts, and so will generate Filtered Scan alerts. This setting may false positive on active hosts (NATs, proxies, DNS caches, etc), so the user may need to deploy the use of Ignore directives to properly tune this directive.
- high - “High” alerts continuously track hosts on a network using a time window to evaluate portscan statistics for that host. A “High” setting will catch some slow scans because of the continuous monitoring, but is very sensitive to active hosts. This most definitely will require the user to tune sfPortscan.

4. watch_ip <ip1,ip2/cidr>

Defines what IPs or networks to watch. IPs or networks not falling into this range are ignored.

5. ignore_scanners <ip_list>

Ignores the source of scan alerts.

6. ignore_scanned <ip_list>

Ignores the destination of scan alerts

7. logfile <file>

This option will output portscan events to the file specified. If `file` does not contain a leading slash, this file will be placed in the Snort config dir.

Format

```
preprocessor sfportscan: proto <protocols> \  
scan_type <portscan|portsweep|decoy_portscan|distributed_portscan|all>\  
sense_level <low|medium|high> watch_ip <IP or IP/CIDR> ignore_scanners <IP list>\  
ignore_scanned <IP list> logfile <path and filename>
```

```
preprocessor flow: stats_interval 0 hash 2
preprocessor sfportscan: proto { all } \
    scan_type { all } \
    sense_level { low }
```

Figure 2.6: sfPortscan Preprocessor Configuration

sfPortscan Alert Output

Unified Output

In order to get all the portscan information logged with the alert, snort generates a pseudo-packet and uses the payload portion to store the additional portscan information of priority count, connection count, IP count, port count, IP range, and port range. The characteristics of the packet are:

Src/Dst MAC Addr == MACDAD IP Protocol == 255 IP TTL == 0

Other than that, the packet looks like the IP portion of the packet that caused the portscan alert to be generated. This includes any IP options, etc. The payload and payload size of the packet is equal to the length of the additional portscan information that is logged. The size tends to be around 100 - 200 bytes.

Open port alerts differ from the other portscan alerts, because open port alerts utilize the tagged packet output system. This means that if an output system that doesn't print tagged packets is used, then the user won't see open port alerts. The open port information is stored in the IP payload and contains the port that is open.

The sfPortscan alert output was designed to work with unified packet logging, so it is possible to extend favorite snort GUIs to display portscan alerts and the additional information in the IP payload using the above packet characteristics.

Log File Output

Logfile output is displayed in the following format, and explained further below:

```
Time: 09/08-15:07:31.603880
event_id: 2
192.168.169.3 -> 192.168.169.5 (portscan) TCP Filtered Portscan
Priority Count: 0
Connection Count: 200
IP Count: 2
Scanner IP Range: 192.168.169.3:192.168.169.4
Port/Proto Count: 200
Port/Proto Range: 20:47557
```

If there are open ports on the target, an additional tagged packet(s) will be appended:

```
Time: 09/08-15:07:31.603881
event_ref: 2
192.168.169.3 -> 192.168.169.5 (portscan) Open Port
Open Port: 38458
```

1. Event_id/Event_ref

These fields are used to link an alert with the corresponding Open Port tagged packet

2. Priority Count

Priority Count keeps track of bad responses (resets, unreachables). The higher the Priority Count, the more bad responses have been received.

3. Connection Count

Connection Count lists how many connections are active on the hosts (src or dst). This is accurate for connection-based protocols, and is more of an estimate for others. Whether or not a portscan was filtered is determined here. High connection count and low priority count would indicate filtered (no response received from target).

4. IP Count

IP Count keeps track of the last IP to contact a host, and increments the count if the next IP is different. For one-to-one scans, this is a low number. For active hosts this number will be high regardless, and one-to-one scans may appear as a distributed scan.

5. Scanned/Scanner IP Range

This field changes depending on the type of alert. Portsweeps (one-to-many) scans display the scanned IP range; Portscans (one-to-one) display the scanner IP.

6. Port Count

Port Count keeps track of the last port contacted and increments this number when that changes. We use this count (along with IP Count) to determine the difference between one-to-one portscans and one-to-one decoys.

Tuning sfPortscan

The most important aspect in detecting portscans is tuning the detection engine for your network(s). Here are some tuning tips:

1. Use the `watch_ip`, `ignore_scanners`, and `ignore_scanned` options.

It's important to correctly set these options. The `watch_ip` option is easy to understand. The analyst should set this option to the list of Cidr blocks and IPs that they want to watch. If no `watch_ip` is defined, sfPortscan will watch all network traffic.

The `ignore_scanners` and `ignore_scanned` options come into play in weeding out legitimate hosts that are very active on your network. Some of the most common examples are NAT IPs, DNS cache servers, syslog servers, and nfs servers. sfPortscan may not generate false positives for these types of hosts, but be aware when first tuning sfPortscan for these IPs. Depending on the type of alert that the host generates, the analyst will know which to ignore it as. If the host is generating portsweep events, then add it to the `ignore_scanners` option. If the host is generating portscan alerts (and is the host that is being scanned), add it to the `ignore_scanned` option.

2. Filtered scan alerts are much more prone to false positives.

When determining false positives, the alert type is very important. Most of the false positives that sfPortscan may generate are of the filtered scan alert type. So be much more suspicious of filtered portscans. Many times this just indicates that a host was very active during the time period in question. If the host continually generates these types of alerts, add it to the `ignore_scanners` list or use a lower sensitivity level.

3. Make use of the Priority Count, Connection Count, IP Count, Port Count, IP range, and Port range to determine false positives.

The portscan alert details are vital in determining the scope of a portscan and also the confidence of the portscan. In the future, we hope to automate much of this analysis in assigning a scope level and confidence level, but for now the user must manually do this. The easiest way to determine false positives is through simple ratio estimations. The following is a list of ratios to estimate and the associated values that indicate a legitime scan and not a false positive.

Connection Count / IP Count: This ratio indicates an estimated average of connections per IP. For portscans, this ratio should be high, the higher the better. For portsweeps, this ratio should be low.

Port Count / IP Count: This ratio indicates an estimated average of ports connected to per IP. For portscans, this ratio should be high and indicates that the scanned host's ports were connected to by fewer IPs. For portsweeps, this ratio should be low, indicating that the scanning host connected to few ports but on many hosts.

Connection Count / Port Count: This ratio indicates an estimated average of connections per port. For portscans, this ratio should be low. This indicates that each connection was to a different port. For portsweeps, this ratio should be high. This indicates that there were many connections to the same port.

The reason that Priority Count is not included, is because the priority count is included in the connection count and the above comparisons take that into consideration. The Priority Count play an important role in tuning because the higher the priority count the more likely it is a real portscan or portsweep (unless the host is firewalled).

4. If all else fails, lower the sensitivity level.

If none of these other tuning techniques work or the analyst doesn't have the time for tuning, lower the sensitivity level. You get the best protection the higher the sensitivity level, but it's also important that the portscan detection engine generates alerts that the analyst will find informative. The low sensitivity level only generates alerts based on error responses. These responses indicate a portscan and the alerts generated by the low sensitivity level are highly accurate and require the least tuning. The low sensitivity level does not catch filtered scans, since these are more prone to false positives.

2.1.8 Telnet Decode

The `telnet_decode` preprocessor allows Snort to normalize Telnet control protocol characters from the session data. In Snort 1.9.0 and above, it accepts a list of ports to run on as arguments. Also in 1.9.0, it normalizes into a separate data buffer from the packet itself so that the raw data may be logged or examined with the `rawbytes` content modifier^{3.5.3}.

By default, `telnet_decode` runs against traffic on ports 21, 23, 25, and 119.

Format

```
preprocessor telnet_decode: <ports>
```

2.1.9 RPC Decode

The `rpc_decode` preprocessor normalizes RPC multiple fragmented records into a single un-fragmented record. It does this by normalizing the packet into the packet buffer. If `stream4` is enabled, it will only process client-side traffic. By default, it runs against traffic on ports 111 and 32771.

Table 2.5: RPC Decoder Options

Option	Description
<code>alert_fragments</code>	Alert on any fragmented RPC record.
<code>no.alert.multiple.requests</code>	Don't alert when there are multiple records in one packet.
<code>no.alert.large.fragments</code>	Don't alert when the sum of fragmented records exceeds one packet.
<code>no.alert.incomplete</code>	Don't alert when a single fragment record exceeds the size of one packet.

Format

```
preprocessor rpc_decode: <ports> [ alert_fragments ] \
```

```
[no_alert_multiple_requests] [no_alert_large_fragments] \  
[no_alert_incomplete]
```

2.1.10 Performance Monitor

This preprocessor measures Snort's real-time and theoretical maximum performance. Whenever this preprocessor is turned on, it should have an output mode enabled, either "console" which prints statistics to the console window or "file" with a file name, where statistics get printed to the specified file name. By default, Snort's real-time statistics are processed. This includes:

- Mbits/Sec
- Drop Rate
- Alerts/Sec
- K-Pkts/Sec
- Avg Bytes/Pkt
- Pat-Matched [percent of data received that Snort processes in pattern matching]
- Syns/Sec
- SynAcks/Sec
- New Sessions Cached/Sec
- Sessions Del fr Cache/Sec
- Current Cached Sessions
- Max Cached Sessions
- Stream Flushes/Sec
- Stream Session Cache Faults
- Stream Session Cache Timeouts
- Frag-Completes/Sec
- Frag-Inserts/Sec
- Frag-Deletes/Sec
- Frag-Flushes/Sec
- Frag-Timeouts
- Frag-Faults
- CPU usage (user)
- CPU usage (sys)
- CPU usage (Idle)
- Mbits/Sec (wire) [average mbits of total traffic]
- Mbits/Sec (ipfrag) [average mbits of IP fragmented traffic]
- Mbits/Sec (ipreass) [average mbits Snort injects after IP reassembly]
- Mbits/Sec (tcprebuilt) [average mbits Snort injects after stream4 reassembly]

- Mbits/Sec (applayer) [average mbits seen by rules and protocol decoders]
- Avg Bytes/Pkt (wire)
- Avg Bytes/Pkt (ipfrag)
- Avg Bytes/Pkt (ipreass)
- Avg Bytes/Pkt (tcprebuilt)
- Avg Bytes/Pkt (applayer)
- K-Pkts/Sec (wire)
- K-Pkts/Sec (ipfrag)
- K-Pkts/Sec (ipreass)
- K-Pkts/Sec (tcprebuilt)
- K-Pkts/Sec (applayer)

The following options can be used with the Performance Monitor:

- **flow** - Prints out statistics about the type of traffic and protocol distributions that Snort is seeing. This option can produce large amounts of output.
- **events** - Turns on event reporting. This prints out statistics as to the number of signatures that were matched by the setwise pattern matcher (*non-qualified events*) and the number of those matches that were verified with the signature flags (*qualified events*). This shows the user if there is a problem with the ruleset that they are running.
- **max** - Turns on the theoretical maximum performance that Snort calculates given the processor speed and current performance. This is only valid for uniprocessor machines, since many operating systems don't keep accurate kernel statistics for multiple CPUs.
- **console** - Prints statistics at the console, this is enabled by default.
- **file** - Prints statistics in a comma-delimited format to the file that is specified. Not all statistics are output to this file. You may also use **snortfile** which will output into your defined Snort log directory. Both of these directives can be overridden on the command line with the **-Z** option.
- **pktcnt** - Adjusts the number of packets to process before checking for the time sample. This boosts performance, since checking the time sample reduces Snort's performance. By default, this is 10000.
- **time** - Represents the number of seconds between intervals.
- **accumulate** or **reset** - Defines what type of drop statistics are kept by the operating system. By default, **accumulate** is used.
- **atexitonly** - Dump stats for entire life of snort.

Examples

```
preprocessor perfmonitor: time 30 events flow file stats.profile max \
    console pktcnt 10000
preprocessor perfmonitor: time 300 file /var/tmp/snortstat pktcnt 10000
```

2.1.11 HTTP Inspect

HTTPInspect is a generic HTTP decoder for user applications. Given a data buffer, HTTPInspect will decode the buffer, find HTTP fields, and normalize the fields. HTTPInspect works on both client requests and server responses.

The current version of HTTPInspect only handles stateless processing. This means that HTTPInspect looks for HTTP fields on a packet-by-packet basis, and will be fooled if packets are not reassembled. This works fine when there is another module handling the reassembly, but there are limitations in analyzing the protocol. Future versions will have a stateful processing mode which will hook into various reassembly modules.

HTTPInspect has a very “rich” user configuration. Users can configure individual HTTP servers with a variety of options, which should allow the user to emulate any type of web server. Within HTTPInspect, there are two areas of configuration: global and server.

Global Configuration

The global configuration deals with configuration options that determine the global functioning of HTTPInspect. The following example gives the generic global configuration format:

Format

```
preprocessor http_inspect: global \  
    iis_unicode_map <map_filename> \  
    codemap <integer> \  
    [detect_anomalous_servers] \  
    [proxy_alert]
```

You can only have a single global configuration, you’ll get an error if you try otherwise.

Configuration

1. `iis_unicode_map <map_filename> [codemap <integer>]`

This is the global `iis_unicode_map` file. The `iis_unicode_map` is a required configuration parameter. The map file can reside in the same directory as `snort.conf` or specified via a fully-qualified path to the map file.

The `iis_unicode_map` file is a Unicode codepoint map which tells HTTPInspect which codepage to use when decoding Unicode characters. For US servers, the codemap is usually 1252.

A Microsoft US Unicode codepoint map is provided in the snort source *etc* directory by default. It is called *unicode.map* and should be used if no other codepoint map is available. A tool is supplied with Snort to generate custom Unicode maps—`ms_unicode_generator.c`, which is available at <http://www.snort.org/dl/contrib/>.

⚠NOTE

Remember that this configuration is for the global IIS Unicode map—individual servers can reference their own IIS Unicode map.

2. `detect_anomalous_servers`

This global configuration option enables generic HTTP server traffic inspection on non-HTTP configured ports, and alerts if HTTP traffic is seen. Don’t turn this on if you don’t have a default server configuration that encompasses all of the HTTP server ports that your users might access. In the future, we want to limit this to specific networks so it’s more useful, but for right now, this inspects all network traffic.

3. proxy-alert

This enables global alerting on HTTP server proxy usage. By configuring HTTPInspect servers and enabling `allow-proxy-use`, you will only receive proxy use alerts for web users that aren't using the configured proxies or are using a rogue proxy server.

Please note that if users aren't required to configure web proxy use, then you may get a lot of proxy alerts. So, please only use this feature with traditional proxy environments. Blind firewall proxies don't count.

Example Global Configuration

```
preprocessor http_inspect: global iis_unicode_map unicode.map 1252
```

Server Configuration

There are two types of server configurations: default and by IP address.

Default This configuration supplies the default server configuration for any server that is not individually configured. Most of your web servers will most likely end up using the default configuration.

Example Default Configuration

```
preprocessor http_inspect_server: server default profile all ports { 80 }
```

Configuration by IP Address This format is very similar to “default,” the only difference being that specific IPs can be configured.

Example IP Configuration

```
preprocessor http_inspect_server: server 10.1.1.1 profile all ports { 80 }
```

Server Configuration Options

Important: Some configuration options have an argument of ‘yes’ or ‘no’. This argument specifies whether the user wants the configuration option to generate an httpinspect alert or not. the ‘yes/no’ argument does not specify whether the configuration option itself is on or off, only the alerting functionality. in other words, whether set to ‘yes’ or ‘no’, HTTP normalization will still occur, and rules based off HTTP traffic will still trigger.

1. profile <all|apache|iis>

Users can configure HTTPInspect by using pre-defined HTTP server profiles. Profiles allow the user to easily configure the preprocessor for a certain type of server, but are not required for proper operation.

There are three profiles available: all, apache, and iis.

1-A. all

The “all” profile is meant to normalize the URI using most of the common tricks available. We alert on the more serious forms of evasions. This is a great profile for detecting all types of attacks, regardless of the HTTP server. “`profile all`” sets the configuration options described in Table 2.6.

Table 2.6: Options for the “all” Profile

Option	Setting
flow_depth	300
chunk encoding	alert on chunks larger than 500000 bytes
iis_unicode_map	codepoint map in the global configuration
ascii decoding	on, alert off
multiple slash	on, alert off
directory normalization	on, alert off
apache whitespace	on, alert off
double decoding	on, alert on
%u decoding	on, alert on
bare byte decoding	on, alert on
iis unicode codepoints	on, alert on
iis backslash	on, alert off
iis delimiter	on, alert off
webroot	on, alert on
non_strict URL parsing	on
tab_uri_delimiter	is set

1-B. apache

The “**apache**” profile is used for Apache web servers. This differs from the “iis” profile by only accepting UTF-8 standard Unicode encoding and not accepting backslashes as legitimate slashes, like IIS does. Apache also accepts tabs as whitespace. “**profile apache**” sets the configuration options described in Table 2.7.

Table 2.7: Options for the “apache” Profile

Option	Setting
flow_depth	300
chunk encoding	alert on chunks larger than 500000 bytes
ascii decoding	on, alert off
multiple slash	on, alert off
directory normalization	on, alert off
webroot	on, alert on
apache whitespace	on, alert on
utf_8 encoding	on, alert off
non_strict url parsing	on
tab_uri_delimiter	is set

1-C. iis

The “**iis**” profile mimics IIS servers. So that means we use IIS Unicode codemaps for each server, %u encoding, bare-byte encoding, double decoding, backslashes, etc. “**profile iis**” sets the configuration options described in Table 2.8.

The default options used by HTTP Inspect do not use a profile and are described in Table 2.9.

Table 2.8: Options for the “iis” Profile

Option	Setting
flow_depth	300
chunk encoding	alert on chunks larger than 500000 bytes
iis_unicode_map	codepoint map in the global configuration
ascii decoding	on, alert off
multiple slash	on, alert off
directory normalization	on, alert off
webroot	on, alert on
double decoding	on, alert on
%u decoding	on, alert on
bare byte decoding	on, alert on
iis unicode codepoints	on, alert on
iis backslash	on, alert off
iis delimiter	on, alert on
apache whitespace	on, alert on
non_strict URL parsing	on

Table 2.9: Default HTTP Inspect Options

Option	Setting
port	80
flow_depth	300
chunk encoding	alert on chunks larger than 500000 bytes
ascii decoding	on, alert off
utf_8 encoding	on, alert off
multiple slash	on, alert off
directory normalization	on, alert off
webroot	on, alert on
iis backslash	on, alert off
apache whitespace	on, alert off
iis delimiter	on, alert off
non_strict URL parsing	on

Profiles must be specified as the first server option and cannot be combined with any other options except:

- `ports`
- `iis_unicode_map`
- `allow_proxy_use`
- `flow_depth`
- `no_alerts`
- `inspect_uri_only`
- `oversize_dir_length`

These options must be specified after the 'profile' option.

Example

```
preprocessor http_inspect_server: server 1.1.1.1 profile all ports { 80 3128 }
```

2. `ports {<port> [<port>< ... >]}`

This is how the user configures which ports to decode on the HTTP server. Encrypted traffic (SSL) cannot be decoded, so adding port 443 will only yield encoding false positives.

3. `iis_unicode_map <map_filename> codemap <integer>`

The IIS Unicode map is generated by the program `ms_unicode_generator.c`. This program is located on the Snort.org web site at <http://www.snort.org/dl/contrib/> directory. Executing this program generates a Unicode map for the system that it was run on. So, to get the specific Unicode mappings for an IIS web server, you run this program on that server and use that Unicode map in this configuration.

When using this option, the user needs to specify the file that contains the IIS Unicode map and also specify the Unicode map to use. For US servers, this is usually 1252. But the `ms_unicode_generator` program tells you which codemap to use for your server, it's the ANSI code page. You can select the correct code page by looking at the available code pages that the `ms_unicode_generator` outputs.

4. `flow_depth <integer>`

This specifies the amount of server response payload to inspect. This option significantly increases IDS performance because we are ignoring a large part of the network traffic (HTTP server response payloads). A small percentage of Snort rules are targeted at this traffic and a small `flow_depth` value may cause false negatives in some of these rules. Most of these rules target either the HTTP header, or the content that is likely to be in the first hundred or so bytes of non-header data. Headers are usually under 300 bytes long, but your mileage may vary.

This value can be set from -1 to 1460. A value of -1 causes Snort to ignore all server side traffic for ports defined in `ports`. Inversely, a value of 0 causes Snort to inspect all HTTP server payloads defined in `ports` (note that this will likely slow down IDS performance). Values above 0 tell Snort the number of bytes to inspect in the first packet of the server response.

5. `ascii <yes|no>`

The `ascii` decode option tells us whether to decode encoded ASCII chars, a.k.a `%2f = /`, `%2e = .`, etc. It is normal to see ASCII encoding usage in URLs, so it is recommended that you disable HTTPInspect alerting for this option.

6. `utf_8 <yes|no>`

The `utf_8` decode option tells HTTPInspect to decode standard UTF-8 Unicode sequences that are in the URI. This abides by the Unicode standard and only uses % encoding. Apache uses this standard, so for any Apache servers, make sure you have this option turned on. As for alerting, you may be interested in knowing when you have a UTF-8 encoded URI, but this will be prone to false positives as legitimate web clients use this type of encoding. When `utf_8` is enabled, ASCII decoding is also enabled to enforce correct functioning.

7. `u_encode` <yes|no>

This option emulates the IIS %u encoding scheme. How the %u encoding scheme works is as follows: the encoding scheme is started by a %u followed by 4 characters, like %uxxxx. The xxxx is a hex-encoded value that correlates to an IIS Unicode codepoint. This value can most definitely be ASCII. An ASCII character is encoded like %u002f = /, %u002e = ., etc. If no `iis_unicode_map` is specified before or after this option, the default codemap is used.

You should alert on %u encodings, because we are not aware of any legitimate clients that use this encoding. So it is most likely someone trying to be covert.

8. `bare_byte` <yes|no>

Bare byte encoding is an IIS trick that uses non-ASCII characters as valid values when decoding UTF-8 values. This is not in the HTTP standard, as all non-ASCII values have to be encoded with a %. Bare byte encoding allows the user to emulate an IIS server and interpret non-standard encodings correctly.

The alert on this decoding should be enabled, because there are no legitimate clients that encode UTF-8 this way since it is non-standard.

9. `base36` <yes|no>

This is an option to decode base36 encoded chars. This option is based off of info from http://www.yk.rim.or.jp/~shikap/patch/spp_http_decode.patch.

If %u encoding is enabled, this option will not work. You have to use the base36 option with the `utf_8` option. Don't use the %u option, because base36 won't work. When base36 is enabled, ASCII encoding is also enabled to enforce correct behavior.

10. `iis_unicode` <yes|no>

The `iis_unicode` option turns on the Unicode codepoint mapping. If there is no `iis_unicode_map` option specified with the server config, `iis_unicode` uses the default codemap. The `iis_unicode` option handles the mapping of non-ASCII codepoints that the IIS server accepts and decodes normal UTF-8 request.

You should alert on the `iis_unicode` option, because it is seen mainly in attacks and evasion attempts. When `iis_unicode` is enabled, ASCII and UTF-8 decoding are also enabled to enforce correct decoding. To alert on UTF-8 decoding, you must enable also enable `utf_8 yes`.

11. `double_decode` <yes|no> The `double_decode` option is once again IIS-specific and emulates IIS functionality. How this works is that IIS does two passes through the request URI, doing decodes in each one. In the first pass, it seems that all types of iis encoding is done: utf-8 unicode, ascii, bare byte, and %u. In the second pass, the following encodings are done: ascii, bare byte, and %u. We leave out utf-8 because I think how this works is that the % encoded utf-8 is decoded to the Unicode byte in the first pass, and then UTF-8 is decoded in the second stage. Anyway, this is really complex and adds tons of different encodings for one character. When `double_decode` is enabled, so ASCII is also enabled to enforce correct decoding.

12. `non_rfc_char` {<byte> [<byte ...>]}

This option lets users receive an alert if certain non-RFC chars are used in a request URI. For instance, a user may not want to see null bytes in the request URI and we can alert on that. Please use this option with care, because you could configure it to say, alert on all '/' or something like that. It's flexible, so be careful.

13. `multi_slash` <yes|no>

This option normalizes multiple slashes in a row, so something like: "foo/////////bar" get normalized to "foo/bar."

If you want an alert when multiple slashes are seen, then configure with a `yes`, otherwise, use `no`.

14. `iis_backslash` <yes|no>

Normalizes backslashes to slashes. This is again an IIS emulation. So a request URI of "/foo\bar" gets normalized to "/foo/bar."

15. directory <yes|no>

This option normalizes directory traversals and self-referential directories.

The directory:

```
/foo/fake\_dir/../bar
```

gets normalized to:

```
/foo/bar
```

The directory:

```
/foo/./bar
```

gets normalized to:

```
/foo/bar
```

If you want to configure an alert, specify “yes”, otherwise, specify “no.” This alert may give false positives, since some web sites refer to files using directory traversals.

16. apache.whitespace <yes|no>

This option deals with the non-RFC standard of using tab for a space delimiter. Apache uses this, so if the emulated web server is Apache, enable this option. Alerts on this option may be interesting, but may also be false positive prone.

17. iis_delimiter <yes|no>

This started out being IIS-specific, but Apache takes this non-standard delimiter as well. Since this is common, we always take this as standard since the most popular web servers accept it. But you can still get an alert on this option.

18. chunk.length <non-zero positive integer>

This option is an anomaly detector for abnormally large chunk sizes. This picks up the Apache chunk encoding exploits, and may also alert on HTTP tunneling that uses chunk encoding.

19. no_pipeline_req

This option turns HTTP pipeline decoding off, and is a performance enhancement if needed. By default, pipeline requests are inspected for attacks, but when this option is enabled, pipeline requests are not decoded and analyzed per HTTP protocol field. It is only inspected with the generic pattern matching.

20. non_strict

This option turns on non-strict URI parsing for the broken way in which Apache servers will decode a URI. Only use this option on servers that will accept URIs like this: "get /index.html alsjdfk alsj lj aj la jsj s\n". The non_strict option assumes the URI is between the first and second space even if there is no valid HTTP identifier after the second space.

21. allow_proxy_use

By specifying this keyword, the user is allowing proxy use on this server. This means that no alert will be generated if the proxy_alert global keyword has been used. If the proxy_alert keyword is not enabled, then this option does nothing. The allow_proxy_use keyword is just a way to suppress unauthorized proxy use for an authorized server.

22. no_alerts

This option turns off all alerts that are generated by the HTTPInspect preprocessor module. This has no effect on HTTP rules in the ruleset. No argument is specified.

23. oversize_dir_length <non-zero positive integer>

This option takes a non-zero positive integer as an argument. The argument specifies the max char directory length for URL directory. If a url directory is larger than this argument size, an alert is generated. A good argument value is 300 characters. This should limit the alerts to IDS evasion type attacks, like whisker -i 4.

24. inspect_uri_only

This is a performance optimization. When enabled, only the URI portion of HTTP requests will be inspected for attacks. As this field usually contains 90-95% of the web attacks, you'll catch most of the attacks. So if you need extra performance, enable this optimization. It's important to note that if this option is used without any uricontent rules, then no inspection will take place. This is obvious since the URI is only inspected with uricontent rules, and if there are none available, then there is nothing to inspect.

For example, if we have the following rule set:

```
alert tcp any any -> any 80 ( msg:"content"; content: "foo"; )
```

and then we inspect the following URI:

```
get /foo.htm http/1.0\r\n\r\n
```

No alert will be generated when 'inspect_uri_only' is enabled. The 'inspect_uri_only' configuration turns off all forms of detection except uricontent inspection.

25. webroot

This option generates an alert when a directory traversal traverses past the web server root directory. This generates much fewer false positives than the directory option, because it doesn't alert on directory traversals that stay within the web server directory structure. It only alerts when the directory traversals go past the web server root directory, which is associated with certain web attacks.

26. tab_uri_delimiter

This option turns on the use of the tab character (0x09) as a delimiter for a URI. Apache accepts tab as a delimiter, IIS does not. For IIS, a tab in the URI should be treated as any other character. Whether this option is on or not, a tab is treated as whitespace if a space character (0x20) precedes it. No argument is specified.

Examples

```
preprocessor http_inspect_server: server 10.1.1.1 \  
    ports { 80 3128 8080 } \  
    flow_depth 0 \  
    ascii no \  
    double_decode yes \  
    non_rfc_char { 0x00 } \  
    chunk_length 500000 \  
    non_strict \  
    no_alerts
```

```
preprocessor http_inspect_server: server default \  
    ports { 80 3128 } \  
    non_strict \  
    non_rfc_char { 0x00 } \  
    flow_depth 300 \  
    apache_whitespace yes \  
    directory no \  
    iis_backslash no \  
    no_alerts
```

```

        u_encode yes \
        ascii no \
        chunk_length 500000 \
        bare_byte yes \
        double_decode yes \
        iis_unicode yes \
        iis_delimiter yes \
        multi_slash no

preprocessor http_inspect_server: server default \
        profile all \
        ports { 80 8080 }

```

2.1.12 ASN.1 Detection

The `asn.1` detection plugin decodes a packet or a portion of a packet, and looks for various malicious encodings.

The general configuration of the `asn.1` rule option is as follows:

```
asn1: [keyword [argument]], . . .
```

Multiple keywords can be used in an `'asn1'` option and the implied logic is boolean OR. So if any of the arguments evaluate as true, the whole option evaluates as true.

ASN.1 Keywords

The ASN.1 keywords provide programmatic detection capabilities as well as some more dynamic type detection. Most of the keywords don't have arguments as the detection is looking for non-configurable information. If a keyword does have an argument, the keyword is followed by a comma and the argument is the next word. If a keyword has multiple arguments, then a comma separates each.

1. **bitstring_overflow**

The `bitstring_overflow` option detects invalid bitstring encodings that are known to be remotely exploitable.

2. **double_overflow**

The `double_overflow` detects a double ASCII encoding that is larger than a standard buffer. This is known to be an exploitable function in Microsoft, but it is unknown at this time which services may be exploitable.

3. **oversize_length**

This detection keyword compares ASN.1 type lengths with the supplied argument. The syntax looks like, "`oversize_length 500`". This means that if an ASN.1 type is greater than 500, then this keyword is evaluated as true. This keyword must have one argument which specifies the length to compare against.

4. **absolute_offset**

This is the absolute offset from the beginning of the packet. For example, if you wanted to decode snmp packets, you would say "`absolute_offset, 0`". `absolute_offset` has one argument—the offset value. Offset may be positive or negative.

5. **relative_offset**

This is the relative offset from the last content match or `byte.test/jump`. `relative_offset` has one argument—the offset number. So if you wanted to start decoding an ASN.1 sequence right after the content "foo", you would specify `'content:"foo"; asn1: bitstring_overflow, relative_offset, 0'`. Offset values may be positive or negative.

ASN.1 Examples

The following rules use ASN.1 decoding options:

```
alert udp any any -> any 161 (msg:"Oversize SNMP Length"; \
    asn1: oversize_length, 10000, absolute_offset, 0;)

alert tcp any any -> any 80 (msg:"ASN1 Relative Foo"; content:"foo"; \
    asn1: bitstring_overflow, print, relative_offset, 0;)
```

2.1.13 X-Link2State Mini-Preprocessor

This mini-preprocessor is designed to detect the X-Link2State (www.microsoft.com/technet/security/bulletin/MS05-021.mspx) vulnerability in Microsoft Exchange Server.



NOTE

The X-Link2State functionality will be rolled up into an SMTP preprocessor in the Snort v2.5 timeframe.

Format

```
preprocessor xlink2state: ports { <port> [<port> <...>] } [drop]
```

Example Configuration

```
preprocessor xlink2state: ports { 25 691 }
```

Table 2.10: X-Link2State Options

Option	Description
ports	List of space-separated ports to inspect
drop	Drop this connection (Inline-mode only)

2.2 Event Thresholding

You can use event thresholding to reduce the number of logged alerts for noisy rules. This can be tuned to significantly reduce false alarms, and it can also be used to write a newer breed of rules. Thresholding commands limit the number of times a particular event is logged during a specified time interval. See Section 3.8 for more information.

2.3 Output Modules

Output modules are new as of version 1.6. They allow Snort to be much more flexible in the formatting and presentation of output to its users. The output modules are run when the alert or logging subsystems of Snort are called, after the preprocessors and detection engine. The format of the directives in the rules file is very similar to that of the preprocessors.

Multiple output plugins may be specified in the Snort configuration file. When multiple plugins of the same type (log, alert) are specified, they are stacked and called in sequence when an event occurs. As with the standard logging and alerting systems, output plugins send their data to /var/log/snort by default or to a user directed directory (using the -l command line switch).

Output modules are loaded at runtime by specifying the output keyword in the rules file:

```
output <name>: <options>
```

```
output alert_syslog: log_auth log_alert
```

Figure 2.7: Output Module Configuration Example

2.3.1 alert_syslog

This module sends alerts to the syslog facility (much like the -s command line switch). This module also allows the user to specify the logging facility and priority within the Snort rules file, giving users greater flexibility in logging alerts.

Available Keywords

Facilities

- log_auth
- log_authpriv
- log_daemon
- log_local0
- log_local1
- log_local2
- log_local3
- log_local4
- log_local5
- log_local6
- log_local7
- log_user

Priorities

- log_emerg
- log_alert
- log_crit
- log_err
- log_warning
- log_notice
- log_info
- log_debug

Options

- log_cons
- log_ndelay
- log_perror
- log_pid

Format

alert_syslog: <facility> <priority> <options>

NOTE

As WIN32 does not run syslog servers locally by default, a hostname and port can be passed as options. The default host is 127.0.0.1. The default port is 514.

output alert_syslog: [host=<hostname[:<port>],] <facility> <priority> <options>

output alert_syslog: 10.1.1.1:514, <facility> <priority> <options>

Figure 2.8: Syslog Configuration Example

2.3.2 alert_fast

This will print Snort alerts in a quick one-line format to a specified output file. It is a faster alerting method than full alerts because it doesn't need to print all of the packet headers to the output file

Format

alert_fast: <output filename>

output alert_fast: alert.fast

Figure 2.9: Fast Alert Configuration

2.3.3 alert_full

This will print Snort alert messages with full packet headers. The alerts will be written in the default logging directory (/var/log/snort) or in the logging directory specified at the command line.

Inside the logging directory, a directory will be created per IP. These files will be decoded packet dumps of the packets that triggered the alerts. The creation of these files slows Snort down considerably. This output method is discouraged for all but the lightest traffic situations.

Format

alert_full: <output filename>

```
output alert_full: alert.full
```

Figure 2.10: Full Alert Configuration

2.3.4 alert_unixsock

Sets up a UNIX domain socket and sends alert reports to it. External programs/processes can listen in on this socket and receive Snort alert and packet data in real time. This is currently an experimental interface.

Format

```
alert_unixsock
```

```
output alert_unixsock
```

Figure 2.11: UNIXSock Alert Configuration

2.3.5 log_tcpdump

The log_tcpdump module logs packets to a tcpdump-formatted file. This is useful for performing post-process analysis on collected traffic with the vast number of tools that are available for examining tcpdump-formatted files. This module only takes a single argument: the name of the output file. Note that the file name will have the UNIX timestamp in seconds appended the file name. This is so that data from separate Snort runs can be kept distinct.

Format

```
log_tcpdump: <output filename>
```

```
output log_tcpdump: snort.log
```

Figure 2.12: Tcpdump Output Module Configuration Example

2.3.6 database

This module from Jed Pickel sends Snort data to a variety of SQL databases. More information on installing and configuring this module can be found on the [91]incident.org web page. The arguments to this plugin are the name of the database to be logged to and a parameter list. Parameters are specified with the format parameter = argument. see Figure 2.13 for example usage.

Format

```
database: <log | alert>, <database type>, <parameter list>
```

The following parameters are available:

host - Host to connect to. If a non-zero-length string is specified, TCP/IP communication is used. Without a host name, it will connect using a local UNIX domain socket.

port - Port number to connect to at the server host, or socket filename extension for UNIX-domain connections.

dbname - Database name

user - Database username for authentication

password - Password used if the database demands password authentication

sensor_name - Specify your own name for this Snort sensor. If you do not specify a name, one will be generated automatically

encoding - Because the packet payload and option data is binary, there is no one simple and portable way to store it in a database. Blobs are not used because they are not portable across databases. So i leave the encoding option to you. You can choose from the following options. Each has its own advantages and disadvantages:

hex (default) - Represent binary data as a hex string.

Storage requirements - 2x the size of the binary

Searchability - very good

Human readability - not readable unless you are a true geek, requires post processing

base64 - Represent binary data as a base64 string.

Storage requirements - ~1.3x the size of the binary

Searchability - impossible without post processing

Human readability - not readable requires post processing

ascii - Represent binary data as an ASCII string. This is the only option where you will actually lose data. Non-ASCII Data is represented as a '.'. If you choose this option, then data for IP and TCP options will still be represented as hex because it does not make any sense for that data to be ASCII.

Storage requirements - slightly larger than the binary because some characters are escaped (&,<,>)

Searchability - very good for searching for a text string impossible if you want to search for binary

human readability - very good

detail - How much detailed data do you want to store? The options are:

full (default) - Log all details of a packet that caused an alert (including IP/TCP options and the payload)

fast - Log only a minimum amount of data. You severely limit the potential of some analysis applications if you choose this option, but this is still the best choice for some applications. The following fields are logged: **timestamp**, **signature**, **source ip**, **destination ip**, **source port**, **destination port**, **tcp flags**, and **protocol**)

Furthermore, there is a logging method and database type that must be defined. There are two logging types available, **log** and **alert**. Setting the type to log attaches the database logging functionality to the log facility within the program. If you set the type to log, the plugin will be called on the log output chain. Setting the type to alert attaches the plugin to the alert output chain within the program.

There are five database types available in the current version of the plugin. These are **mssql**, **mysql**, **postgresql**, **oracle**, and **odbc**. Set the type to match the database you are using.

NOTE

The database output plugin does not have the ability to handle alerts that are generated by using the **tag** keyword. See section 3.7.5 for more details.

```
output database: log, mysql, dbname=snort user=snort host=localhost password=xyz
```

Figure 2.13: Database Output Plugin Configuration

2.3.7 csv

The csv output plugin allows alert data to be written in a format easily importable to a database. The plugin requires 2 arguments: a full pathname to a file and the output formatting option.

The list of formatting options is below. If the formatting option is default, the output is in the order the formatting option is listed.

- timestamp
- sig_generator
- sig_id
- sig_rev
- msg
- proto
- src
- srcport
- dst
- dstport
- ethsrc
- ethdst
- ethlen
- tcpflags
- tcpseq
- tcpack
- tcplen
- tcpwindow
- ttl
- tos
- id
- dgmlen
- iplen
- icmptype
- icmpcode
- icmpid
- icmpseq

Format

```
output alert_csv: <filename> <format>
```

```
output alert_csv: /var/log/alert.csv default
```

```
output alert_csv: /var/log/alert.csv timestamp, msg
```

Figure 2.14: CSV Output Configuration

2.3.8 unified

The unified output plugin is designed to be the fastest possible method of logging Snort events. The unified output plugin logs events in binary format, allowing another programs to handle complex logging mechanisms that would otherwise diminish the performance of Snort.

The name *unified* is a misnomer, as the unified output plugin creates two different files, an *alert* file, and a *log* file. The alert file contains the high-level details of an event (eg: IPs, protocol, port, message id). The log file contains the detailed packet information (a packet dump with the associated event ID). Both file types are written in a binary format described in *spo_unified.h*.

NOTE

Files have the file creation time (in Unix Epoch format) appended to each file when it is created.

Format

```
output alert_unified: <base file name> [, <limit <file size limit in MB>]
output log_unified: <base file name> [, <limit <file size limit in MB>]
```

```
output alert_unified: snort.alert, limit 128
output log_unified: snort.log, limit 128
```

Figure 2.15: Unified Configuration Example

2.3.9 alert_prelude

NOTE

support to use alert_prelude is not built in by default. To use alert_prelude, snort must be built with the `-enable-prelude` argument passed to `./configure`.

The alert_prelude output plugin is used to log to a Prelude database. For more information on Prelude, see <http://www.prelude-ids.org/>.

format

```
output alert_prelude: profile <name of prelude profile>
    [, info <priority number for info priority alerts>]
    [, low <priority number for low priority alerts>]
    [, medium <priority number for medium priority alerts>]
    [, high <priority number for high priority alerts>]
```

```
output alert_prelude: profile snort, info 4, low 3, medium 2, high 1
```

Figure 2.16: alert_prelude configuration example

2.3.10 log null

Sometimes it is useful to be able to create rules that will alert to certain types of traffic but will not cause packet log entries. In Snort 1.8.2, the log_null plugin was introduced. This is equivalent to using the -n command line option but it is able to work within a ruletype.

Format

```
output log_null
```

```
output log_null # like using snort -n
```

```
ruletype info {
    type alert
    output alert_fast: info.alert
    output log_null
}
```

Figure 2.17: Log Null Usage Example

Chapter 3

Writing Snort Rules

How to Write Snort Rules and Keep Your Sanity

3.1 The Basics

Snort uses a simple, lightweight rules description language that is flexible and quite powerful. There are a number of simple guidelines to remember when developing Snort rules.

Most Snort rules are written in a single line. This was required in versions prior to 1.8. In current versions of Snort, rules may span multiple lines by adding a backslash \ to the end of the line.

Snort rules are divided into two logical sections, the rule header and the rule options. The rule header contains the rule's action, protocol, source and destination IP addresses and netmasks, and the source and destination ports information. The rule option section contains alert messages and information on which parts of the packet should be inspected to determine if the rule action should be taken.

Figure 3.1 illustrates a sample Snort rule.

```
alert tcp any any -> 192.168.1.0/24 111 (content:"|00 01 86 a5|"; msg:"mountd access");
```

Figure 3.1: Sample Snort Rule

The text up to the first parenthesis is the rule header and the section enclosed in parenthesis contains the rule options. The words before the colons in the rule options section are called option *keywords*.

NOTE

Note that the rule options section is not specifically required by any rule, they are just used for the sake of making tighter definitions of packets to collect or alert on (or drop, for that matter).

All of the elements in that make up a rule must be true for the indicated rule action to be taken. When taken together, the elements can be considered to form a logical AND statement. At the same time, the various rules in a Snort rules library file can be considered to form a large logical OR statement.

3.2 Rules Headers

3.2.1 Rule Actions

The rule header contains the information that defines the who, where, and what of a packet, as well as what to do in the event that a packet with all the attributes indicated in the rule should show up. The first item in a rule is the rule action. The rule action tells Snort what to do when it finds a packet that matches the rule criteria. There are 5 available default actions in Snort, alert, log, pass, activate, and dynamic. In addition, if you are running Snort in inline mode, you have additional options which include drop, reject, and sdrops.

1. alert - generate an alert using the selected alert method, and then log the packet
2. log - log the packet
3. pass - ignore the packet
4. activate - alert and then turn on another dynamic rule
5. dynamic - remain idle until activated by an activate rule , then act as a log rule
6. drop - make iptables drop the packet and log the packet
7. reject - make iptables drop the packet, log it, and then send a TCP reset if the protocol is TCP or an ICMP port unreachable message if the protocol is UDP.
8. sdrops - make iptables drop the packet but does not log it.

You can also define your own rule types and associate one or more output plugins with them. You can then use the rule types as actions in Snort rules.

This example will create a type that will log to just tcpdump:

```
ruletype suspicious
{
    type log
    output log_tcpdump: suspicious.log
}
```

This example will create a rule type that will log to syslog and a MySQL database:

```
ruletype redalert
{
    type alert
    output alert_syslog: LOG_AUTH LOG_ALERT
    output database: log, mysql, user=snort dbname=snort host=localhost
}
```

3.2.2 Protocols

The next field in a rule is the protocol. There are four protocols that Snort currently analyzes for suspicious behavior – tcp, udp, icmp, and ip. In the future there may be more, such as ARP, IGRP, GRE, OSPF, RIP, IPX, etc.

3.2.3 IP Addresses

The next portion of the rule header deals with the IP address and port information for a given rule. The keyword `any` may be used to define any address. Snort does not have a mechanism to provide host name lookup for the IP address fields in the rules file. The addresses are formed by a straight numeric IP address and a CIDR[3] block. The CIDR block indicates the netmask that should be applied to the rule's address and any incoming packets that are tested against the rule. A CIDR block mask of `/24` indicates a Class C network, `/16` a Class B network, and `/32` indicates a specific machine address. For example, the address/CIDR combination `192.168.1.0/24` would signify the block of addresses from `192.168.1.1` to `192.168.1.255`. Any rule that used this designation for, say, the destination address would match on any address in that range. The CIDR designations give us a nice short-hand way to designate large address spaces with just a few characters.

In Figure 3.1, the source IP address was set to match for any computer talking, and the destination address was set to match on the `192.168.1.0` Class C network.

There is an operator that can be applied to IP addresses, the negation operator. This operator tells Snort to match any IP address except the one indicated by the listed IP address. The negation operator is indicated with a `!`. For example, an easy modification to the initial example is to make it alert on any traffic that originates outside of the local net with the negation operator as shown in Figure 3.2.

```
alert tcp !192.168.1.0/24 any -> 192.168.1.0/24 111 \
  (content: "|00 01 86 a5|"; msg: "external mountd access");)
```

Figure 3.2: Example IP Address Negation Rule

This rule's IP addresses indicate any tcp packet with a source IP address not originating from the internal network and a destination address on the internal network.

You may also specify lists of IP addresses. An IP list is specified by enclosing a comma separated list of IP addresses and CIDR blocks within square brackets. For the time being, the IP list may not include spaces between the addresses. See Figure 3.3 for an example of an IP list in action.

```
alert tcp ![192.168.1.0/24,10.1.1.0/24] any -> \
  [192.168.1.0/24,10.1.1.0/24] 111 (content: "|00 01 86 a5|"; \
  msg: "external mountd access");)
```

Figure 3.3: IP Address Lists

3.2.4 Port Numbers

Port numbers may be specified in a number of ways, including any ports, static port definitions, ranges, and by negation. Any ports are a wildcard value, meaning literally any port. Static ports are indicated by a single port number, such as `111` for portmapper, `23` for telnet, or `80` for http, etc. Port ranges are indicated with the range operator `..`. The range operator may be applied in a number of ways to take on different meanings, such as in Figure 3.4.

Port negation is indicated by using the negation operator `!`. The negation operator may be applied against any of the other rule types (except `any`, which would translate to none, how Zen...). For example, if for some twisted reason you wanted to log everything except the X Windows ports, you could do something like the rule in Figure 3.5.

```
log udp any any -> 192.168.1.0/24 1:1024 log udp
traffic coming from any port and destination ports ranging from 1 to 1024

log tcp any any -> 192.168.1.0/24 :6000

log tcp traffic from any port going to ports less than or equal to 6000

log tcp any :1024 -> 192.168.1.0/24 500:

log tcp traffic from privileged ports less than or equal to 1024 going to ports greater than or equal to 500
```

Figure 3.4: Port Range Examples

```
log tcp any any -> 192.168.1.0/24 !6000:6010
```

Figure 3.5: Example of Port Negation

3.2.5 The Direction Operator

The direction operator `->` indicates the orientation, or direction, of the traffic that the rule applies to. The IP address and port numbers on the left side of the direction operator is considered to be the traffic coming from the source host, and the address and port information on the right side of the operator is the destination host. There is also a bidirectional operator, which is indicated with a `<>` symbol. This tells Snort to consider the address/port pairs in either the source or destination orientation. This is handy for recording/analyzing both sides of a conversation, such as telnet or POP3 sessions. An example of the bidirectional operator being used to record both sides of a telnet session is shown in Figure 3.6.

Also, note that there is no `<-` operator. In Snort versions before 1.8.7, the direction operator did not have proper error checking and many people used an invalid token. The reason the `<-` does not exist is so that rules always read consistently.

```
log tcp !192.168.1.0/24 any <> 192.168.1.0/24 23
```

Figure 3.6: Snort rules using the Bidirectional Operator

3.2.6 Activate/Dynamic Rules

NOTE

Activate and Dynamic rules are being phased out in favor of a combination of tagging (3.7.5) and flowbits (3.6.10).

Activate/dynamic rule pairs give Snort a powerful capability. You can now have one rule activate another when it's action is performed for a set number of packets. This is very useful if you want to set Snort up to perform follow on recording when a specific rule goes off. Activate rules act just like alert rules, except they have a `*required*` option field: `activates`. Dynamic rules act just like log rules, but they have a different option field: `activated_by`. Dynamic rules have a second required field as well, `count`.

Activate rules are just like alerts but also tell Snort to add a rule when a specific network event occurs. Dynamic rules are just like log rules except are dynamically enabled when the activate rule id goes off.

Put 'em together and they look like Figure 3.7.

```
activate tcp !$HOME_NET any -> $HOME_NET 143 (flags: PA; \
  content: "|E8COFFFFFF/bin"; activates: 1; \
  msg: "IMAP buffer overflow!");
dynamic tcp !$HOME_NET any -> $HOME_NET 143 (activated_by: 1; count: 50;)
```

Figure 3.7: Activate/Dynamic Rule Example

These rules tell Snort to alert when it detects an IMAP buffer overflow and collect the next 50 packets headed for port 143 coming from outside `$HOME_NET` headed to `$HOME_NET`. If the buffer overflow happened and was successful, there's a very good possibility that useful data will be contained within the next 50 (or whatever) packets going to that same service port on the network, so there's value in collecting those packets for later analysis.

3.3 Rule Options

Rule options form the heart of Snort's intrusion detection engine, combining ease of use with power and flexibility. All Snort rule options are separated from each other using the semicolon (;) character. Rule option keywords are separated from their arguments with a colon (:) character.

There are four major categories of rule options.

meta-data These options provide information about the rule but do not have any affect during detection

payload These options all look for data inside the packet payload and can be inter-related

non-payload These options look for non-payload data

post-detection These options are rule specific triggers that happen after a rule has "fired."

3.4 Meta-Data Rule Options

3.4.1 msg

The msg rule option tells the logging and alerting engine the message to print along with a packet dump or to an alert. It is a simple text string that utilizes the \ as an escape character to indicate a discrete character that might otherwise confuse Snort's rules parser (such as the semi-colon ; character).

Format

msg: "<message text>";

3.4.2 reference

The reference keyword allows rules to include references to external attack identification systems. The plugin currently supports several specific systems as well as unique URLs. This plugin is to be used by output plugins to provide a link to additional information about the alert produced.

Make sure to also take a look at <http://www.snort.org/pub-bin/sigs-search.cgi/> for a system that is indexing descriptions of alerts based off of the sid (See Section 3.4.3).

Table 3.1: Supported Systems

System	URL Prefix
bugtraq	http://www.securityfocus.com/bid/
cve	http://cve.mitre.org/cgi-bin/cvename.cgi?name=
nessus	http://cgi.nessus.org/plugins/dump.php3?id=
arachnids	(currently down) http://www.whitehats.com/info/IDS
mcafee	http://vil.nai.com/vil/dispVirus.asp?virus_k=
url	http://

Format

reference: <id system>,<id>; [reference: <id system>,<id>;]

```

alert tcp any any -> any 7070 (msg:"IDS411/dos-realaudio"; \
    flags:AP; content:"|fff4 fffd 06|"; reference:arachnids,IDS411;)

alert tcp any any -> any 21 (msg:"IDS287/ftp-wuftp260-venglin-linux"; \
    flags:AP; content:"|31c031db 31c9b046 cd80 31c031db|"; \
    reference:arachnids,IDS287; reference:bugtraq,1387; \
    reference:cve,CAN-2000-1574;)

```

Figure 3.8: Reference Usage Examples

3.4.3 sid

The sid keyword is used to uniquely identify Snort rules. This information allows output plugins to identify rules easily. This option should be used with the rev keyword. (See section 3.4.4)

- <100 Reserved for future use
- 100-1,000,000 Rules included with the Snort distribution
- >1,000,000 Used for local rules

The file sid-msg.map contains a mapping of alert messages to Snort rule IDs. This information is useful when post-processing alert to map an ID to an alert message.

Format

```
sid: <snort rules id>;
```

Example

This example is a rule with the Snort Rule ID of 1000983.

```
alert tcp any any -> any 80 (content:"BOB"; sid:1000983; rev:1;)
```

3.4.4 rev

The sid keyword is used to uniquely identify revisions of Snort rules. Revisions, along with Snort rule id's, allow signatures and descriptions to be refined and replaced with updated information. This option should be used with the sid keyword. (See section 3.4.3)

Format

```
rev: <revision integer>
```

Example

This example is a rule with the Snort Rule Revision of 1.

```
alert tcp any any -> any 80 (content:"BOB"; sid:1000983; rev:1;)
```

3.4.5 classtype

The classtype keyword categorizes alerts to be attack classes. By using the and prioritized. The user can specify what priority each type of rule classification has. Rules that have a classification will have a default priority set.

Format

```
classtype: <class name>;
```

Rule classifications are defined in the `classification.config` file. The config file uses the following syntax:

```
config classification: <class name>,<class description>,<default priority>
```

The standard classifications included with Snort are listed in Table 3.2. The standard classifications are ordered with 3 default priorities currently. A priority 1 is the most severe priority level of the default rule set and 4 is the least severe.

Table 3.2: Snort Default Classifications

Classtype	Description	Priority
attempted-admin	Attempted Administrator Privilege Gain	high
attempted-user	Attempted User Privilege Gain	high
shellcode-detect	Executable code was detected	high
successful-admin	Successful Administrator Privilege Gain	high
successful-user	Successful User Privilege Gain	high
trojan-activity	A Network Trojan was detected	high
unsuccessful-user	Unsuccessful User Privilege Gain	high
web-application-attack	Web Application Attack	high
attempted-dos	Attempted Denial of Service	medium
attempted-recon	Attempted Information Leak	medium
bad-unknown	Potentially Bad Traffic	medium
denial-of-service	Detection of a Denial of Service Attack	medium
misc-attack	Misc Attack	medium
non-standard-protocol	Detection of a non-standard protocol or event	medium
rpc-portmap-decode	Decode of an RPC Query	medium
successful-dos	Denial of Service	medium
successful-recon-largescale	Large Scale Information Leak	medium
successful-recon-limited	Information Leak	medium
suspicious-filename-detect	A suspicious filename was detected	medium
suspicious-login	An attempted login using a suspicious username was detected	medium
system-call-detect	A system call was detected	medium
unusual-client-port-connection	A client was using an unusual port	medium
web-application-activity	access to a potentially vulnerable web application	medium
icmp-event	Generic ICMP event	low
misc-activity	Misc activity	low
network-scan	Detection of a Network Scan	low
not-suspicious	Not Suspicious Traffic	low
protocol-command-decode	Generic Protocol Command Decode	low
string-detect	A suspicious string was detected	low
unknown	Unknown Traffic	low

```

alert tcp any any -> any 80 (msg:"EXPLOIT ntpdx overflow"; \
    dsize: >128; classtype:attempted-admin; priority:10 );

alert tcp any any -> any 25 (msg:"SMTP expn root"; flags:A+; \
    content:"expn root"; nocase; classtype:attempted-recon;)

```

Figure 3.9: Example Classtype Rules

Warnings

classtype uses classifications defined by the classification config option. The classifications used by the rules provided with Snort are defined in `etc/classification.config`

3.4.6 Priority

The priority tag assigns a severity level to rules. A classtype rule assigns a default priority that may be overridden with a priority rule. For an example in conjunction with a classification rule refer to Figure 3.9. For use by itself, see Figure 3.10

Format

```
priority: <priority integer>;
```

```

alert TCP any any -> any 80 (msg: "WEB-MISC phf attempt"; flags:A+; \
    content: "/cgi-bin/phf"; priority:10;)

```

Figure 3.10: Example Priority Rule

3.5 Payload Detection Rule Options

3.5.1 content

The content keyword is one of the more important features of Snort. It allows the user to set rules that search for specific content in the packet payload and trigger response based on that data. Whenever a content option pattern match is performed, the Boyer-Moore pattern match function is called and the (rather computationally expensive) test is performed against the packet contents. If data exactly matching the argument data string is contained anywhere within the packet's payload, the test is successful and the remainder of the rule option tests are performed. Be aware that this test is case sensitive.

The option data for the content keyword is somewhat complex; it can contain mixed text and binary data. The binary data is generally enclosed within the pipe (|) character and represented as bytecode. Bytecode represents binary data as hexadecimal numbers and is a good shorthand method for describing complex binary data. Figure 3.11 contains an example of mixed text and binary data in a Snort rule.

Note that multiple content rules can be specified in one rule. This allows rules to be tailored for less false positives.

If the rule is preceded by a `!`, the alert will be triggered on packets that do not contain this content. This is useful when writing rules that want to alert on packets that do not match a certain pattern

⚠ NOTE

Also note that the following characters must be escaped inside a content rule:
: ; \ "

Format

```
content: [!] "<content string>";
```

Example

```
alert tcp any any -> any 139 (content:"|5c 00|P|00|I|00|P|00|E|00 5c|");
```

Figure 3.11: Mixed Binary Bytecode and Text in a 'content' keyword

```
alert tcp any any -> any 80 (content:! "GET");
```

Figure 3.12: Negation Example

Changing content behavior

The **content** keyword has a number of modifier keywords. The modifier keywords change how the previously specified content works. These modifier keywords are:

1. depth
2. offset
3. distance
4. within
5. nocase
6. rawbytes

3.5.2 nocase

The **nocase** keyword allows the rule writer to specify that the Snort should look for the specific pattern, ignoring case. **nocase** modifies the previous 'content' keyword in the rule.

Format

```
nocase;
```

Example

```
alert tcp any any -> any 21 (msg:"FTP ROOT"; content:"USER root"; nocase;)
```

Figure 3.13: Content rule with nocase modifier

3.5.3 rawbytes

The rawbytes keyword allows rules to look at the raw packet data, ignoring any decoding that was done by preprocessors. This acts as a modifier to the previous content 3.5.1 option.

format

```
rawbytes;
```

Example

This example tells the content pattern matcher to look at the raw traffic, instead of the decoded traffic provided by the telnet decoder.

```
alert tcp any any -> any 21 (msg: "Telnet NOP"; content: "|FF F1|"; rawbytes;)
```

3.5.4 depth

The depth keyword allows the rule writer to specify how far into a packet Snort should search for the specified pattern. depth modifies the previous ‘content’ keyword in the rule.

A depth of 5 would tell Snort to only look for the specified pattern within the first 5 bytes of the payload.

As the depth keyword is a modifier to the previous ‘content’ keyword, there must be a content in the rule before ‘depth’ is specified.

See Figure 3.14 for an example of a combined content, offset, and depth search rule.

Format

```
depth: <number>;
```

3.5.5 offset

The offset keyword allows the rule writer to specify where to start searching for a pattern within a packet. offset modifies the previous ‘content’ keyword in the rule.

An offset of 5 would tell Snort to start looking for the specified pattern after the first 5 bytes of the payload.

As this keyword is a modifier to the previous ‘content’ keyword, there must be a content in the rule before ‘offset’ is specified.

See Figure 3.14 for an example of a combined content, offset, and depth search rule.

Format

```
offset: <number>;
```

```
alert tcp any any -> any 80 (content: "cgi-bin/phf"; offset:4; depth:20;)
```

Figure 3.14: Combined Content, Offset and Depth Rule. Skip the first 4 bytes, and look for cgi-bin/phf in the next 20 bytes

3.5.6 distance

The distance keyword allows the rule writer to specify how far into a packet Snort should ignore before starting to search for the specified pattern relative to the end of the previous pattern match.

This can be thought of as exactly the same thing as depth (See Section 3.5.5), except it is relative to the end of the last pattern match instead of the beginning of the packet.

Format

```
distance: <byte count>;
```

Example

The rule listed in Figure 3.15 maps to a regular expression of `/ABCDE.{1}EFGH/`.

```
alert tcp any any -> any any (content:"ABC"; content: "DEF"; distance:1;)
```

Figure 3.15: distance usage example

3.5.7 within

The within keyword is a content modifier that makes sure that at most N bytes are between pattern matches using the Content (See Section 3.5.1). It's designed to be used in conjunction with the distance (Section 3.5.6) rule option.

The rule listed in Figure 3.16 constrains the search to not go past 10 bytes past the ABCDE match.

Format

```
within: <byte count>;
```

Examples

```
alert tcp any any -> any any (content:"ABC"; content: "EFG"; within:10;)
```

Figure 3.16: within usage example

3.5.8 uricontent

The uricontent parameter in the Snort rule language searches the NORMALIZED request URI field. This means that if you are writing rules that include things that are normalized, such as `%2f` or directory traversals, these rules will not alert. The reason is that the things you are looking for are normalized out of the URI buffer.

For example, the URI:

```
/scripts/..%c0%af../winnt/system32/cmd.exe?/c+ver
```

will get normalized into:


```
/winnt/system32/cmd.exe?/c+ver
```

Another example, the URI:

```
\begin{verbatim} /cgi-bin/aaaaaaaaaaaaaaaaaaaaaa/..%252fp%68f?
```

will get normalized into:

```
/cgi-bin/phf?
```

When writing a uricontent rule, write the content that you want to find in the context that the URI will be normalized. For example, if Snort normalizes directory traversals, do not include directory traversals.

You can write rules that look for the non-normalized content by using the content option. (See Section 3.5.1)

For a description of the parameters to this function, see the content rule options in Section 3.5.1.

This option works in conjunction with the HTTP Inspect preprocessor specified in Section 2.1.11.

Format

```
uricontent:[!]<content string>;
```

3.5.9 isdataat

Verify that the payload has data at a specified location, optionally looking for data relative to the end of the previous content match.

Format

```
isdataat:<int>[,relative];
```

Example

```
alert tcp any any -> any 111 (content:"PASS"; isdataat:50,relative; \
  content:"!|0a|"; distance:0;)
```

This rule looks for the string PASS exists in the packet, then verifies there is at least 50 bytes after the end of the string PASS, then verifies that there is not a newline character within 50 bytes of the end of the PASS string.

3.5.10 pcre

The pcre keyword allows rules to be written using perl compatible regular expressions. For more detail on what can be done via a pcre regular expression, check out the PCRE web site <http://www.pcre.org>

Format

```
pcre:[!]"(</regex>/|m<delim><regex><delim>)[ismxAEGRUB]";
```

The post-re modifiers set compile time flags for the regular expression.

Table 3.3: Perl compatible modifiers

i	case insensitive
---	------------------

s	include newlines in the dot metacharacter
m	By default, the string is treated as one big line of characters. <code>^</code> and <code>\$</code> match at the beginning and ending of the string. When <code>m</code> is set, <code>^</code> and <code>\$</code> match immediately following or immediately before any newline in the buffer, as well as the very start and very end of the buffer.
x	whitespace data characters in the pattern are ignored except when escaped or inside a character class

Table 3.4: PCRE compatible modifiers

A	the pattern must match only at the start of the buffer (same as <code>^</code>)
E	Set <code>\$</code> to match only at the end of the subject string. Without <code>E</code> , <code>\$</code> also matches immediately before the final character if it is a newline (but not before any other newlines).
G	Inverts the "greediness" of the quantifiers so that they are not greedy by default, but become greedy if followed by <code>"?"</code> .

Table 3.5: Snort specific modifiers

R	Match relative to the end of the last pattern match. (Similar to <code>distance:0</code> ;))
U	Match the decoded URI buffers (Similar to <code>uricontent</code>)
B	Do not use the decoded buffers (Similar to <code>rawbytes</code>)

The modifiers `R` and `B` should not be used together.

Example

This example performs a case-insensitive search for the string `BLAH` in the payload.

```
alert ip any any -> any any (pcre:"/BLAH/i");)
```

3.5.11 byte_test

Test a byte field against a specific value (with operator). Capable of testing binary values or converting representative byte strings to their binary equivalent and testing them.

For a more detailed explanation, please read Section 3.11.5.

Format

```
byte_test: <bytes to convert>, [!]<operator>, <value>, <offset> \
    [,relative] [,<endian>] [,<number type>, string];
```

Option	Description
<code>bytes_to_convert</code>	Number of bytes to pick up from the packet
<code>operator</code>	Operation to perform to test the value: <ul style="list-style-type: none"> • <code><</code> - less than • <code>></code> - greater than • <code>=</code> - equal • <code>!</code> - not • <code>&</code> - bitwise AND • <code>^</code> bitwise OR
<code>value</code>	Value to test the converted value against
<code>offset</code>	Number of bytes into the payload to start processing
<code>relative</code>	Use an offset relative to last pattern match
<code>endian</code>	Endian type of the number being read: <ul style="list-style-type: none"> • <code>big</code> - Process data as big endian (default) • <code>little</code> - Process data as little endian
<code>string</code>	Data is stored in string format in packet
<code>number type</code>	Type of number being read: <ul style="list-style-type: none"> • <code>hex</code> - Converted string data is represented in hexadecimal • <code>dec</code> - Converted string data is represented in decimal • <code>oct</code> - Converted string data is represented in octal

Any of the operators can also include `!` to check if the operator is not true. If `!` is specified without an operator, then the operator is set to `=`.

NOTE

Snort uses the C operators for each of these operators. If the `&` operator is used, then it would be the same as using `if (data & value) { do-something(); }`

3.5.12 byte_jump

The `byte_jump` option allows rules to be written for length encoded protocols trivially. By having an option that reads the length of a portion of data, then skips that far forward in the packet, rules can be written that skip over specific portions of length-encoded protocols and perform detection in very specific locations.

The `byte_jump` option does this by reading some number of bytes, convert them to their numeric representation, move that many bytes forward and set a pointer for later detection. This pointer is known as the detect offset end pointer, or `doe_ptr`.

For a more detailed explanation, please read Section 3.11.5.

Format

```
byte_jump: <bytes_to_convert>, <offset> \
    [,relative] [,multiplier <multiplier value>] [,big] [,little][,string]\
```

```

alert udp $EXTERNAL_NET any -> $HOME_NET any \
(msg:"AMD procedure 7 plog overflow "; \
content: "|00 04 93 F3|"; \
content: "|00 00 00 07|"; distance: 4; within: 4; \
byte_test: 4,>, 1000, 20, relative;)

alert tcp $EXTERNAL_NET any -> $HOME_NET any \
(msg:"AMD procedure 7 plog overflow "; \
content: "|00 04 93 F3|"; \
content: "|00 00 00 07|"; distance: 4; within: 4; \
byte_test: 4, >,1000, 20, relative;)

alert udp any any -> any 1234 \
(byte_test: 4, =, 1234, 0, string, dec; \
msg: "got 1234!");)

alert udp any any -> any 1235 \
(byte_test: 3, =, 123, 0, string, dec; \
msg: "got 123!");)

alert udp any any -> any 1236 \
(byte_test: 2, =, 12, 0, string, dec; \
msg: "got 12!");)

alert udp any any -> any 1237 \
(byte_test: 10, =, 1234567890, 0, string, dec; \
msg: "got 1234567890!");)

alert udp any any -> any 1238 \
(byte_test: 8, =, 0xdeadbeef, 0, string, hex; \
msg: "got DEADBEEF!");)

```

Figure 3.17: Byte Test Usage Example

```
[,hex] [,dec] [,oct] [,align] [,from_beginning];
```

Option	Description
bytes_to_convert	Number of bytes to pick up from the packet
offset	Number of bytes into the payload to start processing
relative	Use an offset relative to last pattern match
multiplier <value>	Multiply the number of calculated bytes by <value> and skip forward that number of bytes.
big	Process data as big endian (default)
little	Process data as little endian
string	Data is stored in string format in packet
hex	Converted string data is represented in hexadecimal
dec	Converted string data is represented in decimal
oct	Converted string data is represented in octal
align	Round the number of converted bytes up to the next 32-bit boundary
from_beginning	Skip forward from the beginning of the packet payload instead of from the current position in the packet.

```
alert udp any any -> any 32770:34000 (content: "|00 01 86 B8|"; \
    content: "|00 00 00 01|"; distance: 4; within: 4; \
    byte_jump: 4, 12, relative, align; \
    byte_test: 4, >, 900, 20, relative; \
    msg: "statd format string buffer overflow";)
```

Figure 3.18: byte jump Usage Example

3.5.13 ftpbounce

The ftpbounce keyword detects FTP bounce attacks.

Format

```
ftpbounce;
```

Example

```
alert tcp $EXTERNAL_NET any -> $HOME_NET 21 (msg:"FTP PORT bounce attempt"; \
    flow:to_server,established; content:"PORT"; nocase; ftpbounce; pcre:"/^PORT/smi"; \
    classtype:misc-attack; sid:3441; rev:1;)
```

3.5.14 regex

The regex keyword has been superseded by PCRE. See Section 3.5.10.

3.5.15 content-list

The content-list keyword is broken and should not be used.

3.6 Non-payload Detection Rule Options

3.6.1 fragoffset

The fragoffset keyword allows one to compare the IP fragment offset field against a decimal value. To catch all the first fragments of an IP session, you could use the fragbits keyword and look for the More fragments option in conjunction with a fragoffset of 0.

Format

```
fragoffset:[<|>]<number>
```

```
alert ip any any -> any any \
    (msg: "First Fragment"; fragbits: M; fragoffset: 0;)
```

Figure 3.19: Fragoffset Usage Example

3.6.2 ttl

The ttl keyword is used to check the IP time-to-live value. This option keyword was intended for use in the detection of traceroute attempts.

Format

```
ttl:[[<number>-]><=]<number>;
```

Example

This example checks for a time-to-live value that is less than 3.

```
ttl:<3;
```

This example checks for a time-to-live value that between 3 and 5.

```
ttl:3-5;
```

3.6.3 tos

The tos keyword is used to check the IP TOS field for a specific value.

Format

```
tos:[!]<number>;
```

Example

This example looks for a tos value that is not 4

```
tos:!4;
```

3.6.4 id

The id keyword is used to check the IP ID field for a specific value. Some tools (exploits, scanners and other odd programs) set this field specifically for various purposes, for example, the value 31337 is very popular with some hackers.

Format

```
id:<number>;
```

Example

This example looks for the IP ID of 31337.

```
id:31337;
```

3.6.5 ipopts

The ipopts keyword is used to check if a specific IP option is present.

The following options may be checked:

rr - Record route

eol - End of list

nop - No op

ts - Time Stamp

sec - IP security option

lsrr - Loose source routing

ssrr - Strict source routing

satid - Stream identifier

any - any IP options are set

The most frequently watched for IP options are strict and loose source routing which aren't used in any widespread internet applications.

Format

```
ipopts:<rr|eol|nop|ts|sec|lsrr|ssrr|satid|any>;
```

Example

This example looks for the IP Option of Loose Source Routing.

```
ipopts:lsrr;
```

Warning

Only a single ipopts keyword may be specified per rule.

3.6.6 fragbits

The **fragbits** keyword is used to check if fragmentation and reserved bits are set in the IP header.

The following bits may be checked:

M - More Fragments

D - Don't Fragment

R - Reserved Bit

The following modifiers can be set to change the match criteria:

+ match on the specified bits, plus any others

***** match if any of the specified bits are set

! match if the specified bits are not set

Format

fragbits: [**+*********!**] <[MDR]>

Example

This example checks if the More Fragments bit and the Do not Fragment bit are set.

```
fragbits:MD+;
```

3.6.7 dsize

The **dsize** keyword is used to test the packet payload size. This may be used to check for abnormally sized packets. In many cases, it is useful for detecting buffer overflows.

Format

dsize: [<>] <number> [<> <number>];

Example

This example looks for a dsize that is between 300 and 400 bytes.

```
dsize:300<>400;
```

Warning

dsize will fail on stream rebuilt packets, regardless of the size of the payload.

3.6.8 flags

The flags keyword is used to check if specific TCP flag bits are present.

The following bits may be checked:

F - FIN (LSB in TCP Flags byte)

S - SYN

R - RST

P - PSH

A - ACK

U - URG

1 - Reserved bit 1 (MSB in TCP Flags byte)

2 - Reserved bit 2

0 - No TCP Flags Set

The following modifiers can be set to change the match criteria:

+ - match on the specified bits, plus any others

***** - match if any of the specified bits are set

! - match if the specified bits are not set

To handle writing rules for session initiation packets such as ECN where a SYN packet is sent with the previously reserved bits 1 and 2 set, an option mask may be specified. A rule could check for a flags value of S,12 if one wishes to find packets with just the syn bit, regardless of the values of the reserved bits.

Format

```
flags: [!|*|+]<FSRPAU120>[,<FSRPAU120>];
```

Example

This example checks if just the SYN and the FIN bits are set, ignoring reserved bit 1 and reserved bit 2.

```
alert tcp any any -> any any (flags:SF,12;)
```

3.6.9 flow

The flow rule option is used in conjunction with TCP stream reassembly (see Section 2.1.3). It allows rules to only apply to certain directions of the traffic flow.

This allows rules to only apply to clients or servers. This allows packets related to \$HOME_NET clients viewing web pages to be distinguished from servers running the \$HOME_NET.

The established keyword will replace the **flags: A+** used in many places to show established TCP connections.

Options

Option	Description
to_client	Trigger on server responses from A to B
to_server	Trigger on client requests from A to B
from_client	Trigger on client requests from A to B
from_server	Trigger on server responses from A to B
established	Trigger only on established TCP connections
stateless	Trigger regardless of the state of the stream processor (useful for packets that are designed to cause machines to crash)
no_stream	Do not trigger on rebuilt stream packets (useful for dsize and stream4)
only_stream	Only trigger on rebuilt stream packets

Format

```
flow: [(established|stateless)]  
      [, (to_client|to_server|from_client|from_server)]  
      [, (no_stream|only_stream)]
```

```
alert tcp !$HOME_NET any -> $HOME_NET 21 (msg:"cd incoming detected"; \  
      flow:from_client; content:"CWD incoming"; nocase;)
```

```
alert tcp !$HOME_NET 0 -> $HOME_NET 0 (msg: "Port 0 TCP traffic"; \  
      flow:stateless;)
```

Figure 3.20: Flow usage examples

3.6.10 flowbits

The `flowbits` rule option is used in conjunction with conversation tracking from the Flow preprocessor (see Section 2.1.4). It allows rules to track states across transport protocol sessions. The `flowbits` option is most useful for TCP sessions, as it allows rules to generically track the state of an application protocol.

There are seven keywords associated with `flowbits`. Most of the options need a user-defined name for the specific state that is being checked. This string should be limited to any alphanumeric string including periods, dashes, and underscores.

Option	Description
set	Sets the specified state for the current flow.
unset	Unsets the specified state for the current flow.
toggle	Sets the specified state if the state is unset, otherwise unsets the state if the state is set.
isset	Checks if the specified state is set.
isnotset	Checks if the specified state is not set.
noalert	Cause the rule to not generate an alert, regardless of the rest of the detection options.

Format

```
flowbits: [set|unset|toggle|isset,reset,noalert] [, <STATE_NAME>];
```

```

alert tcp any 143 -> any any (msg:"IMAP login";
  content:"OK LOGIN"; flowbits:set,logged_in;
  flowbits:noalert;)

alert tcp any any -> any 143 (msg:"IMAP LIST"; content:"LIST";
  flowbits:isset,logged_in;)

```

Figure 3.21: Flowbits Usage Examples

3.6.11 seq

The seq keyword is used to check for a specific TCP sequence number.

Format

```
seq:<number>;
```

Example

This example looks for a TCP sequence number of 0.

```
seq:0;
```

3.6.12 ack

The ack keyword is used to check for a specific TCP acknowledge number.

Format

```
ack: <number>;
```

Example

This example looks for a TCP acknowledge number of 0.

```
ack:0;
```

3.6.13 window

The ack keyword is used to check for a specific TCP window size.

Format

```
window:[!]<number>;
```

Example

This example looks for a TCP window size of 55808.

```
window:55808;
```

3.6.14 itype

The itype keyword is used to check for a specific ICMP type value.

Format

```
itype: [<|>]<number> [<><number>];
```

Example

This example looks for an ICMP type greater than 30.

```
itype:>30;
```

3.6.15 icode

The itype keyword is used to check for a specific ICMP code value.

Format

```
icode: [<|>]<number> [<><number>];
```

Example

This example looks for an ICMP code greater than 30.

```
code:>30;
```

3.6.16 icmp_id

The itype keyword is used to check for a specific ICMP ID value.

This is useful because some covert channel programs use static ICMP fields when they communicate. This particular plugin was developed to detect the stacheldraht DDoS agent.

Format

```
icmp_id:<number>;
```

Example

This example looks for an ICMP ID of 0.

```
icmp_id:0;
```

3.6.17 icmp_seq

The itype keyword is used to check for a specific ICMP sequence value.

This is useful because some covert channel programs use static ICMP fields when they communicate. This particular plugin was developed to detect the stacheldraht DDoS agent.

Format

```
icmp_seq: <number>;
```

Example

This example looks for an ICMP Sequence of 0.

```
icmp_seq:0;
```

3.6.18 rpc

The rpc keyword is used to check for a RPC application, version, and procedure numbers in SUNRPC CALL requests.

Wildcards are valid for both version and procedure numbers by using '*';

Format

```
rpc: <application number>, [<version number>|*], [<procedure number>|*]>;
```

Example

The following example looks for an RPC portmap GETPORT request.

```
alert tcp any any -> any 111 (rpc: 100000,*,3);
```

Warning

Because of the fast pattern matching engine, the RPC keyword is slower than looking for the RPC values by using normal content matching.

3.6.19 ip_proto

The ip_proto keyword allows checks against the IP protocol header. For a list of protocols that may be specified by name, see /etc/protocols.

Format

```
ip_proto:[!><] <name or number>;
```

Example

This example looks for IGMP traffic.

```
alert ip any any -> any any (ip_proto:igmp;)
```

3.6.20 sameip

The sameip keyword allows rules to check if the source ip is the same as the destination IP.

Format

```
sameip;
```

Example

This example looks for any traffic where the Source IP and the Destination IP is the same.

```
alert ip any any -> any any (sameip;)
```

3.7 Post-Detection Rule Options

3.7.1 logto

The logto option tells Snort to log all packets that trigger this rule to a special output log file. This is especially handy for combining data from things like NMAP activity, HTTP CGI scans, etc. It should be noted that this option does not work when Snort is in binary logging mode.

Format

```
logto:"filename";
```

3.7.2 session

The session keyword is built to extract user data from TCP Sessions. There are many cases where seeing what users are typing in telnet, rlogin, ftp, or even web sessions is very useful.

There are two available argument keywords for the session rule option, printable or all. The printable keyword only prints out data that the user would normally see or be able to type.

The all keyword substitutes non-printable characters with their hexadecimal equivalents.

Format

```
session: [printable|all];
```

Example

The following example logs all printable strings in a telnet packet.

```
log tcp any any <> any 23 (session:printable;)
```

Warnings

Using the session keyword can slow Snort down considerably, so it should not be used in heavy load situations. The session keyword is best suited for post-processing binary (pcap) log files.

3.7.3 resp

The `resp` keyword is used attempt to close sessions when an alert is triggered. In Snort, this is called flexible response.

Flexible Response supports the following mechanisms for attempting to close sessions:

Option	Description
<code>rst_snd</code>	Send TCP-RST packets to the sending socket
<code>rst_rcv</code>	Send TCP-RST packets to the receiving socket
<code>rst_all</code>	Send TCP_RST packets in both directions
<code>icmp_net</code>	Send a ICMP_NET_UNREACH to the sender
<code>icmp_host</code>	Send a ICMP_HOST_UNREACH to the sender
<code>icmp_port</code>	Send a ICMP_PORT_UNREACH to the sender
<code>icmp_all</code>	Send all above ICMP packets to the sender

These options can be combined to send multiple responses to the target host.

Format

```
resp: <resp_mechanism>[,<resp_mechanism>[,<resp_mechanism>]];
```

Warnings

This functionality is not built in by default. Use the `--enable-flexresp` flag to configure when building Snort to enable this functionality.

Be very careful when using Flexible Response. It is quite easy to get Snort into an infinite loop by defining a rule such as:

```
alert tcp any any -> any any (resp:rst_all;)
```

It is easy to be fooled into interfering with normal network traffic as well.

Example

The following example attempts to reset any TCP connection to port 1524.

```
alert tcp any any -> any 1524 (flags:S; resp:rst_all;)
```

3.7.4 react

This keyword implements an ability for users to react to traffic that matches a Snort rule. The basic reaction is blocking interesting sites users want to access: New York Times, slashdot, or something really important - napster and porn sites. The React code allows Snort to actively close offending connections and/or send a visible notice to the browser. The notice may include your own comment. The following arguments (basic modifiers) are valid for this option:

- `block` - close connection and send the visible notice
- `warn` - send the visible, warning notice (will be available soon)

The basic argument may be combined with the following arguments (additional modifiers):

- `msg` - include the `msg` option text into the blocking visible notice

- proxy: <port_nr> - use the proxy port to send the visible notice (will be available soon)

Multiple additional arguments are separated by a comma. The react keyword should be placed as the last one in the option list.

Format

```
react: <react_basic_modifier[, react_additional_modifier]>;

alert tcp any any <> 192.168.1.0/24 80 (content: "bad.htm"; \
    msg: "Not for children!"; react: block, msg;)
```

Figure 3.22: React Usage Example

Warnings

React functionality is not built in by default. This code is currently bundled under Flexible Response, so enabling Flexible Response (`-enable-flexresp`) will also enable React.

Be very careful when using react. Causing a network traffic generation loop is very easy to do with this functionality.

3.7.5 tag

The tag keyword allow rules to log more than just the single packet that triggered the rule. Once a rule is triggered, additional traffic involving the source and/or destination host is *tagged*. Tagged traffic is logged to allow analysis of response codes and post-attack traffic. *tagged* alerts will be sent to the same output plugins as the original alert, but it is the responsibility of the output plugin to properly handle these special alerts. Currently, the database output plugin, described in Section 2.3.6, does not properly handle *tagged* alerts.

Format

tag: <type>, <count>, <metric>, [direction]

type

- **session** - Log packets in the session that set off the rule
- **host** - Log packets from the host that caused the tag to activate (uses [direction] modifier)

count - Count is specified as a number of units. Units are specified in the <metric> field.

metric

- **packets** - Tag the host/session for <count> packets
- **seconds** - Tag the host/session for <count> seconds

Note, any packets that generate an alert will not be tagged. For example, it may seem that the following rule will tag the first 600 seconds of any packet involving 10.1.1.1.

```
alert tcp any any <> any 10.1.1.1 (tag:host,600,seconds,src;)
```

However, since the rule will fire on every packet involving 10.1.1.1, no packets will get tagged. The *flowbits* option would be useful here.

```
alert tcp any any <> any 10.1.1.1 (flowbits:isnotset,tagged;
    flowbits:set,tagged; tag:host,600,seconds,src;)
```


Example

This example logs the first 10 seconds of any telnet session.

```
alert tcp any any -> any 23 (flags:s,12; tag:session,10,seconds;)
```

3.8 Event Thresholding

Event thresholding can be used to reduce the number of logged alerts for noisy rules. This can be tuned to significantly reduce false alarms, and it can also be used to write a newer breed of rules. Thresholding commands limit the number of times a particular event is logged during a specified time interval.

There are 3 types of thresholding:

- **limit**
Alerts on the 1st m events during the time interval, then ignores events for the rest of the time interval.
- **threshold**
Alerts every m times we see this event during the time interval.
- **both**
Alerts once per time interval after seeing m occurrences of the event, then ignores any additional events during the time interval.

Thresholding commands can be included as part of a rule, or you can use standalone threshold commands that reference the generator and SID they are applied to. There is no functional difference between adding a threshold to a rule, or using a separate threshold command applied to the same rule. There is a logical difference. Some rules may only make sense with a threshold. These should incorporate the threshold command into the rule. For instance, a rule for detecting a too many login password attempts may require more than 5 attempts. This can be done using the ‘limit’ type of threshold command. It makes sense that the threshold feature is an integral part of this rule.

In order for rule thresholds to apply properly, these rules must contain a SID.

Only one threshold may be applied to any given generator and SID pair. If more than one threshold is applied to a generator and SID pair, Snort will terminate with an error while reading the configuration information.

3.8.1 Standalone Options

This format supports 6 threshold options as described in Table 3.6—all are required.

Table 3.6: Standalone Options

Option	Arguments
gen_id	<generator ID>
sig_id	<Snort signature ID>
type	limit, threshold, or both
track	by_src or by_dst
count	<number of events>
seconds	<time period over which count is accrued>

3.8.2 Standalone Format

```
threshold gen_id <gen-id>, sig_id <sig-id>, \  
    type <limit|threshold|both>, \  
    track <by_src|by_dst>, count <s>, seconds <m>
```

3.8.3 Rule Keyword Format

This format supports 4 threshold options as described in Table 3.7—all are required.

Table 3.7: Rule Keyword Options

Option	Arguments
type	limit, threshold, or both
track	by_src or by_dst
count	<number of events>
seconds	<time period over which count is accrued>

3.8.4 Rule Keyword Format

```
threshold: type <limit|threshold|both>, track <by_src|by_dst>, \  
    count <n>, seconds <m>;
```

For either standalone or rule format, all tracking is by src or by dst ip, ports or anything else are not tracked.

Thresholding can also be used globally, this allows you to specify a threshold for every rule. Standard thresholding tests are applied first to an event, if they do not block a rule from being logged, and then the global thresholding test is applied—thresholds in a rule will override a global threshold. Global thresholds do not override what’s in a signature or a more specific stand-alone threshold.

The global threshold options are the same as the standard threshold options with the exception of the ‘sig_id’ field. The sig_id field must be set to 0 to indicate that this threshold command applies to all sig_id values with the specified gen_id. To apply the same threshold to all gen_id’s at the same time, and with just one command specify a value of gen_id=0.

The format for global threshold commands is as such:

```
threshold gen_id <gen-id>, sig_id 0, \  
    type <limit|threshold|both>, \  
    track <by_src|by_dst>, \  
    count <n>, \  
    seconds <m>
```

This applies a threshold to every event from <gen-id>.

or

```
threshold gen_id 0 , sig_id 0, \  
    type <limit|threshold|both>, \  
    track <by_src|by_dst>, \  
    count <n>, \  
    seconds <m>
```

This applies a threshold to every event from every gen-id.

3.8.5 Examples

Standalone Thresholds

Limit logging to 1 event per 60 seconds:

```
threshold gen_id 1, sig_id 1851, \  
    type limit, track by_src, \  
    count 1, seconds 60
```

Limit logging to every 3rd event:

```
threshold gen_id 1, sig_id 1852, \  
    type threshold, track by_src, \  
    count 3, seconds 60
```

Limit logging to just 1 event per 60 seconds, but only if we exceed 30 events in 60 seconds:

```
threshold gen_id 1, sig_id 1853, \  
    type both, track by_src, \  
    count 30, seconds 60
```

Rule Thresholds

This rule logs the first event of this SID every 60 seconds.

```
alert tcp $external_net any -> $http_servers $http_ports \  
    (msg:"web-misc robots.txt access"; flow:to_server, established; \  
    uricontent:"/robots.txt"; nocase; reference:nessus,10302; \  
    classtype:web-application-activity; threshold: type limit, track \  
    by_src, count 1 , seconds 60 ; sid:1000852; rev:1;)
```

This rule logs every 10th event on this SID during a 60 second interval. So if less than 10 events occur in 60 seconds, nothing gets logged. Once an event is logged, a new time period starts for type=threshold.

```
alert tcp $external_net any -> $http_servers $http_ports \  
    (msg:"web-misc robots.txt access"; flow:to_server, established; \  
    uricontent:"/robots.txt"; nocase; reference:nessus,10302; \  
    classtype:web-application-activity; threshold: type threshold, \  
    track by_dst, count 10 , seconds 60 ; sid:1000852; rev:1;)
```

This rule logs at most one event every 60 seconds if at least 10 events on this SID are fired.

```
alert tcp $external_net any -> $http_servers $http_ports \  
    (msg:"web-misc robots.txt access"; flow:to_server, established; \  
    uricontent:"/robots.txt"; nocase; reference:nessus,10302; \  
    classtype:web-application-activity; threshold: type both , track \  
    by_dst, count 10 , seconds 60 ; sid:1000852; rev:1;)
```

Global Thresholds

Limit to logging 1 event per 60 seconds per IP triggering each rule (rule gen_id is 1):

```
threshold gen_id 1, sig_id 0, type limit, track by_src, count 1, seconds 60
```

Limit to logging 1 event per 60 seconds per IP, triggering each rule for each event generator:

```
threshold gen_id 0, sig_id 0, type limit, track by_src, count 1, seconds 60
```

Events in Snort are generated in the usual way, thresholding is handled as part of the output system. Read `gen-msg.map` for details on gen ids.

Users can also configure a memcap for threshold with a “config:” option:

```
config threshold: memcap <bytes>
```

3.9 Event Suppression

Event suppression stops specified events from firing without removing the rule from the rule base. Suppression uses a CIDR block notation to select specific networks and users for suppression. Suppression tests are performed prior to either standard or global thresholding tests.

Suppression commands are standalone commands that reference generators, SIDs, and IP addresses via a CIDR block. This allows a rule to be completely suppressed, or suppressed when the causative traffic is going to or coming from a specific IP or group of IP addresses.

You may apply multiple suppression commands to a SID. You may also combine one threshold command and several suppression commands to the same SID.

3.9.1 Format

The suppress command supports either 2 or 4 options, as described in Table 3.8.

Table 3.8: Suppression Options

Option	Argument	Required?
<code>gen_id</code>	<generator id>	required
<code>sig_id</code>	<Snort signature id>	required
<code>track</code>	<code>by_src</code> or <code>by_dst</code>	optional, requires ip
<code>ip</code>	<code>ip[/mask]</code>	optional, requires track

```
suppress gen_id <gen-id>, sig_id <sig-id>, \  
    track <by_src|by_dst>, ip <ip|mask-bits>
```

3.9.2 Examples

Suppress this event completely:

```
suppress gen_id 1, sig_id 1852:
```

Suppress this event from this IP:

```
suppress gen_id 1, sig_id 1852, track by_src, ip 10.1.1.54
```

Suppress this event to this CIDR block:

```
suppress gen_id 1, sig_id 1852, track by_dst, ip 10.1.1.0/24
```

3.10 Snort Multi-Event Logging (Event Queue)

Snort supports logging multiple events per packet/stream that are prioritized with different insertion methods, such as max content length or event ordering using the event queue.

The general configuration of the event queue is as follows:

```
config event_queue: [max_events [size]] [log [size]] [order_events [TYPE]]
```

3.10.1 Event Queue Configuration Options

There are three configuration options to the configuration parameter 'event_queue'.

1. **max_queue**

This determines the maximum size of the event queue. For example, if the event queue has a max size of 8, only 8 events will be stored for a single packet or stream.

The default value is 8.

2. **log**

This determines the number of events to log for a given packet or stream. You can't log more than the `max_event` number that was specified.

The default value is 3.

3. **order_events**

This argument determines the way that the incoming events are ordered. We currently have two different methods:

- **priority** - The highest priority (1 being the highest) events are ordered first.
- **content_length** - Rules are ordered before decode or preprocessor alerts, and rules that have a longer content are ordered before rules with shorter contents.

The method in which events are ordered does not affect rule types such as pass, alert, log, etc.

The default value is `content_length`.

3.10.2 Event Queue Configuration Examples

The default configuration:

```
config event_queue: max_queue 8 log 3 order_events content_length
```

Example of a reconfigured event queue:

```
config event_queue: max_queue 10 log 3 order_events content_length
```

Use the default event queue values, but change event order:

```
config event_queue: order_events priority
```

Use the default event queue values but change the number of logged events:

```
config event_queue: log 2
```

3.11 Writing Good Rules

There are some general concepts to keep in mind when developing Snort rules to maximize efficiency and speed.

3.11.1 Content Matching

The 2.0 detection engine changes the way Snort works slightly by having the first phase be a setwise pattern match. The longer a content option is, the more *exact* the match. Rules without *content* (or *uricontent*) slow the entire system down.

While some detection options, such as *pcre* and *byte_test*, perform detection in the payload section of the packet, they do not use the setwise pattern matching engine. If at all possible, try and have at least one *content* option if at all possible.

3.11.2 Catch the Vulnerability, Not the Exploit

Try to write rules that target the vulnerability, instead of a specific exploit.

For example, look for a the vulnerable command with an argument that is too large, instead of shellcode that binds a shell.

By writing rules for the vulnerability, the rule is less vulnerable to evasion when an attacker changes the exploit slightly.

3.11.3 Catch the Oddities of the Protocol in the Rule

Many services typically send the commands in upper case letters. FTP is a good example. In FTP, to send the username, the client sends:

```
user username_here
```

A simple rule to look for FTP root login attempts could be:

```
alert tcp any any -> any any 21 (content:"user root";)
```

While it may *seem* trivial to write a rule that looks for the username root, a good rule will handle all of the odd things that the protocol might handle when accepting the user command.

For example, each of the following are accepted by most FTP servers:

```
user root
user root
user root
user      root
user<tab>root
```

To handle all of the cases that the FTP server might handle, the rule needs more smarts than a simple string match.

A good rule that looks for root login on ftp would be:

```
alert tcp any any -> any 21 (flow:to_server,established; content:"root";
    pcre:"/user\s+root/i";)
```

There are a few important things to note in this rule:

- The rule has a *flow* option, verifying this is traffic going to the server on an established session.
- The rule has a *content* option, looking for *root*, which is the longest, most unique string in the attack. This option is added to allow Snort's setwise pattern match detection engine to give Snort a boost in speed.
- The rule has a *pcrc* option, looking for user, followed at least one space character (which includes tab), followed by root, ignoring case.

3.11.4 Optimizing Rules

The content matching portion of the detection engine has recursion to handle a few evasion cases. Rules that are not properly written can cause Snort to waste time duplicating checks.

The way the recursion works now is if a pattern matches, and if any of the detection options after that pattern fail, then look for the pattern again after where it was found the previous time. Repeat until the pattern is not found again or the opt functions all succeed.

On first read, that may not sound like a smart idea, but it is needed. For example, take the following rule:

```
alert ip any any -> any any (content:"a"; content:"b"; within:1;)
```

This rule would look for “a”, immediately followed by “b”. Without recursion, the payload “aab” would fail, even though it is obvious that the payload “aab” has “a” immediately followed by “b”, because the first “a” is not immediately followed by “b”.

While recursion is important for detection, the recursion implementation is not very smart.

For example, the following rule options are not optimized:

```
content:"|13|"; dsize:1;
```

By looking at this rule snippet, it is obvious the rule looks for a packet with a single byte of 0x13. However, because of recursion, a packet with 1024 bytes of 0x13 could cause 1023 too many pattern match attempts and 1023 too many dsize checks. Why? The content 0x13 would be found in the first byte, then the dsize option would fail, and because of recursion, the content 0x13 would be found again starting after where the previous 0x13 was found, once it is found, then check the dsize again, repeating until 0x13 is not found in the payload again.

Reordering the rule options so that discrete checks (such as dsize) are moved to the beginning of the rule speed up Snort.

The optimized rule snippet would be:

```
dsize:1; content:"|13|";
```

A packet of 1024 bytes of 0x13 would fail immediately, as the dsize check is the first option checked and dsize is a discrete check without recursion.

The following rule options are discrete and should generally be placed at the beginning of any rule:

- dsize
- flags
- flow
- fragbits

- icmp_id
- icmp_seq
- icode
- id
- ipopts
- ip_proto
- itype
- seq
- session
- tos
- ttl
- ack
- window
- resp
- sameip

3.11.5 testing numerical values

The rule options *byte_test* and *byte_jump* were written to support writing rules for protocols that have length encoded data. RPC was the protocol that spawned the requirement for these two rule options, as RPC uses simple length based encoding for passing data.

In order to understand *why* *byte_test* and *byte_jump* are useful, let's go through an exploit attempt against the sadmind service.

This is the payload of the exploit:

```

89 09 9c e2 00 00 00 00 00 00 02 00 01 87 88 .....
00 00 00 0a 00 00 00 01 00 00 00 01 00 00 20 .....
40 28 3a 10 00 00 00 0a 4d 45 54 41 53 50 4c 4f @(:.....metasplo
49 54 00 00 00 00 00 00 00 00 00 00 00 00 00 00 it.....
00 00 00 00 00 00 00 00 40 28 3a 14 00 07 45 df .....@(:...e.
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
00 00 00 00 00 00 00 00 06 00 00 00 00 00 00 00 .....
00 00 00 00 00 00 00 00 04 00 00 00 00 00 00 04 .....
7f 00 00 01 00 01 87 88 00 00 00 0a 00 00 00 04 .....
7f 00 00 01 00 01 87 88 00 00 00 0a 00 00 00 11 .....
00 00 00 1e 00 00 00 00 00 00 00 00 00 00 00 00 .....
00 00 00 00 00 00 00 00 3b 4d 45 54 41 53 50 4c 4f .....;metasplo
49 54 00 00 00 00 00 00 00 00 00 00 00 00 00 00 it.....
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 .....
00 00 00 00 00 00 00 00 06 73 79 73 74 65 6d 00 00 .....system..
00 00 00 15 2e 2e 2f 2e 2e 2f 2e 2e 2f 2e 2e 2f ...../.../.../
2e 2e 2f 62 69 6e 2f 73 68 00 00 00 00 00 04 1e ../bin/sh.....
<snip>

```

Let's break this up, describe each of the fields, and figure out how to write a rule to catch this exploit.

There are a few things to note with RPC:

- Numbers are written as uint32s, taking four bytes. The number 26 would show up as 0x0000001a.
- Strings are written as a uint32 specifying the length of the string, the string, and then null bytes to pad the length of the string to end on a 4 byte boundary. The string "bob" would show up as 0x00000003626f6200.

```
89 09 9c e2    - the request id, a random uint32, unique to each request
00 00 00 00    - rpc type (call = 0, response = 1)
00 00 00 02    - rpc version (2)
00 01 87 88    - rpc program (0x00018788 = 100232 = sadmind)
00 00 00 0a    - rpc program version (0x0000000a = 10)
00 00 00 01    - rpc procedure (0x00000001 = 1)
00 00 00 01    - credential flavor (1 = auth\_unix)
00 00 00 20    - length of auth\_unix data (0x20 = 32)

## the next 32 bytes are the auth\_unix data
40 28 3a 10    - unix timestamp (0x40283a10 = 1076378128 = feb 10 01:55:28 2004 gmt)
00 00 00 0a    - length of the client machine name (0x0a = 10)
4d 45 54 41 53 50 4c 4f 49 54 00 00    - metasploit

00 00 00 00    - uid of requesting user (0)
00 00 00 00    - gid of requesting user (0)
00 00 00 00    - extra group ids (0)

00 00 00 00    - verifier flavor (0 = auth\_null, aka none)
00 00 00 00    - length of verifier (0, aka none)
```

The rest of the packet is the request that gets passed to procedure 1 of sadmind.

However, we know the vulnerability is that sadmind trusts the uid coming from the client. sadmind runs any request where the client's uid is 0 as root. As such, we have decoded enough of the request to write our rule.

First, we need to make sure that our packet is an RPC call.

```
content:"|00 00 00 00|"; offset:4; depth:4;
```

Then, we need to make sure that our packet is a call to sadmind.

```
content:"|00 01 87 88|"; offset:12; depth:4;
```

Then, we need to make sure that our packet is a call to the procedure 1, the vulnerable procedure.

```
content:"|00 00 00 01|"; offset:16; depth:4;
```

Then, we need to make sure that our packet has auth_unix credentials.

```
content:"|00 00 00 01|"; offset:20; depth:4;
```

We don't care about the hostname, but we want to skip over it and check a number value after the hostname. This is where byte_test is useful. Starting at the length of the hostname, the data we have is:

```
00 00 00 0a 4d 45 54 41 53 50 4c 4f 49 54 00 00
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00
```

We want to read 4 bytes, turn it into a number, and jump that many bytes forward, making sure to account for the padding that RPC requires on strings. If we do that, we are now at:

```
00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
00 00 00 00
```

which happens to be the exact location of the uid, the value we want to check.

In english, we want to read 4 bytes, 36 bytes from the beginning of the packet, and turn those 4 bytes into an integer and jump that many bytes forward, aligning on the 4 byte boundary. To do that in a Snort rule, we use:

```
byte_jump:4,36,align;
```

then we want to look for the uid of 0.

```
content:"|00 00 00 00|"; within:4;
```

Now that we have all the detection capabilities for our rule, let's put them all together.

```
content:"|00 00 00 00|"; offset:4; depth:4;
content:"g00 01 87 88|"; offset:12; depth:4;
content:"|00 00 00 01|"; offset:16; depth:4;
content:"|00 00 00 01|"; offset:20; depth:4;
byte_jump:4,36,align;
content:"|00 00 00 00|"; within:4;
```

The 3rd and fourth string match are right next to each other, so we should combine those patterns. We end up with:

```
content:"|00 00 00 00|"; offset:4; depth:4;
content:"|00 01 87 88|"; offset:12; depth:4;
content:"|00 00 00 01 00 00 00 01|"; offset:16; depth:8;
byte_jump:4,36,align;
content:"|00 00 00 00|"; within:4;
```

If the sadmind service was vulnerable to a buffer overflow when reading the client's hostname, instead of reading the length of the hostname and jumping that many bytes forward, we would check the length of the hostname to make sure it is not too large.

To do that, we would read 4 bytes, starting 36 bytes into the packet, turn it into a number, and then make sure it is not too large (let's say bigger than 200 bytes). In Snort, we do:

```
byte_test:4,>,200,36;
```

Our full rule would be:

```
content:"|00 00 00 00|"; offset:4; depth:4;
content:"|00 01 87 88|"; offset:12; depth:4;
content:"|00 00 00 01 00 00 00 01|"; offset:16; depth:8;
byte_test:4,>,200,36;
```

Chapter 4

Making Snort Faster

4.1 MMAPed pcap

On Linux, a modified version of libpcap is available that implements a shared memory ring buffer. Phil Woods (cpw@lanl.gov) is the current maintainer of the libpcap implementation of the shared memory ring buffer. The shared memory ring buffer libpcap can be downloaded from his website at <http://public.lanl.gov/cpw/>.

Instead of the normal mechanism of copying the packets from kernel memory into userland memory, by using a shared memory ring buffer, libpcap is able to queue packets into a shared buffer that Snort is able to read directly. This change speeds up Snort by limiting the number of times the packet is copied before Snort gets to perform its detection upon it.

Once Snort linked against the shared memory libpcap, enabling the ring buffer is done via setting the environment variable *PCAP_FRAMES*. *PCAP_FRAMES* is the size of the ring buffer. According to Phil, the maximum size is 32768, as this appears to be the maximum number of iovecs the kernel can handle. By using *PCAP_FRAMES=max*, libpcap will automatically use the most frames possible. On Ethernet, this ends up being 1530 bytes per frame, for a total of around 52 Mbytes of memory for the ring buffer alone.

Chapter 5

Snort Development

Currently, this chapter is here as a place holder. It will someday contain references on how to create new detection plugins and preprocessors. End users don't really need to be reading this section. This is intended to help developers get a basic understanding of whats going on quickly.

If you are going to be helping out with Snort development, please use the HEAD branch of cvs. We've had problems in the past of people submitting patches only to the stable branch (since they are likely writing this stuff for their own IDS purposes). Bugfixes are what goes into STABLE. Features go into HEAD.

5.1 Submitting Patches

Patches to Snort should be sent to the `snort-devel@lists.sourceforge.net` mailing list. Patches should be done with the command `diff -nu snort-orig snort-new`.

5.2 Snort dataflow

First, traffic is acquired from the network link via libpcap. Packets are passed through a series of decoder routines that first fill out the packet structure for link level protocols then are further decoded for things like TCP and UDP ports.

Packets are then sent through the registered set of preprocessors. Each preprocessor checks to see if this packet is something it should look at.

Packets are then sent through the detection engine. The detection engine checks each packet against the various options listed in the Snort rules files. Each of the keyword options is a plugin. This allows this to be easily extensible.

5.2.1 Preprocessors

For example, a TCP analysis preprocessor could simply return if the packet does not have a TCP header. It can do this by checking:

```
if (p->tcph==null)
    return;
```

Similarly, there are a lot of `packet_flags` available that can be used to mark a packet as "reassembled" or logged. Check out `src/decode.h` for the list of `pkt_*` constants.

5.2.2 Detection Plugins

Basically, look at an existing output plugin and copy it to a new item and change a few things. Later, we'll document what these few things are.

5.2.3 Output Plugins

Generally, new output plugins should go into the barnyard project rather than the Snort project. We are currently cleaning house on the available output options.

5.3 The Snort Team

Creator and Lead Architect	Marty Roesch
Lead Snort Developers	Marc Norton Andrew Mullican Steve Sturges
Snort Rules Maintainer	Brian Caswell
Snort Rules Team	Nigel Houghton Alex Kirk Judy Novak Matt Watchinski
Win32 Maintainer	Chris Reid
RPM Maintainers	JP Vossen Daniel Wittenberg
Inline Developers	Victor Julien Rob McMillen William Metcalf
Major Contributors	Daniel Roelker Andrew Baker Erek Adams Scott Campbell Roman D. Michael Davis Chris Green Jed Haile Jeremy Hewlett Glenn Mansfield Keeni Chad Kreimendahl Jeff Nathan Andreas Ostling Dragos Ruiu Fyodor Yarochkin Phil Wood

Bibliography

- [1] <http://packetstorm.securify.com/mag/phrack/phrack49/p49-06>
- [2] <http://www.nmap.org>
- [3] <http://public.pacbell.net/dedicated/cidr.html>
- [4] <http://www.whitehats.com>
- [5] <http://www.incident.org/snortdb>
- [6] <http://www.pcre.org>