

SnortTM Users Manual

2.2.0RC1

The Snort Project

28th June 2004

Copyright ©1998-2003 Martin Roesch

Copyright ©2001-2003 Chris Green

Copyright ©2003 Sourcefire, Inc.

Contents

1	Snort Overview	6
1.1	Getting Started	6
1.2	Sniffer Mode	6
1.3	Packet Logger Mode	7
1.4	Network Intrusion Detection Mode	8
1.4.1	NIDS Mode Output Options	8
1.4.2	High Performance Configuration	9
1.4.3	Changing Alert Order	10
1.5	Miscellaneous	10
1.6	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	portscan detector	16
2.1.2	portscan ignorehosts	17
2.1.3	frag2	17
2.1.4	stream4	18
2.1.5	flow	19
2.1.6	flow-portscan	20
2.1.7	telnet decode	25
2.1.8	rpc decode	25
2.1.9	performance monitor	26
2.1.10	http inspect	27
2.2	event thresholding	35

2.2.1	standalone options	35
2.2.2	standalone format	35
2.2.3	rule keyword format	36
2.2.4	rule keyword format	36
2.2.5	examples	37
2.3	event suppression	39
2.3.1	format	39
2.3.2	examples	39
2.4	output modules	40
2.4.1	alert_syslog	40
2.4.2	alert_fast	41
2.4.3	alert_full	42
2.4.4	alert_unixsock	42
2.4.5	log_tcpdump	42
2.4.6	database	43
2.4.7	csv	44
2.4.8	unified	46
2.4.9	log null	46
3	Writing Snort Rules	
	How to Write Snort Rules and Keep Your Sanity	47
3.1	The Basics	47
3.2	Rules Headers	48
3.2.1	Rule Actions	48
3.2.2	Protocols	48
3.2.3	IP Addresses	48
3.2.4	Port Numbers	49
3.2.5	The Direction Operator	49
3.2.6	Activate/Dynamic Rules	50
3.3	Rule Options	52
3.4	Meta-Data Rule Options	52
3.4.1	msg	52
3.4.2	reference	52
3.4.3	sid	53
3.4.4	rev	53
3.4.5	classtype	54

3.4.6	Priority	55
3.5	Payload Detection Rule Options	56
3.5.1	content	56
3.5.2	nocase	56
3.5.3	rawbytes	57
3.5.4	depth	57
3.5.5	offset	58
3.5.6	distance	58
3.5.7	within	59
3.5.8	uricontent	59
3.5.9	isdataat	60
3.5.10	pcre	60
3.5.11	byte_test	61
3.5.12	byte_jump	62
3.5.13	regex	63
3.5.14	content-list	63
3.6	Non-payload Detection Rule Options	63
3.6.1	Fragoffset	63
3.6.2	ttl	64
3.6.3	tos	64
3.6.4	id	64
3.6.5	ipopts	65
3.6.6	Fragbits	66
3.6.7	dsize	66
3.6.8	flags	67
3.6.9	flow	67
3.6.10	flowbits	68
3.6.11	seq	69
3.6.12	ack	69
3.6.13	window	69
3.6.14	itype	70
3.6.15	icode	70
3.6.16	icmp_id	70
3.6.17	icmp_seq	71
3.6.18	rpc	71

3.6.19	ip_proto	71
3.6.20	sameip	72
3.7	Post-Detection Rule Options	72
3.7.1	logto	72
3.7.2	session	72
3.7.3	sest	73
3.7.4	React	74
3.7.5	tag	74
3.8	writing good rules	76
3.8.1	content matching	76
3.8.2	catch the vulnerability, not the exploit	76
3.8.3	catch the oddities of the protocol in the rule	76
3.8.4	optimizing rules	77
3.8.5	testing numerical values	78
4	Making Snort Faster	82
4.1	MMAPed pcap	82
5	Snort Development	83
5.1	submitting patches	83
5.2	snort dataflow	83
5.2.1	preprocessors	83
5.2.2	detection plugins	84
5.2.3	output plugins	84

Chapter 1

Snort Overview

This manual is based off of *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> and Jeremy Hewlett <jh@snort.org>. If you have a better way to say something or something in the documentation is outdated, drop us a line and we will update it. The documentation is now in L^AT_EX format in the `doc/snortman.tex` file if you would like to submit patches for this document. Small documentation updates are the easiest way to help the Snort Project out.

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. There are three main modes in which Snort can be configured: sniffer, packet logger, and network intrusion detection system. Sniffer mode simply reads the packets off of the network and displays them for you in a continuous stream on the console. Packet logger mode logs the packets to the disk. Network intrusion detection mode is the most complex and configurable configuration, allowing Snort to analyze network traffic for matches against a user defined rule set and perform several actions based upon what it sees.

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.



if both hosts are on the home network, then they are recorded based upon the higher of the two's 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 that is really required to place Snort into

logger mode is the specification of a logging directory at the command line with the `-l` switch, the `-b` binary logging switch merely provides a modifier to tell it 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 set 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 it's not necessary to specify the `-e` switch either.

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

This will configure Snort to run in it's most basic NIDS form, logging packets that the rules tell it to in plain ASCII to 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:

- A fast** fast alert mode, write the alert in a simple format with a timestamp, alert message, source and destination IPs/ports
- A full** this is also the default alert mode, so if you specify nothing this will automatically be used
- A unsock** send alerts to a UNIX socket that another program can listen on
- A none** turn off alerting
- A console** send "fast-style" alerts to the console (screen)
- A cmg** generate "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. If you wish to disable packet logging all together, 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 2.4.1 for more details on configuring syslog output.

Here are some output configuration examples:

- Log to default (decoded ASCII) facility and send alerts to syslog

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

- 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 High Performance Configuration

If you want Snort to go *fast* (like keep up with a 1000 Mbps connect), 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 and have another program performing the slow actions, such as writing into a database.

If you want a text file that's easily parsable, but still be 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
```

1.4.3 Changing Alert Order

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

For people who know what they're doing, the `-o` switch has been provided to change the default rule application behavior to Pass rules, then Alert, then Log:

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

1.5 Miscellaneous

If you want to run snort in daemon mode, you can add `-D` switch to any combination above. Please NOTICE that if you want to be able to restart snort by sending SIGHUP signal to the daemon, you will need to use full path to snort binary, when you start it, i.g.:

```
/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.

If you're going to be posting packet logs to public mailing lists you might want to try out the `-O` switch. This switch obfuscates your the IP addresses in the packet printouts. This is handy if you don't want the 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:

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

This will read the packets from a log file and dump the packets to the screen, obfuscating only the addresses from the 192.168.1.0/24 class C network.

1.6 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 get Snort running in several different modes.

⚠NOTE

Often `\?` is needed to escape the `?` in many shells.

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 putting them in place in the file in the place where the include appears.

Format

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



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 Variables (2.1) 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 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

The 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 -. * \$var - define meta variable * \$(var) - replace with the contents of variable var * \$(var:-default) - replace with the contents of the variable var or with default if var is undefined. * \$(var:?message) - replace with the contents of variable var or print out the error message message and exit

See Figure 2.2 for an example of these rules modifiers in action.

```
var MY_NET 192.168.1.0/24
log tcp any any -> $MY_NET 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	explanation
order	config order: pass alert log activation dynamic	Change the order that rules are evaluated
alertfile	config alertfile: alerts	Set the alerts output file
classification	config classification: misc-activity,Misc activity,3	See 3.2
decode_arp	config decode_arp	Turn on arp decoding (snort -a)
dump_chars_only	config dump_chars_only	Turn on character dumps (snort -C)
dump_payload	config dump_payload	Dump application layer (snort -d)
decode_data_link	config decode_data_link	Decode Layer2 headers (snort -e)
bpf_file	config bpf_file: filters.bpf	Specify BPF filters (snort -F)
set_gid	config set_gid: 30	Change to GID to specified GID (snort -g)
daemon	config daemon	Fork as a daemon (snort -D)
interface	config interface: xl0	Set the network interface (snort -i)

alert_with_interface_name	config alert_with_interface_name	Append interface name to alert (snort -I)
logdir	config logdir: /var/log/snort	Set the logdir (snort -l)
umask	config umask: 022	Set umask when running (snort -m)
pkt_count	config pkt_count: 13	Exit after N packets (snort -n)
nolog	config nolog	Disable Logging. Note: Alerts will still occur. (snort -N)
obfuscate	config obfuscate	Obfuscate IP Addresses (snort -O)
no_promisc	config no_promisc	Disable promiscuous mode (snort -p)
quiet	config quiet	Disable banner and status reports (snort -q)
chroot	config chroot: /home/snort	Chroot to specified dir (snort -t)
checksum_mode	config checksum_mode : all	Types of packets to calculate checksums. Values: none, noip, notcp, noicmp, noudp, or all
set_uid	set_uid: snort_user	Set UID to <id> (snort -u)
utc	config utc	Use UTC instead of local time for timestamps (snort -U)
verbose	config verbose	Use Verbose logging to stdout (snort -v)
dump_payload_verbose	config dump_payload_verbose	Dump raw packet starting at link layer (snort -X)
show_year	config show_year	show year in timestamps (snort -y)
stateful	config stateful	set 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	turn off the alerts generated by the decode phase of snort
disable_tcptopt_experimental_alerts	config disable_tcptopt_experimental_alerts	turn off alerts generated by experimental tcp options
disable_tcptopt_obsolete_alerts	config disable_tcptopt_obsolete_alerts	turn off alerts generated by obsolete tcp options

disable_tcptopt_ttcp_alerts	config disable_tcptopt_ttcp_alerts	turn off alerts generated by T/TCP options
disable_tcptopt_alerts	config disable_tcptopt_alerts	disable option length validation alerts
disable_ipopt_alerts	config disable_ipopt_alerts	disable IP option length validation alerts
detection	config detection: search-method ac no_stream_inserts max_queue_events	Make changes to the detection engine.
reference	config reference: www http://	add a new reference system to snort

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 through 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 portscan detector

the snort portscan preprocessor is developed by patrick mullen.

what the snort portscan preprocessor does

- log the start and end of portscans from a single source ip to the standard logging facility.
- if a log file is specified, logs the destination ips and ports scanned as well as the type of scan.

a portscan is defined as tcp connection attempts to more than p ports in t seconds or udp packets sent to more than p ports in t seconds. ports can be spread across any number of destination ip addresses, and may all be the same port if spread across multiple ips. this version does single->single and single->many portscans. the next full release will do distributed portscans (multiple->single or multiple->multiple). a portscan is also defined as a single stealth scan packet, such as null, fin, syn-fin, xmas, etc. this means that from scan-lib in the standard distribution of snort you should comment out the section for stealth scan packets. the benefit is with the portscan module these alerts would only show once per scan, rather than once for each packet. if you use the external logging feature you can look at the technique and type in the log file.

the arguments to this module are:

- network to monitor the network/cidr block to monitor for portscans
- number of ports number of ports accessed in the detection period
- detection period number of seconds to count that the port access threshold is considered for
- logdir/filename the directory/filename to place alerts in. alerts are also written to the standard alert file

format

```
portscan: <monitor network> <number of ports> <detection period> <file path>
```

```
preprocessor portscan: 192.168.1.0/24 5 7 /var/log/portscan.log
```

Figure 2.4: portscan preprocessor configuration example

2.1.2 portscan ignorehosts

another module from patrick mullen that modifies the portscan detection system's operation. if you have servers which tend to trip off the portscan detector (such as ntp, nfs, and dns servers), you can tell portscan to ignore tcp syn and udp portscans from certain hosts. the arguments to this module are a list of ips/cidr blocks to be ignored.

format

```
portscan-ignorehosts: <host list>
```

```
preprocessor portscan-ignorehosts: 192.168.1.5/32 192.168.3.0/24
```

Figure 2.5: portscan ignorehosts module configuration example

2.1.3 frag2

frag2, introduced in snort 1.8, is a new IP defragmentation preprocessor. frag2 is designed to replace the defrag preprocessor. this defragmenter is designed to be memory efficient and use the same memory management routines that are in use in other parts of snort.

frag2 has configurable memory usage and fragment timeout options. given no arguments, frag2 uses the default memory limit of 4194304 bytes (4MB) and a timeout period of 60 seconds. the timeout period is used to determine a length of time that a unassembled fragment should be discarded.

in snort 1.8.7, several options were added to help catch the use of evasion techniques such as fragroute.

format

```
preprocessor frag2: [memcap <xxx>], [timeout <xx>], [min_ttl <xx>], \
    [detect_state_problems], [ttl_limit <xx>]
```

timeout <seconds> amount of time to keep an inactive stream in the state table, sessions that are flushed will automatically be picked up again if more activity is seen, default is 30 seconds

memcap <bytes> number of bytes to set the memory cap at, if this limit is exceeded frag2 will aggressively prune inactive reassemblers, default is 4mb

detect_state_problems turns on alerts for events such as overlapping fragments

min_ttl sets the minimum ttl that frag2 will accept

ttl_limit sets the delta value that will set off an evasion alert. (initial fragment ttl +/- ttl limit)

```
preprocessor frag2: memcap 16777216, timeout 30
```

Figure 2.6: frag2 preprocessor configuration

2.1.4 stream4

the stream4 module provides tcp stream reassembly and stateful analysis capabilities to snort. robust stream reassembly capabilities allow snort to ignore "stateless" attacks such as stick and snot produce. stream4 also gives large scale users the ability to track more than 256 simultaneous tcp streams. stream4 should be able to scale to handle 32,768 simultaneous tcp connections in its default configuration.

stream4 contains two configurable modules, the stream4 preprocessor and the associated stream4 reassemble plugin. the stream4_reassemble options are listed below.

stream4 format

```
preprocessor stream4: [noinspect], [keepstats], [timeout <seconds>], \  
    [memcap <bytes>], [detect_scans], [detect_state_problems], \  
    [disable_evasion_alerts], [ttl_limit <count>]
```

noinspect disable stateful inspection

keepstats record session summary information in <logdir>/session.log

timeout <seconds> amount of time to keep an inactive stream in the state table, sessions that are flushed will automatically be picked up again if more activity is seen, default is 30 seconds

memcap <bytes> number of bytes to set the memory cap at, if this limit is exceeded stream4 will aggressively prune inactive sessions, default is 8mb

detect_scans turns on alerts for portscan events

detect_state_problems turns on alerts for stream events of note, such as evasive rst packets, data on the syn packet, and out of window sequence numbers

disable_evasion_alerts turns off alerts for events such as tcp overlap

ttl_limit sets the delta value that will set off

stream4_reassemble format

```
preprocessor stream4_reassemble: [clientonly], [serveronly],\  
    [noalerts], [ports <portlist>]
```

clientonly provide reassembly for the client side of a connection only

serveronly provide reassembly for the server side of a connection only

noalerts don't alert on events that may be insertion or evasion attacks

ports <portlist> - a whitespace separated list of ports to perform reassembly for, all provides reassembly for all ports, default provides reassembly for ports 21 23 25 53 80 110 111 143 and 513

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.

Table 2.2: stream4 defaults

option	default
session timeout	30 seconds
session memory cap	8388608 bytes
stateful inspection	active
stream stats	inactive
state problem alerts	inactive
portscan alerts	inactive

Table 2.3: stream4_reassemble defaults

option	default
reassemble client	active
reassemble server	inactive
reassemble ports	21 23 25 53 80 143 110 111 513 1433
reassembly alerts	active

2.1.5 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 obsoletes the conversation preprocessor.

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

memcap	number of bytes to allocate
rows	number of rows for the flow hash table. ¹
stats_interval	interval (in seconds) to dump statistics to stdout. set this to 0 to disable.
hash	hashing method to use. ²

example configuration

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

2.1.6 flow-portscan

this module is designed to detect portscans based off flow creation in the flow preprocessors. the goal is to catch one->many hosts and one->many ports scans.

the flow preprocessor to portscan recognizer is taken from experience with spp_conversation/portscan2 by jason larsen & jed haile and ipaudit by jon rifkin.

this subsystem became a bit more complicated than originally intended but it does a good job of mitigating false positives from devices such as squid proxies. the new design is also a lot more memory consistent than portscan1 or 2. it also ignores single port syn floods as they are a dos, not a portscan.

memory requirements should be way down from portscan2 architecture though but there's slightly less information saved off. the new architecture operates similarly to a ring buffer. when a scanner has not been active in a long time, it's only reclaimed when there is no more memory to use.

all of the prior methods for portscan detection in snort are deprecated and will be removed in the near future. if you have custom code against conversation or one of the portscan preprocessors, consider making it a module in flow or portscan.

the flow preprocessor must first be enabled in order for flow-portscan to function properly.

the basic components of flow-portscan are:

1. scoreboards

scoreboards contain information regarding timescales for a single ip address. there are two scoreboards, one for talkers (nodes that are active on your network) and one for scanners (nodes that have talked to a previously unknown port in your server-watch-net)

2. uniqueness tracker

the uniqueness tracker is used to determine if this connection should count as something "new" for a particular ip. it checks if a connection is a new type of connection for a source ip by disregarding the source port. any change in sip, dip, ip_proto, and dport indicates a new unique connection and will be processed further for the server statistics table and scoring. this keeps things like a web page with 15 images from rapidly increasing point scores with lots of accesses to the same web server.

3. server statistics tracker

this is used to track flows destined to the "server-watchnet" and keep "hitcounts" on the number of times a particular service has been requested with unique requests since snort has started. this hitcount is tracked by dip, dport, and protocol.

if a service is very popular, connections can be ignored for scoring by comparing the hitcount to the "server-ignore-limit". if there are more requests to this service than the server-ignore-limit, then flow-portscan will completely ignore this service. similarly, the "server-scanner-limit" controls if a request to a service counts as scanner points or as talker points.

if a request to a service is not in the server-watchnet, it will count as talker points. if no server-watchnet is defined, all alerts will be talker alerts.

execution path of flow-portscan

1. flow-portscan receives a new flow message from the flow module
2. the uniqueness tracker determines if message is a new type of flow by looking for changes in sip, dip, ip_proto, and dport. if this is not unique, and the tcp flags are normal, exit out.
3. if this connection is to an destination ip in the server-watchnet:
during the "server-learning-time", it increments the hitcounts for service popularity.
if it's otherwise just get the stored hitcount. if the hitcount is greater than the server-ignore-limit, exit out. if it's less than the server-scanner-limit, mark the incremented points as scanner points.

4. a connection is marked as either a talker or a scanner by step 3.

there are 4 time scales; 2 each for the ip scanner and ip talker.

the fixed timescales detect n events in m seconds. this is the typical type of portscan alert.

the sliding timescales adjust the "score reset point" on each event after the first. this adjusts the side of the window we're detecting portscan events in by taking

```
end = end + ((end - start) * sliding-scale-factor)
```

each time scale has it's own point tally that is incremented per new flow. each set of points only touches either the talker-fixed-score and talker-sliding-score or scanner-fixed-score and scanner-sliding-score

5. evaluate the score against individual thresholds, either talker or scanner.

```
if(fixed_limit <= fixed_score)
generate_alert()
```

format

```
preprocessor flow-portscan: [scoreboard-memcap-talker <bytes>] \
                           [scoreboard-rows-talker <count>] \
                           [scoreboard-rows-scanner <count>] \
                           [scoreboard-memcap-scanner <bytes>] \
                           [scanner-fixed-threshold <integer>] \
                           [scanner-sliding-threshold <integer>] \
                           [scanner-fixed-window <integer>] \
                           [scanner-sliding-window <integer>] \
                           [scanner-sliding-scale-factor <float>] \
                           [talker-fixed-threshold <integer>] \
                           [talker-sliding-threshold <integer>] \
                           [talker-fixed-window <integer>] \
```

```

[talker-sliding-window <integer>] \
[talker-sliding-scale-factor <float>] \
[unique-memcap <bytes>] \
[unique-rows <integer>] \
[server-memcap <bytes>] \
[server-rows <integer>] \
[server-watchnet <ip list in snort notation>] \
[src-ignore-net <ip list in snort notation>] \
[dst-ignore-net <ip list in snort notation>] \
[tcp-penalties <on|off>] \
[server-learning-time <seconds>] \
[server-ignore-limit <hit count>] \
[server-scanner-limit <hit count>] \
[alert-mode <once|all>] \
[output-mode <msg|pktkludge>] \
[base-score <integer>] \
[dumpall <1>]

```

1. scoreboard-rows-talker (default: 1000000).
number of rows to use for the talker table.
2. scoreboard-rows-scanner (default: 250000).
number of rows to use for the scanner table.
3. unique-rows (default: 1000000)
how many rows to allocate for the uniqueness tracker.
4. server-rows (default: 65536)
how many rows to allocate for server learning

general note about rows: higher row counts will take more memory away from the memory caps for a specific subsystem. in the snort output, this is referred to as "overhead bytes" and the percentage of overhead encountered will be shown. higher row counts provide a larger hash table to minimize collisions and have a faster overall processing time at the expense of memory. the hash tables themselves use a pseudorandom hardening salt that is picked at initialization time.
5. scoreboard-memcap-talker (default: 25165824)
number of bytes to use for the talker table.
6. scoreboard-memcap-scanner (default: 6291456)
number of bytes to use for the scanner table.
7. unique-memcap (default: 25165824)
how many bytes to allocate to the uniqueness tracker. the more memory given, the less that connections to a busy server will appear as a scan target on a popular service.
8. server-memcap (default: 2097152)
how many bytes to allocate for server learning
9. scanner-fixed-threshold (default: 15)
number of points that a scanner must accumulate in the scanner-fixed-window time range. set to 0 to disable this type of alert.

10. talker-fixed-threshold (default: 15)
number of points that a scanner must accumulate in talker-fixed-window time range. set to 0 to disable this type of alert.
11. scanner-sliding-threshold (default: 40)
number of points that a scanner must accumulate in scanner-sliding-window time range. set to 0 to disable this type of alert.
12. talker-sliding-threshold (default: 30)
number of points that a scanner must accumulate in talker-sliding-window time range. set to 0 to disable this type of alert.
13. scanner-fixed-window (default: 15)
how many seconds we should go before resetting the fixed scanner score.
14. talker-fixed-window (default: 30)
how many seconds we should go before resetting the fixed talker score.
15. scanner-sliding-window (default: 20)
how many seconds we should go before resetting the fixed scanner score.
16. talker-sliding-window (default: 30)
how many seconds we should go before resetting the sliding talker score.
17. scanner-sliding-scale-factor (default: 0.5)
how much to increase the sliding window by each time we get a new sliding scanner entry. it's current size + (<scale factor> * current_size).
18. talker-sliding-scale-factor (default: 0.5)
how much to increase the sliding window by each time we get a new sliding talker entry. it's current size + (<scale factor> * current_size).
19. src-ignore-net
the ip list of what source ips to ignore.
20. dst-ignore-net
the ip list of what destination ips to ignore.
21. tcp-penalties (default: on)
if this is enabled, when a new tcp flow enters the portscan detection set, check the tcp flags for non-standard session initiators and assign penalty points for odd combinations such as syn+fin
22. flag mapping
23. server-watchnet
the ip list of what machines to learn services on. busy servers should be placed here to help the portscan detector learn what services are requested on the network.

Table 2.5: flag mapping

syn or syn+ecn bits	base_score (defaults to 1 point)
syn+fin+th_ack and anything else	5 points
syn+fin and anything else without ack	3 points
anything else	2 points

24. server-learning-time (default: 28800)

how many seconds we should keep increment hitcounts of services on ips in the server-watchnet

this does not perform validation that the service is connected correctly. it is possible while learning that someone floods the table with unique connections, causing something to become a service that you do not wish to be a service. it's generally assumed that the learning time will occur at a time where traffic is "typical". future versions of snort should allow this state to be saved and modifiable. if this caveat is a concern in your environment, do not set a server watchnet and rely only on talker scores.

25. server-ignore-limit (default: 500)

how many requests a port on an ip in the server-watchnet must see before it is ignored for the purposes of portscans.

26. server-scanner-limit (default: 500)

how many requests a port on an ip in the server-watchnet must see before it is treated as a talker rather than a scanner. this is a minimum number of requests that must be seen during the server-learning-time for the flow to be treated as a talker connection rather than as a scanner connection.

27. alert-mode (default: once)

Table 2.6: alert modes

once	alert only on the first time we get a scan entry hit. this dramatically reduces clutter because the scan alert in the first place tells one to look for other event types.
all	alert each time the score increases beyond a threshold

28. output-mode (default: msg)

Table 2.7: output modes

msg	a variable text message with the scores included
pktkludge	generate a fake packet and use the logging output system

29. dumpall when snort is exiting, dump the entire contents of the server table, the uniqueness tracker table, and the scoreboard entries. this is useful if you suspect an underlying bug in the algorithms used or if you would just like to see what it has learned. set this to "1" to enable.**30.** base-score (default: 1) default score for a new connection. this is probably only useful for debugging.

example configuration

```
preprocessor flow-portscan: server-watchnet [10.0.0.0/8] \  
                           unique-memcap 5000000 \  
                           unique-rows 50000 \  
                           tcp-penalties on \  
                           server-scanner-limit 50 \  
                           alert-mode all \  
                           output-mode msg \  
                           server-learning-time 3600
```

2.1.7 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}.

it defaults to running on ports 21, 23, 25, and 119.

format

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

2.1.8 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 the packet buffer. if `stream4` is enabled, it will only process client side traffic. it defaults to running on ports 111 and 32771.

Table 2.8: rpc decoder options

option	purpose
<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]
```

2.1.9 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. the default statistics that are processed are snort's real-time statistics. this includes:

1. packets received
2. packets dropped
3. % packets dropped
4. packets received
5. kpackets per second
6. average bytes per packets
7. mbits per second (wire)
8. mbits per second (rebuilt) [this is the average mbits that snort injects after rebuilding packets]
9. mbits per second (total)
10. pattern matching percent [the average percent of data received that snort processes in pattern matching]
11. cpu usage (user time) (system time) (idle time)
12. alerts per second
13. syn packets per second
14. syn/ack packet per second
15. new sessions per second
16. deleted sessions per second
17. total sessions
18. max sessions during time interval
19. stream flushes per second
20. stream faults per second
21. stream timeouts
22. frag completes per second
23. frag inserts per second
24. frag deletes per second
25. frag flushes per second
26. frag timeouts

27. frag faults

when the keyword "flow" is enabled, statistics are printed out about the type of traffic and protocol distributions that snort is seeing. this option can produce large amounts of output.

the keyword "events" turns on event reporting. this prints out statistics as to the number of signatures that were matched by the setwise pattern matcher and the number of those matches that were verified with the signature flags. we call these non-qualified and qualified events. it shows the user if there is a problem with the ruleset that they are running.

the keyword "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.

the keyword "console" prints statistics at the console, this is on by default.

the keyword "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.

the keyword "pktcnt" adjusts the number of pkts 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.

the keyword "time" represents the number of seconds between intervals.

examples

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

2.1.10 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. this always needs to be specified in the global configuration, otherwise you get an error. this map file should live where you keep your snort.conf (or provide a fully qualified path name to it).

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 is available. you can generate your own unicode maps by using the program ms_unicode_generator.c located in the snort contrib directory. 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 go to. in the future we want to limit this to particular 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 the types of attacks regardless of the http server. "profile all" sets the following configuration options.

Table 2.9: profile "all" options

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
looking for null bytes in url	on, alert on
multiple slash	on, alert off
directory normalization	on, alert off
apache whitespace	on
double decoding	on
%u decoding	on
bare byte decoding	on
iis unicode codepoints	alert on
iis backslash	on, alert off
iis delimiter	on

1-B. apache

the "apache" profile is used for apache web servers. this differs from the 'iis' profile by only excepting utf-8 standard unicode encoding and not excepting backslashes as legitimate slashes, like iis does. apache also excepts tabs as whitespace. "profile apache" sets the following configuration options:

Table 2.10: profile "apache" options

flow_depth	300
chunk encoding	alert on chunks larger than 500000 bytes
ascii decoding	on alert off
looking for null bytes in url	on, alert on
multiple slash	on, alert off
directory normalization	on, alert off
apache whitespace	on, alert on
utf_8 encoding	on, alert off
non_strict url parsing	on

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 following configuration options:

Table 2.11: profile "iis" options

flow_depth	300
iis_unicode_map	codepoint map in the global configuration
ascii decoding	on, alert off
multiple slash	on, alert off
directory normalization	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 on
apache whitespace	on, alert on

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

1. ports
2. iis_unicode_map
3. allow_proxy_use
4. flow_depth
5. no_alerts
6. inspect_uri_only
7. 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 what ports to decode on the http server. encrypted traffic (ssl) cannot be decoded, so adding ports 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 in `snort 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 you 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, that we don't really have rules for anyway. most of the http server rules that we do have are for the http header and a few bytes after that, so we can catch those alerts by specifying a `flow_depth` of about 150 - 300. mileage may vary.

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 to not enable `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 an 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 chars, 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 char 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 chars as valid values in 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 encoded 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, so is ascii encoding 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.

users should alert on the iis_unicode option, because it is seen mainly in attacks and evasion attempts. when iis_unicode is enabled, so is ascii and utf-8 decoding to enforce correct decoding. to alert on utf-8 decoding, the user 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 decoded in the second stage. anyway, this is really complex and adds tons of different encodings for one char. when double_decode is enabled, so is ascii to enforce correct decoding.

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

this option let's 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 give an 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 a no.

14. iis_backslash <yes|no>

normalize 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 a user wants to configure an alert, then specify "yes", otherwise "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 non-rfc standard of tab for a space delimiter. apache uses this, so if the emulated web server is apache you need to 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 chars. 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, then 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 less 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.

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 \  
    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 }
```

2.2 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:

1. **limit**
alert on the 1st *m* events during the time interval, then ignore events for the rest of the time interval.
2. **threshold**
alert every *m* times we see this event during the time interval.
3. **both**
alert once per time interval after seeing *m* occurrences of the event, then ignore 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.

you may apply only one threshold to any given sid. if you try to apply more than one threshold command to a sid, snort will terminate while reading the configuration information.

2.2.1 standalone options

this format supports 6 threshold options - all are required.

Table 2.12: standalone options

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>

2.2.2 standalone format

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

2.2.3 rule keyword format

this format supports 4 threshold options - all are required.

Table 2.13: rule keyword options

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

2.2.4 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 1st to an event, if they do not block a rule from being logged then the global thresholding test is applied - thresholds in a rule will override a global 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 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.

2.2.5 examples

standalone thresholds

limit to logging 1 event per 60 seconds:

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

limit to logging every 3rd event:

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

limit to logging 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 1st 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 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>
```

2.3 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.

2.3.1 format

the suppress command supports either 2 or 4 options

Table 2.14: suppression options

gen_id	<generator id>	required
sig_id	<snort signature id>	required
track	by_src or by_dst	optional, requires ip
ip	ip[/mask]	optional, requires track

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

2.3.2 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
```


2.4 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

2.4.1 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>



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

2.4.2 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

```
output alert_fast: alert.fast
```

Figure 2.9: fast alert configuration

format

```
alert_fast: <output filename>
```

2.4.3 alert_full

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 per ip will be created. 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

2.4.4 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

2.4.5 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 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

2.4.6 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 loose 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

2.4.7 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

2.4.8 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

2.4.9 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 2.16: 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 is the rule options. The words before the colons in the rule options section are called option keywords.



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.

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

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[?] 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.

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

```
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

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.

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

```
activate tcp !$HOME_NET any -> $HOME_NET 143 (flags: PA; \
  content: "|E8C0FFFFFF|/bin"; activates: 1; \
  msg: "IMAP buffer overflow!");)
dynamic tcp !$HOME_NET any -> $HOME_NET 143 (activated_by: 1; count: 50;)
```

Figure 3.7: Activate/Dynamic Rule Example

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/snort-db/> <http://www.snort.org/snort-db/> 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



The following characters must be escaped inside a content rule:

`: ; \ "`

Format

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

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`

Example

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.

```
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"; depth:3; nocase;)
```

Figure 3.12: Negation Example

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 3.5.1option.

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 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 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 3.5.4), 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/aaaaaaaaaaaaaaaaaaaaaaaaaa/..%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 ??.

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. ^ and \$ match at the beginning and ending of the string. When m is set, ^ and \$ 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 ^)
E	Set \$ to match only at the end of the subject string. Without E, \$ 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 "?".

Table 3.5: Snort specific modifiers

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

The modifiers R and B should not be used together.

Example

This example looks for the string BLAH, ignoring the case of BLAH.

```
alert ip any any -> any any (pcree:"/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.

Format

```
byte_test: <bytes_to_convert>, <operator>, <value>, <offset> \
          [, [relative],[big],[little],[string],[hex],[dec],[oct]]
```

bytes_to_convert number of bytes to pick up from the packet

operator operation to perform to test the value (<,>,<=,<!,&)

value value to test the converted value against

offset number of bytes into the payload to start processing

relative use an offset relative to last pattern match

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

```
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

3.5.12 byte_jump

The `byte_jump` option is used to grab some number of bytes, convert them to their numeric representation, jump the `doe_ptr` up that many bytes (for further pattern matching/byte_testing). This will allow relative pattern matches to take into account numerical values found in network data.

Format

```
byte_jump: <bytes_to_convert>, <offset> \
[, [relative],[big],[little],[string],[hex],[dec],[oct],[align]]
```

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

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

```

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 regex

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

3.5.14 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 2.1.4). 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

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??). 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.

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 !$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 3.21: Flow 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 sesp

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:

rst_snd send TCP-RST packets to the sending socket

rst_rcv send TCP-RST packets to the receiving socket

rst_all send TCP_RST packets in both directions

icmp_net send a ICMP_NET_UNREACH to the sender

icmp_host send a ICMP_HOST_UNREACH to the sender

icmp_port send a ICMP_PORT_UNREACH to the sender

icmp_all 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

The `react` keyword based on flexible response (Flex Resp) implements flexible reaction 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 Flex Resp code allows Snort to actively close offending connections and/or send a visible notice to the browser (warn modifier available soon). 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

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 `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 2.4.6, 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 writing good rules

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

3.8.1 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.

3.8.2 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.

3.8.3 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 *pcre* option, looking for user, followed at least one space character (which includes tab), followed by root, ignoring case.

3.8.4 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

3.8.5 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, lets 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 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 06 00 00 00 00 00 00 00 00 .....
00 00 00 00 00 00 00 04 00 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 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 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>

```

lets 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, then 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 0x00000003b626f6200.

```

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, lets put them all together.

```
content:"|00 00 00 00|"; offset:4; depth:4;
content:"|00 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 (lets 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.

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>