SNORT RULE WRITING

Pouyan Zamani

THE BASICS

Most Snort rules are written in **a single line.**

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.

SAMPLE RULE 1

The text up to the first parenthesis is the **rule header.**

The section enclosed in parenthesis contains the **rule options.**

The words before the colons in the rule options section are called option ***keywords.***

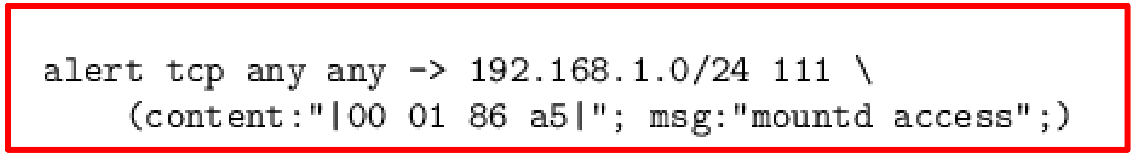

```
alert tcp any any -> 192.168.1.0/24 111 \
    (content:"|00 01 86 a5|"; msg:"mountd access";)
```

When taken together, the elements can be considered to form a **logical AND** statement.

The various rules in a Snort rules library file can be considered to form a large **logical OR** statement.

RULE HEADERS (RULE ACTIONS)

The first item in a rule is the rule action.

- **alert** - generate an alert using the selected alert method, and then log the packet
- **log** - log the packet
- **pass** - ignore the packet
- **drop** - block and log the packet
- **reject** - block the packet, log it, and then send a TCP reset if the protocol is TCP or an ICMP port unreachable message if the protocol is UDP.
- **sdrop** - block the packet but do not log it.

RULE HEADERS (PROTOCOL)

There are four protocols that Snort currently analyzes for suspicious behavior:

- **TCP**
- **UDP**
- **ICMP**
- **IP**

In the future there may be more, such as ARP, IGRP, GRE, OSPF, RIP, IPX, etc.

RULE HEADERS (IP ADDRESS)

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 config file.

The addresses are formed by a straight numeric IP address and a CIDR block.

An IP list is specified by enclosing a ***comma*** separated list of IP addresses and CIDR blocks within ***square brackets***.

IP Operator:

- ! : Not

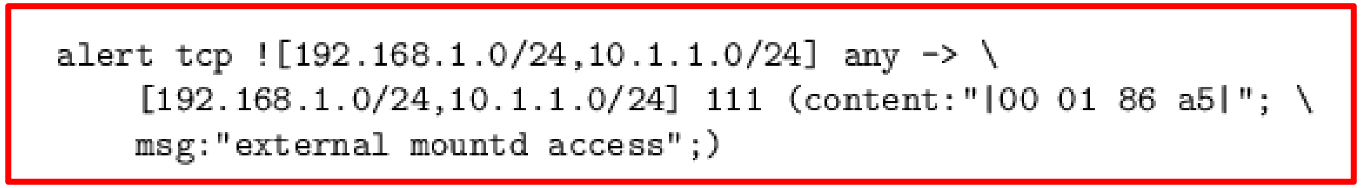

```
alert tcp ![192.168.1.0/24,10.1.1.0/24] any -> \
    [192.168.1.0/24,10.1.1.0/24] 111 (content:"|00 01 86 a5|"; \
    msg:"external mountd access";)
```

RULE HEADERS (PORT NUMBERS)

Any ports are a wildcard value, meaning literally any port.

Static ports are indicated by a single port number.

Port ranges are indicated with the range operator ":"

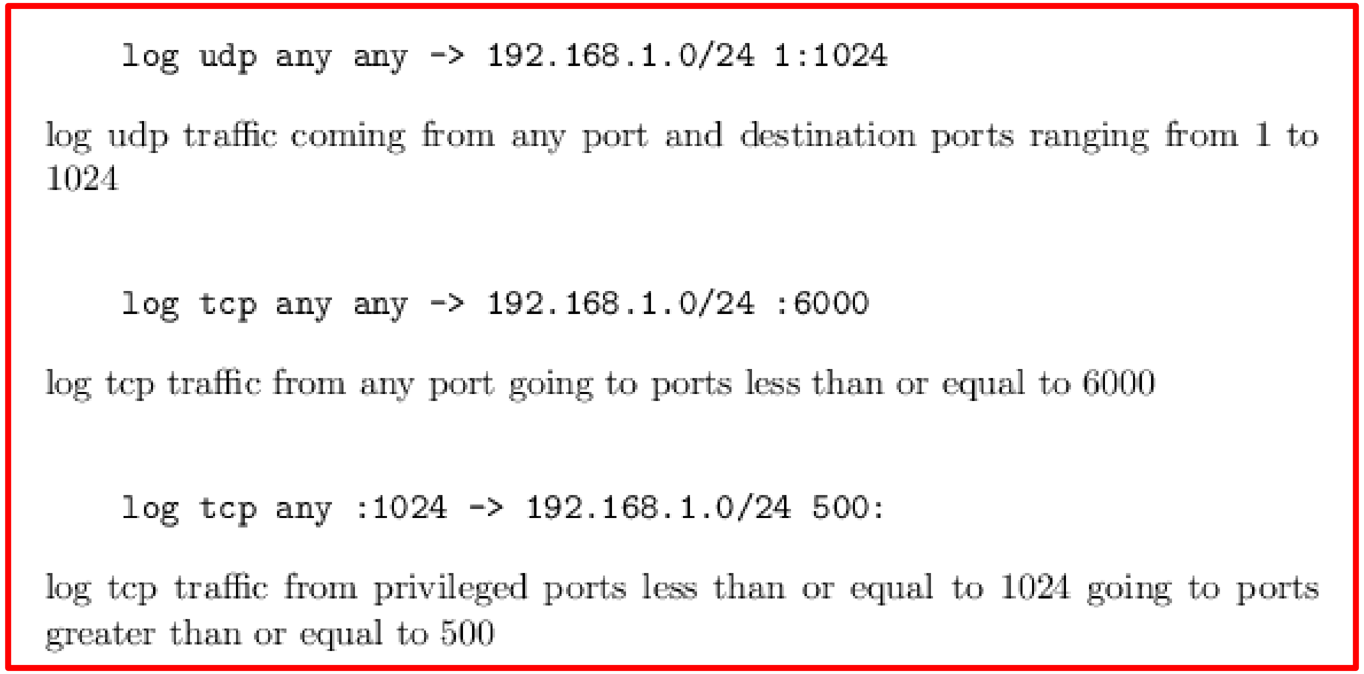

```
log udp any any -> 192.168.1.0/24 1:1024
```

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

RULE HEADERS (DIRECTION OPERATOR)

The direction operator **“->”** indicates the orientation, or direction, of the traffic that the rule applies to.

The IP address and port numbers on the left side of the direction operator is considered to be the traffic coming from the source host.

The IP address and port numbers on the right side of the operator is considered to be the traffic going to the destination host.

There is also a bidirectional operator, which is indicated with a **“<>”** symbol.

Also, note that there is no “<-” operator.

RULE OPTIONS

Rule options form the heart of Snort's intrusion detection engine.

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.

Four major categories of rule options:

- **General**: 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.”

GENERAL RULE OPTIONS

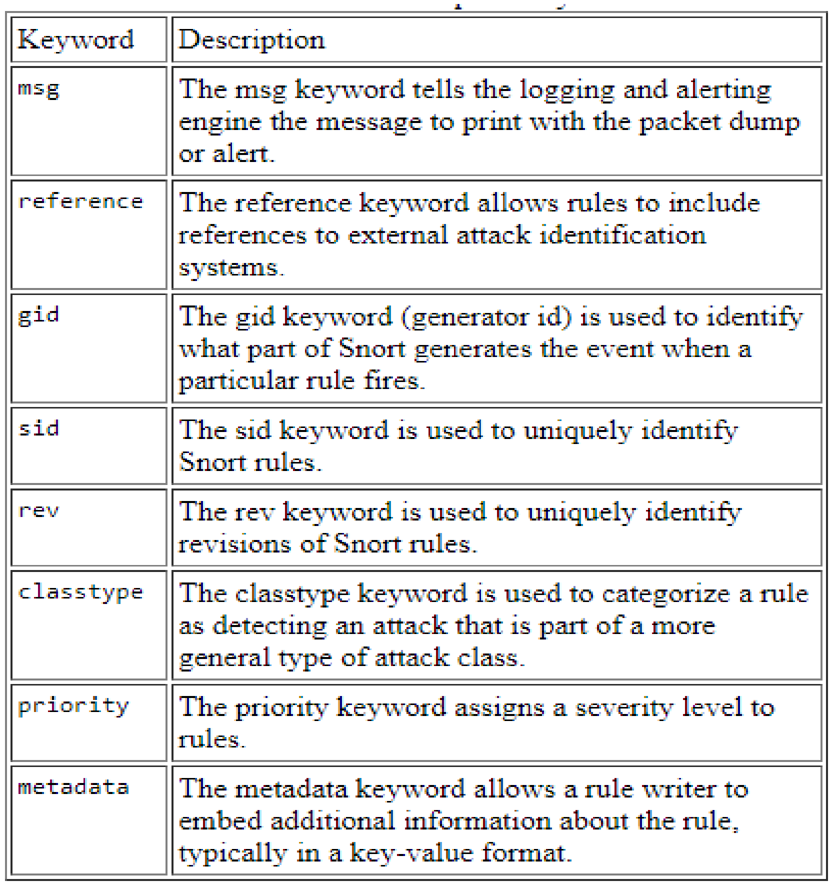

Keyword	Description
msg	The msg keyword tells the logging and alerting engine the message to print with the packet dump or alert.
reference	The reference keyword allows rules to include references to external attack identification systems.
gid	The gid keyword (generator id) is used to identify what part of Snort generates the event when a particular rule fires.
sid	The sid keyword is used to uniquely identify Snort rules.
rev	The rev keyword is used to uniquely identify revisions of Snort rules.
classtype	The classtype keyword is used to categorize a rule as detecting an attack that is part of a more general type of attack class.
priority	The priority keyword assigns a severity level to rules.
metadata	The metadata keyword allows a rule writer to embed additional information about the rule, typically in a key-value format.

HTTP://MANUAL-SNORT-ORG.S3-WEBSITE-US-EAST-1.AMAZONAWS.COM/NODE31.HTML

PAYLOAD DETECTION RULE OPTIONS (CONTENT)

It allows the user to set rules that search for specific content in the packet payload and trigger response based on that data.

Be aware that this test is case sensitive.

The option data for the content keyword can contain mixed text and binary data.

The binary data is generally enclosed within the **pipe** “|” character and represented as bytecode(HEX).

Also note that the following characters must be escaped inside a content rule with a **Slash** “/”:

- ;
- \
- "

PAYLOAD DETECTION RULE OPTIONS (CONTENT)

Offset: Allows the rule writer to specify ***where to start searching*** for a pattern within a packet.

Distance: Allows the rule writer to specify ***how far into a packet*** Snort should ignore before starting to search for the specified pattern relative to the end of the ***previous pattern match.***

Depth: Allows the rule writer to specify ***how far into a packet*** Snort should search for the specified pattern.

Within: A content modifier that makes sure that ***at most N bytes are between pattern matches*** using the ***content*** keyword.

PAYLOAD DETECTION RULE OPTIONS (CONTENT)

Nocase: Allows the rule writer to specify that the Snort should look for the specific pattern, ignoring case.

Fast_Pattern: A content modifier that sets the content within a rule to be used with the fast pattern matcher.

- The default behavior of fast pattern determination is to use the longest HTTP buffer content.
- If no HTTP buffer is present, then the fast pattern is the longest content.
- This option may be specified only once per rule.

PCRE: Allows rules to be written using perl compatible regular expressions.

HTTPS://PERLDOC.PERL.ORG/PERLRE

PAYLOAD DETECTION RULE OPTIONS (CONTENT)

Byte_Test: Test a byte field against a specific value (with operator).

Sample: Byte_Test: 2, &, 64, 4 ;

- How to Read it: at offset 4, consider 2 bytes, & it with 01000000

Byte_Jump: Allows for the ability to select a <num of bytes> from an <offset> and moves the detection pointer to that position.

Sample: Byte_Jump: 2, 4…..

- How to Read it: at offset 4, Skip 2 bytes……

HTTP://MANUAL-SNORT-ORG.S3-WEBSITE-US-EAST-1.AMAZONAWS.COM/NODE32.HTML

NON-PAYLOAD DETECTION RULE OPTIONS

Keyword	Description
fragoffset	allows one to compare the IP fragment offset field against a decimal value.
ttl	used to check the IP time-to-live value.
tos	used to check the IP TOS field for a specific value.
id	used to check the IP ID field for a specific value.
ipopts	used to check if a specific IP option is present.
fragbits	used to check if fragmentation and reserved bits are set in the IP header.
dsize	used to test the packet payload size.
flags	used to check if specific TCP flag bits are present.
flow	allows rules to only apply to certain directions of the traffic flow.
flowbits	allows rules to track states during a transport protocol session.

NON-PAYLOAD DETECTION RULE OPTIONS 2

Keyword	Description
seq	used to check for a specific TCP sequence number.
ack	used to check for a specific TCP acknowledge number.
window	used to check for a specific TCP window size.
itype	used to check for a specific ICMP type value.
icode	used to check for a specific ICMP code value.
icmp_id	used to check for a specific ICMP ID value.
icmp_seq	used to check for a specific ICMP sequence value.
rpc	used to check for a RPC application, version, and procedure numbers in SUNRPC CALL requests.
ip_proto	allows checks against the IP protocol header.
sameip	allows rules to check if the source ip is the same as the destination IP.

POST-DETECTION RULE OPTIONS

Keyword	Description
logto	tells Snort to log all packets that trigger this rule to a special output log file.
session	built to extract user data from TCP Sessions.
resp	used attempt to close sessions when an alert is triggered.
react	implements an ability for users to react to traffic that matches a Snort rule by closing connection and sending a notice.
tag	allow rules to log more than just the single packet that triggered the rule.
replace	Replace the prior matching content with the given string of the same length. Available in inline mode only.
detection_filter	Track by source or destination IP address and if the rule otherwise matches more than the configured rate it will fire

RULE THRESHOLDS

Type

- Limit: alerts on the 1st m events during the time interval, then ignores events for the rest of the time interval.
- Threshold: alerts every m times we see this event during the time interval.
- Both: alerts once per time interval after seeing m occurrences of the event, then ignores any additional events during the time interval.

Track

- By_src: Count is maintained for each unique source IP addresses.
- By_dst: Count is maintained for each unique destination IP addresses

Count c / **Seconds** s : limited number of alerts / limited time period