Rocky Enterprise Linux 9.2 Manual Pages on command 'firewalld.richlanguage.5'

$ man firewalld.richlanguage.5

FIREWALLD.RICHLANG(5) firewalld.richlanguage FIREWALLD.RICHLANG(5)

NAME

firewalld.richlanguage - Rich Language Documentation

DESCRIPTION

With the rich language more complex firewall rules can be created in an easy to understand way. The language uses keywords with values and is an abstract representation of ip*tables rules.

The rich language extends the current zone elements (service, port, icmp-block, icmp-type, masquerade, forward-port and source-port) with additional source and destination addresses, logging, actions and limits for logs and actions.

This page describes the rich language used in the command line client and D-Bus interface. For information about the rich language representation used in the zone configuration files, please have a look at firewalld.zone(5).

A rule is part of a zone. One zone can contain several rules. If some rules interact/contradict, the first rule that matches "wins".

General rule structure

rule

[source]

[destination]

service|port|protocol|icmp-block|icmp-type|masquerade|forward-port|source-port

[log|nflog]

[audit]

[accept|reject|drop|mark]

The complete rule is provided as a single line string. A destination is allowed here as long as it does not conflict with the destination of a service.

Rule structure for source black or white listing

rule

source

[log|nflog]

[audit]

accept|reject|drop|mark

This is used to grant or limit access from a source to this machine or machines that are reachable by this machine. A destination is not allowed here.

Important information about element options: Options for elements in a rule need to be added exactly after the element. If the option is placed somewhere else it might be used for another element as far as it matches the options of the other element or will result in a rule error.

Rule

rule [family="ipv4|ipv6"] [priority="priority"]

If the rule family is provided, it can be either "ipv4" or "ipv6", which limits the rule to IPv4 or IPv6. If the rule family is not provided, the rule will be added for IPv4 and IPv6. If source or destination addresses are used in a rule, then the rule family need to be provided. This is also the case for port/packet forwarding.

If the rule priority is provided, it can be in the range of -32768 to 32767 where lower values have higher precedence. Rich rules are sorted

by priority. Ordering for rules with the same priority value is undefined. A negative priority value will be executed before other firewalld primitives. A positive priority value will be executed after other firewalld primitives. A priority value of 0 will place the rule in a chain based on the action as per the "Information about logging and actions" below.

Source

source [not] address="address[/mask]"|mac="mac-address"|ipset="ipset"

With the source address the origin of a connection attempt can be limited to the source address. An address is either a single IP address, or a network IP address, a MAC address or an IPSet. The address has to match the rule family (IPv4/IPv6). Subnet mask is expressed in either dot-decimal (/x.x.x.x) or prefix (/x) notations for IPv4, and in prefix notation (/x) for IPv6 network addresses. It is possible to invert the sense of an address by adding not before address. All but the specified address will match then.

Destination

destination [not] address="address[/mask]"|ipset="ipset"

With the destination address the target can be limited to the destination address. The destination address is using the same syntax as the source address.

The use of source and destination addresses is optional and the use of a destination addresses is not possible with all elements. This depends on the use of destination addresses for example in service entries.

Service

service name="service name"

The service service name will be added to the rule. The service name is one of the firewalld provided services. To get a list of the supported services, use firewall-cmd --get-services.

If a service provides a destination address, it will conflict with a destination address in the rule and will result in an error. The services using destination addresses internally are mostly services using multicast.

Port

port port="port value" protocol="tcp|udp|sctp|dccp"

The port port value can either be a single port number portid or a port range portid-portid. The protocol can either be tcp, udp, sctp or dccp.

Protocol

protocol value="protocol value"

The protocol value can be either a protocol id number or a protocol name. For allowed protocol entries, please have a look at /etc/protocols.

Tcp-Mss-Clamp

tcp-mss-clamp="value=pmtu|value=number >= 536|None"

The tcp-mss-clamp sets the maximum segment size in the rule.

The tcp-mss-clamp has an optional attribute value can be either be set to "pmtu" or a number greater than or equal to 536. If attribute value is not present then the maximum segment size is automatically set to "pmtu".

ICMP-Block

icmp-block name="icmptype name"

The icmptype is the one of the icmp types firewalld supports. To get a listing of supported icmp types: firewall-cmd --get-icmptypes

It is not allowed to specify an action here. icmp-block uses the action reject internally.

Masquerade

masquerade

Turn on masquerading in the rule. A source and also a destination address can be provided to limit masquerading to this area.

It is not allowed to specify an action here.

Note: IP forwarding will be implicitly enabled.

ICMP-Type

icmp-type name="icmptype name"

The icmptype is the one of the icmp types firewalld supports. To get a listing of supported icmp types: firewall-cmd --get-icmptypes

Forward-Port

forward-port port="port value" protocol="tcp|udp|sctp|dccp" to-port="port value" to-addr="address"

Forward port/packets from local port value with protocol "tcp" or "udp" to either another port locally or to another machine or to another port on another machine.

The port value can either be a single port number or a port range portid-portid. The to-addr is an IP address. The protocol can either be tcp, udp, sctp or dccp.

It is not allowed to specify an action here. forward-port uses the action accept internally.

Note: IP forwarding will be implicitly enabled if to-addr is specified.

Source-Port

source-port port="port value" protocol="tcp|udp|sctp|dccp"

The source-port port value can either be a single port number portid or a port range portid-portid. The protocol can either be tcp, udp, sctp or dccp.

Log

log [prefix="prefix text"] [level="log level"] [limit value="rate/duration"]

Log new connection attempts to the rule with kernel logging for example in syslog. You can define a prefix text with a maximum length of 127 characters that will be added to the log message as a prefix. Log level can be one of "emerg", "alert", "crit", "error", "warning", "notice", "info" or "debug", where default (i.e. if there's no one specified) is "warning". See syslog(3) for description of levels. See Limit section for description of limit tag.

Note: The iptables backend truncates prefix to 29 characters.

NFLog

nflog [group="group id"] [prefix="prefix text"] [queue-size="threshold"] [limit value="rate/duration"]

Log new connection attempts to the rule using kernel logging to pass the packets through a 'netlink' socket to users or applications monitoring the multicast group. The minimum and default value for group is 0, maximum value is 65535. See NETLINK_NETFILTER in netlink(7) man page and NFLOG in both iptables-extensions(8) and nft(8) man pages for a more detailed description.

You can define a prefix text with a maximum length of 127 characters that will be added to the log message as a prefix. The queue-size option can be set to increase the queue threshold which can help limit context switching. The default value for queue-size is 1, maximum value is 65535. See iptables-extensions(8) and nft(8) for more details.

See Limit section for description of limit tag.

Note: The iptables backend truncates prefix to 63 characters.

Audit

audit [limit value="rate/duration"]

Audit provides an alternative way for logging using audit records sent to the service auditd. Audit type will be discovered from the rule action automatically. Use of audit is optional. See Limit section for description of limit tag.

Action

An action can be one of accept, reject, drop or mark.

The rule can either contain an element or also a source only. If the rule contains an element, then new connection matching the element will be handled with the action. If the rule does not contain an element, then everything from the source address will be handled with the action.

accept [limit value="rate/duration"]

reject [type="reject type"] [limit value="rate/duration"]

drop [limit value="rate/duration"]

mark set="mark[/mask]" [limit value="rate/duration"]

With accept all new connection attempts will be granted. With reject they will not be accepted and their source will get a reject ICMP(v6) message. The reject type can be set to specify appropriate ICMP(v6) error message. For valid reject types see --reject-with type in iptables-extensions(8) man page. Because reject types are different for IPv4 and IPv6 you have to specify rule family when using reject type. With drop all packets will be dropped immediately, there is no information sent to the source. With mark all packets will be marked in the PREROUTING chain in the mangle table with the mark and mask

combination. See Limit section for description of limit tag.

Limit

limit value="rate/duration"

It is possible to limit Log, NFLog, Audit and Action. A rule using this tag will match until this limit is reached. The rate is a natural positive number [1, ..] The duration is of "s", "m", "h", "d". "s" means seconds, "m" minutes, "h" hours and "d" days. Maximum limit value is "2/d", which means at maximum two matches per day.

Information about logging and actions

Logging can be done with the log, nflog and audit actions. A new chain is added to all zones: zone_log. This will be jumped into before the deny chain to be able to have a proper ordering.

The rules or parts of them are placed in separate chains according to the priority and action of the rule:

zone_pre

zone_log

zone_deny

zone_allow

zone_post

When priority < 0, the rich rule will be placed in the zone_pre chain.

When priority == 0 Then all logging rules will be placed in the zone_log chain. All reject and drop rules will be placed in the zone_deny chain, which will be walked after the log chain. All accept rules will be placed in the zone_allow chain, which will be walked after the deny chain. If a rule contains log and also deny or allow actions, the parts are placed in the matching chains.

When priority > 0, the rich rule will be placed in the zone_post chain.

EXAMPLES

These are examples of how to specify rich language rules. This format (i.e. one string that specifies whole rule) uses for example firewall-cmd --add-rich-rule (see firewall-cmd(1)) as well as D-Bus interface.

Example 1

Enable new IPv4 and IPv6 connections for protocol 'ah'

rule protocol value="ah" accept

Example 2

Allow new IPv4 and IPv6 connections for service ftp and log 1 per minute using audit

rule service name="ftp" log limit value="1/m" audit accept

Example 3

Allow new IPv4 connections from address 192.168.0.0/24 for service tftp and log 1 per minutes using syslog

rule family="ipv4" source address="192.168.0.0/24" service name="tftp" log prefix="tftp" level="info" limit value="1/m" accept

Example 4

New IPv6 connections from 1:2:3:4:6:: to service radius are all rejected and logged at a rate of 3 per minute. New IPv6 connections from other sources are accepted.

rule family="ipv6" source address="1:2:3:4:6::" service name="radius" log prefix="dns" level="info" limit value="3/m" reject

rule family="ipv6" service name="radius" accept

Example 5

Forward IPv6 port/packets receiving from 1:2:3:4:6:: on port 4011 with protocol tcp to 1::2:3:4:7 on port 4012

rule family="ipv6" source address="1:2:3:4:6::" forward-port to-addr="1::2:3:4:7" to-port="4012" protocol="tcp" port="4011"

Example 6

White-list source address to allow all connections from 192.168.2.2

rule family="ipv4" source address="192.168.2.2" accept

Example 7

Black-list source address to reject all connections from 192.168.2.3

rule family="ipv4" source address="192.168.2.3" reject type="icmp-admin-prohibited"

Example 8

Black-list source address to drop all connections from 192.168.2.4

rule family="ipv4" source address="192.168.2.4" drop

SEE ALSO

firewall-applet(1), firewalld(1), firewall-cmd(1), firewall-config(1), firewalld.conf(5), firewalld.direct(5), firewalld.dbus(5), firewalld.icmptype(5), firewalld.lockdown-whitelist(5), firewall-offline-cmd(1), firewalld.richlanguage(5), firewalld.service(5), firewalld.zone(5), firewalld.zones(5), firewalld.policy(5), firewalld.policies(5), firewalld.ipset(5), firewalld.helper(5)

NOTES

firewalld home page:

http://firewalld.org

More documentation with examples:

http://fedoraproject.org/wiki/FirewallD

AUTHORS

Thomas Woerner <twoerner@redhat.com>

Developer

Jiri Popelka <jpopelka@redhat.com>

Developer

Eric Garver <eric@garver.life>

Developer