

IPTABLES

NAME

iptables --- administration tool for IPv4 packet filtering and NAT

SYNOPSIS

iptables [-t table] {-A|-C|-D} chain rule-specification

iptables [-t table] -I chain [rulenum] rule-specification

iptables [-t table] -R chain rulenum rule-specification

iptables [-t table] -D chain rulenum

iptables [-t table] -S [chain [rulenum]]

iptables [-t table] {-F|-L|-Z} [chain [rulenum]] [options...]

iptables [-t table] -N chain

iptables [-t table] -X [chain]

iptables [-t table] -P chain target

iptables [-t table] -E old-chain-name new-chain-name

rule-specification = [matches...] [target]

`match = -m matchname [per-match-options]`

`target = -j targetname [per-target-options]`

DESCRIPTION

Iptables is used to set up, maintain, and inspect the tables of IPv4 packet filter rules in the Linux kernel. Several different tables may be defined. Each table contains a number of built-in chains and may also contain user-defined chains.

Each chain is a list of rules which can match a set of packets. Each rule specifies what to do with a packet that matches. This is called a 'target', which may be a jump to a user-defined chain in the same table.

TARGETS

A firewall rule specifies criteria for a packet and a target. If the packet does not match, the next rule in the chain is the examined; if it does match, then the next rule is specified by the value of the target, which can be the name of a user-defined chain or one of the special values ACCEPT, DROP, QUEUE or RETURN.

ACCEPT means to let the packet through. DROP means to drop the packet on the floor.

QUEUE means to pass the packet to userspace. (How the packet can be received by a userspace process differs by the particular queue handler. 2.4.x and 2.6.x kernels up to 2.6.13 include the ip_queue queue handler. Kernels 2.6.14 and later additionally include the nfnetlink_queue queue handler. Packets with a target of QUEUE will be sent to queue number '0' in this case. Please also see the NFQUEUE target as described later in this man page.)

RETURN means stop traversing this chain and resume at the next rule in the previous (calling) chain. If the end of a built-in chain is reached or a rule in a built-in chain with target RETURN is matched, the target specified by the chain policy determines the fate of the packet.

TABLES

There are currently five independent tables (which tables are present at any time depends on the kernel configuration options and which modules are present). -t, --table table This option specifies the packet matching table which the command should operate on. If the kernel is configured with automatic module loading, an attempt will be made to load the appropriate module for that table if it is not already there.

The tables are as follows:

filter: This is the default table (if no -t option is passed). It contains the built-in chains INPUT (for packets destined to local sockets), FORWARD (for packets being routed through the box), and OUTPUT (for locally-generated packets). **nat:** This table is consulted when a packet that creates a new connection is encountered. It consists of three built-ins: PREROUTING (for altering packets as soon as they come in), OUTPUT (for altering locally-generated packets before

routing), and POSTROUTING (for altering packets as they are about to go out). **mangle**: This table is used for specialized packet alteration. Until kernel 2.4.17 it had two built-in chains: PREROUTING (for altering incoming packets before routing) and OUTPUT (for altering locally-generated packets before routing). Since kernel 2.4.18, three other built-in chains are also supported: INPUT (for packets coming into the box itself), FORWARD (for altering packets being routed through the box), and POSTROUTING (for altering packets as they are about to go out). **raw**: This table is used mainly for configuring exemptions from connection tracking in combination with the NOTRACK target. It registers at the netfilter hooks with higher priority and is thus called before ip_conntrack, or any other IP tables. It provides the following built-in chains: PREROUTING (for packets arriving via any network interface) OUTPUT (for packets generated by local processes) **security**: This table is used for Mandatory Access Control (MAC) networking rules, such as those enabled by the SECMARK and CONNSECMARK targets. Mandatory Access Control is implemented by Linux Security Modules such as SELinux. The security table is called after the filter table, allowing any Discretionary Access Control (DAC) rules in the filter table to take effect before MAC rules. This table provides the following built-in chains: INPUT (for packets coming into the box itself), OUTPUT (for altering locally-generated packets before routing), and FORWARD (for altering packets being routed through the box). ○

OPTIONS

The options that are recognized by iptables can be divided into several different groups. CO

COMMANDS

These options specify the desired action to perform. Only one of them can be specified on the command line unless otherwise stated below. For long versions of the command and option names, you need to use only enough letters to ensure that iptables can differentiate it from all other options.

- A, --append chain rule-specification** Append one or more rules to the end of the selected chain. When the source and/or destination names resolve to more than one address, a rule will be added for each possible address combination.
- C, --check chain rule-specification** Check whether a rule matching the specification does exist in the selected chain. This command uses the same logic as -D to find a matching entry, but does not alter the existing iptables configuration and uses its exit code to indicate success or failure.
- D, --delete chain rule-specification** Delete one or more rules from the selected chain. There are two versions of this command: the rule can be specified as a number in the chain (starting at 1 for the first rule) or a rule to match.
- I, --insert chain [rulenum] rule-specification** Insert one or more rules in the selected chain as the given rule number. So, if the rule number is 1, the rule or rules are inserted at the head of the chain. This is also the default if no rule number is specified.
- R, --replace chain rulenum rule-specification** Replace a rule in the selected chain. If the source and/or destination names resolve to multiple addresses, the command will fail. Rules are numbered starting at 1.
- L, --list [chain]** List all rules in the selected chain. If no chain is selected, all chains are listed. Like every other iptables command, it applies to the specified table (filter is the default), so NAT rules get listed by `iptables -t nat -n -L`. Please note that it is often used with the -n option, in order to avoid long reverse DNS lookups. It is legal to specify the -Z (zero) option as well, in which case the chain(s) will be atomically listed and zeroed. The exact output is affected by the other arguments given. The exact rules are suppressed until you use `iptables -L -v`.
- S, --list-rules [chain]** Print all rules in the selected chain. If no chain is selected, all chains are printed like `iptables-save`. Like every other iptables command, it applies to the specified table (filter is the default).
- F, --flush**

[chain]Flush the selected chain (all the chains in the table if none is given). This is equivalent to deleting all the rules one by one. -Z, --zero [chain [rulenum]]Zero the packet and byte counters in all chains, or only the given chain, or only the given rule in a chain. It is legal to specify the -L, --list (list) option as well, to see the counters immediately before they are cleared. (See above.) -N, --new-chain chainCreate a new user-defined chain by the given name. There must be no target of that name already. -X, --delete-chain [chain]Delete the optional user-defined chain specified. There must be no references to the chain. If there are, you must delete or replace the referring rules before the chain can be deleted. The chain must be empty, i.e. not contain any rules. If no argument is given, it will attempt to delete every non-builtin chain in the table. -P, --policy chain targetSet the policy for the chain to the given target. See the section TARGETS for the legal targets. Only built-in (non-user-defined) chains can have policies, and neither built-in nor user-defined chains can be policy targets. -E, --rename-chain old-chain new-chainRename the user specified chain to the user supplied name. This is cosmetic, and has no effect on the structure of the table. -hHelp. Give a (currently very brief) description of the command syntax.

PARAMETERS

The following parameters make up a rule specification (as used in the add, delete, insert, replace and append commands). -4, --ipv4This option has no effect in iptables and iptables-restore. -6, --ipv6If a rule using the -6 option is inserted with (and only with) iptables-restore, it will be silently ignored. Any other uses will throw an error. This option allows to put both IPv4 and IPv6 rules in a single rule file for use with both iptables-restore and ip6tables-restore. [!] -p, --protocol protocolThe protocol of the rule or of the packet to check. The specified protocol can be one of tcp, udp, udplite, icmp, esp, ah, sctp or the special keyword "all", or it can be a numeric value, representing one of these protocols or a different one. A protocol name from /etc/protocols is also allowed. A "!" argument before the protocol inverts the test. The number zero is equivalent to all. "all" will match with all protocols and is taken as default when this option is omitted. [!] -s, --source address[/mask][,...]Source specification. Address can be either a network name, a hostname, a network IP address (with /mask), or a plain IP address. Hostnames will be resolved once only, before the rule is submitted to the kernel. Please note that specifying any name to be resolved with a remote query such as DNS is a really bad idea. The mask can be either a network mask or a plain number, specifying the number of 1's at the left side of the network mask. Thus, a mask of 24 is equivalent to 255.255.255.0. A "!" argument before the address specification inverts the sense of the address. The flag --src is an alias for this option. Multiple addresses can be specified, but this will expand to multiple rules (when adding with -A), or will cause multiple rules to be deleted (with -D). [!] -d, --destination address[/mask][,...]Destination specification. See the description of the -s (source) flag for a detailed description of the syntax. The flag --dst is an alias for this option. -m, --match matchSpecifies a match to use, that is, an extension module that tests for a specific property. The set of matches make up the condition under which a target is invoked. Matches are evaluated first to last as specified on the command line and work in short-circuit fashion, i.e. if one extension yields false, evaluation will stop. -j, --jump targetThis specifies the target of the rule; i.e., what to do if the packet matches it. The target can be a user-defined chain (other than the one this rule is in), one of the special builtin targets which decide the fate of the packet immediately, or an extension (see EXTENSIONS below). If this option is omitted in a rule (and -g is not used), then matching the rule will have no effect on the packet's fate, but the counters on the rule will be incremented. -g, --goto chainThis specifies that the processing should continue in a user specified chain. Unlike the --jump option return will not continue

processing in this chain but instead in the chain that called us via `--jump.[!]` `-i`, `--in-interface nameName` of an interface via which a packet was received (only for packets entering the INPUT, FORWARD and PREROUTING chains). When the `"!"` argument is used before the interface name, the sense is inverted. If the interface name ends in a `"+"`, then any interface which begins with this name will match. If this option is omitted, any interface name will match.[!]
`-o`, `--out-interface nameName` of an interface via which a packet is going to be sent (for packets entering the FORWARD, OUTPUT and POSTROUTING chains). When the `"!"` argument is used before the interface name, the sense is inverted. If the interface name ends in a `"+"`, then any interface which begins with this name will match. If this option is omitted, any interface name will match.[!]
`-f`, `--fragment`This means that the rule only refers to second and further fragments of fragmented packets. Since there is no way to tell the source or destination ports of such a packet (or ICMP type), such a packet will not match any rules which specify them. When the `"!"` argument precedes the `"-f"` flag, the rule will only match head fragments, or unfragmented packets.
`-c`, `--set-counters packets bytes`This enables the administrator to initialize the packet and byte counters of a rule (during INSERT, APPEND, REPLACE operations). **OT**

HER OPTIONS

The following additional options can be specified:
`-v`, `--verbose`Verbose output. This option makes the list command show the interface name, the rule options (if any), and the TOS masks. The packet and byte counters are also listed, with the suffix 'K', 'M' or 'G' for 1000, 1,000,000 and 1,000,000,000 multipliers respectively (but see the `-x` flag to change this). For appending, insertion, deletion and replacement, this causes detailed information on the rule or rules to be printed. `-v` may be specified multiple times to possibly emit more detailed debug statements.
`-n`, `--numeric`Numeric output. IP addresses and port numbers will be printed in numeric format. By default, the program will try to display them as host names, network names, or services (whenever applicable).
`-x`, `--exact`Expand numbers. Display the exact value of the packet and byte counters, instead of only the rounded number in K's (multiples of 1000) M's (multiples of 1000K) or G's (multiples of 1000M). This option is only relevant for the `-L` command.
`--line-numbers`When listing rules, add line numbers to the beginning of each rule, corresponding to that rule's position in the chain.
`--modprobe=command`When adding or inserting rules into a chain, use command to load any necessary modules (targets, match extensions, etc).
Sumber : <http://ipset.netfilter.org/>