A look at xtables-addons

Pablo Neira Ayuso <pablo@netfilter.org

GeoIP

geoip match options:

[!] --src-cc, --source-country country[,country...]
Match packet coming from (one of) the specified country(ies)

[!] --dst-cc, --destination-country country[,country…] Match packet going to (one of) the specified country(ies)\n"

NOTE: The country is inputed by its ISO3166 code.

GeoIP (2)

- A few perl scripts:
  - xt_geoip_dl: download and unpack CSV file
  - xt_geoip_build: transform to xt_geoip format

GeoIP (3)

CSV file format:

format:network,geoname_id,registered_country_geoname_id,represented_country_geoname_id,is_anonymous_proxy,is_satellite_provider

1.0.0.0/24,2077456,2077456,,0,0

GeoIP (4)

- Script to autogenerate nftables file:
  - Variables that define location as string (253 countries), eg.
    - define $norway = 3144096
  - Maps
    - Network Address : ISO code
    - ISO code : Continent
    - ISO code : Country name

TARPIT

- “Goal:
  - Allow incoming TCP connections to be established.
  - Passing data should result in the connection being switched to the persist state (0 byte window), in which the remote side stops sending data and asks to continue every 60 seconds.
  - Attempts to shut down the connection should be ignored completely, so the remote side ends up having to time it out.”
- Packet generator to send arbitrary responses?

DELUDE

- Like REJECT…
  - But it replies with SYN+ACK to SYN packets.
- Packet generator to send arbitrary responses?

DHCPMAC

- Matches and mangles dhcp MAC field
- Layer 7 matching and mangling for UDP?

DNETMAP

- 1:1 mapping in IPv4 subnets
- Goal: A single rule can map a private subnet to a shorter public subnet.
- Already supported via nftables maps.

ECHO

- Echo packets back to the origin.
- An example of xtables target according to documentation.

Fuzzy

- rate limit based on a fuzzy logic controller (FLC)

  --lower-limit number
- Specifies the lower limit, in packets per second.

  --upper-limit number
- Specifies the upper limit, also in packets per second.
- “The goal is very similar to that of "limit" match, but using techniques of Fuzzy Control...”

gradm

- Grsecurity RBAC enabled
- Only userspace codebase in xtables-addons tree.

iface

- Match on interface flags, eg. IFF_UP
- Extend nft_meta to support for this.

ipp2p

- ...

ipv4options

- Match in ipv4 options

  --flags [!]symbol[,…]
- Nop
- Security
- lsrr (RFC 791)
- timestamp (RFC 781, 791)
- record-route", /* RFC 791 */
- [9] = "ssrr", /* RFC 791 */
- [11] = "mtu-probe", /* RFC 1063 */
- [12] = "mtu-reply", /* RFC 1063 */
- [18] = "traceroute", /* RFC 1393 */
- [20] = "router-alert", /* RFC 2113 */

length2

- --layer3 Match against layer3 size (e.g. L4 + IPv6 header)\n"
- " --layer4 Match against layer4 size (e.g. L5 + SCTP header)\n" TCP, UDP, UDPLITE, SCTP, DCCP, ICMP, ICMPv6, AH, ESP, IPIP, IPV6-in-IPv6
- " --layer5 Match against layer5 size (e.g. L7 + chunk headers)\n"
- " --layer7 Match against layer7 payload (e.g. SCTP payload)\n"
- "[!] --length n[:n] Match packet length against value or range\n"
- " of values (inclusive)\n"

LOGMARK

- Combines LOG and MARK targets.
- Already supported in nftables.

lscan

- Detects port scans: Could be implemented via ruleset?
- --stealth
- Match if the packet did not belong to any known TCP connection (Stealth/FIN/XMAS/NULL scan).

- --synscan
- Match if the connection was a TCP half-open discovery (SYN scan), i.e. the connection was torn down after the 2nd packet in the 3-way handshake.
- --cnscan
- Match if the connection was a TCP full open discovery (connect scan), i.e. the connection was torn down after completion of the 3-way handshake.
- --grscan
- Match if data in the connection only flew in the direction of the remote side, e.g. if the connection was terminated after a locally running daemon sent its identification. (E.g. openssh, smtp, ftpd.) This may falsely trigger on warranted single-direction data flows, usually bulk data transfers such as FTP DATA connections or IRC DCC. Grab Scan Detection should only be used on ports where a protocol runs that is guaranteed to do a bidirectional exchange of bytes.

psd

- Another portscan detector.
- Derived from Solar Designer's scanlogd
- Hashtable (source address):
  - Timestamp
  - Port
  - Counter
  - Weight (based on priviledged or unpriviledge ports)
- Configurable option (to tune portscan detection heuristics)
  - time threshold
  - Weight theshold

quota2

- Supported by stateful objects (quota)

SYSRQ

- Magic packet to trigger sysrq

condition

- sysctl entry to turn on/off matching on rule
  - /proc/net/nf_condition/Iname