Check Point®
SOFTWARE TECHNOLOGIES LTD.

Workshop in Information Security

Building a Firewall within the Linux Kernel

Implementation Details

Lecturer: Eran Tromer

Teaching assistant: Ariel Haviv

Advisor: Assaf Harel

Our Way to do Things

- Tables:
  - Table row/entry is a `struct`.
  - A table is a bunch of consecutive entries.
  - Each table is `kmalloc`ed, and accessed with `mmap`.
  - Table size in rows/bytes is exposed through sysfs files.
  - Each table is a device, with its own minor#.

- Firewall variables:
  - Setting and getting firewall variables and features are done usinf sysfs.

Rule Base

- A rule base is usually a list of rules.
- Every rule is:
  - Source and destination sockets.
  - A mask to be applied to IPs, to enable handling networks.
  - Verdict of matching packet.
- We will try to match every packet against one of the rules (order matters).
- Usually there will be a default action if no match if found.
  - AKA: Cleanup Rule.

Connection Table

- A connection tables is where we keep track of every active connection. In our case it is TCP.
- Every table row is:
  - Source and destination IPs and ports.
  - TCP state.
  - Expire time of this connection.
- TCP packets will be:
  - Accepted if they conform with the protocol state, or begin a new connection.
  - Dropped otherwise.

Log

- A log is what we show to the user to let him/her see that we are actually doing our job and not randomly accepting and dropping packets.
- A log row is:
  - Source and destination IPs and ports.
  - Time of row modification/creation.
  - The action taken, and why.
  - A counter.
- Usually we will want to group similar decisions and not bloat the log, that’s why we will need a counter.