ctsyncd

past, current and future

Holger Eitzenberger

Sept. 10th 2007

Contents

- Past
- Current Status
- Future Improvements
- Q&A

Past

- kernel implementation
- not in vanilla kernel
- largely unmaintained
- SMP issues
- Astaro Security Gateway V6 (2006)

Current

- idea for userspace implementation about 1st half 2006
- encouragement by Kriztian Kovazs :-)
- using libnetfilter-conntrack (version 0.0.30!)
- first commit in April 7th 2006

Current - Goals

- **copy of kernel implementation in userspace**
  - ACTIVE-PASSIVE
  - possibly multiple passive nodes
- **have a working version until end of 2006 (V7)**
- **performance important**
  - existing ASG V6 user base using kernel implementation
  - first prototype ok
  - sync TCP connections only
- **protocol**
  - TIPC?

Current - Implementation

- **MASTER-SLAVE mode**
- **servant to Astaro HA daemon**
- **simple but effective ctsync protocol between nodes**
  - 8 byte protocol header
- **mcast link-local address**
- **kernel netlink data as payload**
- **sync TCP established connections**
- **hardcoded**
  - most filtering

```
struct cts_msghdr {
    uint8_t ver;
    uint8_t type;
    uint16_t data_len;
    uint32_t seq;
    unsigned char data[];
    /* dword aligned netlink
       data */
};
```

Current - Caching

Caches

- **conntrack refcounted**
- **tuple cache**
  - CT stays there for duration of connection
- **sequence cache**
  - used to check for dead entries
  - shorter lifetime
- **conntrack on either tuple cache or both of them**

Current - Conclusion

- simple & straightforward implementation
- few bugs after all
- performance ok
- stable codebase
- room for improvements

Future

Improvements

- generalize protocol (TCP, UDP, SCTP, ...) handling interface
- implement generic filtering capability on top
- make Netlink part better
  - considered libnl
    - implemented halfway through, stopped
  - libnetfilter-queue update
- which filter model?
  - BPF alike?
  - Netlink is different
  - Kernel-side filtering?
- generic HA daemon interface
  - keepalived, ...
- Publish
  - Bazaar: release early and often
  - GPL

Fin

No Timeline :-)

Q&A