

IPVS Update: v4.1 to date

Netfilter Workshop, Amsterdam

Simon Horman <simon@horms.net>
Horms Solutions Ltd. www.horms.net

June 2016

- Bug fixes with an occasional new feature

- Reset XPS's sender_cpu on forwarding
- Correct crash in sync daemon for v0 of sync protocol with FTP
- Call skb_orphan for forwarding to militate against confusion if bound to a local socket in early_demux
- Fix crash if scheduler is changed resulting from RCU conversion in v3.10
- Do not use uninitialised value for local source address of tunnels, also introduced in v3.10
- Fix panic on unreachable ipv6 next hop
- Various IPv6 fixes

v4.2: Bug Fix of Note

Fix by Alex Gartrell for this gem:

```
unshare -n /bin/bash <<EOF
ip addr add dev lo face::1/128
ipvsadm -A -t [face::1]:15213
ipvsadm -a -t [face::1]:15213 -r b00c::1
echo boom | nc face::1 15213
EOF
```

- Add `sync_maxlen` parameter to sync daemon
- Add multicast support to sync daemon
- Add OVF scheduler

- Avoid illegal access to `inet_sk(skb->sk)` for SYNACK packets attached to request sockets
- netns handling rework
- Add `sysctl` to ignore tunneled packets: Avoid routing loops between IPVSeS
- Scheduling of ICMP and inverse packets

- Compilation fixes

- Correct Call-ID header parsing in SIP PE if it is not the first header
- Allow rescheduling after RST
- Drop first packet to redirect conntrack
- Minor cleanups

- SIP PE enhancements from Marco Angaroni:
 - Update SIP PE to handle outgoing connections
 - Add drop_entry protection to SIP PE
 - Don't update conntrack in OPS mode
 - Optimize release of connections in OPS mode
 - Handle connections started by real-servers

- Count pre-established TCP states as active