

# 2.7.x Networking

David S. Miller  
davem@redhat.com

# SKB DMA Mapping

- Noticed by Anton Blanchard
- SKB has scatterlist, but not a true one from `asm/scatterlist.h`
- Devices map each scatterlist element one by one.
- Slow, especially on IOMMU platforms
- Need to migrate over, or do it all at once
- Consultation with Jeff Garzik needed

# MPLS Stack

- Two implementations exist
- One by MPLS hardware hacker and one by Jamal and myself
- Differences being worked out (netlink for configuration, packet handler design, state machine for MPLS actions, etc.)
- Shooting for 2.7.x integration
- Backporting to 2.6.x is possible

# Unified Flow Cache

- Initially explored during IPSEC development
- One lookup leads to socket, or firewalling rule, or forwarding path
- Tricky issues:
  - things like raw and multicast sockets
  - light SMP synchronization
  - garbage collection
  - keeping rule update costs low
  - Keeping key size from growing out of control

# IPSEC

- Better PMTU handling
  - Have to remember inner packet keys in order to patch up properly when ICMP comes back
- Better tunnel handling
  - Herbert Xu working on this already
- Perhaps async crypto (ask James)
- HW acceleration IPSEC support
- Improve MSS calculation, still a hack

# Misc. and CRAZY ideas

- More Arnaldo style cleanups
- Maybe put sockets in page cache
- Some kind of receive zero-copy mechanism, ask Rusty.
- Integrate CIPE somehow, maybe even a 2.6.x task
- Flesh out rest of IPV6 support, ask Yoshifuji.
- Anyone have other ideas?