

# BSDCan 2011 FreeBSD Developer Summit

Network Stack Working Group Summary

Bjoern A. Zeeb, 2011-05-12

- SMP scalability
- User space network stack
- IPv6
- Future of Firewalls
- The evolving NIC

- Virtual network stack
- DNSsec in base
- net daemons in base
- netstack event notifications
- DDoS

# SMP scalability

- Eliminate cache missing.
- Reaching limit w/o affinity awareness.
- Wait for pcbgroups/RSS being in HEAD before benchmarking/PMC your workloads.
- Only Chelsio driver currently.
- Doing it in software gets about 2/3 performance of hardware.
- Socket API.

# User space netstack

- Tor project, RTEMS, kernel bypass, proxy things, keep kernel clean, regression testing and devel, restart process !machine.
- NetBSD rump: performance concerns.
- Same code base as kernel, symbol munging, LD\_PRELOAD, ... customization.
- See it as a library -> fork/exec/...rt sharing
- Problems wrt. hardware and virtualization.

# IPv6

- "I don't care as long as it's there".
- Pass TAHI & compliance.
- Performance we need to look into.
- Vendors upstreaming some patches.
- Some issues raised - carp, ipfw /32.

# Future of Firewalls

- Once again - future of ipfilter?
- How do we scale pf?
- Drop OpenBSD upstream.
- Users like it, unique features like pfsync.
- Configuration conversion?
- OpenBSD has docs and tutorials we need to continue to leverage.
- Offloading simple things to hardware (from all firewalls).

# The evolving NIC

- Finish offloading/csum for v6 as for v4.
- Offloading not always works out best.
- Eliminate code duplication in drivers.
- Come up with KPIs between drivers and L3.
- Polling vs. interrupts.



# Virtual network stack

- Framework and if cloner changes for 9.0?
- Segregating IP spaces.
- Need better user land APIs to deal with multiple stacks from a single process.
- Middle thing between VNETs and MRT?
- Need to fix MRT (limited to 16 instances).

# DNSSEC in base

- Would need to ship the root key.
- Do not want libc to do all crypto.
- nss\_dnssec?
- A lightweight caching resolver?
- SSH finger print verification out of DNS ...
- People will look into signing freebsd.org.

# net daemons in base

- routed/route6d needs a maintainer.
- What should be in base?  
RIP(ng) vs. OSPF vs. BGP, .. vs. maintaing.
- Some people want mpd in base  
(depends on kernel things).
- RFC on the linux routing sockets?  
RFC3549 (informational)

# Netevent notifications

- rtsock needs a successor. Look at netlink?
- Keep virtualization in mind. Recursion?
- Which notifications do we have and want?  
=> net@
- Can we fix ifIndex?

# DDoS

- Something at L2 or netisr?
- Per protocol base?
- Currently no work ongoing.
- Does direct dispatch handle the problem?