

# BSDCan 2011 FreeBSD Developer Summit

Network Stack Working Group

Bjoern A. Zeeb, 2011-05-11

# > whoarewe

Bjoern A. Zeeb

Robert N. M. Watson

Ryan Stone

George V. Neville-Neil

Philip Paeps

Rui Paulo

Mike Silbersack

Navdeep Parhar

Ermal Luci

Zack Kirsch

Brooks Davis

Hiroki Sato

Sam Leffler

Peter Grehan

Max Laier

Rene Ladan

James Gritton

Hariprasad Shanmugam

Peter Losher

Jeff Prem

Scott Ulrich

Anuranjan Shukla

Andrew Boyer

Pramod Srinivasan

Kristof Provost

Anil Madhavapeddy

Chris Buechler

John Gemignani

Mike Karels

# Overview

90 minutes, you do the maths.

- (1) SMP scalability > 32core, affinity, ... (*Robert Watson*)
- (2) User Space Network Stack (*George Neville-Neil*)
- (3) IPv6 (*Bjoern Zeeb*)
- (4) Future Of Firewalls (*Bjoern Zeeb*)
- (5) The evolving NIC: stack & offload features? (*Bjoern Zeeb*)
- (6) Network Stack Virtualization (*Bjoern Zeeb*)
- (7) DNSSEC integration into base (*Brooks Davis*)
- (8) Network related daemons in base (*Hiroki Sato*)
- (9) Network event notification (*Hiroki Sato*)

# SMP scalability > 32core, affinity, ...

Robert Watson

- (1) models of affinity for sockets, flows, etc.
- (2) integration with network card features, and especially, anticipating forthcoming features
- (3) making sure we take into account highly multithreaded mips vs intel-style differences
- (4) APIs for applications
- (5) per protocol work delivery models  
-- what makes sense?

# User Space Network Stack

George Neville-Neil / Bjoern Zeeb

- (1) Network stack regression testing, Tor, RTEMS, NIC-to-Userspace packet reception.
- (2) Port it, get it compile, symbol munging, LD\_PRELOAD?, simple tap(4) sample, testing
- (3) Customization for various use cases

# IPv6

Bjoern Zeeb

- (1) RFC compliance and regression tests
- (2) Cleanup
- (3) Security
- (4) Performance
- (5) New(er) RFCs and Features
- (6) High-Availability and user space management  
{SLAAC, DHCPv6, carp, ...}

# Future Of Firewalls

Bjoern Zeeb

- (1) Future of pf(4) and FreeBSD?  
{another import, parallel versions, TCP internals, statistics, config stability, config translator, ...}
- (2) What will happen to ipfilter?
- (3) Getting to 10Gbit/s full-featured?  
Offload simple port filters to HW, etc.
- (4) Config file translators.
- (5) ``One real firewall'' (and clang)

# The evolving NIC: what should the stack do with offload features?

Bjoern Zeeb

- (1) Cleaner APIs, less impact on TCP stack?
- (2) Add TSO6, User space TSO, CSUM\_\* and IPv6, cleanup LRO
- (3) Factor out a lot of copy and paste from NIC drivers
- (4) line-by-line review from bottom up
- (5) ifnet, etc. KPI rework for 10.



# Network Stack Virtualization

Bjoern Zeeb

(1) Where we are

(2) What we should get done for 9.0-R.

(3) What do you want / need?

# DNSSec integration into base

Brooks Davis

- (1) \$subject says it all
- (2) APIs.
- (3) Do we need to ship a verifying caching resolver default config?
- (4) Integrate into our install? Implications?
- (5) Consider the future of bind in base along the same time?

# Network related daemons in base

Hiroki Sato

- (1) mpd in base
- (2) routed route6d outdated, lack of features, partly broken
- (3) Cases for small CPE builders

# Network Event Notification

Hiroki Sato

- (1) rtsock, devctl present
- (2) Some events duplicated, some limited
- (3) NIC drivers might need to implement some one-by-one
- (4) What kind of events do we have and need?
- (5) Which mechanism (rtsock vs. devctl) should we use for those?
- (6) Do we need another mechanism? netlink?