



DPDK, VPP & pfSense 3.0

Jim Thompson

DPDK Summit Userspace - Dublin- 2017



%whoami



- ❑ Co-founded company in 1992
 - ❑ Focused on network security
- ❑ Originally “Netgate” was name an open source firewall for SunOS 4
 - ❑ Ported to Solaris (2.3), Windows NT, BSDi
 - ❑ Added IPsec (then quite new)
 - ❑ Wrote CiscoSecure (TACACS+ / RADIUS server)
- ❑ Might know us best for DPDK in a Box


DPDK in a Box



- ❑ Simple, easy intro to running DPDK
- ❑ Minnowboard “Turbot” w/ 4C E3845 Atom, 2GB RAM, 32GB M.2 SSD, 2 x i210 Ethernet
- ❑ CentOS + pre-compiled / installed DPDK + source code + testpmd



DPDK in a Box, next-gen



%whoami



- ❓ Co-founded company in 1992
- ❓ Originally “Netgate” was name an open source firewall for SunOS 4
 - ❓ Also ported to Solaris (2.3), Windows NT, BSDi
 - ❓ Also wrote CiscoSecure (TACACS+ server)
- ❓ Might know us best for DPDK in a Box
- ❓ **We are the company behind the pfSense project**

- ❑ FreeBSD-based router/firewall distribution
 - ❑ “Like a Cisco ASA”
 - ❑ 100,000 lines of PHP
 - ❑ Orchestration + WebGUI
- ❑ Started in 2004
- ❑ First release 4 October 2006
- ❑ 2.4 release 2 October 2017
- ❑ 1.375 million **current** installs
 - ❑ 960,000+ on 64-bit Intel
 - ❑ 8,000 32-bit ARM
 - ❑ ARM64 underway

- ❓ pfSense under strain after 10 years
 - ❓ Increasing use of 10Gbps and higher Ethernet
 - ❓ Increasing packet rate requirements
 - ❓ No central 'status' db, so most config changes require restarting stack
 - ❓ No API — makes automated test & interfacing to orchestration difficult
- ❓ Offsite meeting in 2014
 - ❓ Simple goals
 - ❓ 10Gbps IP4/IP6 forwarding tinygrams, with ACLs
 - ❓ 10Gbps IPsec (large frames)
 - ❓ Eliminate PHP
 - ❓ Add API
 - ❓ Project named

Pennybacker



Experiments in kernel-bypass



- ❓ Libuinet
 - ❓ Integrated IPsec
- ❓ Netmap-fw
 - ❓ August 2015
 - ❓ FreeBSD ip_forward() over netmap
 - ❓ DXR for IPv4 routing
 - ❓ 5Mpps per core (Haswell), but single-threaded (netmap bugs)
- ❓ Packet-journey
 - ❓ October 2015
 - ❓ Native DPDK app
 - ❓ Has ACLs, No IPsec
 - ❓ 20Mpps (6C Sandy Bridge)
 - ❓ We integrated Poptrie (not much gain)
 - ❓ We were adding IPsec, when ...

- ❓ VPP open sourced by Cisco February 2016
 - ❓ Went snowboarding in Japan in March (AsiaBSDCon)
- ❓ Much better performance
 - ❓ Vector processing —> Instruction cache remains hot
 - ❓ First packet pays toll, second is cheap
 - ❓ Production-grade
 - ❓ Has 'C' API, tons of instrumentation, tons of features
- ❓ Changed direction!
- ❓ Contributed back ;-)

Pennybacker components



- ❑ DPDK
- ❑ [FD.io](#)'s VPP
 - ❑ IPv4/V6 interfaces w/VLANs, static & dynamic routing, ACLS, NAT, DHCP client, IPsec
- ❑ Clixon
 - ❑ Generated NETCONF / RESTCONF / CLI
 - ❑ <http://www.clixon.org>
- ❑ Strongswan
 - ❑ IKEv1/IKEv2
- ❑ FRRouting
 - ❑ Formerly Qagga
- ❑ Glue code
- ❑ Test harness
 - ❑ Tested to 40Gbps IPsec on i7-6950x w/ 8955 QAT

Pennybacker



- ❑ Launch on AWS EC2 end of October
 - ❑ Compatible with “VPC Transit”
- ❑ C3000 / Xeon HW w/QAT November
- ❑ Azure “soon”

Next-steps



- ❑ Test to 100Gbps IPsec.
- ❑ Fill-out Clixon for VPP components
 - ❑ GRE, VXLAN, PPPoE, MPLS, L2 features
- ❑ NAPALM
- ❑ WebGUI
- ❑ OpenVPN implementation
- ❑ ARM64
 - ❑ NXP LS1043A port underway
 - ❑ Cavium CN8100 hw on my desk next week
- ❑ ...

Questions?

Jim Thompson

jim@netgate.com