

This one goes up to 11!  
(11.1 actually)

# FreeBSD for networks

---

MUHAMMAD MOINUR RAHMAN  
BOFH@FREEBSD.ORG

# Who am I?

---

Ports developer and conference hopper repeat-offender

Consultant –

- Network Systems
- Large scale FreeBSD deployments
- Professional paranoid

# What is FreeBSD?

---

Complete Operating System

Tools and source code

More than 24,000 3rd party open source software packages

Complete documentation

An open source community



# Who uses FreeBSD?

---

NetApp

Dell/EMC/Isilon

Dell/KACE

Panasas

Apple

Limelight Networks

Swisscom

Sentex

Microsoft

WhatsApp

Juniper Networks

Verisign

Perseus Telecom

Sony

XipLink

McAfee

NYI

Yahoo

# Why use FreeBSD?

---

Innovation

Great tools

Mature release model

Excellent documentation in many languages

○ [https://www.freebsd.org/doc/zh\\_CN/books/handbook/](https://www.freebsd.org/doc/zh_CN/books/handbook/)

Business friendly licence

Open community

# Produce a whole system

---

Operating system

Device drivers

Compilers and associated tools

Debugging tools

Editors

Packaging system

Ready for coding when install is done

# Changes in (recent?) years

---

FreeBSD 11.1 (2017) is not FreeBSD 4.11 (2005)!

- New package manager: `pkg(8)`
- Easy to use package building tool: `poudriere(8)`
- Binary system updates: `freebsd-update(8)`
- Many performance improvements (SMP, jemalloc, etc...)
- Many new features (ZFS, Capsicum, pf, etc...)
- Many improvements to old favourites (`jail(8)`, `rc.conf(5)`, etc...)

# Improvements to filesystems

---

FreeBSD now includes two very mature and time-proven filesystems

## UFS

- Traditional Unix filesystem
- High performance
- Snapshots
- Journalized Soft Updates

## ZFS

- Zetabyte File System (originally from Sun)
- Filesystem and volume manager
- RAID (many options)
- Fully up to date in FreeBSD!



# Jails

---

Light-weight virtualisation: run multiple tenants on a single kernel

- Separate filesystem namespace
- ZFS delegation features
- VIMAGE network stacks

# Jail use cases

---

- Web-based virtual hosting
- Email hosting
- Service isolation with micro-services

# The FreeBSD network stack

---

TCP/IP was originally developed on BSD and FreeBSD.

FreeBSD is still the reference implementation for many network protocols.

- Full support for IPv4 and IPv6
- Active development on TCP with pluggable congestion control
- Reference implementation of SCTP

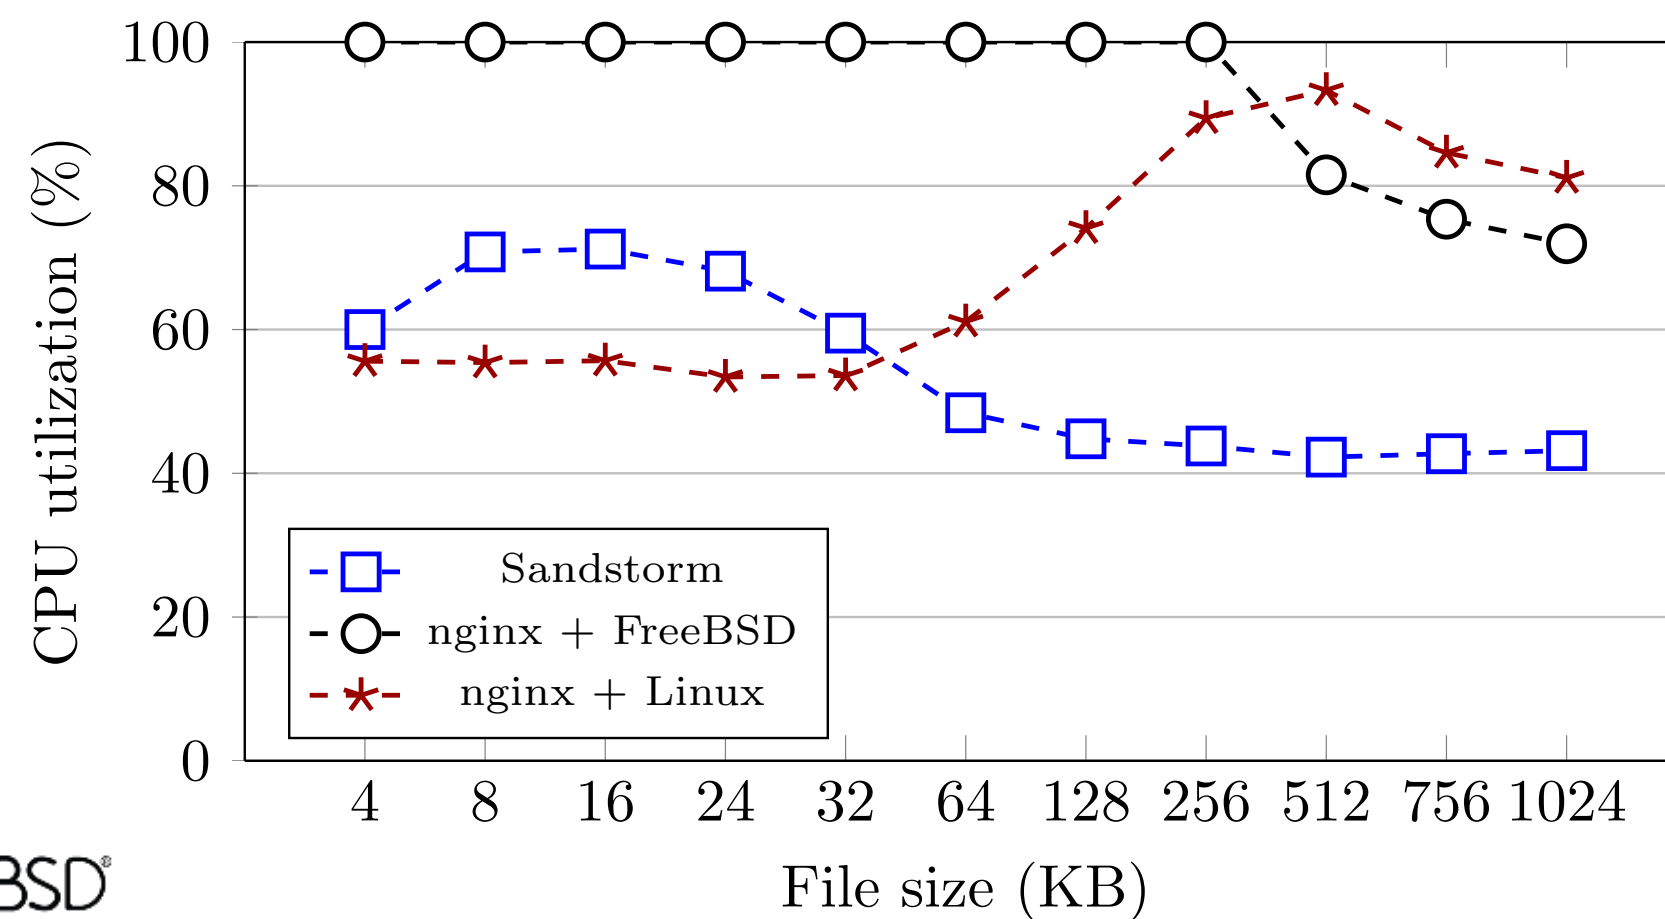
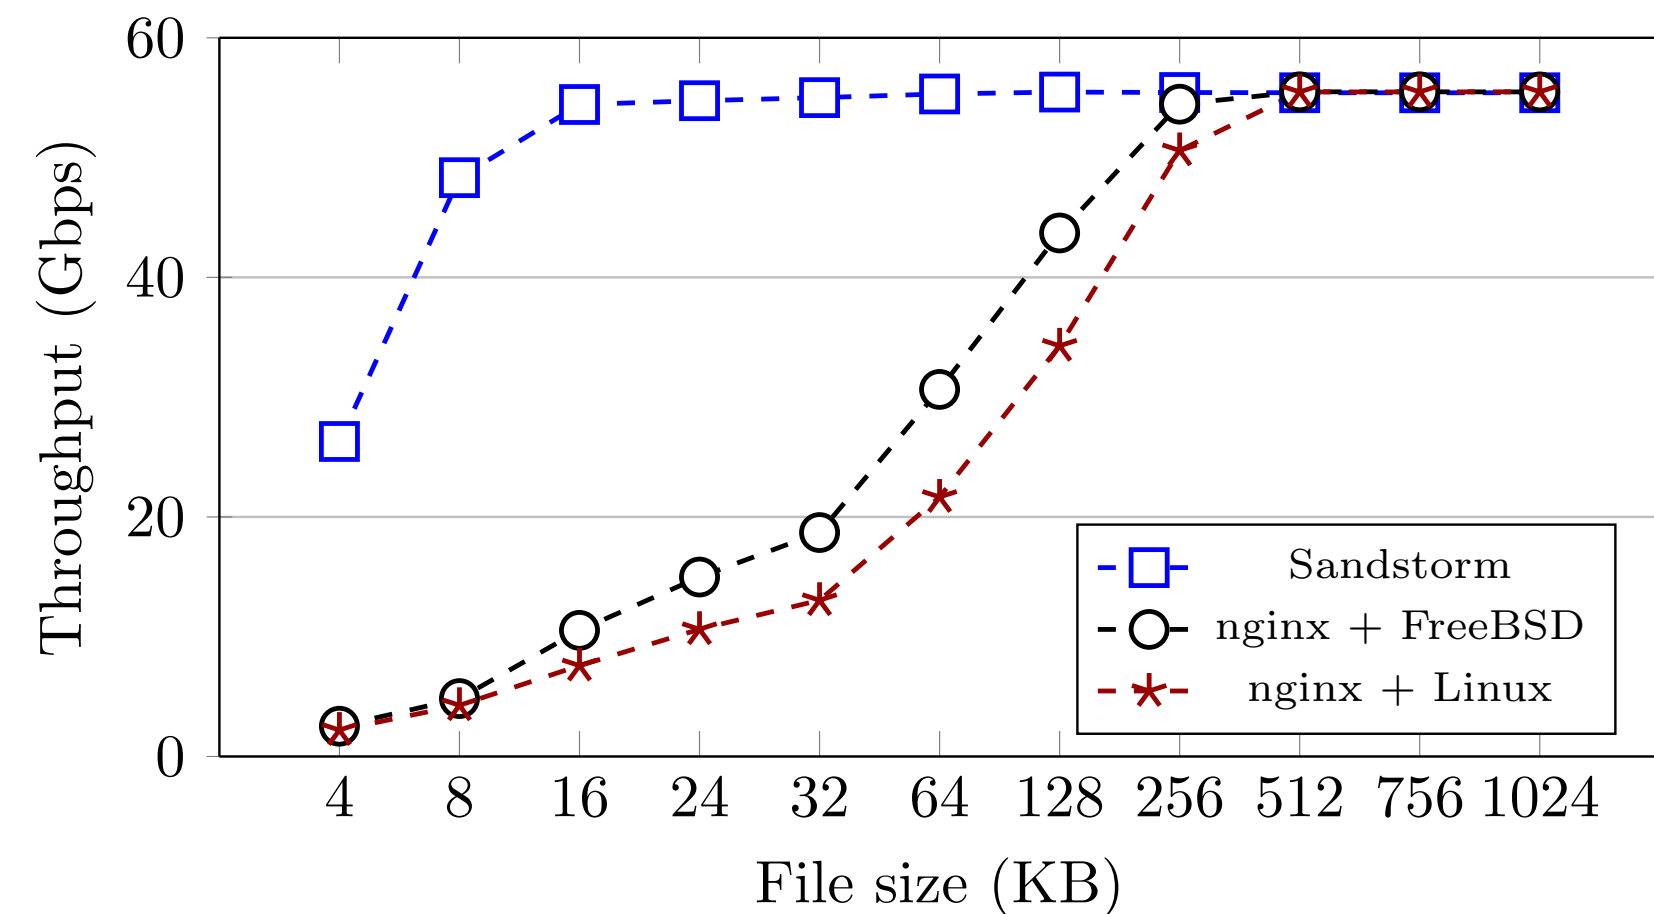
# Pluggable TCP stacks

---

Your choice of congestion control:

- BBR (in `-CURRENT` ... coming to 11.x **Soon™**)
- RACK
- CUBIC
- NewReno

# Performance improvements in networking



- 30 years since the network-stack design developed
- Massive changes in architecture, micro-architecture, memory...
  - Optimising compilers
  - Cache-centered CPUs
  - Multiprocessing, NUMA
  - DMA, multiqueue
  - 10 Gigabit/s Ethernet
- Performance lost to 'generality' throughout stack
- Revisit fundamentals through clean-slate stack
- Orders-of-magnitude performance gains

# Performance improvements in networking

Year	Version	Feature
1983	4.2BSD	BSD sockets, TCP/IP implementation
1986	4.3BSD	VJ/Karels congestion control
1999	FreeBSD 3.1	sendfile(2)
2000	FreeBSD 4.2	TCP accept filters
2001	FreeBSD 4.4	TCP ISN randomisation
2002	FreeBSD 4.5	TCP SYN cache/cookies
2003	FreeBSD 5.0-5.1	IPv6, TCP TIMEWAIT state reduction
2004	FreeBSD 5.2-5.3	TCP host cache, SACK, fine-grained locking
2008	FreeBSD 6.3	TCP LRO, TSO
2008	FreeBSD 7.0	T/TCP removed, socket-buffer autosizing
2009	FreeBSD 7.1	Read-write locking, full TCP offload (TOE)
2009	FreeBSD 8.0	TCP ECN
2012	FreeBSD 9.0	Pluggable TCP congestion control, connection groups

# Active transport community

---

FreeBSD network stack developers are active members of the transport community.

- Developing and testing new congestion control algorithms
- Performance improvements on different workloads
- Tie-ins with security folks (bump in the wire / line-rate encryption)
- Some work on various multi-path TCP implementations

# Firewalls

---

- IPFW: "native" FreeBSD firewall
- pf: fork of the OpenBSD packet filter
- ipfilter: for fans of legacy firewalls

All three firewalls are well-documented in the FreeBSD Handbook and online manual pages included with the operating system.



# More networking

---

- Multi-IP jails (IPv4 and IPv6)
- VIMAGE for multi-tenant routers
- Your choice of firewalls: ipfw, pf, (ipfilter)
- Multiple FIBs for complex routing
- Zebra, Quagga, BIRD, OpenBGPD, OpenOSPFd packages

# VIMAGE

---

- Multiple network stacks for multi-tenant systems
- Combine with jails for very light-weight virtualisation
- Each VIMAGE jail gets (among other things):
  - Choice of firewall
  - Multiple FIBs
  - All the security features of jails

# Even more networking

---

- IPSEC, IKEv2, etc ...
- Layer 2: bridge (dot1d, dot1q), lagg, vlans, spanning tree
- Very nearly working MSTP support (\*)
- Very active "transport community"

# Virtualisation

---

## Ready to use images

- VMWare
- Virtual Box
- qemu
- HyperV

## bhyve

- Native hypervisor
- Runs Linux, Windows and FreeBSD images
- Also used on Mac OS (xhyve)
- BSD Licensed

# Other security features

---

In addition to jails, FreeBSD sports many other exciting security features

## MAC and Audit frameworks

- Who did what and when?
- Much more in-depth than merely logging
- Send audit trails to remote machines

## Capsicum

- Better than privilege separation
- Capabilities for UNIX
- Sandboxing

# System call translation

---

"Linux personality disorder" / "Linuxolator"

Natively run a substantial subset of Linux ELF binaries

Often runs Linux binaries faster than Linux [\*]

Use cases: not invented here binaries for Linux, databases, CAD tools,...

Known to work: Oracle, Eagle CAD, Mentor, many others!

[\*] Usual disclaimers apply. Batteries not included. May contain traces of nuts. Etc.

# Some highlights of 11.1-RELEASE

---

- Many improvements to ZFS
- Broadcom Wi-Fi driver improvements
- bhyve features for ARMv7
- Ported bhyve to ARMv8

# You too can join the FreeBSD community!

---

Join the mailing lists

Clone or checkout the code

- [svn.freebsd.org](http://svn.freebsd.org)
- [github/freebsd](https://github.com/freebsd)

Submit patches

- [reviews.freebsd.org](http://reviews.freebsd.org)

Get a mentor

Get proposed to core@

Granted a commit bit (all commits ReviewedBy)

Be freed from mentorship

Find a mentee



# Learn more about FreeBSD

---

Website: [www.freebsd.org](http://www.freebsd.org)

FreeBSD Foundation: [www.freebsd.foundation.org](http://www.freebsd.foundation.org)

GitHub: [github.com/freebsd](https://github.com/freebsd)

Mailing Lists

Forums

FreeBSD Handbook

IRC

