



*The FreeBSD build option, OpenZFS, bhyve,  
compat\_linux, and jail.conf.d nexus*

*A powerful whole that is far greater than the sum of its parts*

**Michael Dexter**

**@michaeldexter**  
**editor@callfortesting.org**

**EuroBSDcon 2022**

# Thank you for your Patience

- Been describing building blocks since EuroBSDcon 2008
- 49.95 years old and it's better late than never
- I trust you love BSD Unix and want to *share* that love
- Many of these themes came up at this event!
- Shifting focus appropriately...

# I wish to make five simple points

- The Importance Standards and Compliance
- The Power of ELI5
- Owning the Stack, Virtualizing the REST
- The Venn Diagram from Hell
- BSD to the rescue thanks to new features

# Importance of Standards and Compliance

*Putting on a  
different hat...*

# Importance of Standards and Compliance



- Data Protection Best Practices White Paper
- Help the Standards sausage get made
- Promote ZFS at the Software Developers Summit

# SNIA TWG Software Development Guidelines

“15.2 Policies and Guidelines for all software development within TWGs (Technical Working Groups)

6. In general, commercially friendly software licenses, such as BSD, are preferred as inbound and outbound licenses.”



**SDC** 

STORAGE DEVELOPER CONFERENCE

SNIA  SANTA CLARA, 2017

# **By the Book: Open Source Reference Implementations for Key SNIA Terminology**

**Michael Dexter  
Gainframe, SNIA DPCO**

# Importance of Standards

- SNIA Standards
  - LTFS
  - NDMP
  - Redfish/Swordfish...
- RFCs/Open Standards
  - NFS
  - iSCSI
  - Fiber Channel...



# Importance of Standards: Get Involved!

- Fun with UEFI/GPT Partitioning
  - Any users?
  - Great computer science refresher
  - “Reserved” is your friend
- Standards need your help
  - Typos! Poor formatting!
  - Ambiguous statements!
  - Questionable implementations!

# Importance of Standards Validation

- Standards range from ad hoc to strict
- Validation tools range from zero cost to expensive
- Ideally a 1:1 relationship exists between them
- Continuous validation is essential and never enough
- NVMe validation = “Price of a midsize car every year”
- Great opportunity for the Foundations

# Importance of Compliance

## ISO/IEC DIS 27040

Information technology – Security techniques – Storage security

- Storage security risks: Data breaches, loss...
- Organizational controls for storage
- People controls for storage
- Physical controls for storage
- Technological controls for storage

***Buckle Up!***

Explain it to me like  
I am five | fifty years old

*EQ 5!*

# Explain it to me like I am five years old

- Originated in the US version of The Office?
- Prior to that...

*“If you can't explain it to a six-year-old,  
you don't understand it yourself.”*

Albert Einstein

# Explain it to me like I am five years old

- This is on all of us, BSD tools are rarely to blame
- MySQL and its fifteen minute rule
- Transparency and discovery, implying mentorship
- Equally a technical and non-technical function!
- Embrace that fact!
- Consider OS, Ports, Docs, and... Solutions? Policies?
  - Note “Turnkey Linux”

Own the Stack  
Virtualize the REST

*Freedom!*

# Own the Stack, Virtualize the REST


- If it's proprietary, you cannot fix it, full stop
- Fewer repos, fewer rat holes
- Appliance experience brings appliance transparency
- Operating system unity fosters scriptability
- Is your REST API and GUI larger than your OS?
- The BSDs, illumos, and *maybe* Fuschia/React/Haiku



# The Venn Diagram From Hell

*How did we get here?*


# The Venn Diagram From Hell



- Switches and Routers
- Storage Devices
- Security Cameras
- IoT Appliances
- Phones, PBXs
- Cloud Servers

**ALL ON THE  
SAME KERNELS**


# The Venn Diagram From Hell



- Proprietary FSs
- Proprietary ASICs
- Proprietary GUIs
- Proprietary CLIs
- Proprietary GPUs
- Support...

**ALL ON THE  
SAME KERNELS**

# The Venn Diagram From Hell Themes



- VC Funded
- GPL Violating
- Updated?
- Weak Support
- Expensive...
- IPO or Acquisition!

**ALL ON THE  
SAME KERNELS**

# The Venn Diagram From Hell

- Where's THAT in the GNU manifesto?
- The Open Source Definition?
- Is this the software freedom we asked for?
- The Cloud only compounds this
- Computer Science != Azure Credentials
- Not a single school teaches Storage...

# Theo was right in more ways than one...

*“Linux people do what they do because they hate Microsoft. We do what we do because we love Unix.”*

*– Theo de Raadt*

- Does the DOS CLI described over the phone?
- What would 20 years of bitter ex-Windows users developing Unix products look like?

# Case in Point: OpenGPU


- AMD launched the [gpuopen.com](https://gpuopen.com) initiative!
- Works reasonably well with the S7150 and S7150 x2/Proxmox
- [github.com/GPUOpen-LibrariesAndSDKs/MxGPU-Virtualization](https://github.com/GPUOpen-LibrariesAndSDKs/MxGPU-Virtualization)
- [github.com/kasperlewau/MxGPU-Virtualization](https://github.com/kasperlewau/MxGPU-Virtualization)
- Kasper to the rescue! Six commits ahead of upstream!
- Post-pandemic... only cloud vendors get access to newer code
- Cards are available, licenses are not 😞

# Case in Point: “Open Networking” Routers

- Open Network Install Environment (ONIE)
- [opencomputeproject.github.io/onie](https://opencomputeproject.github.io/onie)

## Released Versions

---

The released version are available for download  [opencomputeproject/onie](https://opencomputeproject.github.io/onie).



**There aren't any releases here**

You can create a release to package software, along with release notes and links to binary files, for other people to use. Learn more about releases in [our docs](#).



## Cases in Point...

How did we get here?

Old Habits Die Hard?

# Congratulations!

*BSD appliances often do a  
better job at GPL compliance  
despite zero obligations*

# Combating The Venn Diagram From Hell

- Hardware plays a critical role
- Foundations must be first in line with HW vendors
- Foundations must come before cloud vendors
- What does a secure management API look like?
  - Ubiquitous scriptability is a great start
  - Has the dust settled enough for a web GUI?

# BSD to the Rescue

*Can we have  
nice things?*

# BSD to the Rescue

- Decades of production-quality open standard implementations
- Decades of ties to academia, very good learning materials
- Unified structures, conventions, and code bases
- Many new virtualization options
- Decades of solid, steady progress including...

# BSD to the Rescue

- OpenBSD with PF, “the Human-readable firewall”
  - OpenSSH “p” SHIPS IN WINDOWS
  - Does not require a support session to install
  - Absolute success by every measurement
- Olivier Cochard-Labbé: BSDRP and FreeNAS
- pfSense/OPNsense

# Enough Context: Something has to change

## FreeBSD Build Options and Standards

- Official build-level feature `WITH_/WITHOUT_` switches
- I see an Open Standards and Specs toy chest!
- What bhyve? Nested Intel Instruction Set Architecture!
- What JAIL? Nested FreeBSD ABI!
- `WITHOUT_ISCSI`? Without RFC 3720!

# FreeBSD Build Options in Detail

- I tested, reached out, encouraged, begged...
- Significantly broken for years, all fixed for 13.0/13.1
- Broken on three branches to celebrate EuroBSDcon
- Hopefully joined by packaged base!
- I smoke test them weekly
- Foundation for OccamBSD



# FreeBSD Build Option Smoke Tests

- 234 options on 13.1 described in `man src.conf`
- Traditionally validated with the Build Option Survey
- `/usr/src/tools/tools/build_option_survey`
- Better approach: BOS Lite
  - Exclude three required options and enable the rest
  - `WITHOUT_AUTO_OBJ` `WITHOUT_UNIFIED_OBJDIR`  
`WITHOUT_INSTALLLIB`

# FreeBSD Build Option Smoke Tests

- Talked to Li-Wen and hope to have an official test
- Sweeping all branches takes only a few minutes
- Eyes on the prize: pre-commit testing. It's quick & easy!

Promise me I am not the only one  
who cares about them

# FreeBSD Build Options: OccamBSD

- [github.com/michaeldexter/occambsd](https://github.com/michaeldexter/occambsd)
- Updated for EuroBSDcon!
- Educational tool → system build tool with profiles
- `src` and `obj` directory overrides
- Physical boot drive support: ZFS root: 240 MB no lz4
- Option to keep built world and kernel
- Around 5 minutes builds on an 8-thread i7 (T480)

# FreeBSD Build Options: OccamBSD

```
build_options="WITHOUT_AUTO_OBJ  
WITHOUT_UNIFIED_OBJDIR WITHOUT_INSTALLLIB  
WITHOUT_BOOT WITHOUT_LOADER_LUA WITHOUT_LOCALES  
WITHOUT_ZONEINFO WITHOUT_EFI WITHOUT_VI"  
kernel_modules="virtio"  
kernel_options="SCHED_ULE GEOM_PART_GPT FFS  
GEOM_LABEL CD9660 TSLOG"  
kernel_devices="pci loop ether acpi uart ahci scbus  
cd virtio virtio_pci virtio_blk vtnet virtio_scsi  
virtio_balloon"  
packages=""
```

# FreeBSD Build Options: OccamBSD

But, but that's no longer FreeBSD!

- Reproducible Builds to the rescue!
- Significant progress has been made
- It's FreeBSD, just less of it

# OpenZFS

- No introduction necessary but...
- Dynamic “readonly” property is revolutionary!
  - No remounting of entire file systems!
  - Could have a significantly read-only OS!
  - Could go read-only in response to threats!
- Nearly-ready native encryption!
- Under-appreciated unprivileged user admin delegation

# OpenZFS

*OpenZFS is an under-appreciated  
regulatory compliance tool*

“Hardening OpenZFS to Further Mitigate Ransomware”

SNIA 2021 Storage Developer Conference

# bhyve Hypervisor, alternatively Xen

- Own the Stack, Virtualize the REST
- My number one request? Higher vCPU count
  - In progress! Thank you John Baldwin!
  - Enables Isolated Build Environments!
  - Allows “Type One” hypervisor-like reduction
    - Shift most resources to a VM
    - Serial-managed embedded host if you want



# bhyve Continued: FreeBSD devctl (8)

- Have a serial connection to your VM host?
- *Dynamically* detach the NIC! Audio! GPU! USB!
- Redistribute them to virtual machines!
- Available as the ACS patch on Linux
- Reduce hardware and virtual machine differences!

Yes, you can do that.

# bhyve Continued: Configuration File!

- Thank you John Baldwin!
- `man 5 bhyve_config`
- `bhyve -o config.dump=1 -m 8G ...`
- Supports “local” parameters that are ignored by `bhyve`
- Designed to be the foundation for higher-level tools

# FreeBSD LinuxKPI Linux ABI Compatibility

- Called `compat_linux` in my day
- Been receiving love!
- Reportedly can perform a `buildworld`!
- Can be supplanted by native FreeBSD binaries
- Joined by WINE and Proton for *some* Windows compatibility

# FreeBSD jail.conf.d

- World's Most Proven NoSQL Database (WMPNSQLD)
  - Check out Mark Atwood's Ignite talk on that...
- An elegantly-simple container management system
- Antranig has a few outstanding issues to address
- Discussing nested directories to leverage `zfs unmount`
- BastilleBSD is a great balance of automagic and manual
- jailhyve for jailed virtual machines

# The Nexus: All *near-base* tools

- FreeBSD provides unprecedented infrastructure
- One-button myth: Complexity never reduces complexity
- Everything described here is a handful of shell scripts
- Shell scripts *driving* in-base components in “real” languages
- THAT’s what human-readability looks like
- From today: BastilleBSD lets Jails push policy... needs help?
- If we need an in-base process supervisor... import one!

# The Nexus: All *near-base* tools

- Standards and Compliance with them provide guidance
- The Power of ELI5: We all have local and global roles in it
- Owning the Stack, Virtualizing the rest liberates you
- BSD to the rescue with a steady flow of features

Let's retire the Venn Diagram from Hell!

# Thank you EuroBSDcon!

Questions?

Time for a Demo?

Thank You!

@michaeldexter  
editor@callfortesting.org

