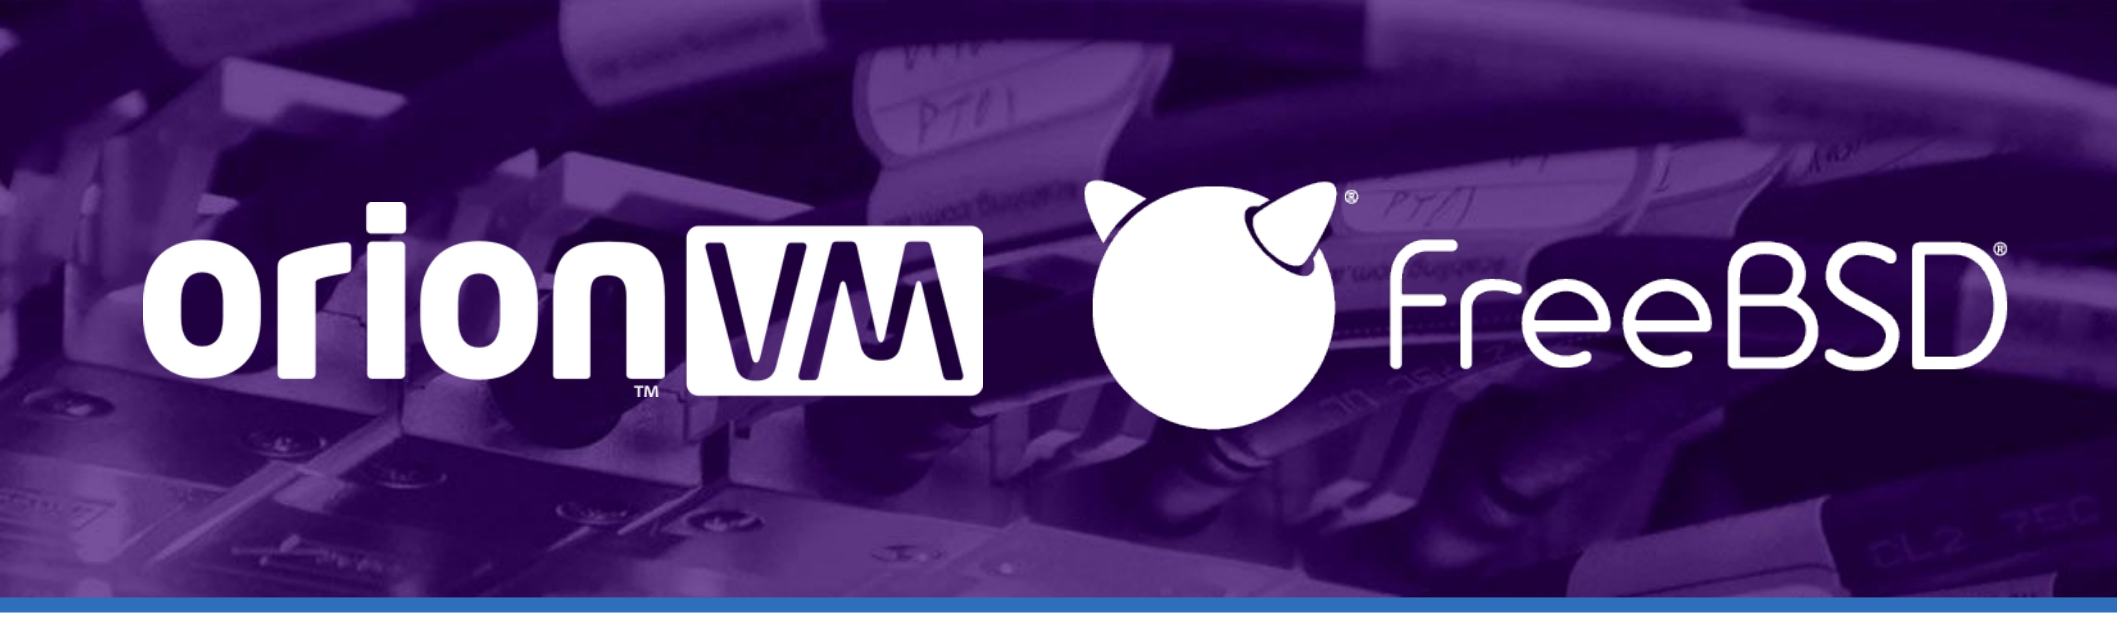

How an Australian IaaS uses FreeBSD

Ruben Schade

FreeBSD Miniconf at Linux.conf.au

Dedication

- NSW Rural Fire Service
- https://www.rfs.nsw.gov.au/volunteer/support-your-local-brigade

Digression #1: Presentation Style

Agenda

- My history with FreeBSD, and introducing OrionVM
- Cloud and VPC storage with FreeBSD, OpenZFS and FreeNAS
- TidalScale for Multi-TiB memory deployments
- Securing cloud and bare metal through PFSense and FreeBSD
- Resources and Thanks
- Q&A

% whoami

- **Ruben Schade**
  Solutions Architect at OrionVM
  Since 2014
- **Role**
  Advise clients on cloud deployments
  FreeBSD template maintainer
- **shutdown -h now**
  Blogger and aspiring author
  Hat and coffee aficionado
  Weeb, and traveller


iBook

My BSD background

- Technically was Haruhi, hah
- FreeBSD user since 2006, on servers and workstations
- Originally on 32 and 64-bit Mac PowerPC
- Used FreeBSD professionally since 2015
- Regular AsiaBSDCon, bhyvecon attendee
- Aim to be a FreeBSD contributor

Digression #2: Gnomenclature

Gnomenclature

- **OrionVM**
  Like the constellation, not the Toyota (Aurion vrrrrm)
- **geli(8)**
  You Jelly of our storage subsystem?
- **gnop(8)**
  Gee-Nop, not Gee-*Nope!*
- **zfs(8), zpool(8)**
  Zed-eff-ess, Zed-pool

Quick OrionVM Overview

Summary

FOUNDED February 2010
OFFICES Sydney, Auckland, San Francisco
BUSINESS Wholesale cloud infrastructure

Locations

Sydney **SY2**
Sydney **SY3**
Melbourne **ME1**

Santa Clara **SV2**
Virginia **DC1**

More PoPs in development.

Features

High-performance Xen IaaS, with hyperconverged VSAN built on a self-healing InfiniBand fabric.

VMs, object storage, backups, desktops, hosted voice, hybrid cloud, bare metal as a service, all on the same platform.

Whitelabel platform with lower rates, greatly simplified interface than other clouds.

OrionVM's Architecture

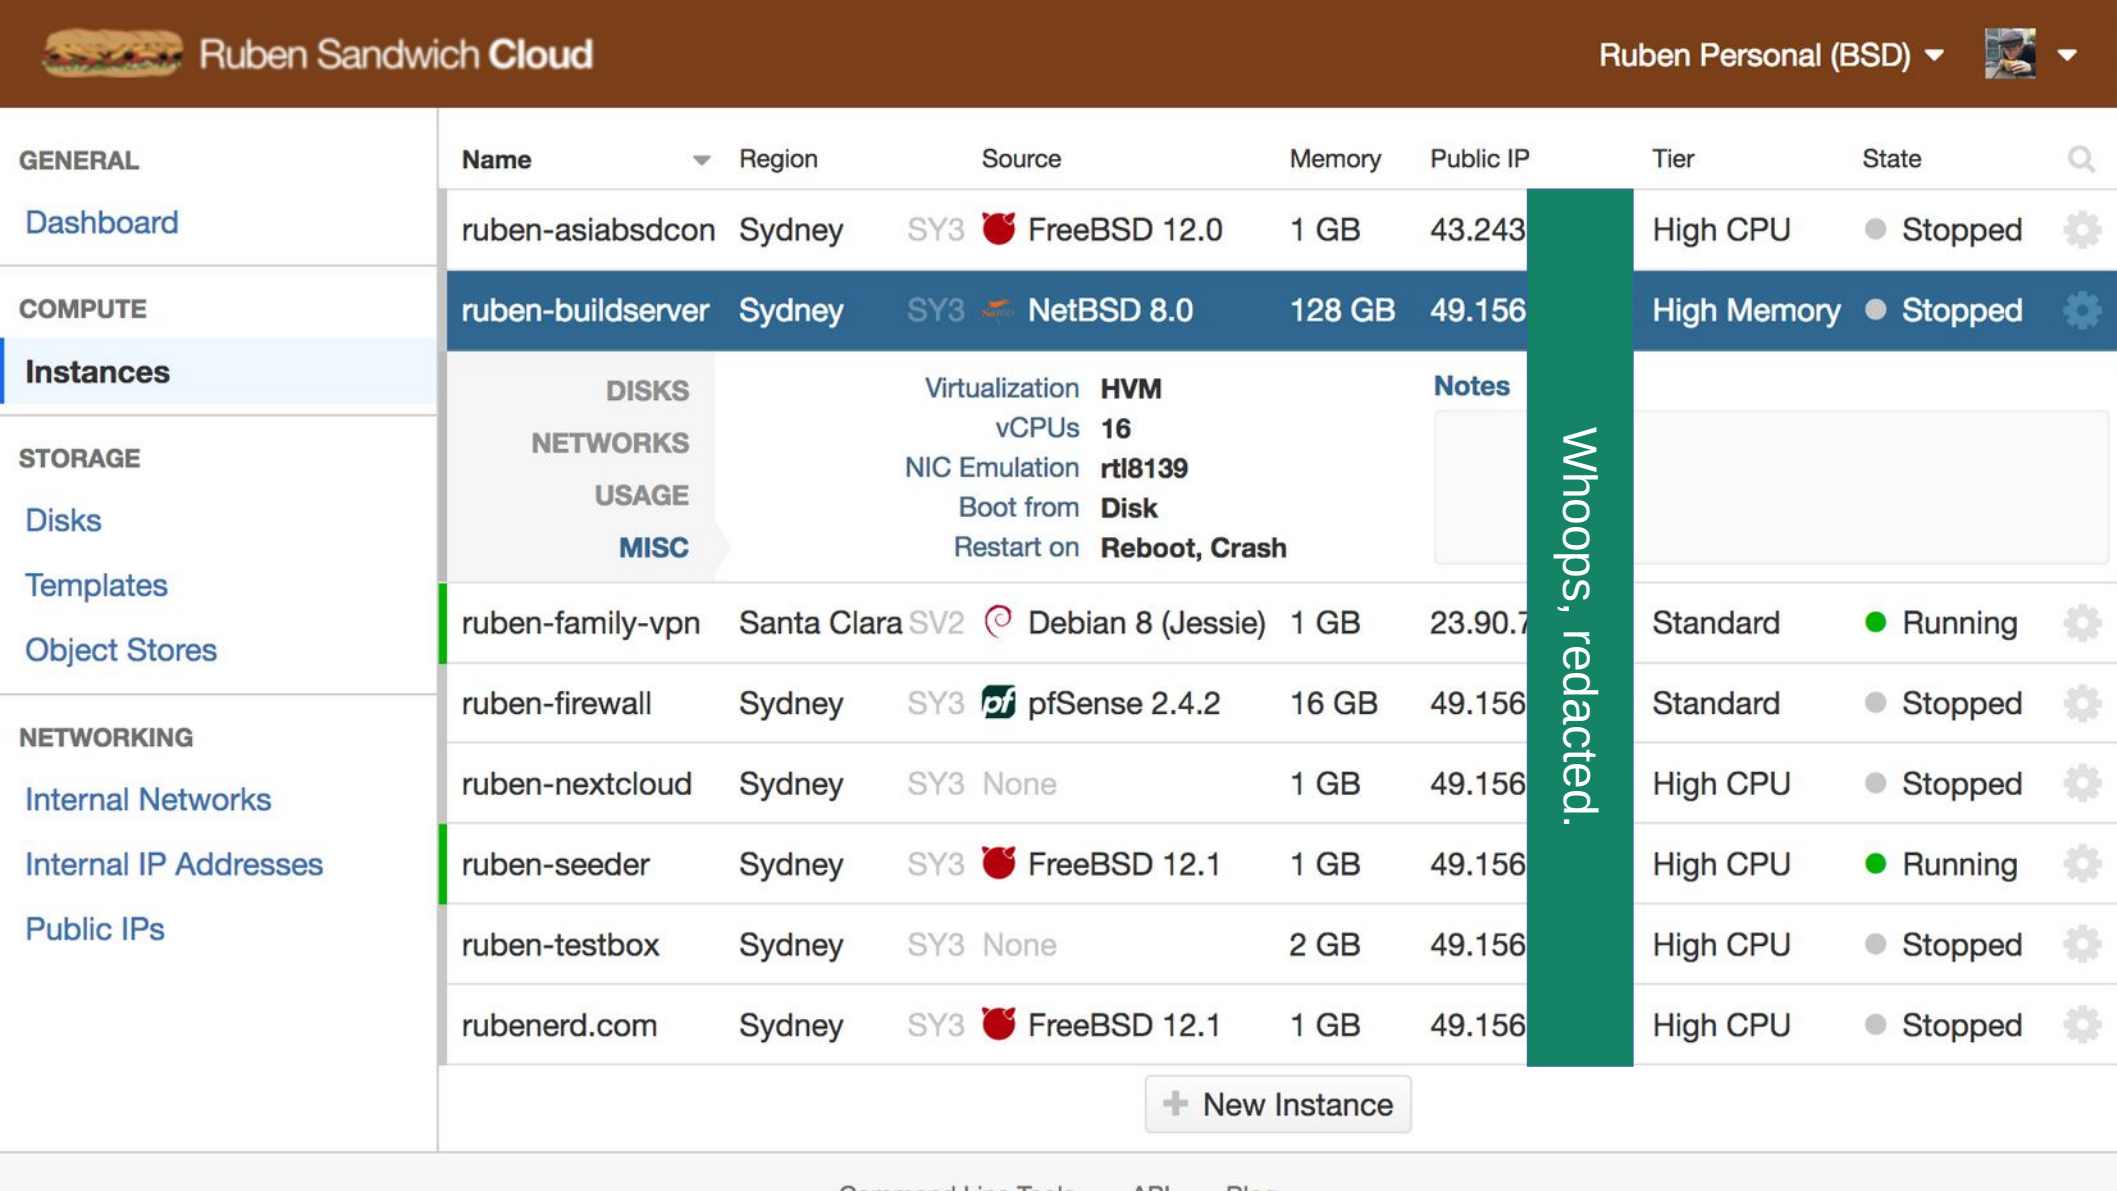
Ruben Sandwich Cloud
Ruben Personal (BSD)
GENERAL
Dashboard
COMPUTE
Instances
STORAGE
Disks
Templates
Object Stores
NETWORKING
Internal Networks
Internal IP Addresses
Public IPs
Name
Region
Source
Memory
Public IP
Tier
State
ruben-asiabsdcon Sydney SY3 FreeBSD 12.0 1 GB 43.243 High CPU Stopped
ruben-buildserver Sydney SY3 NetBSD 8.0 128 GB 49.156 High Memory Stopped
DISKS
NETWORKS
USAGE
MISC
Virtualization HVM
vCPUs 16
NIC Emulation rtl8139
Boot from Disk
Restart on Reboot, Crash
Notes
ruben-family-vpn Santa Clara SV2 Debian 8 (Jessie) 1 GB 23.90.7 Standard Running
ruben-firewall Sydney SY3 pfSense 2.4.2 16 GB 49.156 Standard Stopped
ruben-nextcloud Sydney SY3 None 1 GB 49.156 High CPU Stopped
ruben-seeder Sydney SY3 FreeBSD 12.1 1 GB 49.156 High CPU Running
ruben-testbox Sydney SY3 None 2 GB 49.156 High CPU Stopped
rubenerd.com Sydney SY3 FreeBSD 12.1 1 GB 49.156 High CPU Stopped
Whoops, redacted.
New Instance

FreeBSD Guests

- **OrionVM was the first Australian cloud to offer FreeBSD guests #boom**
- First-class template, same as Linux, Windows Server, and network appliances
- Resellers come with it turned on by default
- Thanks to Colin Percival for his AWS EMI work and inspiration. Sorry about the bump

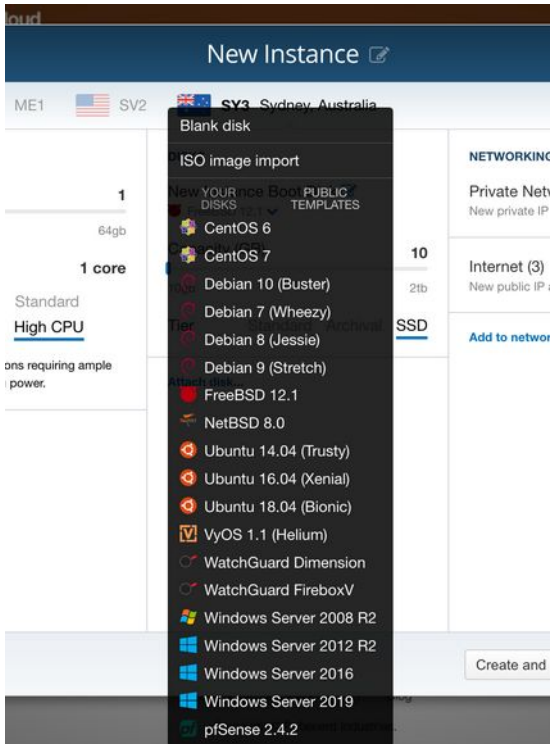

FreeBSD Guests

- RootOnZFS by default
  - https://wiki.freebsd.org/RootOnZFS
- UFS2 version can be enabled by request
- 12.1-RELEASE and 11.3-RELEASE
- Discussed at AsiaBSDCon 2019 WIP session
  - https://2019.asiabsdcon.org/

FreeBSD Guests

AsiaBSDCon 2019 Template Update

- devd(8) integrations for live attaching/detaching:
  - Block devices/disks
  - ISO images
  - NICs and sub-interfaces
  - IP addresses, gateways
  - SSH pubkeys
- Future plans
  - Inline ZFS pool expansion
  - More build automation
  - Binary tools

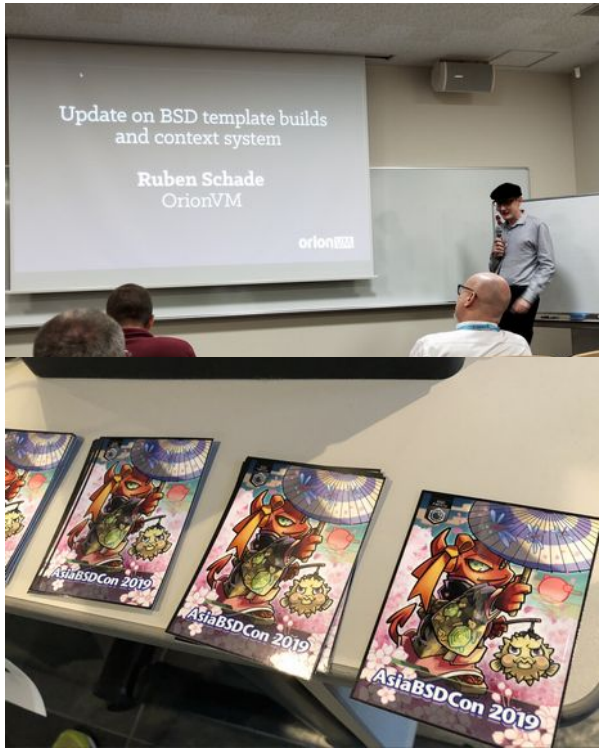

Distraction #3: Why Use It?

Why Use a FreeBSD Guest?

“FreeBSD can do pretty much everything that Linux can do; and most of the things it does better.”

~ Allan Jude

FreeBSD Committer
OpenZFS developer and writer
Chief Architect at Klara Systems

(Un)helpful Advice

- **The “Bucket Problem”**
  - FreeBSD is for compatibility
  - NetBSD is for portability
  - OpenBSD is for security
- **Linux is just a kernel**
  - FreeBSD is a complete system
  - GNU/Linux needs a distribution
- ***Great, what do these mean in practice?***

Why I Use and Recommend It

- **Technical Reasons**
  - System maturity and predictability
  - Toolchain quality (and BSD userland helps)
  - Ease of deployment
  - Documentation
- **Squishy Stuff**
  - Community and events
  - Transparency
- ***Benefits from the illumos, NetBSD, OpenBSD, Linux, macOS communities!***

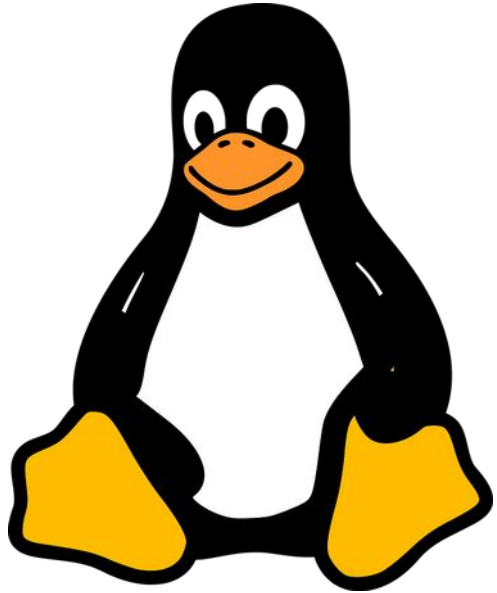

FreeBSD for Cloud Storage

- **OpenZFS is the industry standard**
  - Combined effort from ZoL and illumos (see Matthew Ahrens on freebsd-current)
  - Compression, dedupe, snapshots, shipping out of the box
  - Transparent development, join a call!
- **FreeBSD has excellent tooling and integration**
  - Jails, Capsicum, HAST, RoZ
  - Keep ports, DBs, logs, etc in their own datastore

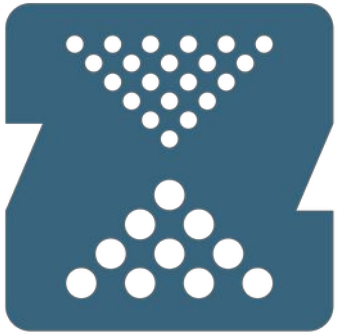

OpenZFS

FreeBSD for Cloud Storage

- **FreeBSD's Integrations with ZFS**
  - Poudriere uses for building ports for pkg(8) by creating an isolated pool that can be blown away
  - iocage(8) uses it by default, ezjail-*(8) can too
  - Allan Jude demonstrated using for atomic upgrades at AsiaBSDCon 2019

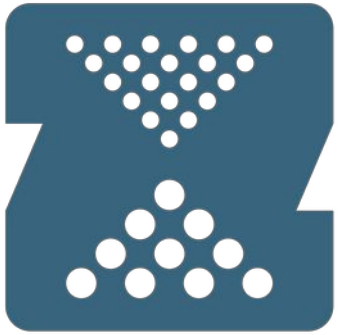

OpenZFS

FreeBSD for Cloud Storage

- **FreeBSD, FreeNAS officially supported templates**
  - Live-attach disks for expanding ZFS pools
  - Use GPT to have appear in */dev/gpt*
- **Can easily export** NFS, Samba to other VMs for free over private network interface
- hastd(8), hastctl(8), and ZFS send/receive for multi-site replication depending on SLA

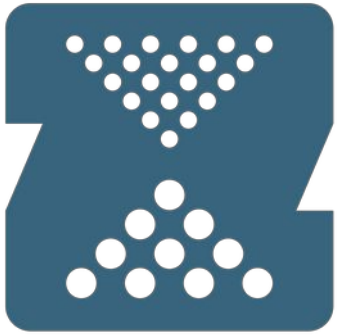

OpenZFS

FreeBSD for Cloud Storage

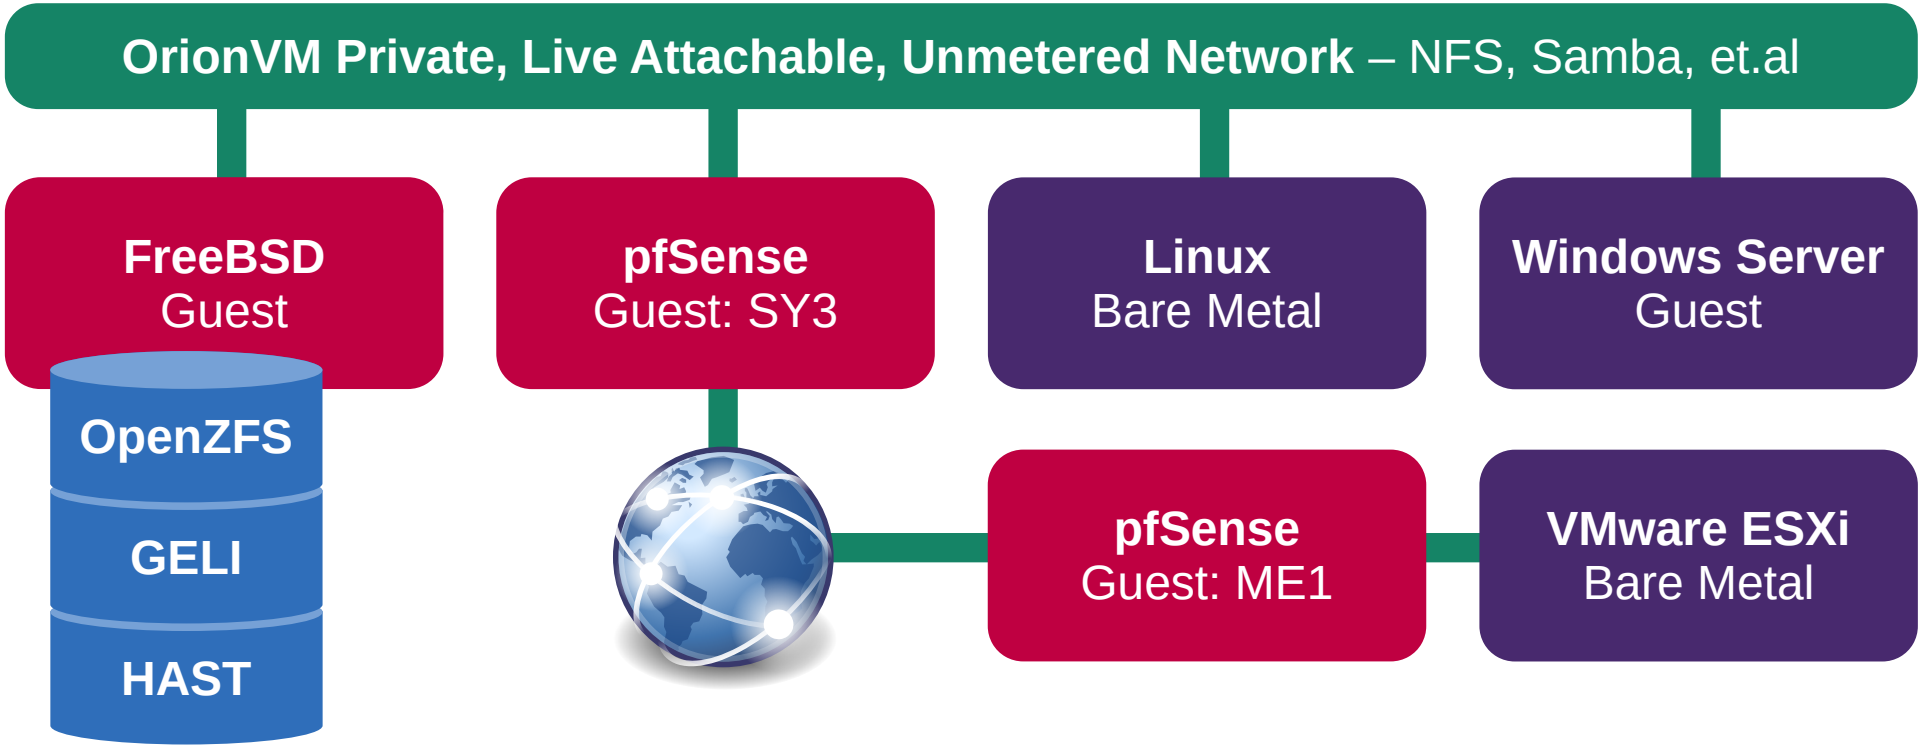

FreeBSD for VPC Storage: FreeNAS

- Perfect for VPC (Virtual Private Cloud) deployments
- For people who don’t require specific SAN hardware (HP 2050s, etc)
- Can be dedicated SAN per deployment, or our multi-tenant cluster
- All the data integrity, snapshotting, and shipping benefits of ZFS
- geli(8) for encryption, hastd(8) for replication

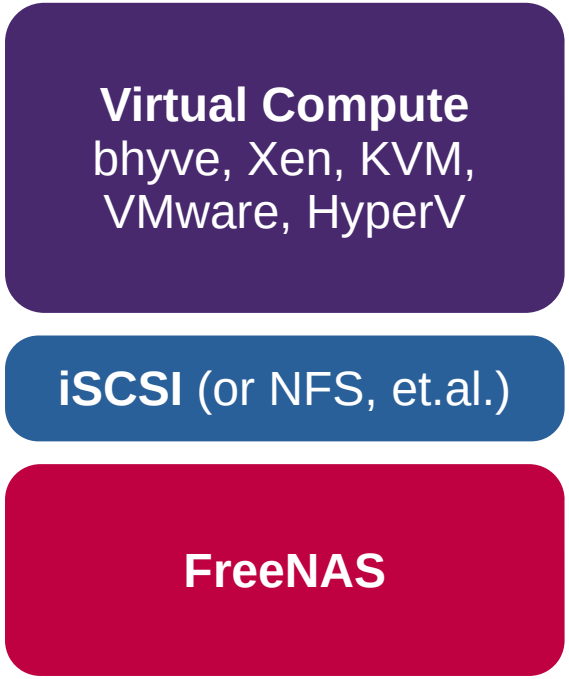

TidalScale

Multi-TiB Memory Deployments

- TidalScale on OrionVM launched 2017
  - https://blog.tidalscale.com/news/topic/orionvm
- Originally based on bhyve to consolidate multiple hypervisors into one large guest over 10G
- Used by big data, analytics, in-memory databases, and former mainframe workloads
- 1.3 TiB to 13 TiB memory, 100+ vCPUs

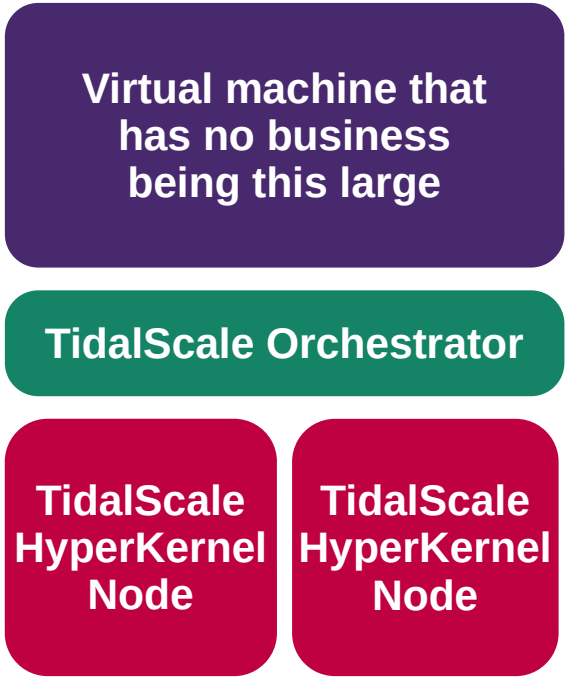

FreeBSD for Securing Deployments

- OrionVM meshes workloads over Layer 2. Guests, colocated hardware, bare metal cluster, MPLS, X-Connects...
- No double NATting like other clouds. NICs get Internet routable addresses directly
- **pfSense** routers sit on the cloud, and protects assets *without extra routers!*
- Can also directly use **FreeBSD's** excellent ipfw(8) in a dedicated guest

FreeBSD for Cloud Storage

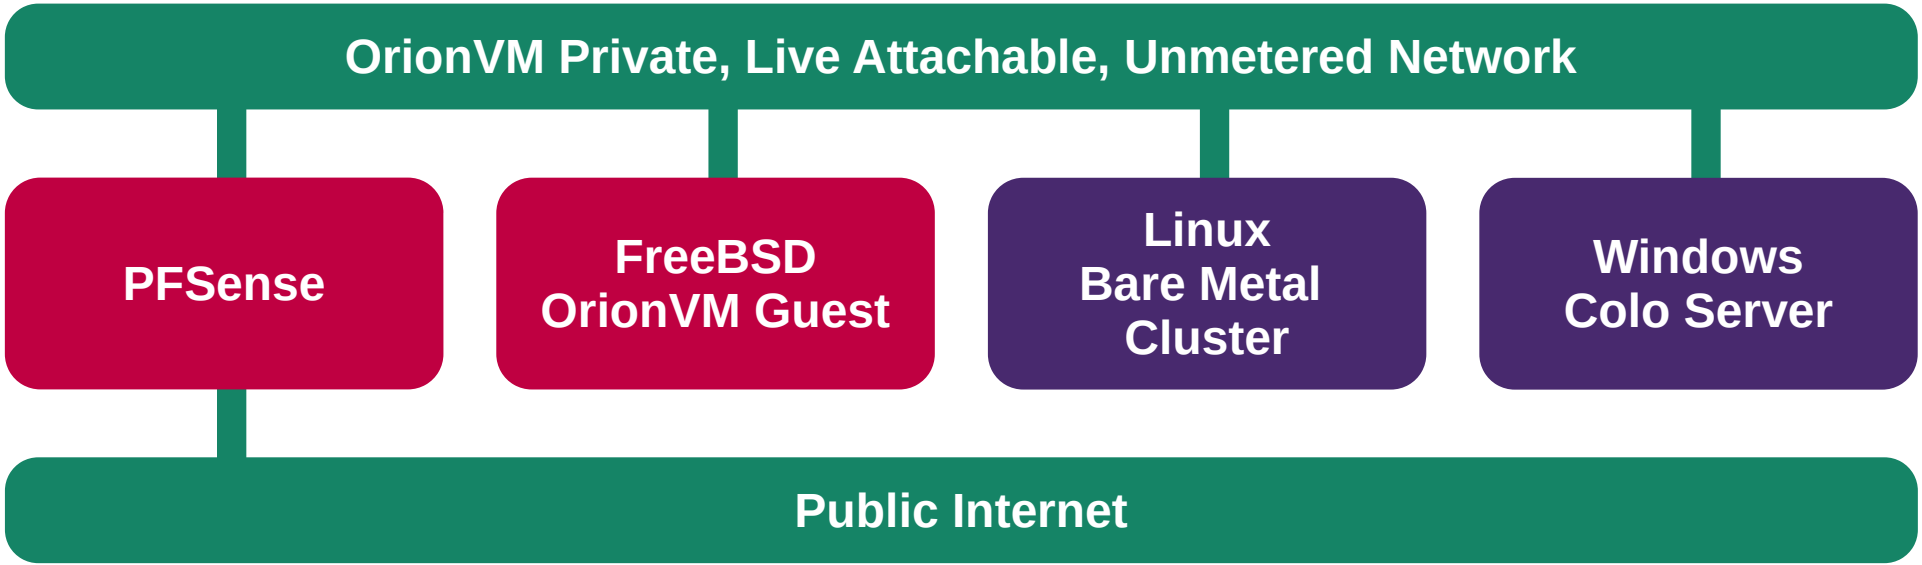

Summary

- You should try FreeBSD, you might like it
- OrionVM's cloud supports FreeBSD, FreeNAS, and pfSense... and Ruben recommends their use
- OpenZFS and FreeNAS perfect for cloud and VPC storage
- FreeBSD lets OrionVM's TidalScale cluster scale to multi TiBs of memory
- pfSense and FreeBSD protects critical client workloads through the cloud, even on dedis

FreeBSD®

Resources and Thanks

- **FreeBSD Foundation** and **Linux.conf.au** organisers
- BSDNow.tv with allanjude@ and bcr@
- bhyve.org and BSDFund.org, Michael Dexter tirelessly coordinating bhyve and other BSD virtualisation efforts
- OpenZFS.org
- PFSense.org
- **Get a reseller demo account!**
  OrionVM.com/contact-us

Legals

- **This presentation and photos**
  © OrionVM, CC BY-SA 4.0
- **OrionVM mark and name**
  ™ OrionVM
- **BSD Daemon “Beastie”**
  © Marshall Kirk McKusick, free to use for BSD projects if done tastefully
- **FreeBSD name, marks**
  FreeBSD Foundation name, marks
  ® FreeBSD Foundation
- **FreeNAS name, marks**
  ® iXSystems, Inc
- **Linux name**
  ® Linus Torvalds
- **OpenZFS name, marks**
  ™ Matt Ahrens, CC BY-SA 3.0
- **pfSense name, marks**
  © Electric Sheep Fencing, LLC
- **Tango Iconset**
  © Tango Desktop Project, CC BY-SA 2.5
- **TidalScale name, marks**
  ® TidalScale, Inc
- **Windows Server name**
  ® Microsoft Corporation
- **Xen Panda Mascot**
  © Xen Project, CC BY-NC 3.0

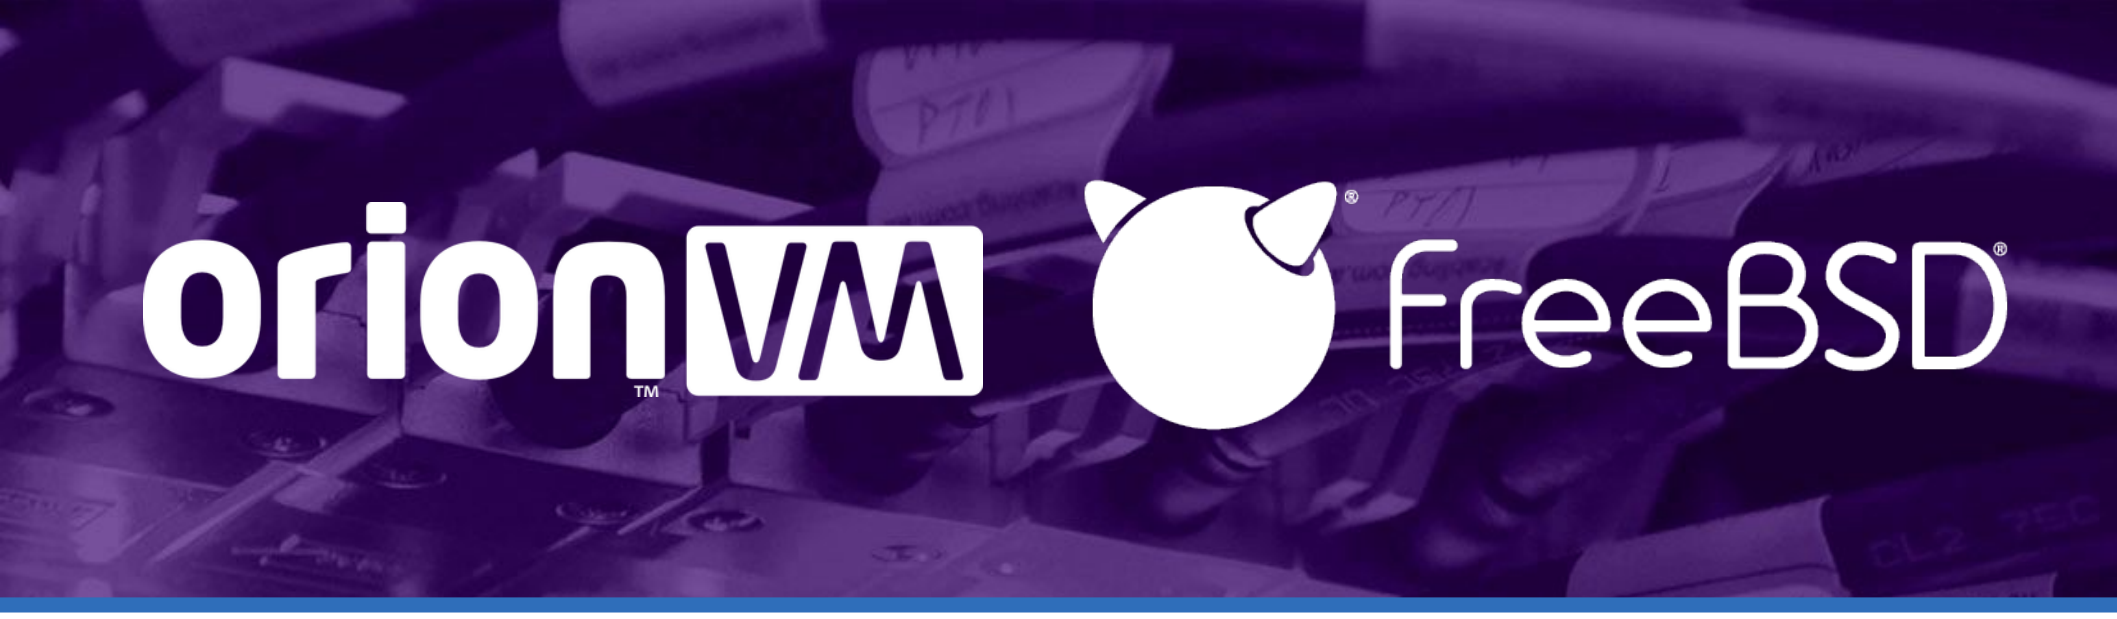

Thanks!

@rubenerd

FreeBSD Miniconf at Linux.conf.au