



State of the OSU Open Source Lab

Lance Albertson, Director
lance@osuosl.org / @ramereth



Session Overview

1. Overview
2. OSL Students
3. Current and New Services
4. Infrastructure Enhancements
5. Goals for this year

Overview

Summary

- FOSS “Hosting Company”
 - Free or low cost hosting services for FOSS projects
 - Co-Location
 - Virtual Machines and other private cloud services
 - Access to a wide array of architectures
 - Software distribution and mirroring
- Mentor undergraduate students in DevOps
 - Gain real-world experience on production systems
 - Past graduates include co-founders of CoreOS
 - Many have key roles in high profile tech companies
- Staff
 - Director
 - 6-10 undergraduate students

History

- Started in 2003
 - Scott Kveton & Jason McKerr co-founders
 - OSU Information Services
- Offered Co-Location hosting for FOSS
 - Gentoo, Debian, and Freenode were early projects
 - Growth spread via word of mouth among projects
 - kernel.org, ASF, Drupal, Linux Foundation
- Seed funding
 - OSU initial funding
 - Google & RealMedia initial sponsors
- Organizational Changes
 - Moved to College of Engineering in 2013
 - Unit of the Center for Applied Systems & Software (CASS)
 - Partnered with another experiential program on campus whose focus was software development and hardware testing

Funding Model

- Corporate donations
 - IBM
 - Google
 - Facebook
- Hosting contracts
 - Linux Foundation
 - Drupal
 - Open Source Robotics Foundation
- In-Kind donations
 - TDS (bandwidth for FTP mirroring)
 - Intel, Facebook, Hudson Trading (hardware donations)
- Support from you! (Thank you!)
- No direct funding from OSU / State of Oregon

Our role in the FOSS Ecosystem

- Provide a neutral hosting facility
 - Foster relationships between FOSS and businesses
- Stable physical home for core FOSS
 - Free hosting
 - Flexible to the needs of projects
- Access to unique hardware and architectures
 - OpenPOWER, ARM, RISC-V, MIPS, etc
- Compute and storage resources
 - Software mirroring
 - CI/CD compute resources
- Systems Engineering expertise
 - Helping projects focus on improving their project
- Training the next generation of FOSS leaders

New Projects in 2018/2019

General Hosting

- Mutt
- FreeBSD Foundation
- MX Linux
- Pentoo Linux
- GNU Radio
- Ohio LinuxFest
- LyX
- CVPM (Computer Vision Package Manager)
- Academic Torrents
- Personal Telco Project
- F-Droid

OpenPOWER Hosting

- OpenWhisk
- Blosc
- OpenBlas
- PyTorch
- PostgreSQL

OpenPOWER Hosting (Cont.)

- Istio
- TensorFlow
- LLVM
- Docker
- OpenCV
- Tesseract OCR
- gcc-gfortran
- x265
- Apple Swift
- LSB
- Fedora
- Mozilla
- FFMpeg
- GHC
- WebM
- NumPy
- Go Lang
- Envoy
- RocksDB

Current Projects (A-M)

Amahi	CRIU	Enlightenment	Goy.Chat	LDC
Anaconda	crostool-NG	F-Droid	Grml	LFS (Linux From Scratch)
Apache Software Foundation	CVPM	Fedora Linux	GSB (Gnome.SlackBuild)	libjpeg-turbo
Apereo	Cygwin	Fedora RISCv	Hackage	libpng
Arch Linux	Darcs	FFmpeg	Hadoop Spark	libvpx
BearSSL	Debian Linux	Finnix	Haiku	LineageOS
Blockchain	Debian RISCv	Firefox	Harmony Agreements	linhes
Blosc	Deeplearning4j	FOSDEM	HHVM	LinuxChix
Bro	Deluge	FOSS Translation	Hortonworks Data Platform	Linux Driver Project
Buildbot	Docker	Freedesktop.org	ICU	Linux Foundation
Buildroot	Drupal	Freedroid	Inkscape	Linux Fund
BusyBox	Eclipse Communication Framework	Funtoo Linux	Istio	Linux Kernel
Cacti	Eclipse	GCC/Clang	Jaws	Linux Plumbers Conference
CentOS Linux	Eclipse OMR	GCC Compile Farm Project	Jellyfish	Linux Standard Base
Checkpoint-restore	Eclipse OpenJ9	GCC Fortran	Jenkins	LLVM
CiviCRM	Elastic Beats	Gentoo Linux	juju-charms	LTTng
Cloud Foundry	Elephant Shed	Ginga Middleware	Jupyter / Base-notebook	LuneOS
CPAN	Elgg	Gnome	K-3D	LyX
Cppcheck	eLinux	GNU Radio	KDE	Mageia
CRAN	ELRepo	Go Language	LAPACK	Mandriva Users

Current Projects (M-Z)

Manjaro Linux	NTPSEC	OpenWRT	Reproducible builds	uClibc
Manulix	NumPy	OrientDB	Robot Operating System	Vector Linux
MariaDB	Nvidia-docker	OSDV	RPM Fusion	VideoLAN
Mesos	Ocaml	OSGeo	RPM	VSXSIMD
Midnight Commander	OCM	PA-RISC	RTEMS	WebDAV
Monkey Server	OFTC	Parrot	Rust	X265 HEVC Encoder
Mozdev	Open Advice	Pentoo Linux	Sahana Foundation	X2go
Mozilla	OpenBLAS	perfbook	Salix Linux	XBMC
Mozillazine	OpenBSD	pgSphere	SHA-3	Xiph
Mulgara	OpenCV	phpBB	Slackware	Yocto
Musescore	open-fcoe	postgresql	Snowpatch	Yum
Musicbrainz	Open Hatch	PowerPC Notebook	Swift	Zarafa
Mutt	OpenID	Presto	Teaching Open Source	
MX Linux	OpenJDK	Privly	Tesseract-OCR	
Mycroft Project	OpenLibm	Python Software Foundation	The Documentation Foundation	
MythTV	open-ldp	PyTorch	Tobin's Linux kernel	
Mythubuntu	OpenMRS	QEMU	Travis-CI	
NAS-Admin	OpenQuake	qiskit-sdk-py	Trinity Rescue Kit	
Node.js	OpenStreetMap	Rainmeter	Trisquel Linux	
NTF	OpenWhisk	Replicant	Ubuntu	

OSL Students

Alumni

- CoreOS
 - Alex Polvi, Brandon Phillips
- Linux Foundation
 - Eric Searcy, Rudy Grigar, Trevor Bramwell
- Microsoft
 - Sarah Cooley
- Many other companies such as:
 - AWS, Tesla, Mozilla, Redhat, Puppet, Zillow, Apple, LinkedIn, Tag1 Consulting, etc

Students Role

- Interact with FOSS projects on daily basis
 - Deploy websites, troubleshoot issues
 - Onboard new projects
- Chef cookbook creation and maintenance
 - Creating new wrapper cookbooks for services
 - Fixing and updating current cookbooks
- Hands-on experience
 - Installing and (un)racking new/old hardware
 - Troubleshooting real-world problems
- Support ticket queue rotation
 - Each student is in charge of every ticket that comes in during that week
 - Ensures each student is knowledgeable about the all of our systems

Hiring process

- Open book quiz
 - Asks basic questions about Linux
 - Simple bash exercises
 - Simple Chef exercise
 - Ensures applicants are passionate enough to complete a quiz
- In person interview
 - Non technical questions ($\frac{1}{3}$)
 - Get a better understanding of their personality
 - See if they would make a good fit with the team
 - Technical questions ($\frac{2}{3}$)
 - Wide range of questions from simple to more complicated
 - Don't expect the students to know all the answers
 - Try and assess their ability to problem solve on the fly

Onboarding process

- Walkthrough guide
 - Account setup
 - How to contribute to internal documentation
 - Basics of Linux
 - Basics of Chef
- Chef onboarding
 - Exercise with a test cookbook
 - Walks them through typical scenarios during cookbook development
 - Create a PR and wait for feedback
 - Senior students provide feedback and work through
- Assign simple tasks initially
 - Within 2-3 months, add them to the ticket rotation schedule

Current and New Services

OSL Managed Platform

- Current and new systems
 - CentOS 7/6 (servers)
 - Debian 8/9 (for student and staff workstations)
 - Managed via Chef
- Chef
 - Wrapper cookbooks using community cookbooks
 - Full unit and integration testing
 - ChefSpec + Test Kitchen + InSpec = Awesome!
 - Jenkins pipeline to automate testing and deployment
- Legacy systems
 - CentOS 6 or Gentoo Linux
 - Managed via CFEngine 2.x

OSL Hardware

- Hardware Budget?
- In-Kind donations
 - Rely mostly on In-Kind donations
 - Intel (Meego) hardware (2012)
 - Dell R610, R710 machines
 - EMC (Neutrino) hardware (2016)
 - Arista 2x 1g/10g switches
 - 8 x 4-node compute nodes
 - Facebook (2016)
 - 3 OCP Racks with total of 90 compute nodes
 - Hudson Trading (2018)
 - “Pallets” of Arista 10g switches
- Wish list
 - 1U/2U compute / storage nodes
 - >3TB SATA HDD's / SSD's
 - 40g End-row switches / 1g ToR switches

Core Infrastructure Services

- Mailing list
 - 200+ lists currently hosted
 - Mailman v2.x based
 - Shared instance
- Email forwarding
 - Includes spam and virus filtering
 - Email store is available if required
- DNS
- Web application hosting
- Systems Engineering consulting for projects

Managed vs. Unmanaged Hosting

- **Managed**
 - Operating System and its updates
 - Configuring and managing services (httpd, etc)
 - Infrastructure design
 - Monitoring and remediation
 - All aspects managed with Chef
- **Unmanaged**
 - We spin up a host and you manage all aspects
 - We only require an account with full sudo for troubleshooting and emergencies

Software Mirroring

- Three server cluster
 - Hosted in Corvallis, Chicago and New York
 - Split via a round-robin DNS
- Stats
 - 1.7Gbps daily average across all three nodes
 - 15 TiB capacity / currently using 12 TiB
 - 100+ repositories/projects hosted
- Hardware Specs
 - 256G RAM
 - IBM POWER8 8286-42A
 - Storage
 - Tiered SAS & SSD RAID6
 - 10Gib uplink
 - Donated by IBM (Thank you!)

Co-Location Hosting

- 300+ Co-location hosts for projects
- Project Racks
 - Gentoo, Linux Foundation, Drupal, Apache Software Foundation, etc
- Projects own their hardware and ship it
- Requirements
 - A reasonable need for having a physical server instead of using a virtual machine
 - Rack mountable server that includes rails
 - Out-of-Band management via IPMI or serial
 - Built by a vendor
 - Exceptions made for special architectures
- OSL Datacenter
 - Shared with University
 - OSL uses around 32 out of 70 racks

Storage: Ceph

- Two clusters
 - 5 node cluster for OpenPOWER OpenStack only
 - 8 node cluster for x86 OpenStack and other OSL services
 - Both deployed in 2018
- x86 Cluster Specs
 - Nautilus (12.x.x)
 - 100 TiB raw capacity
 - 52 x 2 TB SATA drives (6 or 8 per node)
 - 8 x 800 GB NVMe SSD drives (1 per node)
 - Bluestore with SSD's for metadata caching
 - 10g networking
- OpenPOWER Cluster Specs
 - 295 TiB raw capacity
 - 40 x 8 TB SATA drives
 - 40g networking

Storage: Ceph

- Current uses
 - Block storage for OpenStack
 - CephFS (replaced GlusterFS)
- Future plans
 - Object storage
 - Expanding / Upgrading hardware
 - Block storage for Ganeti (maybe?)
 - Geo-Replication?

Private Cloud Platforms

Tale of two platforms...

- Ganeti
 - Stable and easy to maintain
 - Little to no public API
 - Deployed since 2009
 - Poor at self service

- OpenStack
 - Mostly stable and difficult to maintain
 - Excellent public API
 - Deployed since 2013
 - Awesome at self service

Private Cloud: Ganeti

- Been using Ganeti since 2009
 - Powered with KVM and DRBD on local storage
 - Primarily CLI driven, minimal GUI and no public API
- Production cluster
 - 120+ VMs (~15 VMs per node)
 - 8 nodes
 - Storage: 25 TB total / 10.6 TB used
 - Ram: 500 GB total / 350 TB used
 - Projects:
 - phpBB, Busybox/Buildroot, ROS, Jenkins, qemu, etc
- Project Clusters
 - PSF, OSGeo, CiviCRM, OSL Internal
- Current uses
 - “Pet” long-running VMs
 - Important traditional services

Private Cloud: OpenStack

- OSL's adventure with OpenStack
 - Been using it internally since around 2013
 - Test-Kitchen (Integration testing with Chef)
 - Created ppc64le based cluster first
 - x86 cluster opened for projects in 2018
- Specs (x86 cluster)
 - Powered with KVM and Ceph for storage
 - 50+ VMs on 8 compute nodes
 - 128 GB RAM, 10G networking
- Projects
 - Academic Torrents, Ohio LinuxFest, GNU Radio, Foreman, glibc, SharkLinux, AnthonLinux, CVPM, Linux Foundation, GNOME, FreedroidRPG, GCC

OpenPOWER

- Collaboration with IBM for over ten years
- FOSS access to ppc64/ppc64le architecture
- OpenStack powered cluster on POWER9/8
 - 5 x POWER8 systems (~225 VMs)
 - 3 x POWER9 systems (~22 VMs)
- 100+ projects using the cluster
 - Projects:
 - Docker, LLVM, Travis CI, Alpine, GHC, VideoLan, Go, Node.js, OpenJDK, LTTng, RockDB, Nvidia Docker, PostgreSQL, etc
 - Many of the ppc64/ppc64le binaries you use were likely built on this cluster!
- Funded and supported by IBM
- Also host bare metal machines for:
 - GCC Compile Farm, Debian, FreeBSD

OpenPOWER + GPU

- GPU hardware access for FOSS projects
 - OSL + OSU Center for Genome Research and Biocomputing (CGRB)
 - CGRB manages the hardware
- How projects access the GPU hardware
 - Sun of Grid Engine
 - HPC scheduling software managed by the CGRB
 - OSL managed Jenkins portal (<https://powerci.osuosl.org>)
 - Uses Nvidia Docker to get access to the hardware
- Upcoming changes
 - Incorporating OpenStack Zun to provide shell access to GPU resources via containers
 - Sharing via VM not feasible due to pci-passthru limitations

Other IBM resources

- IBM-Z (s390x)
 - Two LPAR's hosted by Marist College in NY
 - Jenkins CI Portal (<https://ibmz-ci.osuosl.org>)
 - Docker images
- AIX
 - Provide hosting but don't manage the resources
 - Access to select FOSS projects for building/testing on AIX

Infrastructure Enhancements

Major projects in the past year

- Chef upgrades and cleanup
- CFEngine to Chef migration
- Ceph
- OpenStack upgrades
- Open Compute compile farm
- IPv6
- Monitoring and metrics
- LetsEncrypt

Chef projects

- 130+ hosts managed
- Upgrades
 - Chef 13 upgrade completed last year
 - Chef 14 upgrade starting soon
 - Updated to latest community cookbooks
- Testing
 - Switched to using cookstyle instead of rspec
 - Starting migration towards InSpec from ServerSpec
- Improvements
 - Proper peer review
 - Better test coverage
 - Using better coding standards

CFEngine to Chef Migration

- ~20 Remaining Legacy systems
- Started migration to Chef in 2013
- Major services remaining that need to be migrated
 - Email relays (4)
 - Mailman (1)
 - Mysql cluster (2)
 - PostgreSQL server (2)
 - Misc managed project VMs (~5-7)
 - VPN/NAT gateway server (1)
 - Other misc OSL hosts (~2)
- Hope to finish the migration by the end of 2019!!

OpenStack upgrades

- Software
 - Completed upgrade to Ocata in December 2018
 - Pike upgrade scheduled for the end of March
 - Plan to continue until running latest stable release
- Storage
 - Migrated from local storage to Ceph
 - Big performance jump and also allowed for live migration
 - Expanded storage capability
- Upcoming changes
 - Adding support for IPv6
 - Enabling additional services for K8s support
 - Improve DNS integration

Open Compute Compile Farm

- Hardware donated by Facebook
- Compute node specs:
 - 140G RAM, single 3TB SATA disk, 10g NIC
- GCC Compile Farm
 - Connected with other projects who needed raw CPU power
- Projects using this:
 - OpenWRT, LineageOS, VLC, Reproducible Builds, RISC-V (Debian/Fedora), F-Droid, GNOME, Freedesktop.org, Buildroot, OpenBSD, GCC Compile Farm
- 59 of 90 nodes allocated

Open Compute Compile Farm

- OCP logistics
 - Rack height
 - 7ft (newer models are taller)
 - Elevator and network cage issues
 - Power
 - 277V 3-Phase AC / 48V DC
 - Network switches with DC PSU's
 - Firmware upgrades
 - BIOS chip removable and burnable
- HVAC limitations
 - Seeking \$150k to upgrade cooling
 - Need your help to fund this!
- Networking logistics
 - On OSU network instead of OSL network
 - Limited IPv4 space - No IPv6 (yet)

IPv6

- Deployed in 2016
 - 2605:bc80:3010::/48
- OSL Public services using IPv6:
 - DNS
 - FTP Mirroring
 - Shared websites (using our load balancer)
- Deployment
 - Static IP address assignment
 - ISP Recommended this configuration
 - Dual-Stack
 - Each subnet gets their own /64
 - Per-project /56 reservations

Monitoring and metrics

- Prometheus & Grafana
 - Deployed in Jan 2019
 - <https://dashboard.osuosl.org>
- Setup exporters for the following:
 - node_exporter (all the CentOS 7 systems)
 - snmp_exporter
 - Switches, PDU and Environmental sensors
 - openstack_exporter
- Working on deploying more exporters
 - apache, haproxy, ipmi, etc
 - Integration with OpenStack Ceilometer / Gnocchi
 - Custom metrics (WIP)
 - FTP project disk usage
 - Database disk usage
 - Backup project disk usage

LetsEncrypt

- Managed with Chef
 - Easy for single hosts
- HAProxy with failover
 - Currently using an NFS-based solution
 - Works but not happy with the implementation
- LetsEncrypt support on ftp.osuosl.org?
 - Not yet (sadly)
 - Currently has a wildcard *.osuosl.org cert
 - Looking at refactoring HAProxy solution

Other misc projects

- RANCID
 - Network switch configs in git
- 10g backend network upgrades
 - Ganeti DRBD
 - Ceph
 - OpenStack
 - New VPN server
 - Backup server
- Migrated from djbdns to Bind
- Razor Server (Netboot installations)
- Increased FTP storage capacity
- Firmware upgrades
 - Thank you Meltdown/Spectre/etc

Goals for the next year

- Complete legacy system migrations
- Upgrade OpenStack & Ceph to latest stable
- Upgrade to Chef 15
- Migrate Ganeti to Ceph
- Deploy a proper ELK stack
 - Integrate more metrics into Grafana
- Start replacing aging OSL Network core
- Start working on CentOS 8 migration?
- Migrate CentOS 6 hosts
 - Upgrade to 7 (or 8 when ready)
- Deploy an ARM OpenStack cluster?



Thank you!

Lance Albertson

lance@osuosl.org / @ramereth

