

TOSS - A RHEL-based Operating System for HPC Clusters

Supercomputing 2017
Red Hat Booth

Ned Bass
System Software Development Group Leader
Livermore Computing Division

November 14, 2017



TOSS is a critical component of LLNL's commodity Linux cluster strategy

What is it?

- A common operating system and computing environment for Tri-Lab Linux clusters.
- A software stack for HPC – large interconnected clusters.
- A methodology for building, quality assurance, integration, and configuration management

Why do it?

- Reduce total cost of ownership and enable application portability
- Consistent source and software across architectures: x86_64, ppc64le, and aarch64
- Install same software on all commodity hardware at the Tri-Labs (LLNL, LANL, Sandia)

RHEL provides important benefits to TOSS maintainers and users

- Access to RH engineers and expertise in various areas
- Software provenance
- Timely access to security patches
- Consistent source for x86_64, ppc64le, and aarch64
- Partner Early Access Program to get early hardware working with RHEL
- Consistent Application Binary Interface within a major release – codes don't need to recompile
- TOSS releases track RHEL releases - TOSS 3.x is based on RHEL 7.y

LLNL software expertise is critical to success of TOSS



- Software support for new hardware sometimes lags behind system deliveries.
- Having in-house software developers is necessary to quickly resolve problems.
- Our solutions are submitted back to Red Hat, benefitting the entire HPC community.



LLNL develops open source systems software for HPC clusters



<https://software.llnl.gov>



./ ConMan

The Console Manager

./ MUNGE

MUNGE Uid 'N' Gid Emporium

363 Software Repositories

Filter Repositories

zfs | zfs/linux | C

ZFS on Linux - the official OpenZFS implementation for Linux.

🔄 ★ 2819 📄 613 🌐

spack | LLNL | Python

A flexible package manager that supports multiple versions, configurations, platforms, and compilers.

🔄 ★ 385 📄 312 🌐

zfs-auto-snapshot | zfs/linux | Shell

ZFS Automatic Snapshot Service for Linux

🔄 ★ 299 📄 92

spl | zfs/linux | C

A shim layer which adds the core interfaces required for OpenZFS.

🔄 ★ 273 📄 150 🌐

pkg-zfs | zfs/linux | -

Native ZFS packaging for Debian and Ubuntu

🔄 ★ 248 📄 33 🌐

dcp | hpc | Shell

dcp is a distributed file copy program that automatically distributes and dynamically balances work equally across nodes in a large distributed system without centralized state.

🔄 ★ 146 📄 10 🌐

rose | rose-compiler | C

A release version of ROSE that passes all regression tests

🔄 ★ 143 📄 56 🌐

pdsh | chaos | C

A high performance, parallel remote shell utility

🔄 ★ 123 📄 24

diod | chaos | C

Distributed I/O Daemon - a 9P file server

🔄 ★ 120 📄 25

TOSS adds system management tools, Lustre, user tools, hardware drivers, and more

Cluster Management Tools

- Pdsh – parallel remote shell
- Powerman – remote power management
- Conman – remote console management
- FreeIPMI – out-of-band systems management
- MUNGE – scalable authentication
- OMS/SMT – Infiniband diagnostics
- Whatsup – node up/down detection
- Genders – cluster configuration database
- CFEngine – configuration management
- SLURM – job scheduling
- Mrsh – remote shell with munge authentication
- Netroot – diskless boot support
- LDMS – lightweight runtime collection of high fidelity data

User Tools

- Compilers (PGI, Intel, GCC, clang)
- Debuggers (Totalview, Allinea)
- MPI libraries (OpenMPI)
- I/O libraries (NetCDF, HDF5)
- Visualization & Graphics (Paraview, VisIt, mplayer, vlc)

Kernel Modules and Patches

- Lustre & ZFS
- Nvidia
- Network drivers (i40e, ixgbe)
- MSR-safe
- NFS support for > 16 groups
- Assorted bug fixes and enhancements





We use as much stock RHEL and EPEL software as we can. Additional TOSS packages are built using the Koji build system from the Fedora project.

New TOSS releases are tagged about once a month to address bugs and security issues

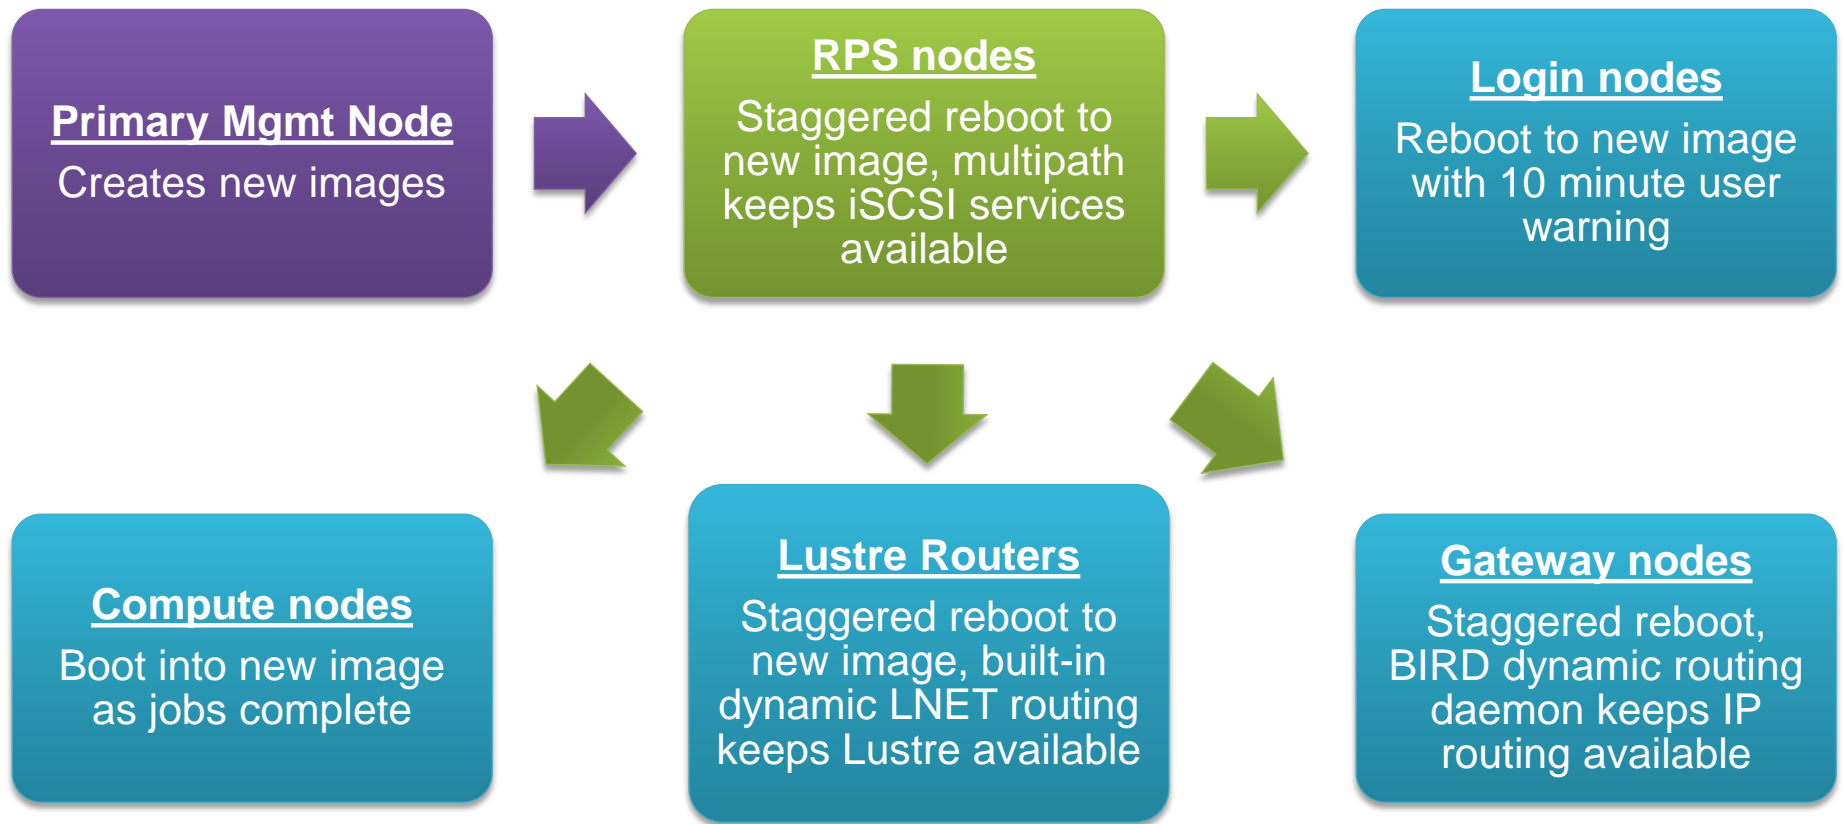
- ❑ Every release candidate is tested end-to-end on compute/Lustre clusters
- ❑ The Synthetic Workload (SWL) test suite is used to validate stability and correctness
 - Contains a mix of MPI applications, benchmarks, and I/O tests
 - Runs jobs in an uncoordinated fashion to simulate real-world compute workloads
 - This approach has been very effective at finding bugs throughout the software stack
before they impact users
- ❑ Not every TOSS release is deployed to production
- ❑ We use Atlassian JIRA for issue tracking and release management



TOSS supports several installation and boot methods for state-full and state-less nodes

Image Type	Use Case
State-full 	<ul style="list-style-type: none">• Primary management nodes• Everything installed on local disk• PXEboot kickstart install from center-wide management node• Image updated during maintenance outage – longer downtimes
NFS root 	<ul style="list-style-type: none">• Compute, login, router, gateway, and Lustre server nodes• Image built ahead of maintenance outage – shorter downtimes• Image can be updated live• Can use local disk for /tmp, /var/spool/cron, etc. (for login nodes)
iSCSI boot target 	<ul style="list-style-type: none">• Similar to nfsroot, but uses iSCSI target as the root device• Benefit: allows for multipath support for fault tolerance and rolling updates – node can survive one server failure• Drawback: live images are immutable
Squashfs 	<ul style="list-style-type: none">• Node copies image into memory and runs it from there• Once booted, nodes no longer requires boot server to stay up• Used for non-primary management (RPS) nodes

The TOSS Update Tool (TUT) minimizes user impact with automated rolling updates.



The primary management node is the only state-full node in a cluster. It runs a Redis server to manage the update process.

We run TOSS on our compute clusters, Lustre servers, and infrastructure systems.

