

Linux OS Plans at The SDCC

Chris Hollowell <hollowec@bnl.gov>
Scientific Data and Computing Center (SDCC)

HEPiX Fall 2021 - 10/25/2021

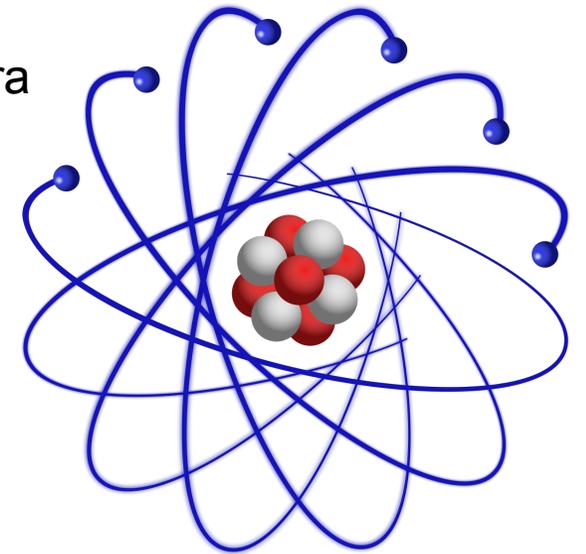
History of Linux at the SDCC

- Have been using Linux at the Scientific Data and Computing Center (SDCC) at BNL since our group's inception in the mid-1990s
 - SDCC was the RHIC Computing Facility (RCF) at the time
- Have been exclusively running Red Hat Linux distributions/derivatives
 - Some discussions about transitioning to Debian in the late 1990s and early 2000s, but never materialized
- Had been running RH (not RHEL) 7.3 and 8.0 in 2003 when Red Hat announced the the end of the Red Hat Linux distribution
 - It was replaced with Fedora Linux, and Red Hat Enterprise Linux (RHEL)
 - Fedora would be a fast-changing, community-driven OS with an approx. 1-year lifecycle, and would remain freely available
 - RHEL was targeted for production/enterprise environments, with a 10-year lifecycle, but only available for purchase



History of Linux at the SDCC (Cont.)

- The end of the Red Hat Linux distribution posed an issue
 - We were not interested in upgrading the OS on our machines yearly - as required with Fedora
 - Significant administrator labor involved, and also potentially difficult for supported experiments to adapt to major yearly OS changes
 - Red Hat Linux generally had a longer support window than Fedora
 - Had concerns over the stability of Fedora Linux with its fast-paced development model
 - Full commercial support for the OS on our compute nodes, which would account for the majority of our RHEL licensing, was not necessary
- Red Hat continued to publish all the sources for their RHEL distribution, even to non-customers
 - Various components did not have GPL licensing, so Red Hat was *not* obligated to do this
- *The publication of the source allowed for the creation of clone/rebuild distributions of RHEL, including Scientific Linux (SL) in 2004*
 - Developed by Fermilab and CERN



Scientific Linux

Scientific Linux at The SDCC

- We upgraded our processor farms to SL 3.0.2 in Fall 2004
 - Presented on [our experience with Scientific Linux](#) at the October 2004 HEPiX
- SL worked quite well for us, and we have continued running new releases of it on our compute farms/clusters
 - Even after Red Hat acquired CentOS in 2014, and dramatically improved the development for that project
 - Additional integrated packages in SL and not in stock CentOS like OpenAFS have simplified our deployments
 - We don't require CERN site-specific packages in CERN CentOS 7 (CC7) like the CERN phonebook
- While primarily developed by FNAL/CERN, SL is also a success story for HEPiX
 - Its existence was motivated in part by the needs of our overall community, as laid out at various HEPiX meetings

BNL HEPiX Fall 2004 Presentation on SL Experience

Scientific Linux Experience at RHIC Computing Facility

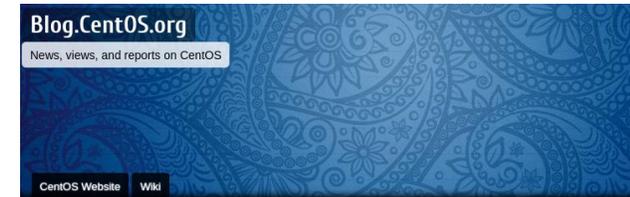
Christopher Hollowell
Physics Department
Brookhaven National Laboratory
hollowec@bnl.gov

Red Hat Linux End Of Life Options (Cont.)

- Scientific Linux
 - In use at several HEP/NP sites
 - No cost
 - Requires few changes to RCF's current installation/custom package management model
 - Built from RHEL3 sources --> compatibility with Enterprise --> few third party application support concerns

History Repeats Itself?

- In 2019, Fermilab announced that there would be no future major releases of Scientific Linux
 - They would continue to support SL 6 and SL 7 through the remainder of their lifecycles
 - CERN had already switched to a CentOS 7-based distribution a few years earlier
- This was logical given the success of CentOS, and how much of a quality distribution Red Hat's acquisition had helped make it
 - But effectively made CentOS 8 the only viable freely available RHEL 8 clone option, and we had some concerns about that...
- In December 2020, the CentOS project announced they were EOL'ing CentOS 8 early (Dec 2021), and replacing this distribution with CentOS Stream 8
 - A slightly forward distribution of RHEL 8
 - Not quite as dramatic a change as the Red Hat -> Fedora/RHEL move, but still significant
 - Particularly because its OS lifecycle support also changed from 10 to 5 years



CentOS Project shifts focus to CentOS Stream

Tuesday, 8, December 2020 Rich Bowen Uncategorized 709 Comments

The future of the CentOS Project is CentOS Stream, and over the next year we'll be shifting focus from CentOS Linux, the rebuild of Red Hat Enterprise Linux (RHEL), to CentOS Stream, which tracks just *ahead* of a current RHEL release. CentOS Linux 8, as a rebuild of RHEL 8, will end at the end of 2021. CentOS Stream continues after that date, serving as the upstream (development) branch of Red Hat Enterprise Linux.

Meanwhile, we understand many of you are deeply invested in CentOS Linux 7, and we'll continue to produce that version through the remainder of the [RHEL 7 life cycle](#).

CentOS Stream will also be the centerpiece of a major shift in collaboration among the CentOS Special Interest Groups (SIGs). This ensures SIGs are developing and testing against what becomes the next version of RHEL. This also provides SIGs a clear single goal, rather than having to build and test for two releases. It gives the CentOS contributor community a great deal of influence in the future of RHEL. And it removes confusion around what "CentOS" means in the Linux distribution ecosystem.

When CentOS Linux 8 (the rebuild of RHEL8) ends, your best option will be to migrate to CentOS Stream 8, which is a small delta from CentOS Linux 8, and has regular updates like traditional CentOS Linux releases. If you are using CentOS Linux 8 in a production environment, and are concerned that CentOS Stream will not meet your needs, we encourage you to contact Red Hat about options.

We have an [FAQ](#) to help with your information and planning needs, as you figure out how this shift of project focus might affect you.

[See also: [Red Hat's perspective on this](#).]

Current Status

- Running RHEL 7 on our critical infrastructure hosts
 - SSH gateways, webservers, databases, etc.
 - Prefer to have a fully commercially supported OS on these hosts
 - Satellite/RHN also provides a value-add for these systems
 - Many of these are VMs on our RHEV clusters
 - No additional cost for RHEL licenses for VMs on RHEV (hypervisors running RHEL server edition)
- Running SL 7 on our HTC farms and HPC clusters
- Supporting these two different distributions has not entailed significant additional effort
 - It has helped that SL 7 is a RHEL 7 derivative/clone
 - We've been operating with this hybrid approach for a number of years without issue
- Initially planned on moving our farms/clusters to CentOS 8 in late 2022 or early 2023
 - Now that CentOS 8 will be early-EOL'd in December, what are our plans?

Available CentOS 8 Alternative Options

- Move all hosts to RHEL 8
 - Potentially use RedHat's academic licensing but need to know details before serious consideration
 - SDCC has 2500+ compute hosts in support of many programs
- Move to another RHEL 8 clone/derivative distribution
 - A few new RHEL 8 clones have become available since the CentOS 8 EOL announcement
 - Rocky Linux
 - Alma Linux
- Migrate to CentOS Stream 8
 - This is the path advocated by CERN and FNAL

CentOS Stream 8 Concerns

- 5-year lifecycle for CentOS stream, instead of the 10-year lifecycle or CentOS/SL
 - Typically perform a major release OS update on our compute nodes every 3-4 years
 - Tightens the schedule
- CentOS Stream will be slightly ahead of RHEL
 - Package updates (potentially including new features) will be available in CentOS Stream somewhat before RHEL
 - Essentially a rolling “preview” of the next RHEL minor release
 - ***While we’ve heard that updated packages in Stream will all have passed Red Hat/CentOS QA before inclusion, the reality is Stream users are still essentially “beta testing” them before they are included in RHEL***
 - Some concerns using Stream will expose our systems to more bugs/issues than we’ve experienced with RHEL/SL
- The kABI/KMI is subject to change, even within a RHEL release to some extent
 - Particularly for kernel symbols not whitelisted
 - <https://access.redhat.com/solutions/444773>
 - Concerns over support for 3rd party kernel modules
 - NVIDIA/CUDA, GPFS, OpenAFS, Lustre
 - Can likely be made functional on RHEL-forward CentOS Stream 8 kernels, but are we certain?
 - Attempts to use pre-built RHEL 8 binary modules more likely a problem
 - Will the vendors for these modules officially support them if there are issues?

CentOS Stream 8 Concerns (Cont.)

- As of today, NVIDIA at least is *not* specifically listing CentOS Stream 8 native support for the latest CUDA release (11.5)
 - <https://docs.nvidia.com/cuda/cuda-installation-guide-linux/index.html>
- This is a concern, as we provide our users with a number of HPC clusters with NVIDIA GPUs

Table 1. Native Linux Distribution Support in CUDA 11.5

Distribution	Kernel ¹	Default GCC	GLIBC	GCC ^{2,3}	ICC ³	NVHPC ³	XLC ³	CLANG	Arm C/C++
x86_64									
RHEL 8.y (y <= 4)	4.18.0-305	8.4.1	2.28	11	2021	19.x, 20.x	NO	12	NO
CentOS 8.y (y <= 4)	4.18.0-305	8.4.1	2.28						
RHEL 7.y (y <= 9)	3.10.0-1160	6.x	2.17						
CentOS 7.y (y <= 9)	3.10.0-1160	6.x	2.17						
OpenSUSE Leap 15.y (y <= 3)	5.3.18-57	7.5.0	2.31						
SUSE SLES 15.y (y <= 3)	5.3.18-57	7.5.0	2.31						
Ubuntu 20.04.3	5.11.0-27	9.3.0	2.31						
Ubuntu 18.04.z (z <= 5)	5.3.0	7.5.0	2.27						
Debian 11.0	5.10.0-8	10.2.1	2.31						
Fedora 34	5.11	11	2.33						

Future Plans

- Will migrate to RHEL 8 on infrastructure hosts running RHEL 7
 - Have already started the process for some systems
 - Completed migration of a number of our Puppet manifests to RHEL 8
- ***For compute hosts we will likely adopt Rocky Linux 8 - but the decision is not finalized***
 - Have successfully tested Rocky Linux 8.4
 - Rocky Linux was created by an original founder of the CentOS and Singularity projects
 - Stated goal of the project is to be 100% RHEL 8 compatible - even from kernel space
 - Eliminates some of the concerns we have with CentOS Stream
 - 10-year lifecycle vs 5-year lifecycle for CentOS Stream 8
 - This is a relatively new project, and its ultimate longevity is unknown
 - But there is also no guarantee the CentOS project will not change CentOS again in another fundamental way, or even eliminate this OS, at some point in the future
- What about the potential implications of diverging from what CERN/FNAL are doing?
 - We have already diverged from the distribution CERN has been running for the past few years
 - We've been running SL 7, they've been running CC 7
 - Has not been a problem since they are ABI/API compatible RHEL 7 clones/derivatives
 - ***But are RHEL 8 and CentOS Stream 8 API/ABI compatible?***



Stream & RHEL API/ABI Compatibility

- Excellent talk given on this subject by Pat Riehecky (an SL developer) from FNAL at the [May 2021 CentOS Dojo](#):
 - [“Thinking About Binary Compatibility and CentOS Stream”](#)
- From the **userland (not kernel)** perspective, RHEL 8 ABI/API compatibility is governed by the published [“Red Hat Enterprise Linux 8: Application Compatibility GUIDE”](#)
 - Component libraries classified for level 1 or 2 compatibility must maintain compatibility between RHEL 8 minor releases
- As a slightly forward-distribution of RHEL 8, CentOS Stream 8 is still beholden to this document
 - New packages in CentOS Stream 8 are eventually destined for the next RHEL 8 minor release, and new RHEL 8 minor releases must satisfy the requirements of this document
- Therefore, RHEL 8 compatible applications using level 1 or 2 classified libraries will also run without issue on CentOS Stream 8
 - Level 2 compatibility is the default for libraries in RHEL 8
- Applications compatible with CentOS Stream 8 will also run on RHEL 8, assuming they do not make use of new features/symbols in CentOS Stream 8
 - Similar situation to running applications built on a newer RHEL minor release on an older minor release (for example a RHEL 8.4 application on RHEL 8.3)

Containers & Linux Distribution Divergence

- When Scientific Linux was created in 2004, having a single shared Linux distribution was critical in the HEP/NP and WLCG community in order to provide a consistent environment for jobs and services in the grid
- ***The world has changed***
 - 17 years is a long time
 - Now container engines like Singularity, Podman, Docker, and orchestration tools like k8s and Openshift are everywhere, and heavily used by our community
 - We are at a point where to first order it does not matter what Linux distribution/version we run on our bare metal
 - Advances in allowing nested Singularity containers, in part thanks to kernel user namespace support, has helped here, permitting even pilots to run in containers
 - At some point in the future there may be (or already are?) HEP/NP sites running Arch, Gentoo or other non-HEP-standard Linux distributions on their bare metal
- Some of the experiments we support at SDCC already run all of their compute jobs in containers
 - For those that don't, there is essentially nothing stopping us from having the batch system run their jobs in Singularity containers automatically with our SL 7 OS image



Conclusions

- SDCC will upgrade our RHEL 7 infrastructure hosts/VMs to RHEL 8 over the next few years
- We will likely migrate from SL 7 to Rocky Linux 8 on our HTC/HPC compute nodes in the late 2022 or early 2023 time frame
 - **Not a final decision**
 - The use of Rocky Linux 8 eliminates some concerns we have with CentOS Stream 8
 - Stability
 - Support for 3rd party kernel modules
 - 10-year lifecycle vs 5-year lifecycle in CentOS Stream
- RHEL 8, CentOS Stream 8, and Rocky Linux 8 are API/ABI compatible from a userland perspective
 - All required to adhere to the “Red Hat Enterprise Linux 8: Application Compatibility GUIDE”
 - With respect to user application compatibility, it effectively doesn’t matter which of these distributions sites ultimately choose to run
- Linux container technology has made significant advances and has become widespread
 - Allows for further divergence in the bare metal Linux distributions/versions used at sites
 - Expect to see continued movement in this direction