



# HEPiX Autumn 2021 – Linux BoF

CERN & FNAL Linux Distro Developments

2021-10-25

Ben Morrice

On behalf of CERN IT

IT-CM




# Agenda

- Situation today
- Recent changes
  - Enterprise Linux Clones
  - Red Hat academic licensing offer
  - Joint statement
  - Proposal (Straw man)
  - CentOS Linux vs CentOS Stream
- Roadmap / timeline
- Distribution landscape
- Questions

# Situation today

- CC7 (CERN CentOS 7) and Scientific Linux 7 continue to be supported operating systems, until 30.06.2024
- C8 (CentOS Linux 8) is available as a supported operating system, until **31.12.2021**



- CS8 (CentOS Stream 8) is available as a supported operating system, until 31.05.2024
- Since becoming available as a supported operating system (01.05.2021), CS8 is now being actively used for production workloads

# Enterprise Linux Clones

- Continued analysis of Enterprise Linux Clones (ELC)
  - Rocky Linux & AlmaLinux being the two most prominent players
  - Both have viable products (stable 8.4 releases)
  - The time taken to releasing an update once available from Red Hat is similar between both projects
  - Rocky Linux is more transparent with its build process, being based on koji versus a custom blackbox on the AlmaLinux side
  - AlmaLinux supports additional architectures, such as powerpc64le

# Red Hat: academic license offer

- After several months of discussion, an academic offer from Red Hat has been made to CERN
  - This offer has not yet been accepted by CERN
  - Please do not disclose this information publicly

# Red Hat: academic license offer

## CERN

- Site license for RHEL (Red Hat Enterprise Linux) at CERN (current and future versions)
- Additional RedHat products are NOT included (Extended Lifecycle Support, RedHat Virtualization, etc)
- Support is self support only

## "Extended research network"

- Up to 1000 free RHEL hosts (per site) – depending on how the site/institution is classified
  - Academic institutions – RHEL can be used for any non-commercial "research"
  - Non-Academic institutions (non-profit) – RHEL can be used, only if the majority of the "research" is focused on CERN (> 50%)
- Other sites are also able to negotiate directly with Red Hat

# Red Hat: What is 'research'

- We don't have the full legal definition yet, but something along the lines of:

*Research: servers exclusively doing computations, analysing and storing research data. HPC, Big data*

*Non-research: computing capacity required to run organisations backend systems such as email servers, file servers storing various documents, HR, ERP, backup servers etc.*



# Red Hat: What this means in practice (CERN)

- The Red Hat offer will satisfy some CERN use cases
- CERN also have systems that require a support contract (eg: Oracle database hosts)
  - We will likely have two contracts at CERN. Self support versus support
- CERN also has needs that can only be satisfied by CentOS Stream
- CERN will be using both distributions (RHEL/CentOS Stream) depending on the use case

# Red Hat: What this means in practice (CERN affiliated sites)

- First 1000 systems per site are completely zero cost (free)
  - For more than 1000 RHEL systems, there would be charges associated
- Onboarding of the site would be performed by the Red Hat CERN reseller
- CERN is not involved except for potential validation that the site is affiliated
- Host auditing is the responsibility of the site
- Host registration to Red Hat is NOT a hard requirement
- It would be up to the site to organise registration to Red Hat, or to manage localised content mirroring
  - CERN will not host Red Hat content for use outside of CERN
- Sites are under no obligation to join this offer

# Red Hat: Summary

- Red Hat's offer will help CERN in *some* use cases, but not all ...
- RHEL will NOT be a silver bullet for HEP
  - The offer is on the table for other sites to potentially use, but it is completely optional

# CERN is not working alone

- CERN has been working closely with FNAL and other HEPiX sites
  - Analysis of scientific use cases/scenarios for CERN/FNAL/HEP
  - A joint statement on our findings/proposal has been published on <https://linux.cern.ch> and to the [scientific-linux-users@fnal.gov](mailto:scientific-linux-users@fnal.gov) mailing list
- Linux Distribution landscape is complicated
  - We've outgrown the "*1 size fits all*" approach
  - Different use cases call for different distributions
    - It's inevitable that multiple distributions will need to be supported

# CERN / Fermilab Joint statement

CERN and Fermilab have been closely evaluating the Linux distribution landscape. We observe that national cyber infrastructure organizations are increasingly supporting more science domains, so in addition to LHC- or HEP-specific considerations, it will be useful to have a choice that is widely recognized and meets the needs of broader science research.

Red Hat has made a proposal to CERN regarding an academic licensing scheme. Ultimately this would require significant overhead at external sites, and therefore we have worries on this proposal's attractiveness for other sites.

Going forward, we propose to target CentOS Stream as the standard distribution for experiments. We feel that deploying CentOS Stream 8 is low risk, and we now have months of experience running IT services and experiment offline workloads on CentOS Stream 8 without any significant issues.

We feel that should issues arise with the adoption of CentOS Stream 8, it would be straightforward to reevaluate other options before CentOS Stream 8 support ends. CentOS Stream 8 is a supported distribution until May 2024. Trivial migration paths are provided by the various ELC (Enterprise Linux Clone) communities.

Continued support for existing workloads on Scientific Linux 7 and CERN CentOS 7 will be maintained as previously planned.

# Proposal (Straw man)


*"Going forward, we propose to target CentOS Stream as the standard distribution for experiments"*

- Access to the latest software, tools and kernel from Red Hat
- We have months of successfully operating CentOS Stream 8 in production, across diverse environments
  - Fully supported at CERN, both puppet managed and unmanaged hosts
  - Rate of change (system updates) are not as scary as we initially thought
  - System update stability
  - Try for yourself: `ssh lxplus8.cern.ch`
- CentOS Stream 8 is the distribution used for RedHat open source products (OpenStack, OpenShift, Ceph, etc)
- CentOS Stream 8 is free to download, distribute, and does not require any host registration


*"We feel that should issues arise with the adoption of CentOS Stream 8, it would be straightforward to reevaluate other options before CentOS Stream 8 support ends"*

- CentOS Stream 8 has an end of life of 31.05.2024
- Migration (without re-install) to RHEL or an ELC in the 8 family is trivial (ref: <https://linux.web.cern.ch/centos8/docs/migration>)

# CentOS Stream 8: update frequency



# CentOS Stream 8: iteration of change



- CentOS Stream 8 has more version history (# rpms)
- Faster iterations of change

```
dnf repoquery \  
--repofrompath=tmprepo,http://linuxsoft.cern.ch/cern/centos/$dist/$repo/x86_64/os \  
--repo=tmprepo --arch x86_64
```



# Proposal: Why CentOS Stream, and not an ELC (Enterprise Linux Clone)?

- CentOS Stream is still a Red Hat product
  - Red Hat is heavily invested in its success
- Latest kernel, software, bug fixes (released before RHEL, and ELCs)
- Strong open community, which allows the ability to influence the development of the OS (and downstream RHEL/ELC)
  - ELCs are a simple clone of RHEL, no real ability to contribute back to the community
- ELCs (Rocky/Alma) are still relatively 'new'
  - Will their traction continue with the release of RHEL8.5, and RHEL 9.0?

# Proposal: Why CentOS Stream, and not RHEL?

- This option is nice for CERN, but could prove awkward during collaboration with CERN affiliated sites
  - Sites may be subject to pay
  - Extra burdens for registering / auditing RHEL systems
  - Sites would need their own content mirroring
- Many upstream projects are not (yet) built against RHEL
  - OpenStack, OpenShift, Ceph, etc
- Using RHEL exclusively is not realistic, but it might be the best option for some use-cases
  - Those needing longer distribution lifecycles (10 years versus 5 years)
  - Those that can't virtualise or containerise

# Roadmap / timeline

- CentOS Linux 8 support ends on **31.12.2021**
  - C8 users are strongly recommended to migrate to CentOS Stream 8 (<https://linux.web.cern.ch/centos8/docs/migration>)
- CentOS Stream 9 is due to be released Q4 2021
  - CERN and FNAL will be supporting this distribution (Q1 2022)
- Red Hat academic licensing offer potentially available Q1 2022
- CERN and FNAL will continue to evaluate the distribution landscape, especially once RHEL9 and associated ELC (Enterprise Linux Clones) for 9 are released (2022)
- CentOS Stream 8 support ends on 31.05.2024
- CERN CentOS 7 / Scientific Linux 7 support ends on 30.06.2024

# Distribution landscape

	2021 H1	2021 H2	2022 H1	2022 H2	2023 H1	2023 H2	2024 H1	2024 H2	2025 H1	2025 H2	2026 H1	2026 H2	2027 H1
<b>LHC Schedule</b>	LS2		Run 3						LS3				
1. CC7 → CS8	CC7	CC7						Future OS					
		CS8											
2. CC7 → CS9	CC7			CC7			Future OS						
				CS9			CS9						
3. CC7 → RHEL8	CC7		CC7				RHEL8						
			RHEL8										
4. CC7 → RHEL9	CC7				CC7		RHEL9						
					RHEL9								

Note: existing C8 systems are assumed to migrate to CS8 before C8 EOL

Option #2 is the strawman proposal for WLCG sites

# Summary

- CERN CentOS 7 and Scientific Linux 7 are still supported operating systems
- CentOS Stream 8 is now a recommended operating system for new installations
- CentOS Linux 8 end of life is **31.12.2021**, migration from C8->CS8 is strongly recommended
- CentOS Stream 9 will be supported by CERN and FNAL, once released
  - Due to timings, the migration path from CC7->CS9 (or CC7->RHEL9) may make more sense than migration from CC7->CS8 (or CC7->RHEL8)
- RHEL site-license option is being reviewed by CERN management



[www.cern.ch](http://www.cern.ch)



# C8 versus CS8: CVE response time

