



Update on the Linux Strategy for CERN (and WLCG)

HEPiX Spring 2023

2023-03-28

Ben Morrice

On behalf of CERN IT

IT-CD



Agenda

- Recap: Why are we here?
- Change of recommended Linux distributions at CERN
 - Why / How / When?
 - CS8/CS9 decommissioning / migration
- CC7 migration path / recommendations

Why are we here?

- Dec 2020: CentOS Project deprecates CentOS Linux to focus on Stream
- Jan 2021: CentOS Stream 8 available for testing at CERN
- Feb 2021: 1st Linux Future Committee meeting to discuss options
- Oct 2021: Joint FNAL-CERN statement on supporting CentOS Stream 8
- Dec 2021: WLCG Management Board approves GDB's recommendation: Red Hat Enterprise Linux, CentOS Stream and RHEL rebuilds should be treated equally
- Feb 2022: CentOS Stream 9 available at CERN
- Mar 2022: Red Hat Academic Site License available at CERN
- Oct 2022: CERN starts reevaluating the original recommendation of CentOS Stream as the default Linux Operating System
- Oct 2022: CERN gathers feedback during HEPiX Autumn 2022 'Linux BoF'
- Dec 2022: CERN / FNAL announce change of recommended Linux Distributions

Reevaluation of CentOS Stream?

- CVE handling with Stream has been complicated
 - Hard to map RH Advisories to Stream packages
 - Several bugs have been encountered in Stream, which were not present in RHEL
- The CS8 release process is... not great.
 - Though should get better
[\(<https://lists.centos.org/pipermail/centos-announce/2023-February/086371.html>\)](https://lists.centos.org/pipermail/centos-announce/2023-February/086371.html)
- CS9 is in a better place, though Red Hat still hasn't put in place full public integration testing

Reevaluation of CentOS Stream?

Potential Options

RECAP

- Recommend **RHEL**
 - CERN has a site license
 - Not all HEP sites are this lucky, though some may profit from the ERN(*)
 - Recommend a **RHEL rebuild** (AlmaLinux / Rocky Linux)
 - CERN site license would not be taken advantage of
 - All HEP sites could profit (if they choose the same rebuild)
 - Recommend **both RHEL and a RHEL rebuild**
 - CERN site license could be taken advantage of for hosts at CERN
 - Some services could use a RHEL rebuild
 - Such as those that have interactions with other HEP sites/users
 - More distributions would need to be supported by the Linux team
 - A timeline to decommission CentOS Stream would also need to be planned
 - Continuing to use CentOS Stream could still be an option 🤔
- We would like to converge on a decision as soon as possible
 - Internal analysis is ongoing
 - Feedback from HEPIX and the wider CERN community will greatly assist
 - Information on any changes to the recommendation will be announced through a Linux Future Committee meeting
 - For **NEW** deployments at CERN:
 - We currently recommend to pause the deployment of CentOS Stream

Reference: <https://indico.cern.ch/event/1200682/contributions/5094007/>

Reevaluation of CentOS Stream?

Potential Options

RECAP

- Recommend **RHEL**
 - CERN has a site license
 - Not all HEP sites are this lucky, though some may profit from the ERN(*)
 - Recommend a **RHEL rebuild** (AlmaLinux / Rocky Linux)
 - CERN site license would not be taken advantage of
 - All HEP sites could profit (if they choose the same rebuild)
 - Recommend **both RHEL and a RHEL rebuild**
 - CERN site license could be taken advantage of for hosts at CERN
 - Some services could use a RHEL rebuild
 - Such as those that have interactions with other HEP sites/users
 - More distributions would need to be supported by the Linux team
 - A timeline to decommission CentOS Stream would also need to be planned
 - Continuing to use CentOS Stream could still be an option 🤔
- We would like to converge on a decision as soon as possible
 - Internal analysis is ongoing
 - Feedback from HEPIX and the wider CERN community will greatly assist
 - Information on any changes to the recommendation will be announced through a Linux Future Committee meeting
 - For **NEW** deployments at CERN:
 - We currently recommend to pause the deployment of CentOS Stream

Reference: <https://indico.cern.ch/event/1200682/contributions/5094007/>

Deciding on a RHEL rebuild?

- 2 valid contenders
 - AlmaLinux
 - Rocky Linux
- Both are *"1:1 binary compatible clone of RHEL"*



Deciding on a RHEL rebuild?

- 2 valid contenders
 - AlmaLinux
 - Rocky Linux
- Both are *"1:1 binary compatible clone of RHEL"*



Deciding on a RHEL rebuild?

- 2 valid contenders
 - AlmaLinux
 - Rocky Linux
- Both are *"1:1 binary compatible clone of RHEL"*



RHEL rebuild comparison


	AlmaLinux	Rocky Linux	"winner"
Ownership / Governance model	Non-Profit 501(c)(6)	"self-imposed not-for-profit organization" Public Benefits Corporation (PBC)	Too early to determine
Package version policy	Mimics RHEL	Previous package versions are not available (at least on 9)	AlmaLinux
CPU architecture support	x86_64, aarch64, ppc64le, s390x	x86_64, aarch64, ppc64le, s390x (ppc64le, s390x exist only on 9)	AlmaLinux
Community involvement	Strong presence within EPEL, Fedora, CentOS	Not strongly visible	AlmaLinux
Errata (updateinfo.xml)	https://errata.almalinux.org (for both 8 and 9)	https://errata.rockylinux.org (only for 8) As of February 2023, also for 9	AlmaLinux
Transparency	Board meetings minutes are public	Unknown	Too early to determine

RHEL rebuild comparison (continued)

	AlmaLinux	Rocky Linux	"winner"
Community contributions	https://almalinux.org/elevate	Unknown	AlmaLinux
Existing traction within HEP	CMS using AlmaLinux8 since Q1-2022, DESY using Alma	Some other HEP sites using Rocky	No obvious dominance
CVE update delay	"1 business day"	"1 business day"	Subjective – see next slides
Major/minor point release delay	"days"	"about a week"	AlmaLinux
Compatibility with CentOS SIGs	"Fully supported"	"Should work"	AlmaLinux
Build system	ALBS	Koji (for 8), Peridot (for 9)	Both use 'custom' systems
Sponsorship (longevity of the project)	CloudLinux (original 'parent'), aws, MS Azure, AMD, etc	CIQ (parent company), aws, Google cloud, MS Azure, etc	No concerns

CVE delay for the "8" family

Delay of CVE updates




AlmaLinux is 'slower',
however is more
consistent

Reference: <https://gitlab.cern.ch/morrice/pkgcompare/-/blob/master/errata.sh>
\$./errata.sh -r 8 -d 60 -g -a
\$ Rocky wins [26 Rocky wins (**39.2381** average hours, 824 total hours) versus 9 Alma wins (**12.9524** average hours, 272 total hours)]

CVE delay for the "9" family

Delay of CVE updates




AlmaLinux is faster, and is more consistent

Reference: <https://gitlab.cern.ch/morrice/pkgcompare/-/blob/master/errata.sh>
\$./errata.sh -r 9 -d 60 -g -a
\$ Alma wins [20 Alma wins (8.51613 average hours, 264 total hours) versus 15 Rocky wins (35.0968 average hours, 1088 total hours)]



Major/Minor (x.y) delays



AlmaLinux is faster, especially on major releases (eg 9.0)

AlmaLinux / RHEL Status

- Both AlmaLinux and RHEL are feature complete
 - Daily testing / weekly production release (with linux.cern.ch and email updates)
 - Base puppet modules work
 - IT already deploying production services on AlmaLinux & RHEL
 - VM (Alma/RHEL), Docker (Alma) and AIMS (pxeboot) images available
 - Lxplus availability: If you want to validate workflows on AlmaLinux, you may do so via
 - `ssh lxplus8.cern.ch` (AlmaLinux 8)
 - `ssh lxplus9.cern.ch` (AlmaLinux 9)
 - **NOTE:** `lxplus.cern.ch` will move to AlmaLinux 9 on **04.12.2023**
 - `lxplus7.cern.ch` will remain available until CC7 EOL (30.06.2024)
 - `lxplus8.cern.ch` will remain available until ~06.2026 "*End of Run 3 + 6 months*"
 - locmap is also available for desktop installations for RHEL/AlmaLinux
 - `dnf -y install locmap-release && dnf -y install locmap`

Reminder:

RHEL Academic Site License ERN

- As part of CERN's Academic Site License negotiations, we negotiated special conditions for HEP sites as part of our "Extended Research Network" (ERN)
 - Sites can obtain 1000 free RHEL licenses for the **purposes of research in partnership with CERN**
 - Interested sites need a commercial agreement with Red Hat and an "introduction" through CERN
- Larger HEP sites can also apply directly to Red Hat for a Site License agreement like CERN's

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

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

CS8/9 decommissioning

With the addition of AlmaLinux, the Linux team now supports **8** Linux distributions

- CC7: CERN CentOS 7
- CS8: CentOS Stream 8
- CS9: CentOS Stream 9
- RHEL7: Red Hat Enterprise Linux 7
- RHEL8: Red Hat Enterprise Linux 8
- RHEL9: Red Hat Enterprise Linux 9
- ALMA8: AlmaLinux 8
- ALMA9: AlmaLinux 9

CS8/9 decommissioning

With the addition of AlmaLinux, the Linux team now supports **8** Linux distributions

- CC7: CERN CentOS 7
- ~~CS8: CentOS Stream 8~~ (as of 30.09.2023 - see also [OTG0074647](#))
- ~~CS9: CentOS Stream 9~~ (as of 30.06.2023 - see also [OTG0074647](#))
- RHEL7: Red Hat Enterprise Linux 7
- RHEL8: Red Hat Enterprise Linux 8
- RHEL9: Red Hat Enterprise Linux 9
- ALMA8: AlmaLinux 8
- ALMA9: AlmaLinux 9

We plan to drop the support for CentOS Stream, bringing the total number of supported Linux operating systems at CERN to **6**

CS8/9 decommissioning

With the addition of AlmaLinux, the Linux team now supports **8** Linux distributions

- ~~CC7: CERN CentOS 7~~ (as of 30.06.2024)
- ~~CS8: CentOS Stream 8~~ (as of 30.09.2023 - see also [OTG0074647](#))
- ~~CS9: CentOS Stream 9~~ (as of 30.06.2023 - see also [OTG0074647](#))
- ~~RHEL7: Red Hat Enterprise Linux 7~~ (as of 30.06.2024)
- RHEL8: Red Hat Enterprise Linux 8
- RHEL9: Red Hat Enterprise Linux 9
- ALMA8: AlmaLinux 8
- ALMA9: AlmaLinux 9

We plan to drop the support for CentOS Stream, bringing the total number of supported Linux operating systems at CERN to **6**

- Supporting around **4** distributions seems to be the 'sweet spot'
 - We'll get back to this level on 30.06.2024

CS8/9 migration to AlmaLinux 8/9 or RHEL 8/9

- The migration of CentOS Stream 8 or 9 to either AlmaLinux 8/9 or RHEL 8/9 is relatively painless process
 - Provided it's the same family (8->8 or 9->9)
 - Full process (including instructions for puppet centrally managed hosts): <https://linux.cern.ch/migration>
 - **Note:** as Stream is ahead of AlmaLinux/RHEL - migration may involve some package **downgrades**

```
dnf remove epel-next-release
# Migrate from CentOS Stream to AlmaLinux
dnf --nogpgcheck --
repofrompath='tmpcern,https://linuxsoft.cern.ch/cern/alma/$releasever/CERN/$basearch/' swap
centos-stream-release almalinux-release

# Migrate from CentOS Stream to RHEL
dnf --nogpgcheck --
repofrompath='tmpcern,https://linuxsoft.cern.ch/cern/rhel/$releasever/CERN/$basearch/' swap
centos-stream-release redhat-release

# Common steps
dnf distro-sync --allowerase --disablerepo=openafs*
dnf distro-sync
reboot
```

CS8/9 migration to AlmaLinux 8/9 or RHEL 8/9

- The migration of CentOS Stream 8 or 9 to either AlmaLinux 8/9 or RHEL 8/9 is relatively painless process
 - Provided it's the same family (8->8 or 9->9)
 - Full process (including instructions for puppet centrally managed hosts): <https://linux.cern.ch/migration>
 - **Note:** as Stream is ahead of AlmaLinux/RHEL - migration may involve some package **downgrades**



If you are using Rocky Linux, you can (if wanted/needed/required) also migrate to AlmaLinux

CC7 migration path

- **90%** of Linux distribution usage at CERN is CC7
 - General end-of-life is 30.06.2024
 - Users need to move to another distribution, but with so multiple options – what should they do?



CC7 migration path


- **90%** of Linux distribution usage at CERN is CC7
 - General end-of-life is 30.06.2024
 - Users need to move to another distribution, but with so multiple options – what should they do?



- CERN IT recommends CC7 users to migrate to the **9** family (AlmaLinux, Red Hat Enterprise Linux)
 - CERN IT do not currently recommend an in-place upgrade, and still propose to perform a "fresh" reinstallation



Why should I 'jump' to the 9 family, skipping 8 ?

- Previously it was common or best practice to migrate from *major x* to *major x+1*
- Usually, we waited until critical bugs were fixed from the 'early' releases (eg: 7.0-7.2, 8.0-8.2)



Why should I 'jump' to the 9 family, skipping 8 ?


- Previously it was common or best practice to migrate from *major x* to *major x+1*
- Usually, we waited until critical bugs were fixed from the 'early' releases (eg: 7.0-7.2, 8.0-8.2)



- RHEL8 was a 'bad' release
 - Modularity was/is a mess
 - CentOS Stream 8 arrived after RHEL8
 - Stream 8 -> RHEL8 workflow took forever to implement (only recently)
 - CentOS Linux 8 deprecated -> FUD
- Red Hat are now releasing RHEL major releases more frequently (every 3 years compared to every 5 years)
 - We don't need to wait as long to perform upgrades
 - CentOS Stream 10 / RHEL10 are coming :)

Why support 8 if the recommendation is to use 9?

- Some user communities are already using 8 in production (CMS)
- ~10% of CERN IT managed hosts (and others) are currently using CentOS Stream 8
 - The migration path from CentOS Stream 8 to another 8 distribution is feasible
 - See <https://linux.cern.ch/migration> for more details on migration




What about the WLCG?

- AlmaLinux for everyone????
 - Not really. Not at all.
 - Nothing has changed from the original [GDB recommendation](#)

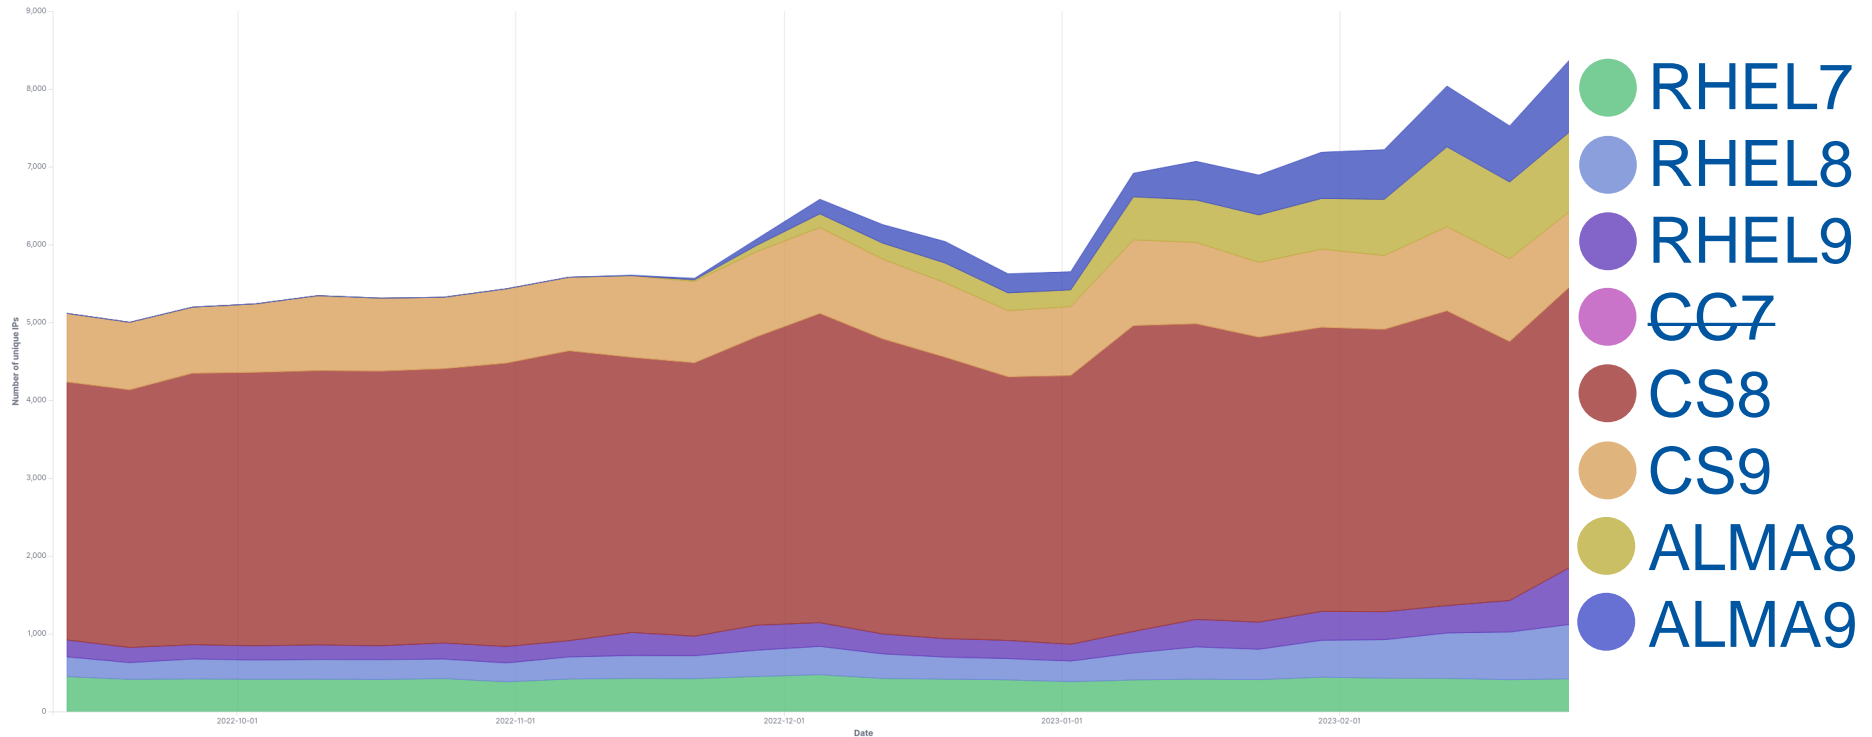
GDB recommends

- *RHEL, CentOS Stream and 100% Enterprise Linux Clones should be **treated equally***
- *Application workloads and middleware can be made to work on any of these (see HEPiX talk) using techniques such as containerization or ABI/API compatibility*
- *If there are issues, these should be reported to the authors (distributions, experiments or middleware)*
- *WLCG sites should plan migrations before 7 EOL (mid 2024)*
- *From **7 to 9 seems promising** during the period 2H 2022-2023 subject to middleware availability*
- *CS9 already available, RHEL/ELC 9s coming in the next few months*
- *8 may also be used already as needed for hardware compatibility or functionality*


Evolution of Linux distribution usage (all distributions)



Evolution of Linux distribution usage (excluding CC7)




Evolution of Linux distribution usage (RHEL/ALMA 8/9)



Questions ? / Discussion





www.cern.ch



Backup slides



Should I use AlmaLinux or RHEL?

- CERN IT supports both AlmaLinux and Red Hat Enterprise Linux
 - AlmaLinux is freely distributable
 - Red Hat Enterprise Linux is a licensed product (though CERN is fortunate enough to hold a site-license) and cannot be distributed outside of CERN
- CERN provided Docker images for RHEL do not (and will not) exist
- CERN provided Docker images for AlmaLinux DO exist (and can be redistributed)

Should I use AlmaLinux or RHEL?

- CERN IT will not dictate which distribution experiments or services use
- Please check <https://linux.cern.ch/which> for full details

Situation / Use-case	Distribution to use
I use software that requires support from a vendor (eg: Oracle, WinCC OA)	RHEL should be used
I need to distribute a virtual machine or docker image to users outside of CERN	AlmaLinux must be used (please see the RHEL page for more details)
My service provides interactive access to non CERN sites (eg: lxplus or lxbatch)	AlmaLinux should be used
My situation is not described above, can I use either AlmaLinux or RHEL?	Yes, you may use either - the choice is yours!

Issues on Stream references

- [0] <https://github.com/systemd/systemd/pull/21558>
- [1] https://bugzilla.redhat.com/show_bug.cgi?id=1966712
- [2] https://bugzilla.redhat.com/show_bug.cgi?id=2123319
- [3] https://bugzilla.redhat.com/show_bug.cgi?id=2055222
- [4] https://bugzilla.redhat.com/show_bug.cgi?id=2037807
- [5] https://bugzilla.redhat.com/show_bug.cgi?id=2046021
- [6] <https://cern.service-now.com/service-portal?id=outage&n=OTG0064504>
- [7] https://bugzilla.redhat.com/show_bug.cgi?id=2119304
- [8] https://bugzilla.redhat.com/show_bug.cgi?id=2043753
- [9] https://bugzilla.redhat.com/show_bug.cgi?id=2138019