



HSF Software Developer Tools & Packaging Working Group #13

The status of Linux distributions for HEP

2023-03-15

Ben Morrice

On behalf of CERN IT

IT-CD-CLI

Agenda

- Change of recommended Linux distributions at CERN
- CS8/CS9 decommissioning / migration
- CC7 migration path
- The CMS experiment: Experiences with AlmaLinux

Future plans (potential options)

- Instead of recommending CentOS Stream, the options **could** be:

1. Recommending RHEL instead

- CERN has a site license
- Not all HEP sites are this lucky

2. Recommending a RHEL rebuild (AlmaLinux / Rocky Linux)

- CERN site license would not be optimised
- All HEP sites can profit

3. Recommending **both** RHEL and a RHEL rebuild

- CERN site license could be optimised 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

- If any of these options were to happen, a timeline to decommission support for CentOS Stream would also need to be planned
- Of course, continuing to use CentOS Stream could still be an option 🤔

RECAP

<https://indico.cern.ch/event/1208501>

Future plans (potential options)

- Instead of recommending CentOS Stream, the options **could** be:

1. Recommending RHEL instead
 - CERN has a site license
 - Not all HEP sites are this lucky
2. Recommending a RHEL rebuild (AlmaLinux / Rocky Linux)
 - CERN site license would not be optimised
 - All HEP sites can profit

RECAP

<https://indico.cern.ch/event/1208501>

3. Recommending **both** RHEL and a RHEL rebuild
 - CERN site license could be optimised 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

- If any of these options were to happen, a timeline to decommission support for CentOS Stream would also need to be planned
- Of course, continuing to use CentOS Stream could still be an option 🤔

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"*



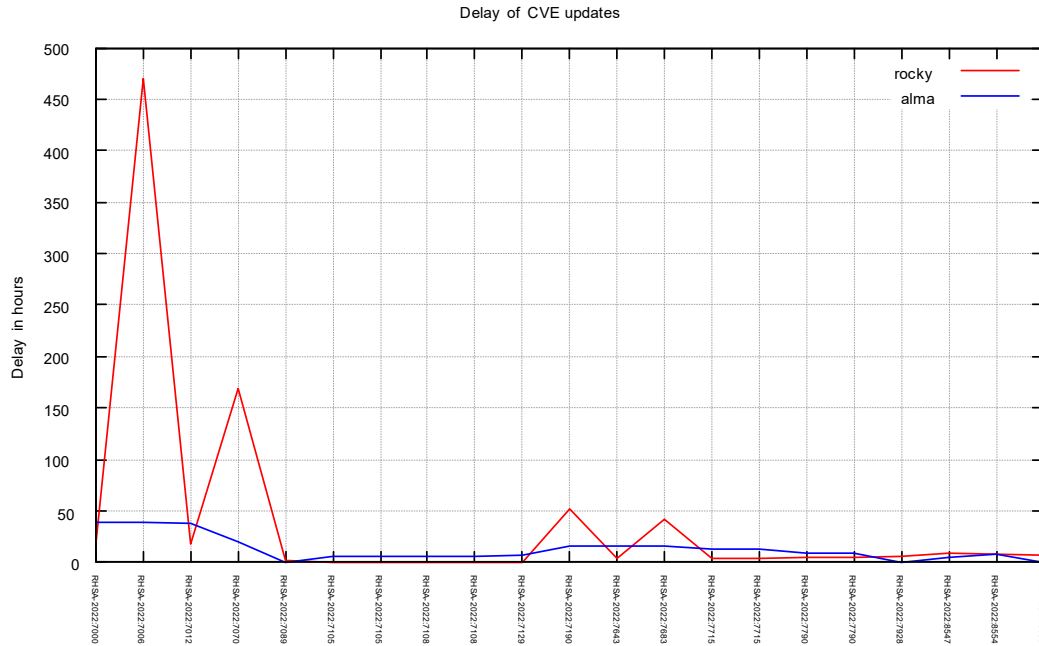
RHEL rebuild comparison

	AlmaLinux	Rocky Linux	"winner"
Ownership / Governance model	Non-Profit 501(c)(6)	<i>"self-imposed not-for-profit organization"</i> Public Benefits Corporation (PBC)	AlmaLinux
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 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

	AlmaLinux	Rocky Linux	"winner"
Community contributions	https://almalinux.org/elevate	Unknown	AlmaLinux
Existing traction within HEP	CMS using Alma 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



AlmaLinux is 'slower',
however is more
consistent

Reference: <https://gitlab.cern.ch/morrice/pkgcompa re/-/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)]

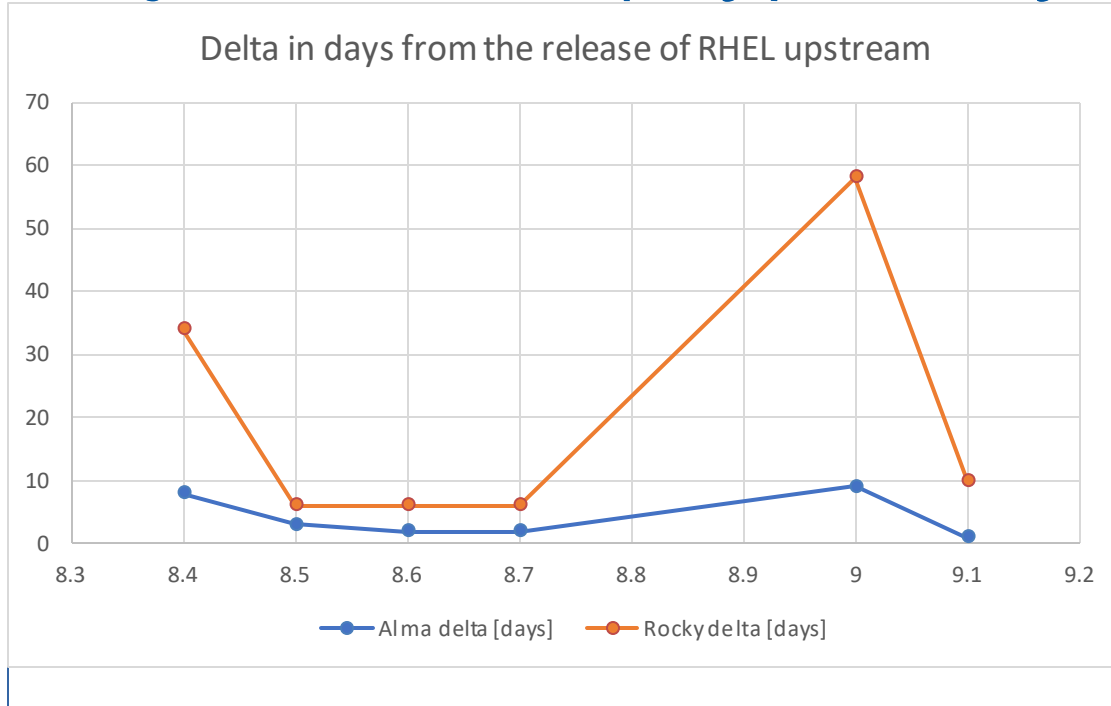
CVE delay for the "9" family



AlmaLinux is faster, and is more consistent

Reference: <https://gitlab.cern.ch/morrice/pkgcompa/re/-/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)

Decision making process / communication

- Internal (CERN Linux) analysis between Rocky Linux and AlmaLinux started in 10.2022 (although this had been watched closely since the inception of both projects)
- LFC #6 explained the potential for change [0]
- CERN gathered HEP community usage at HEPiX Autumn 2022 [1]
- Discussions between CERN and FNAL also took place to ensure we were aligned
- Once a decision had been made 'internally', approval was required from IT management [2], [3]
- After CERN IT management approval, an internal Service Now "outage" was released on 07.12.2022 [4]
- Joint statements from FNAL and CERN were published on 07.12.2022 [5], [6]
- Email notification to Linux Future Committee members sent on 08.12.2022 informing them of the news and invitation for LFC #7 (this presentation)
- LFC #7 explains this process (similar content to this presentation) [7]
- Information was communicated in ITUM-38 [8]

[0] <https://indico.cern.ch/event/1208501> [1] <https://indico.cern.ch/event/1200682/sessions/459106/#20221031>

[2] <https://indico.cern.ch/event/1210018> [3] <https://indico.cern.ch/event/1225366>

[4] <https://cern.service-now.com/service-portal?id=outage&n=OTG0074647>

[5] <https://news.fnal.gov/2022/12/fermilab-cern-recommendation-for-linux-distribution>

[6] <https://linux.web.cern.ch/#fermilabcern-recommendation-for-linux-distribution>

[7] <https://indico.cern.ch/event/1229565> [8] <https://indico.cern.ch/event/1229771>

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
 - ai-bs has Alma/RHEL support (`--alma8`, `--alma9`, `--rhel8`, `--rhel9`)
 - 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`

CS8/9 decommissioning

With the addition of AlmaLinux, the CERN 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 at any one point in time 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 --allowerasing --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
 - Extended Lifecycle Support prolongs this date to 30.06.2026 (with caveats)
 - Users need to move to another distribution, but with so many 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
 - Extended Lifecycle Support prolongs this date to 30.06.2026 (with caveats)
 - Users need to move to another distribution, but with so many 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 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



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
- Docker images for RHEL do not (and will not) exist
- 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!


Evolution of Linux distribution usage (all distributions)



Evolution of Linux distribution usage (excluding CC7)



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




The CMS experiment: *Experiences with Alma Linux*



A summary from the original presentation by Danilo Piparo at the LFC#6 (<https://indico.cern.ch/event/1229565/>)

Questions ? / Discussion





www.cern.ch

