

Future Linux Distribution Scenarios

Tim Bell
tim.bell@cern.ch
CERN IT-CM

Grid Deployment Board
12.01.2021

CentOS 7/8 (was 2019, now out of date)

End of Lifetime (EOL) Dates			
	CentOS-6	CentOS-7	CentOS-8
Full Updates ¹	May 10th, 2017	Q4 2020	May 2024
Maintenance Updates ²	November 30th, 2020	June 30th, 2024	May 31st, 2029
Architecture			
	CentOS-6	CentOS-7	CentOS-8
Latest Release	6.10	7 (2009)	8 (2004)
Supports i386 Systems (32-bit Pentium, AMD, Via)	Yes ³	Yes (1503)	No

- Within the (previous commitment) release
 - Full support (enhancements etc.) first 5 years
 - Maintenance support next 5 years
 - Extended support later available only RHEL(\$)

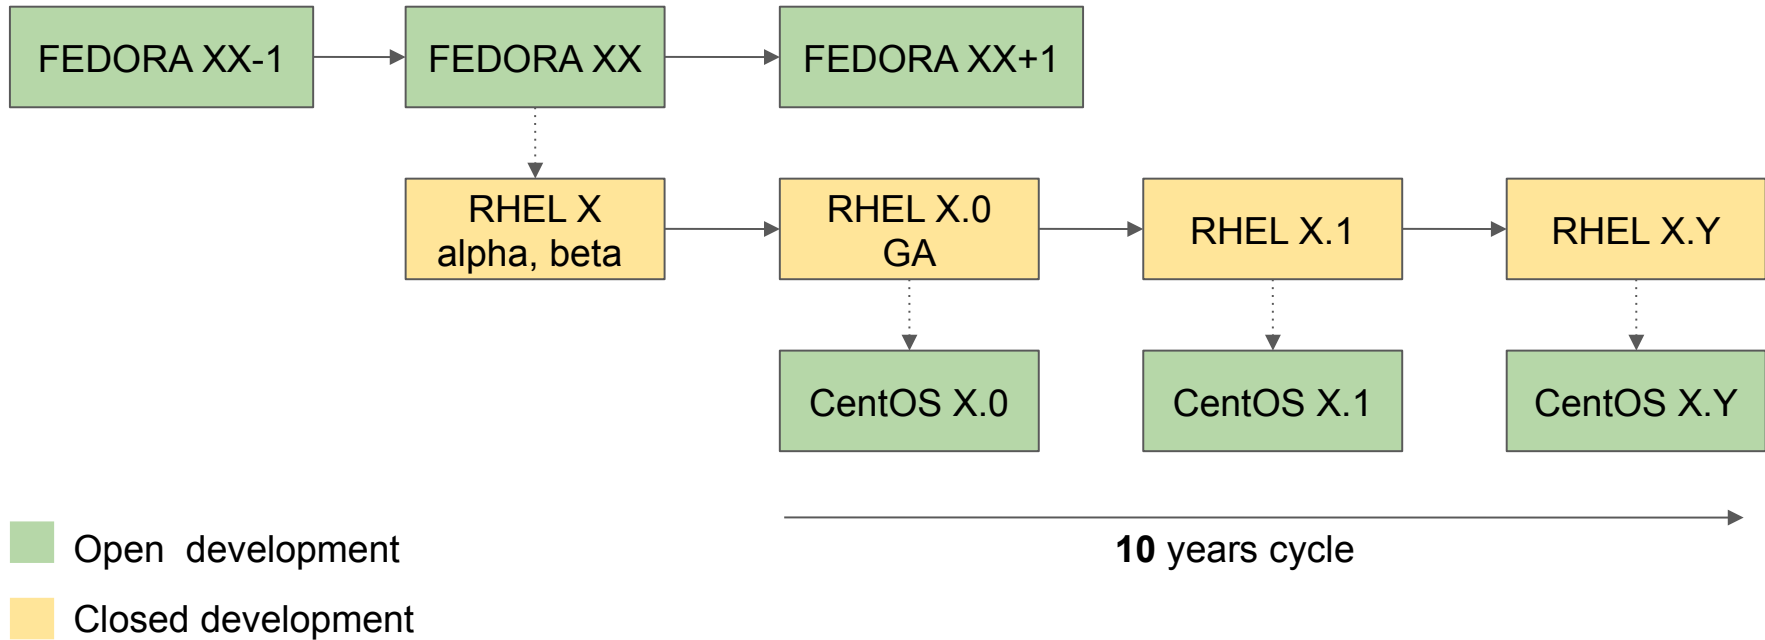
CentOS Governance

- Rebuild mandate as community open source Linux, binary compatible with Red Hat Enterprise Linux
- Red Hat employed all of the community members in 2014
 - And IBM purchased Red Hat in 2019 for **34B USD**
- **Board structure** has 11 members with 3 non-Red Hat including CERN and Fermilab employees
- CERN has hosted and participated in community events
 - e.g. CentOS dojos at CERN - **2018, 2017**
 - FOSDEM

CERN situation before Dec 8th 2020

- Production
 - CERN CentOS 7 – 40K hosts - Majority of OS build work done upstream
 - WLCG physics workloads
 - Online
 - Services (prior to C8 availability Q2 2020)
 - CentOS 8 – 4K hosts – OS release from upstream build only, local automation
 - New services starting with this
 - Would become default towards the end of Run 3 in 2024
- Retired (Nov 2020)
 - Scientific Linux 6
 - No further updates available



Previous Red Hat model



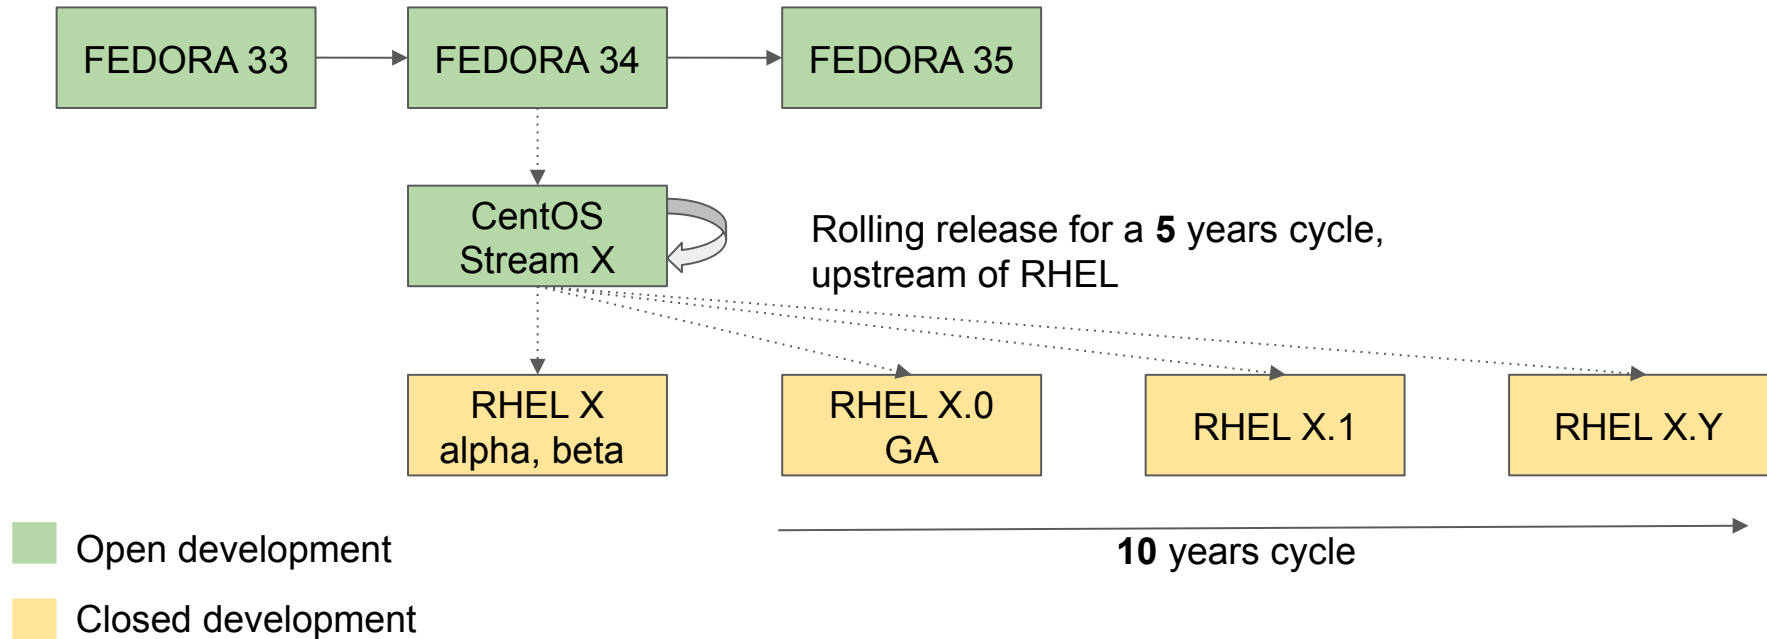
Previous CentOS schedule

Release	Beta avail	Production	Maintenance	End of Life
CentOS 7		2014-07	2019-07	2024-06
CentOS 8 rebuild		2019-09	2024-08	2029-05
CentOS 8 stream		2019-09	2024-08	2029-08
CentOS 9 (est.)	2021	2023		2028

- CentOS stream is a distro derived from the very latest patches for RHEL (i.e. upstream of RHEL) i.e. not 100% compatible with current RHEL release
- CERN has only deployed CentOS 8 rebuild





Updated Red Hat model



Updated CentOS schedule

Release	Beta avail	Production	Maintenance	End of Life
CentOS 7		2014-07	2019-07	2024-06
CentOS 8 rebuild		2019-09	(2024-08)	2021-11 (2029-05)
CentOS 8 stream		2019-09	(2024-08)	2024-08 (2029-05)
CentOS 9 stream (est.)	Q1 2021	2023		2028

- CentOS 8 rebuild will become End Of Life in 2021 ... 8 years shorter than announced
- CentOS 8 stream will become the only CentOS 8 distribution
- CentOS 8/9/... lifetime to move to 5 years maximum, no maintenance support phase
 - Release frequency moves from 5 years to 3 years
 - Beta available prior, development in the open



Implications

- EPEL/ELrepo extended software repository are currently RHEL compatible and used to distribute much of our software
 - EPEL-next is under discussion for stream dynamism
- CentOS SIG community software to base off CentOS stream ?
 - **RDO** OpenStack, Virtualisation
 - FreeIPA, Keycloak, Openshift, Ceph, Puppet, owncloud, Grafana and many more to be seen
- CentOS lifetime from 10 to 5 years (stream), stop rebuild in 1 year
 - **Very unusual change** mid-release by an open source project
 - Unlikely to be aligned with all of the stakeholders windows
 - Or hardware lifetimes (~5 years)
 - Experiment/CERN software certification often 3-4 years after release today
 - Commercial software (e.g. WinCC OA) also need regular certification

Scenario #1

- Follow CentOS 8 stream
 - Being built by the Red Hat teams as an upstream RHEL
 - Aligned with CentOS SIGs for community projects
- Upgrading during the runs would need planning for online or offline software
 - 5 year lifecycle and streaming change would be a different mindset along with significant implications for software suppliers
- WLCG would become multi version environment
 - Sites unlikely to be able to keep 100% in sync
 - Software providers will need to support several releases
 - CentOS stream loses 100% RHEL compatibility
- Containers / CVMFS may help for WLCG but are not a magic bullet yet
 - Lots of work to be done before ALL workloads are ready
 - [Red Hat Universal Base Image](#) may also help
 - Some isolation but not full
 - e.g. latest releases drop old SSL breaking storage access

Scenario #2

- Build a RHEL compatible distro in collaboration with Fermilab (or adopt a community Special Interest Group such as [Rocky](#), [CloudLinux](#) or [Springdale](#))
 - Scientific Linux remains a strong brand, much interest in our plans
- Process is more complex than SL days e.g. Application Streams
- CentOS would not help extend lifetime beyond 5 years
 - [Q13](#) - “We will not be putting hardware, resources, or asking volunteers to work towards that effort, nor will we allow the CentOS brand to be used for such a project.”
- We would need to maintain build/test process
 - And re-engineer this for each release
 - New 3 year cycle for RHEL means more releases
- RHEL source code availability and format is uncertain at 5 years
 - Non-copyleft sources (Apache, BSD, MIT, ...) ?
- Many community packages will be built against stream so we would need to rebuild/re-engineer them too

Scenario #3

- Red Hat may offer a RHEL deal for some scientific sites similar to their [developer program](#), i.e. non-commercial
- Potential 10 year support e.g. for WLCG or online
- Full RHEL compatibility, would therefore need Red Hat OpenStack/Storage/Identity Management (\$)
 - Probably OpenStack RDO/Ceph/FreeIPA Community editions would be stream based
- Would it cover the entire grid (and other sciences)?
- Significant lock in risk with 44,000 machines @ 300CHF/machine/year
 - 13MCHF/year (CERN only 20% of WLCG resources)

Scenario #4

- Select alternative distro, Ubuntu ? Oracle Linux ? openSUSE ?
 - Debian community not huge and only support 3 years
 - Oracle Linux only has a subset of the RHEL packages
- **Ubuntu release cycle** ~ every 2 years alternate short/long support
 - Long Term Support releases open source updates for 5 years
 - Longer, need to purchase **Extended Security Maintenance**
- Significant retraining / reskilling needed for administrators
 - Only some open source tool chains run on both e.g. Puppet
 - RPM, selinux, anaconda ... significant differences
- Would need parallel running until at least 2024
- May limit support from the Red Hat sponsored community edition software e.g. Ceph, FreeIPA, Keycloak, ...
 - Need local builds, testing, documentation, ...

WLCG Thoughts

- Need for prototyping future distribution during lifetime of '7'
 - Stream ? Another rebuild ?
 - Some young communities and uncertainty
- May see more divergence between services and worker nodes
 - Software providers may need to produce more targets
- CVMFS/containers do give some risk reduction
- Many scientific sites are looking to our community for a common solution (HPC+HTC)
- Deadlines
 - End 2021
 - what to do for existing '8' installations ?
 - Mid 2024
 - last 6 months of Run 3 currently after EOL '7' but analysis completion probably into 2025 – '8' migration during the run ?
 - Prepare for Run 4 environment on '9' ? '10' ?

Where now?

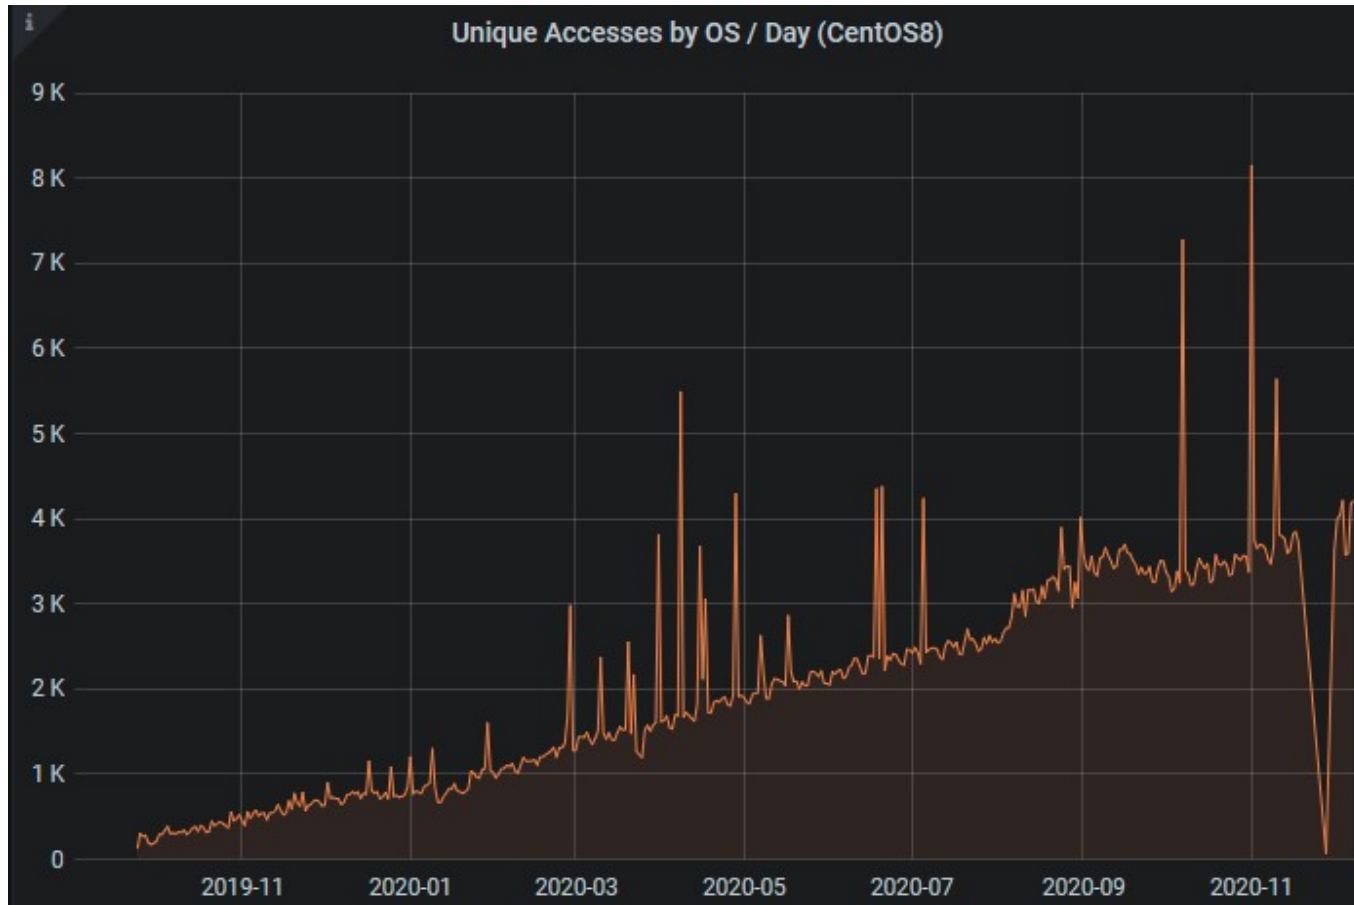


- Review non-urgent migrations to CentOS 8 – especially Pets
 - CentOS 7 support ends around same time as CentOS 8 stream in 2024
- Continue to evangelise 10 year support with the CentOS community
- ARM support will now be more complex and dynamic
- CERN and FNAL communications team issued a statement (slide 20)
- Consult on future approach with Fermilab, HEPiX, WLCG/OSG, accelerator/online teams, Linux Certification committee, ...
- Prepare CentOS 8 stream environment and repositories for new installs and to bridge to 2024 for those already migrated
 - Gain experience on the new model to assess viability for our use cases
- How can we adapt to the more dynamic distro world ?
 - We have had 10 year lifecycle support for a long time
- Thoughts/Comments would be welcome

References

- CentOS Community Announcement (8th December 2020)
 - <https://blog.centos.org/2020/12/future-is-centos-stream/>
- CentOS Community FAQ
 - <https://centos.org/distro-faq/>
- Red Hat Statements
 - <https://www.redhat.com/en/blog/centos-stream-building-innovative-future-enterprise-linux>
 - <http://crunchtools.com/before-you-get-mad-about-the-centos-stream-change-think-about/>
- Feedback from the community – opinionated :-)
 - [Petition](#) signed by over 10000 people
 - A community response [web site](#)
 - [Ycombinator](#)
 - [CentOS-devel](#)

CentOS 8 usage at CERN



Distribution Lifetimes

Distribution	Lifetime (years)	Release Frequency (years)	Link
Red Hat Enterprise	10*	5 (going to 3)	Link
CentOS rebuild	10	5	
CentOS stream	5	3	Link
Fedora	1.1	0.5	Link
Oracle Linux	10	5 (going to 3)	Link
Ubuntu LTS	5*	2	Link
Debian	3	2	Link
openSUSE	3	3	Link
FreeBSD	5	2	Link

* Can be extended a further 5 years at additional cost

Note: Based on various sources and has not been validated with distros

Rebuild Community Initiatives

- Rocky Linux (starting up, no ETA for 8)
 - <https://rockylinux.org/>
- CloudLinux (8 in preparation)
 - <https://blog.cloudlinux.com/announcing-open-sourced-community-driven-rhel-fork-by-cloudlinux>
- Springdale from Princeton (currently no 8 but planned)
 - <http://springdale.math.ias.edu/>

CERN / FNAL Statement (Dec 17 2020)

- “CERN and Fermilab acknowledge the recent decision to shift focus from CentOS Linux to CentOS Stream, and the sudden change of the end of life of the CentOS 8 release. This may entail significant consequences for the worldwide particle physics community. We are currently investigating together the best path forward. We will keep you informed about any developments in this area during Q1 2021. “

<https://listserv.fnal.gov/scripts/wa.exe?A2=LINUX-USERS;f8f017a8.2012>