



Site survey for LINUX options

Julia Andreeva
CERN

Motivation


- Get feedback from the sites regarding their current experience/concerns with various LINUX options
- Understand whether site experience is in line with the recommendations of CERN and FNAL
- Help in preparing WLCG guidance for the sites

Thank you for taking part!

- 90 sites took part in the survey
- T0, 12 T1, 71 T2, 6 T3
- Most of answers we've got are from the European sites

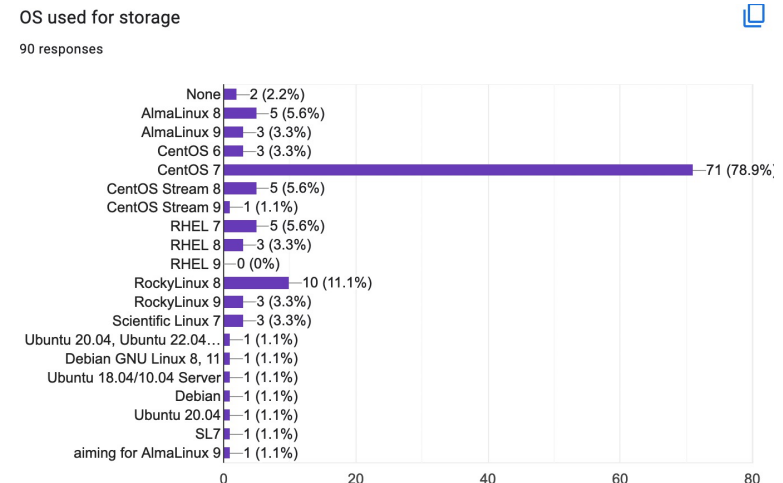
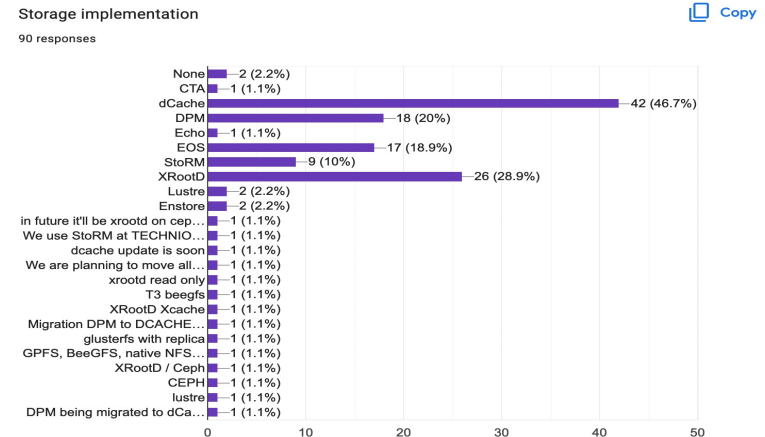
The lowest tier level of your site

90 responses



Storage

- dCache is the most popular solution for storage implementation, in particular taking into account that majority of DPM sites plan to migrate to dCache
- Majority of sites are still on CentOS7, though migration started:
 - 8 sites Alma Linux (8 and 9)
 - 3 sites RHEL8 , 5 RHEL 7
 - 13 Rocky Linux (8 and 9)




Batch system

- Htcondor and Slurm are the most popular solution for batch implementation
- Majority of sites are still on CentOS7, though migration started:
 - 7 sites Alma Linux (8 and 9)
 - 7 sites Rocky Linux (8 and 9)
- Sites which are using RHEL7 might have licence and will upgrade to later RHEL versions

Batch system implementation

 Copy


90 responses



OS used for batch

 Copy

90 responses




Computing Element

- ARC and HTCondor are the most popular solution for batch implementation.
- Majority of sites are still on CentOS7, though migration started:
 - 3 sites Alma Linux (8 and 9)
 - 4 sites Rocky Linux (8 and 9)
- Sites which are using RHEL7 might have licence and will upgrade to later RHEL versions


CE implementation

90 responses



OS used for CE

90 responses



Migration plans from the sites

- 20 sites are planning or already using AL
- 23 sites are planning or already using RL, there are 3 which are already on RL but consider migrating to AL in order to follow CERN/FNAL recommendation
- 6 sites are hesitating between AL and RL
- 11 sites did not make up their mind yet and are waiting for the recommendation from WLCG
- Sites with licences (mainly T1s) plan to have RHEL+RL or RHEL+AL
- There are also sites considering or already using Ubuntu or Debian

Closer look at T1s, current status

Site Name	OS used for storage	OS used for batch	OS used for CE
RRC-KI-T1	CentOS 7	CentOS 7	CentOS 7
NDGF-T1	AlmaLinux 8, CentOS 7, CentOS Stream 8, CentOS Stream 9, RHEL 8, Ubuntu 20.04, Ubuntu 22.04, Gentoo 2.9	AlmaLinux 8, CentOS 7, Ubuntu 20.04, Gentoo 2.9	AlmaLinux 8, CentOS 7, Ubuntu 20.02, Gentoo 2.9
CC-IN2P3	CentOS 7	CentOS 7	RHEL 7
JINR-T1	CentOS 7	CentOS 7	CentOS 7
T1_US_FNAL	Scientific Linux 7	CentOS 7	CentOS 7
KR-KISTI-GSDC-01	CentOS 6, CentOS 7	CentOS 7	CentOS 7
NIKHEF-ELPROD	CentOS 7	CentOS 7	CentOS 7
RAI-LCG2	CentOS 7	RockyLinux 8	CentOS 7
TRIUMF-LCG2	Scientific Linux 7	Scientific Linux 7	Scientific Linux 7
PIC Tier-1	CentOS 7, RockyLinux 8	CentOS 7, RockyLinux 8	CentOS 7, RockyLinux 8
INFN-T1	CentOS 7, CentOS Stream 8, RHEL 8	CentOS 7	CentOS 7
BNL	RHEL 7, RHEL 8	AlmaLinux 8, CentOS 7, SL7	RHEL 7

Closer look at T1s, plans

RRC-KI-T1	Migration to Debian (where it is possible, cor core services, for example) and AlmaLinux
NDGF-T1	CentOS 7 sites going to Alma or Rocky, Ubuntu sites happy and staying. As a distributed tier-1 we're going to be running on whatever the sites chose to run, especially for batch.
CC-IN2P3	Usage of rocky and rhel (depending of the service) is expected (some migration already started)
JINR-T1	AlmaLinux 8
T1_US_FNAL	Short term Centos stream 8, working towards Alma 9 longer term
KR-KISTI-GSDC-01	AlmaLinux (because we do not have RHEL licenses)
NIKHEF-ELPROD	Upgrading to Rocky 8, later Rocky 9
RAI-LCG2	We have already migrated our batch farm to Rocky 8. this was due to using the latest AMD CPUs which didn't have full CentOS7 support. Where possible we will move other services directly to Rocky 9.
TRIUMF-LCG2	RHEL 9 plus AlmaLinux 9
PIC Tier-1	We started to move to Rocky OS 8 as CentOS substitute. However, after the news from CERN and FNAL we are thinking of moving to Alma Linux 9 in future (before June 2024).
INFN-T1	For storage we are considering RHEL9 and Rocky9. For farm nodes we plan to avoid stream, going straight to Rocky/Alma that we consider even.
BNL	RHEL 8/9 on storage, and other critical systems AlmaLinux 8 on our HPC, and AlmaLinux 9 on HTC

Some feedback from the sites regarding AL

- AlmaLinux has fast update release and seem to be the choice of many other sites
- Quick propagation of updates from RHEL downstream;
- Availability and support of third-party software;
- Availability of beta releases for the next minor releases;
- Compatibility with 3rd party kernel modules;
- AlmaLinux Cloud images were available more easily/earlier.
- AlmaLinux was the first EL8 OS (aside from RHEL) to provide GPT partition tables, the others may still be MBR even now.
- AlmaLinux seemed to have better community/governance.
- **No issues observed so far with our choice of Alma, that we would not have had with the commercial upstream distro, or another rebuild;**
- Alma 9: some software not yet supported, and many puppet forge modules we use are not compatible with 9 yet

Some feedback from the sites regarding RL

- Rocky was the first being available
- Entirely positive experience, very fast feedback by developers in case e.g. mirror issues arise.
- Rocky is officially supported by Mellanox/Nvidia OFED drivers
- Great and fast support by the RockyLinux developer community, extensive support in public tooling (Puppet modules etc.) even compared to e.g. AlmaLinux
- Rocky * is supposed to be a faithful rebuilt of RHEL with binary compatibility. It has good hardware vendor support.
- Broad adoption also outside of CERN / HEP community (two times as many hits from Rocky seen on EPEL repos compared to Alma, Rocky even exceeds CentOS stream).
- Longer commitment (in terms of end of security updates dates) than Alma.
- People would like to stick with RL9 because migration to AL9 would result in reinstall of already done services (its just more work).

See more comments on the backup slide 14

RL vs AL

- AlmaLinux offers an easy upgrade from CentOS 7, but official support from hardware vendors for Rocky Linux is better. We consider the latter to be more important.
- Broad adoption of RL also outside of CERN / HEP community (two times as many hits from Rocky seen on EPEL repos compared to Alma, Rocky even exceeds CentOS stream).
- Great and fast support by the RockyLinux developer community, extensive support in public tooling (Puppet modules etc.) even compared to e.g. AlmaLinux
- We would like to stick with RL9 because migration to AL9 would result in reinstall of already done services (its just more work).
- Middleware support, updates for RL appear a few days later than in Alma Linux
- **We would really like to see the working behind why Alma was chosen over Rocky. Is this something those of us who went down the Rocky path should worry about?**

General comments

- Things to be considered for the OS choice:
 - hosting on-site experiments and groups as tier- $\{0,1,2\}$ equivalents (photon, HEP, theory, astro, accelerator communities)
 - WLCG VO requirements need to be consolidated with other communities;
 - grid infrastructure part of an overreaching infrastructure covering also local/national user analyses, simulation, online data analysis,...;
 - GPU support necessary consideration for an OS flavour as well as support for other architectures (ARM)
- There are a lot of ongoing migration (GSI to token, gridftp to webdav, etc.) so a lot of reinstallations are going to happen. It is a good chance to upgrade the OS too, but it would be better to know what OS to install...
- A little more direction early on from CERN might have helped with decision.
- We would also like to ask what tools are used at sites to manage OS vulnerabilities? Pakiti future is unclear.
- In near future we plan to move to kubernetes for CE, storage, other supportive things etc. I would assume picking OS will be less and less relevant to run something.

Conclusions

- So far very positive experience with AL and RL. No bad surprises.
- Sites which already tried Rocky would like to stick with it. Sometimes there is a confusion whether they need to switch to AM because of CERN and FNAL recommendation.
- We need to understand the situation with the tools to manage OS vulnerabilities. Pakiti future is unclear.
- There are sites which are still waiting for guidance from WLCG, so a clear message with recommendations is to be broadcasted.

Backup slide: Some comments regarding RHEL

- Stable and commercial support and major release lifecycle of 10 years. Hardware driver support. Familiarity in ops team because Centos/RHEL are used in the entire infrastructure
- RHEL9 pros: Long support life, kernel 5.x improvements. concerns: Lack of all needed middleware/tools on this OS
- UMD repository missing support for RH8 derivatives, I had to recompile the argus client to install HTCondor-CE on a Centos 8.
- Last time we tried, ARC was unable to run on any RHEL 8 flavor (esp. the APEL parsers)
- RHEL8/9 concern: Availability of cgroups (v2)

Backup slide: Some feedback from the sites regarding RL

- Rocky was the first being available, evaluation of alma based on further consideration on the commercialisation of rocky which is too close to what happen to centos.
- We would like to stick with RL9 because migration to AL9 would result in reinstall of already done services (its just more work). But we can migrate to AL9, if necessary.
- Middleware support, updates appear few days later than in Alma Linux
- Rocky is officially supported by Mellanox/Nvidia OFED drivers
- We already have some experience with Rocky 8 and it works good. Rocky 9 appeared too early, software developers focused on Rocky 8, so many packages are not available yet.
- Great and fast support by the RockyLinux developer community, extensive support in public tooling (Puppet modules etc.) even compared to e.g. AlmaLinux
- Rocky 9 seems to have a weird interaction between podman-compose and systemd, where health checks cause a kernel memory leak through zombie cgroup. Also, Rocky 9 seems to make the use of NetworkManager mandatory. This is probably the same for RHEL9. A change going to Rocky 8 is the introduction of netfilter over iptables, which requires learning an additional skill.
- Rocky * is supposed to be a faithful rebuilt of RHEL with binary compatibility. It has good hardware vendor support.
- Entirely positive, very fast feedback by developers in case e.g. mirror issues arise.
- Broad adoption also outside of CERN / HEP community (two times as many hits from Rocky seen on EPEL repos compared to Alma, Rocky even exceeds CentOS stream). This broad adoption also means broad support in many community Puppet modules. Quick developer responses via Mattermost community, known issues are quickly tracked and solved in due time (e.g. currently missing Errata for Rocky 9) once developers are made aware. Open build system which will also allow for community package building similar to OBS. Longer commitment (in terms of end of security updates dates) than Alma. Upstream packages for our storage implementation (XRootD) are built on Rocky.