

Fedora ELN at Meta: a testbed for fleet upgrades

DevConf.CZ 2023

Davide Cavalca
Production Engineer, Linux Userspace



Agenda

01 Fedora ELN

02 ELN at Meta

03 Results

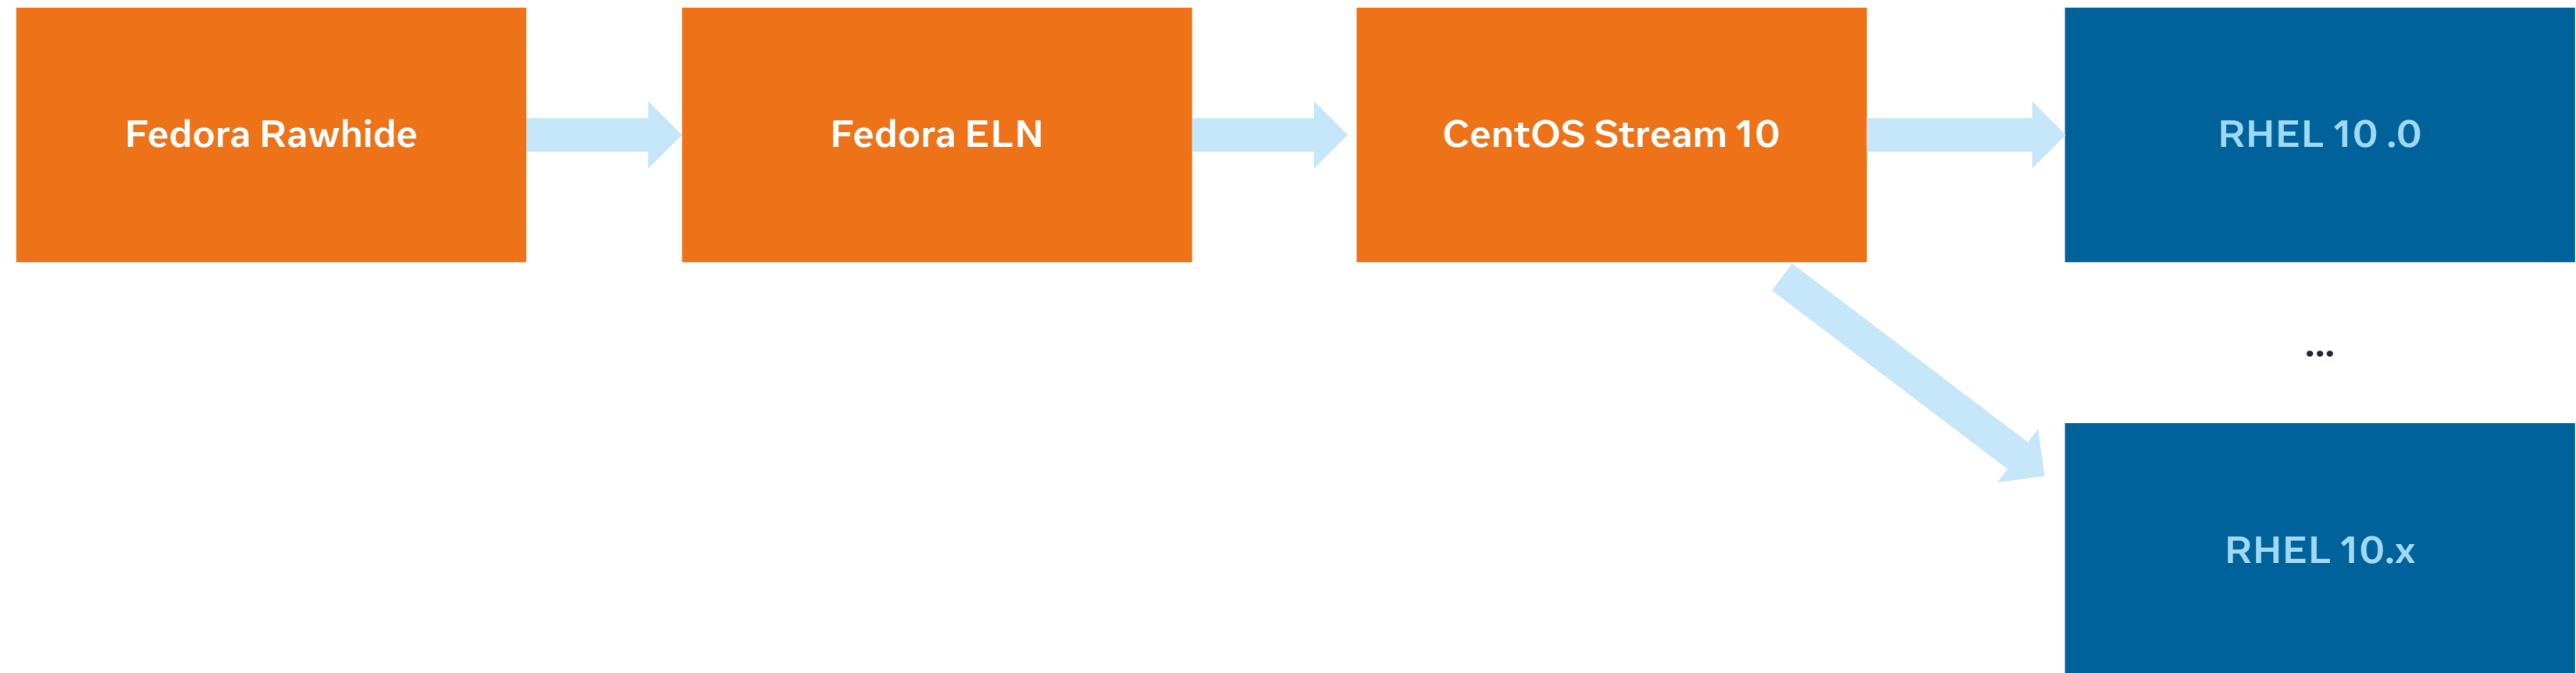
04 Get involved

Fedora ELN

What's ELN

- “Enterprise Linux Next”
- Continuous rebuild of Rawhide with the CentOS macros and toolchain
- Assists in the bringup of the next CentOS Stream major release
- <https://docs.fedoraproject.org/en-US/eln/>

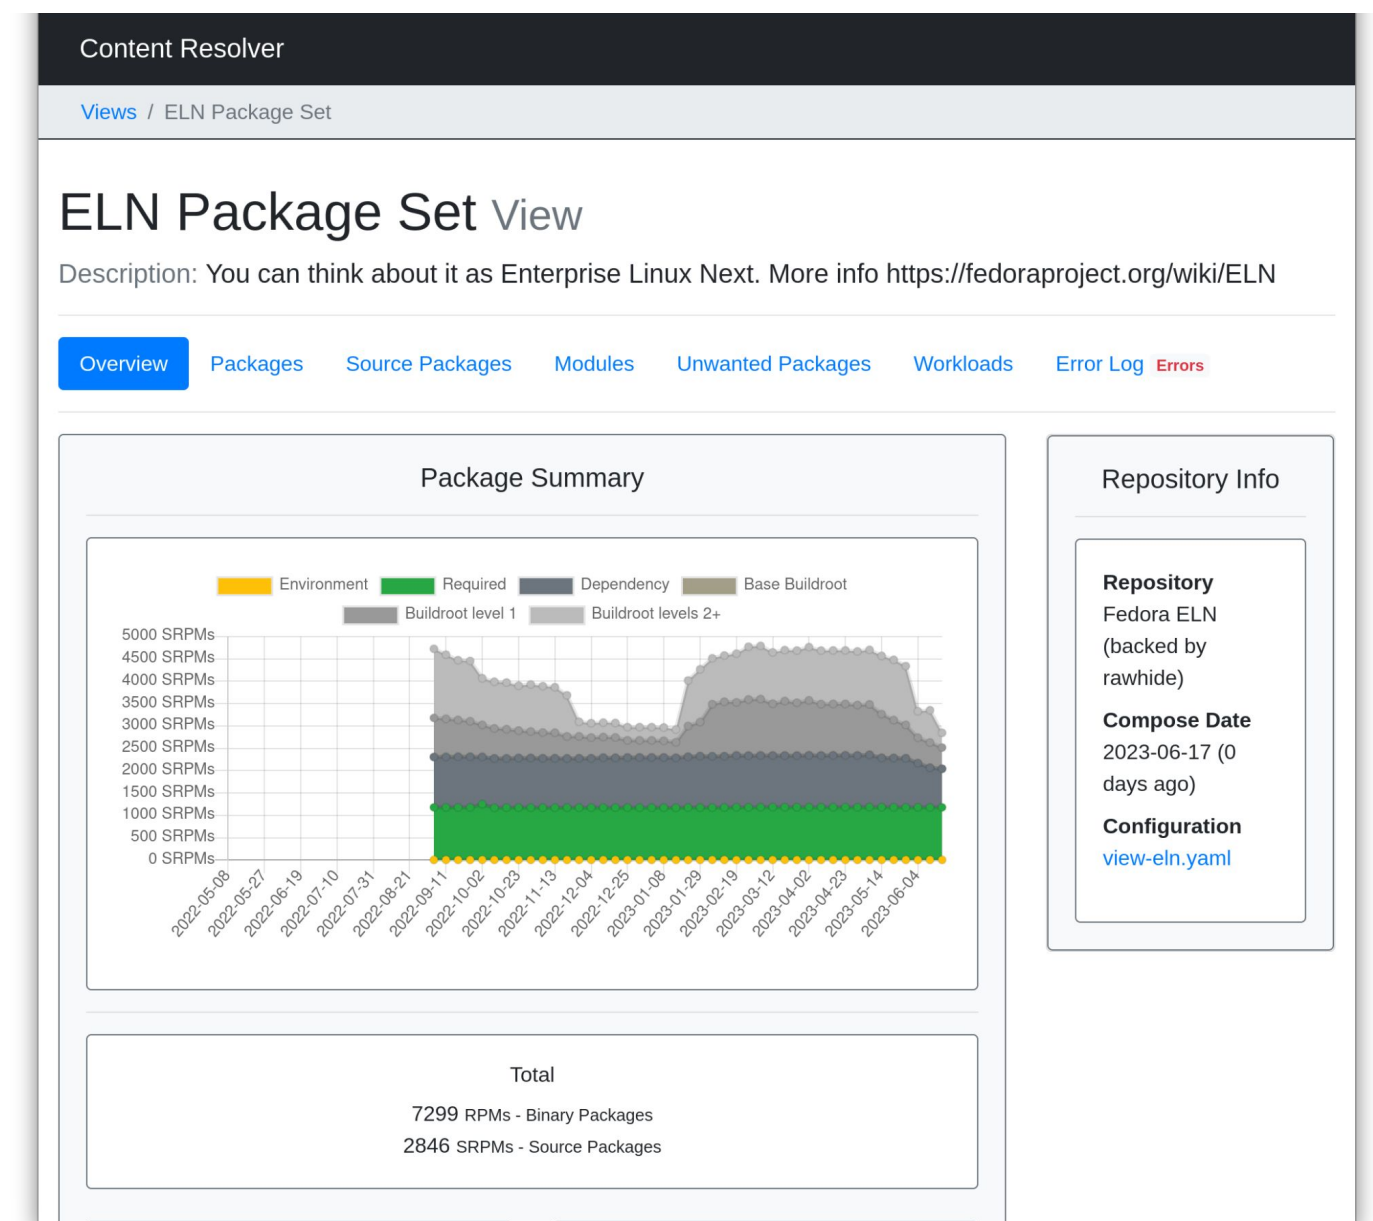




Fedora ELN

How it's made

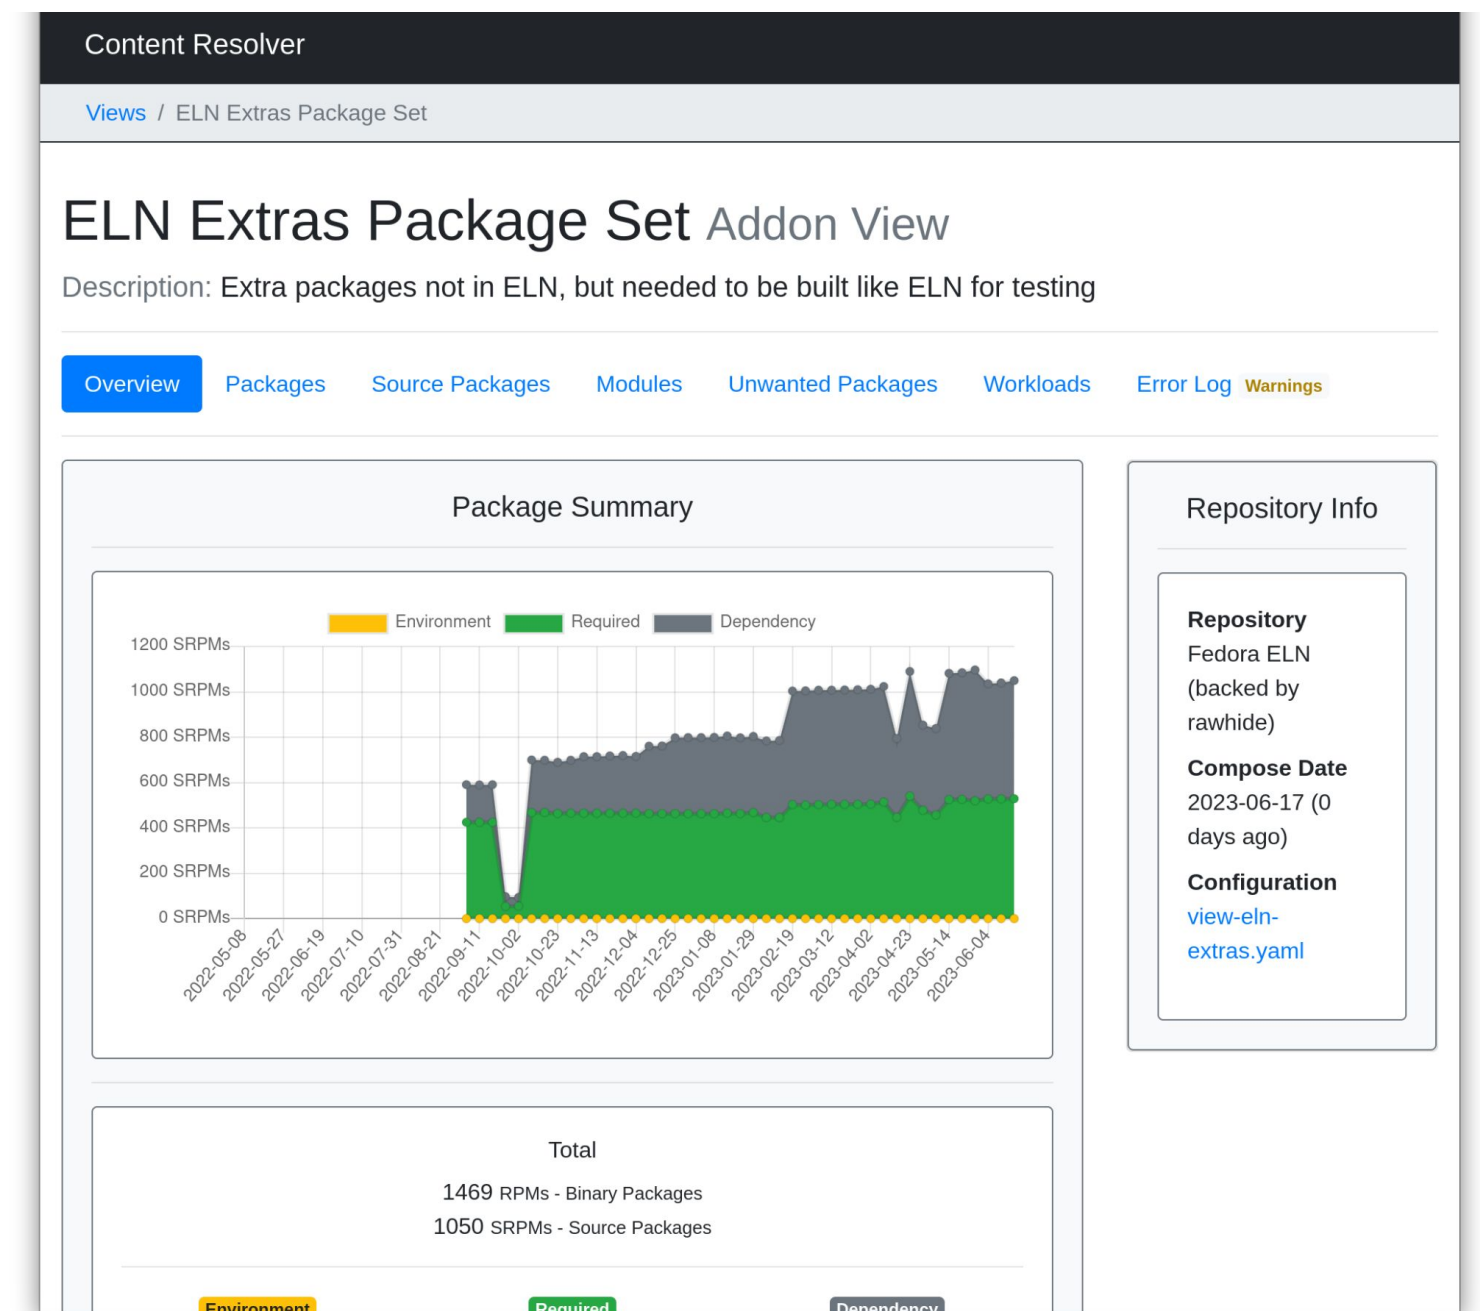
- Frequent composes via ODCS
- Package set defined in Content Resolver
- <https://tiny.distro.builders/view--view-elN.html>



Fedora ELN

ELN Extras

- Additional packages that are built for ELN and composed together
- Expected to be used to bootstrap EPEL 10
- <https://tiny.distro.builders/view--view-elN-extras.html>
- <https://github.com/minimization/content-resolver-input>




Fedora ELN

ELN SIG

- Fedora Special Interest Group responsible for ELN
- #eln:fedoraproject.org on Matrix, #fedora-eln on IRC
- Regular meetings on Fridays at noon, EST
- <https://docs.fedoraproject.org/en-US/eln/sig/>

ELN at Meta



CentOS production deployment

- CentOS Linux 5 -> 6 (~2013-2016)
- CentOS Linux 6 -> 7 (2016-2018)
- CentOS Linux 7 -> CentOS Stream 8 (2018-2022)
- CentOS Stream 8 -> 9 (2022-)
- Current status
 - 86% of the fleet on CentOS Stream 9, 14% on CentOS Stream 8
 - No more CentOS Linux 7 hosts!
 - Majority of containers still on CentOS Stream 8

Problems

- Major OS upgrades are disruptive
- Qualification is bursty, time consuming and expensive
- Last minute discoveries can turn into major issues
 - e.g. the SHA-1 deprecation in CentOS Stream 9
- What if we had a way to start qualification earlier?

Goals

- Streamline bringup of new major OS releases
 - Find and fix bugs long before they even make it into CentOS Stream
 - Identify policy and package changes early on
 - Turn qualification from a point-in-time activity to a continuous one
- Build a continuous testing and integration platform
 - Allow customers to perform testing and validation on their own pace

Bootstrapping

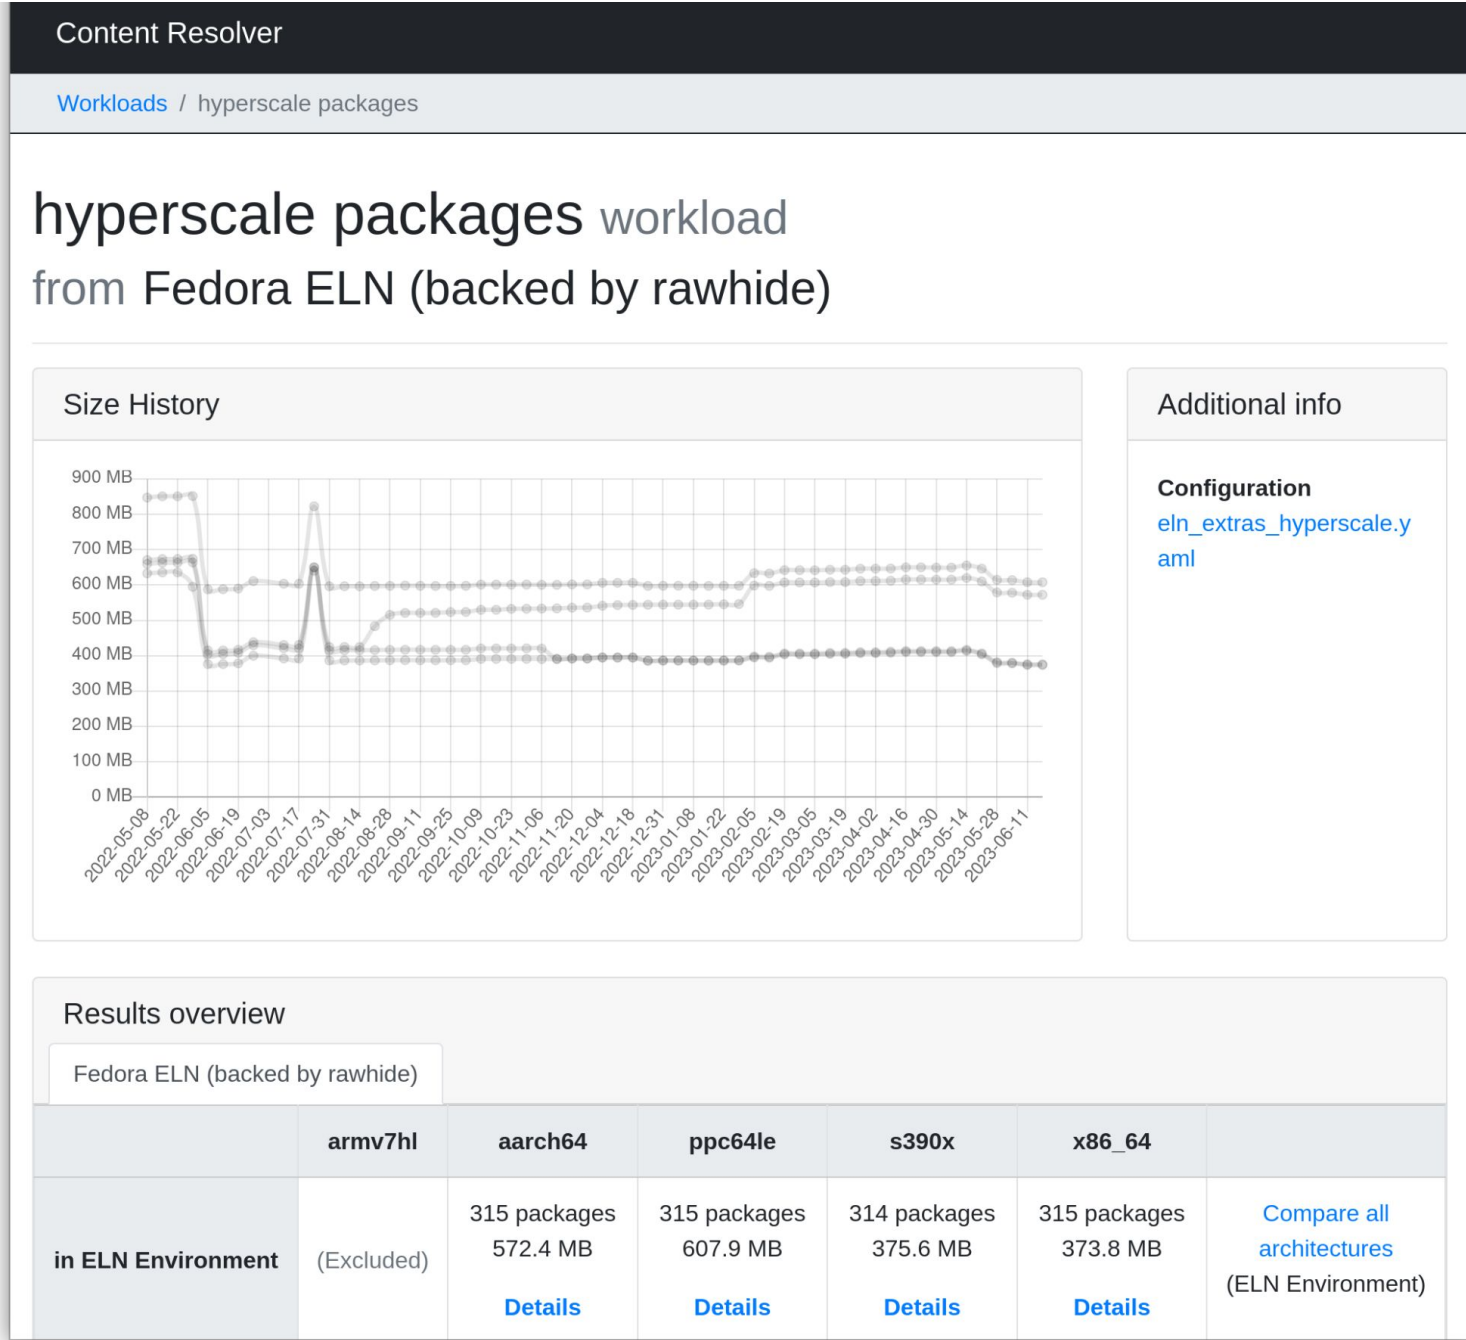
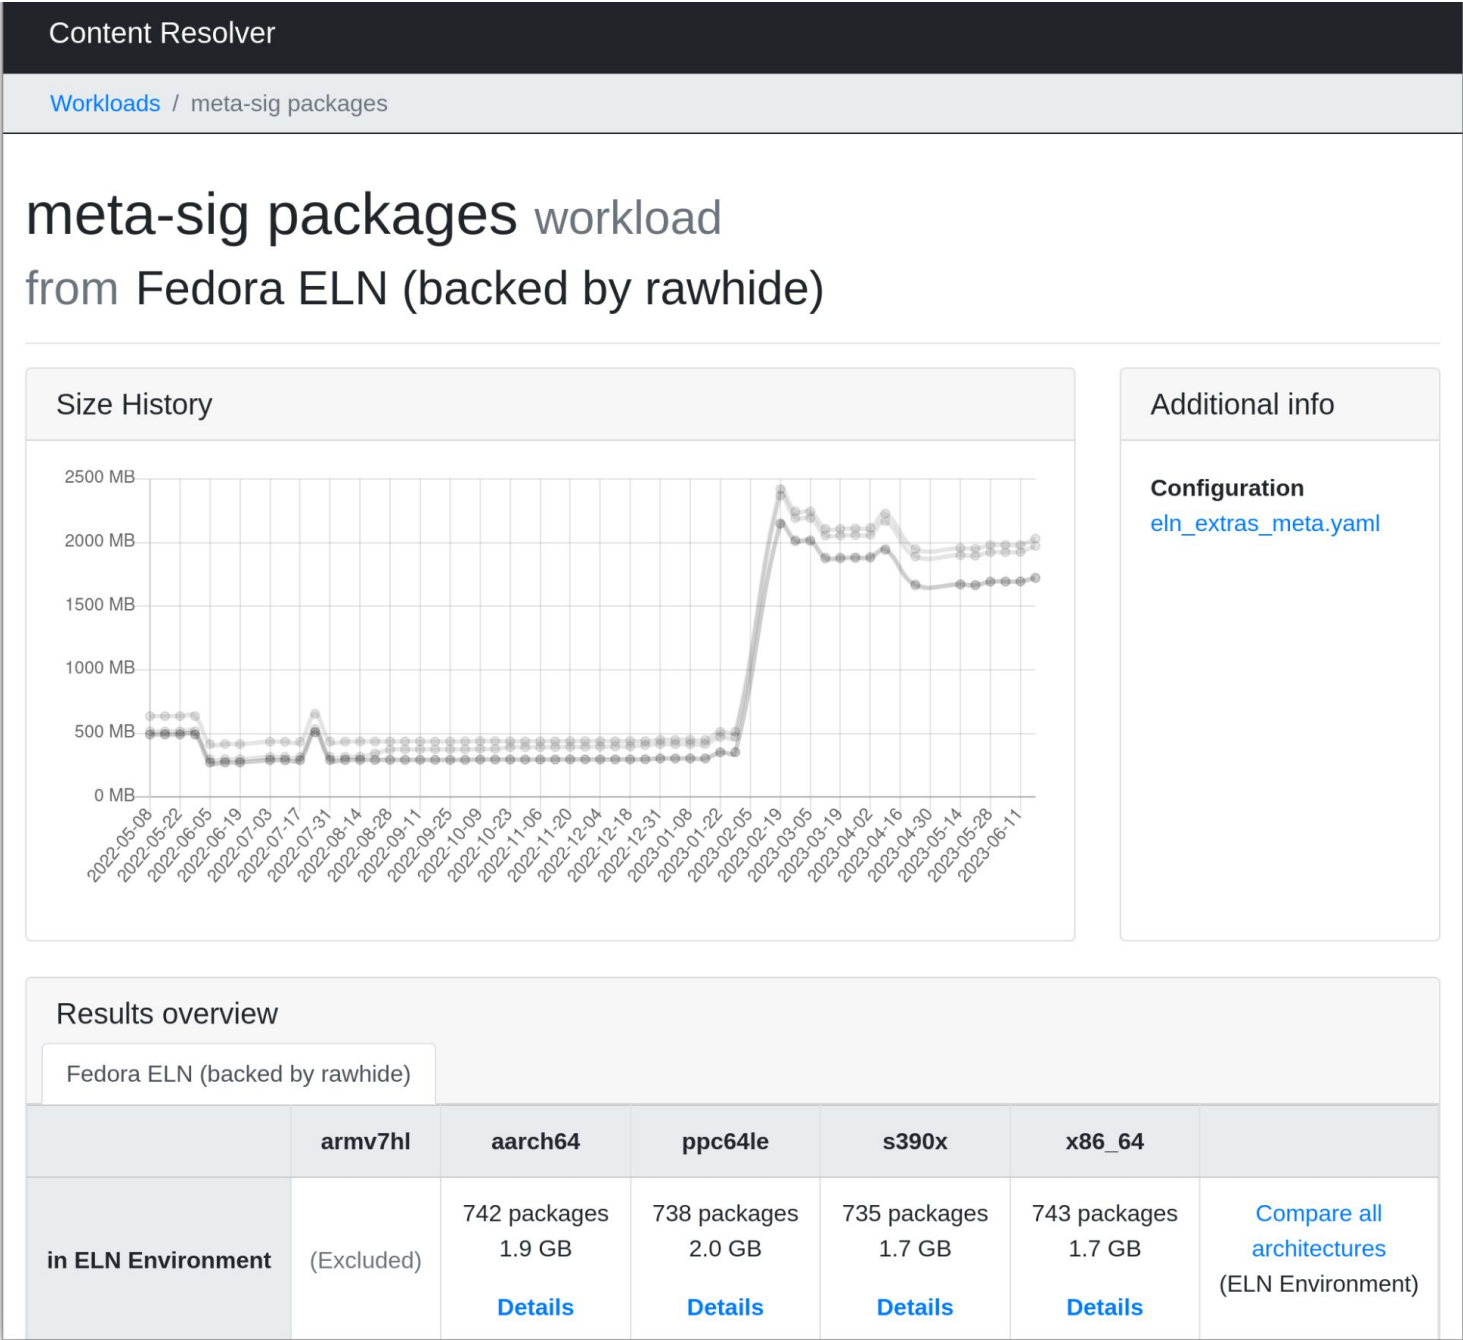
- ELN bringup is akin to a new major OS release bringup
 - Mirror repos
 - Establish the rolling OS updates pipeline
 - Add gating and support logic in Chef
 - Build provisioning and container images
 - Qualify and deploy
- ...and fix all the bugs at every step!

Repos

- ELN composes are built on ODCS
 - <https://odcs.fedoraproject.org/composes/production/latest-Fedora-ELN/>
- Worked with infra to expose them over rsync
 - <https://pagure.io/fedora-infrastructure/issue/9730>
- Published our mirror
 - <http://mirror.facebook.net/fedora-elN/production/latest-Fedora-ELN/>
- Snapshotted periodically for internal use via Rolling OS Updates

What about EPEL?

- Use ELN Extras to provide packages we need currently carried in EPEL
- As a byproduct, this ensures they'll be ready for EPEL 10!
- Two content-resolver workloads
- <https://tiny.distro.builders>



Rolling OS updates

- Periodic repo snapshots for production
- Used to track CentOS Stream updates within a release train
- Rolled out over two weeks every two weeks
- Implemented in Chef
 - DNF configuration on hosts
 - In-place upgrades of existing hosts via `dnf upgrade`

Rolling OS updates

- Setup a new “eln” release train
- Faster snapshot cadence
 - Every day during bootstrapping, then every week
- Streamlined workflow
 - Automated promotion, no need to test for in-place upgrades
- Caveats
 - Composes can be incomplete or broken
 - ELN Extras workflows can fail due to a single broken package

Chef

- ELN identifies itself as Fedora Rawhide
- Added detection logic to our OSS cookbooks leveraging VARIANT_ID
- Monkeypatched it internally to pretend it's CentOS instead
 - So we can minimize ELN-specific logic...
 - ...and actually test the CentOS logic, as that's what'll run in production
- Iterated over missing/broken packages, gating, logic bugs, etc.
- <https://github.com/facebook/chef-cookbooks>

Provisioning

- In-place conversion as a stopgap
 - From existing CentOS Stream 9 hosts
 - Chef recipe to `dnf distro-sync --allowerase`
 - It works!
- Integration in our internal provisioning system
 - Provisioning images
 - Continuous delivery infrastructure
 - Feature parity with CentOS Stream 9

Qualification

- Kernel/hardware testing systems
 - Minimal install and Chef runlist
- Interactive development servers
 - Large variety of packages and usecases
- Container platform hosts
 - Production workloads

Results

Results

Current state

- About four months from zero to fully working provisioning
- Few dozens of test systems running ELN
- Started conversations for a wider ELN deployment
- Began preparations for future CentOS Stream 10 deployment
 - Integrate and communicate expected changes
 - Address discovered issues

Results

RPM signature validation issues

```
error: [...]: Header RSA signature: BAD (package tag 268: invalid OpenPGP signature)
```

```
error: [...]: not an rpm package (or package manifest)
```

- RPM 4.18 has a new OpenPGP implementation based on Sequoia
 - Stricter parser, better validation
 - <https://fedoraproject.org/wiki/Changes/RpmSequoia>
- <https://github.com/rpm-software-management/rpm/issues/2351>
- Resolved by fixing our internal signing service
- Long tail of internal packages to rebuild

Results

RPM signature validation issues

```
Importing GPG key [...]:
```

```
error: Certificate [...]:
```

```
Policy rejects [...]: No binding signature at time 2023-03-21T21:57:39Z
```

```
Key import failed (code 2).
```

- Crypto policy disallows SHA-1 as of CentOS Stream 9
- RPM didn't actually validate the signing key... until Sequoia
- `sq-keyring-linter` can audit and fix existing keys
- Tell your vendors!

Results

util-linux 2.39 regressions

- `nsenter` option parsing changes
 - Broke our image build system
 - <https://github.com/util-linux/util-linux/issues/2143>
- New mount API
 - <http://karelzak.blogspot.com/2023/06/util-linux-v239-improved-mount.html>
 - Remount fails on 5.12 and older kernels
 - <https://github.com/util-linux/util-linux/issues/2283>
 - Independently found in Rawhide by OpenQA

Results

ELN quirks

- auditd ruleset
 - ELN was installing the Fedora default ruleset
 - Broke our Chef logic which assumed the RHEL one
 - <https://src.fedoraproject.org/rpms/audit/pull-request/8>
- systemd presets
 - ELN uses the Fedora presets
 - systemd-networkd was never enabled

Get involved

Get involved

ELN and CentOS Stream 10

- CentOS Stream 10 will begin branching soon
- Fedora Devel discussion: <https://tinyurl.com/bdew9vcs>
- Try out ELN now to get a head start!

Get involved

Participate in the community

- Join the ELN SIG
- Add your EPEL packages to ELN Extras
- <https://docs.fedoraproject.org/en-US/eln/>
- <https://github.com/fedora-eln/>
- <https://tiny.distro.builders/>
- <https://github.com/minimization/content-resolver-input>
- #eln:fedoraproject.org on Matrix, #fedora-eln on IRC

Questions?

THANK YOU FOR YOUR TIME

