

Building community with CentOS Stream

LISA21

Davide Cavalca
Production Engineer

FACEBOOK Infrastructure

Agenda

CentOS at Facebook


Contributing upstream

Hyperscale SIG

Get involved

CentOS at
Facebook

CentOS at Facebook



CentOS at Facebook

Why CentOS?

- Stable releases
- Binary compatibility
- Security updates
- Mature and well understood tooling
- EPEL
- Close relationship with Fedora

CentOS at Facebook

FTL - Fast Thin Layer

- Backports from Fedora Rawhide for stuff we care about
- Mostly plumbing and low-level packages
- GitHub: [facebookincubator/rpm-backports](https://github.com/facebookincubator/rpm-backports)
- %facebook macro to gate internal stuff
- CentOS + FTL = stable distro, moving fast

CentOS at Facebook

Policy deviations

- Upstream kernel
 - cgroup2 by default
 - btrfs on / by default
- iptables: legacy backend instead of nftables
- networking: network-scripts instead of NetworkManager

CentOS at Facebook

Major OS upgrades

- CentOS Linux 5 -> 6 (~2013-2016)
- CentOS Linux 6 -> 7 (2016-2018)
- CentOS Linux 7 -> CentOS Stream 8 (2018-2021)
- Reprovisioning for OS upgrades
 - Clean slate
 - Deprecated unwanted features
 - Policy changes coupling
- Leverage the general host maintenance window
- Tooling and automation for rollouts

CentOS at Facebook

Minor OS upgrades

- Incremental Rolling OS upgrades
- Every two weeks we sync down the latest updates...
- ...and roll them out over two weeks
- ‘dnf upgrade’ kicked off via Chef
- High level monitoring of rollout health
- Easy stop button and opt out for individual packages

CentOS at Facebook

Can we do better?

- FTL
 - Internal backports are forks
 - No clear path to upstreaming
 - Bug fixes get lost
 - Distro updates have to be manually integrated
- Policy deviations
 - Have to be maintained long-term
 - Can impact bug reports and repros
 - No real feedback loop

Contributing
upstream

Contributing upstream

Upstream first

- Community sets the direction
- We move fast; Open Source often moves faster
- We don't need to write everything ourselves
- Sharing our code means sharing the maintenance and having others extend it

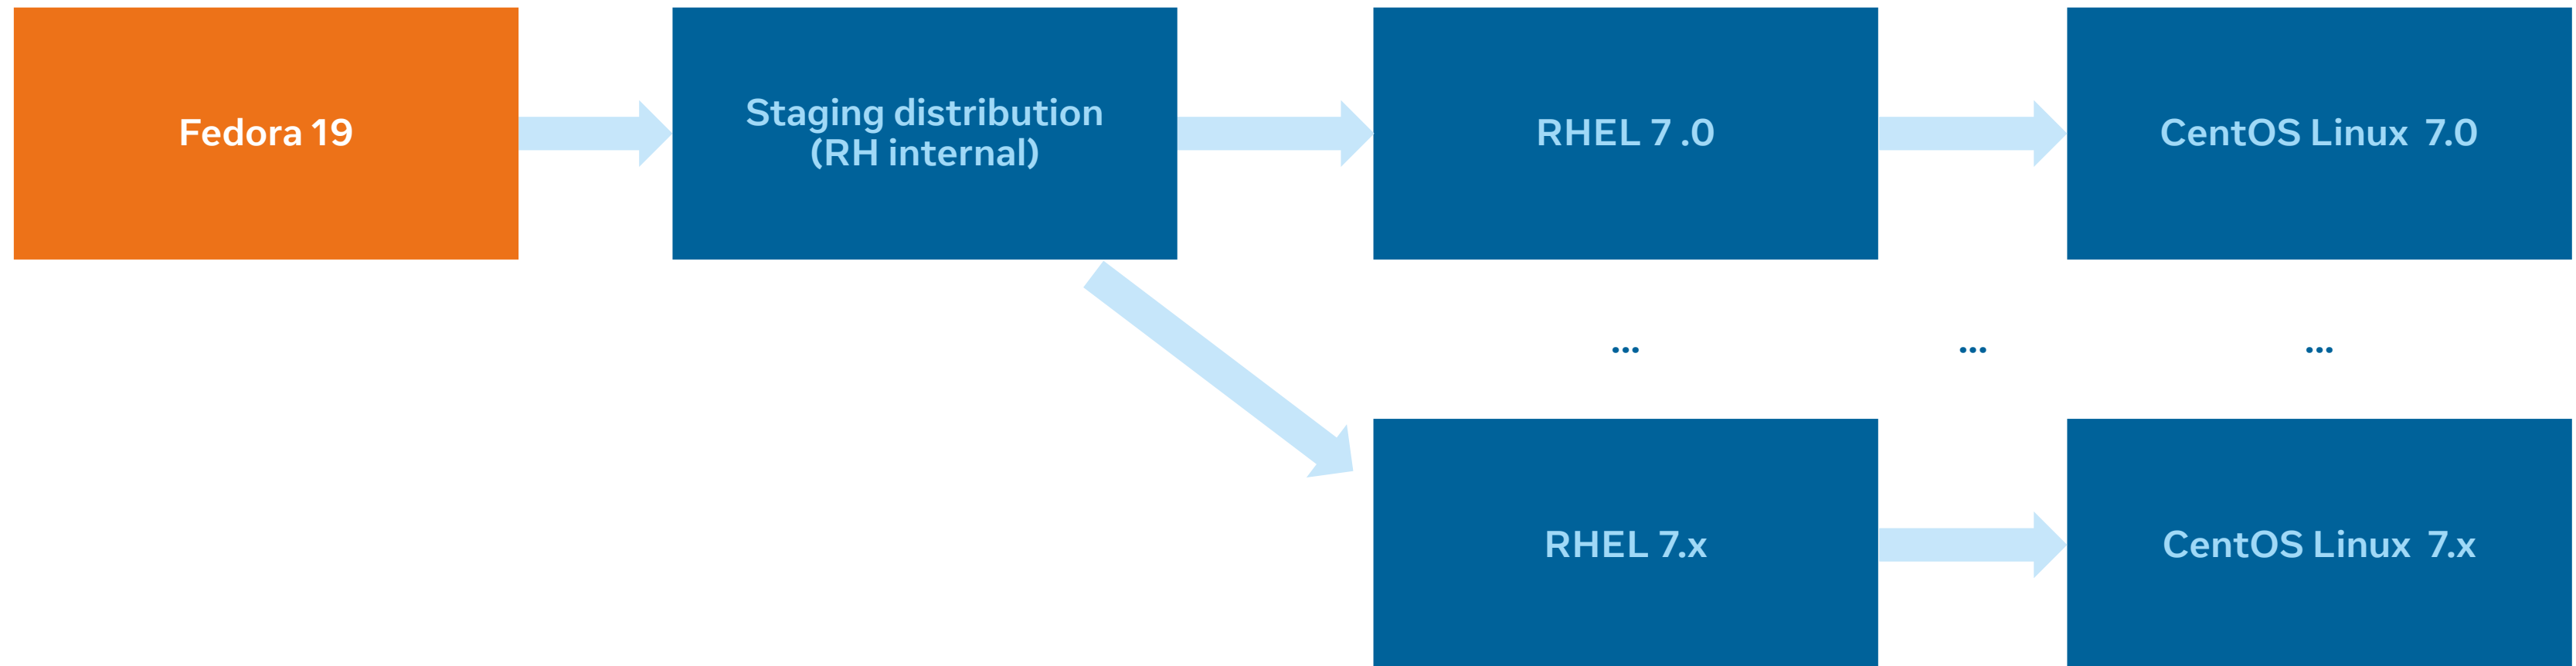
Contributing upstream

How

- Show up
- Engage with the community as a peer
- Solve real problems
- Build trust

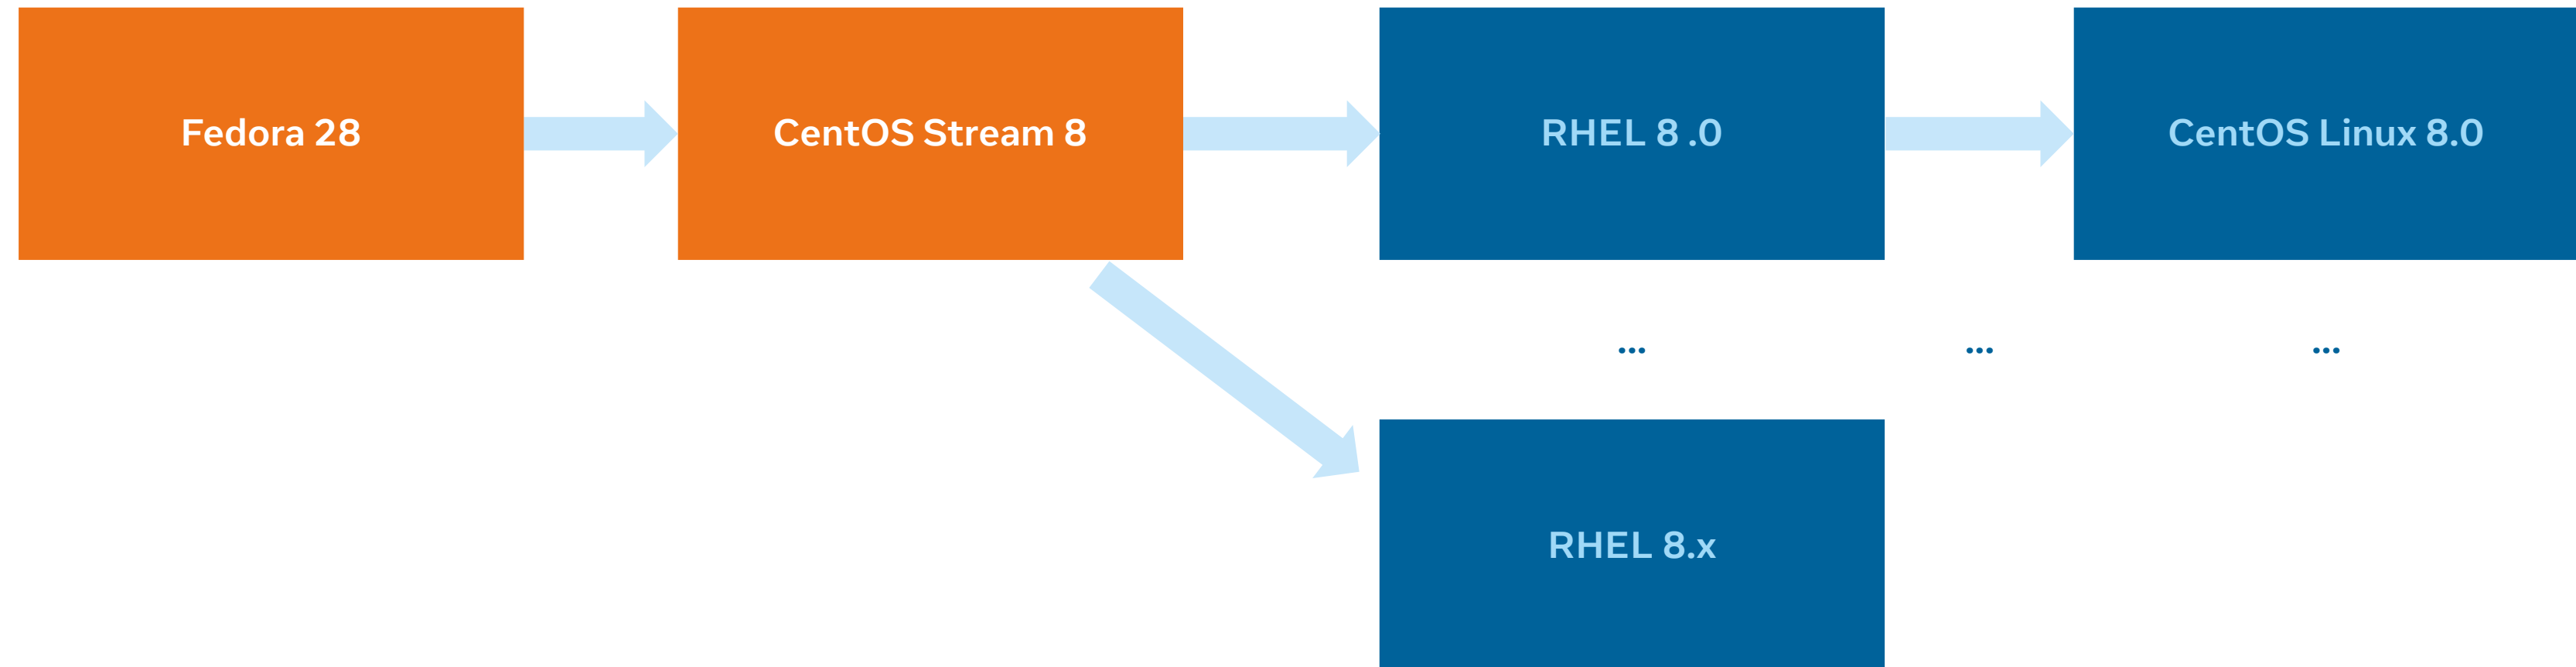
Contributing upstream

CentOS Linux 7



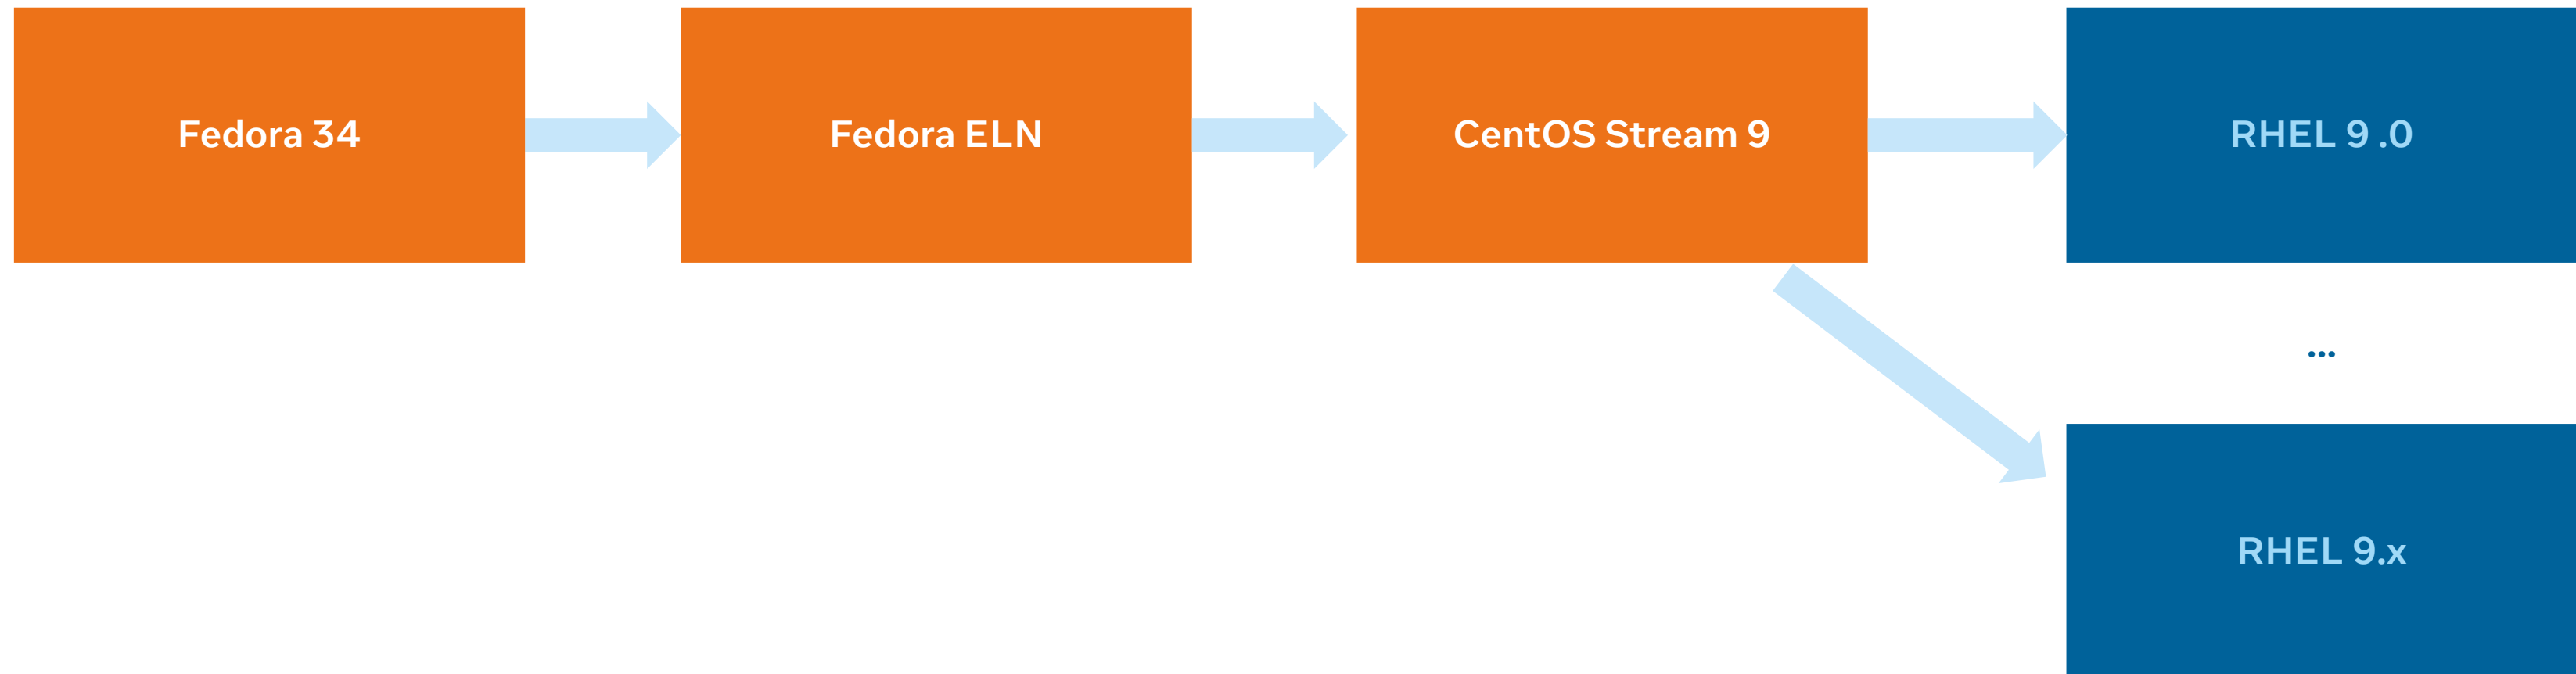
Contributing upstream

CentOS Linux 8 and CentOS Stream 8



Contributing upstream

CentOS Stream 9



CentOS Stream blog post: <https://tinyurl.com/ycn29k2c>

Contributing upstream

Fedora

- Influences the next CentOS Stream major release
- File and fix bugs, maintain packages, drive Changes, etc.
 - <https://src.fedoraproject.org>
 - <https://fedoraproject.org/wiki/Changes>
- Change proposals
 - F33: Btrfs by default
 - F34: Btrfs with zstd compression by default
 - F34: systemd-oomd by default
 - ~~F34~~ F35: DNF RPM Copy-on-Write
 - F35: Btrfs by default for Fedora Cloud
 - F35: fsverity RPM support

Contributing upstream

Fedora EPEL

- Additional packages for RHEL and CentOS based on Fedora
- <https://fedoraproject.org/wiki/EPEL>
- EPEL Packagers SIG
 - Streamline the process to add packages to EPEL
 - Tooling improvements
 - Collective maintenance
 - <https://fedoraproject.org/wiki/EPEL/Packagers>

Contributing upstream

Fedora ELN

- Continuous rebuild of Rawhide with the CentOS macros and toolchain
- Assists in the bringup of the next CentOS Stream major release
- ELN SIG
 - Enablement work to make ELN easier to consume
 - Extending to cover more packages via eln-extra
 - <https://github.com/fedora-eln>
- Continuous testing and integration pipeline
 - Within Facebook: provisioning, Chef, containers
 - Find and fix bugs long before they even make it into CentOS Stream
 - Identify policy and package changes early on

Contributing upstream

CentOS Stream 8

- Continuously delivered distribution tracking the next minor release of RHEL
- File and fix bugs: <https://bugzilla.redhat.com>
 - Product: Red Hat Linux Enterprise 8
 - Version: CentOS Stream
- Follow development and send pull requests
 - <https://git.centos.org>
- Drive change via Special Interest Groups (SIGs)
 - Building blocks of the CentOS community
 - <https://wiki.centos.org/SpecialInterestGroup>

Hyperscale SIG

Hyperscale SIG

What we do

- CentOS Stream focus
- Large scale infrastructure
- Foster cross-company collaboration on packaging and tooling
- Bring in-house development out in the open
- Open to anybody interested in working in this space
- <https://wiki.centos.org/SpecialInterestGroup/Hyperscale>
- #centos-hyperscale on Libera.Chat

Hyperscale SIG

Faster-moving package backports

- Updated backports of distro packages
- Feature enablement, closely tracking upstream development
- Drop in replacements for distro packages
- Stable and targeting production use
- Delivered as a dedicated repository
 - `dnf install centos-release-hyperscale`
- Available packages
 - <https://cbs.centos.org/koji/packages?tagID=2249>
 - dracut, dwarves, grep, less, libvirt, meson, mtr, ninja-build, pykickstart, rasdaemon, systemd, tpm2-tss, tpm2-tools, util-linux, ...

Hyperscale SIG

systemd

- Actively maintained systemd backport
- Running in production at FB
- Tracking latest upstream stable release
 - Staging repo: <https://pagure.io/centos-sig-hyperscale/systemd>
- Based on the Fedora packaging
 - <https://git.centos.org/rpms/systemd/tree/c8s-sig-hyperscale>
- CI/CD pipeline to build and test daily snapshots
 - Keeps the staging repo in sync
 - Builds and tests dailies for the latest git master
 - <https://pagure.io/centos-sig-hyperscale/systemd-releng>

Hyperscale SIG

LLVM 12

- Non-modular LLVM 12 build
- llvm, clang, lld, lldb (+ mesa)
- Modular LLVM 12 is slated for CentOS Stream proper
- https://bugzilla.redhat.com/show_bug.cgi?id=1952248
- <https://reviews.llvm.org/D101972>

Hyperscale SIG

Policy and configuration alternatives

- Modifications of distro packages to enable alternative options
- Meant to be backward compatible and minimize changes
- Example: iptables
 - Only supports nftables in CentOS Linux 8
 - Rebuild to enable the legacy iptables backend as an alternative

Hyperscale SIG

Large-scale testing

- Provide a way to test distro-wide changes in production settings
- Example: DNF/RPM Copy-on-Write
 - <https://fedoraproject.org/wiki/Changes/RPMCoW>
 - Requires patched packaging stack
- Currently deployed in production at FB
- Delivered as a dedicated repository
 - `dnf install centos-release-hyperscale-experimental`

Hyperscale SIG

Work in progress

- Kernel
 - LTS-tracking upstream kernel
 - Feature enablement: btrfs, cgroup2, BPF
- Btrfs
 - Installer support
 - Userspace enablement
 - Transactional updates
- Container images: <https://quay.io/centos hyperscale/centos>
- Cloud images
- <https://pagure.io/centos-sig-hyperscale/sig/issues>

Get involved

Get involved

Participate in the CentOS community

- Read and contribute to the blog: <https://blog.centos.org>
- Join the mailing list: centos-devel@centos.org
- Attend a meeting: <https://www.centos.org/community/calendar>
- Join a SIG: <https://wiki.centos.org/SpecialInterestGroup>
- Report or fix a bug: <https://bugzilla.redhat.com>
- Maintain a package in EPEL: <https://fedoraproject.org/wiki/EPEL/Packagers>
- Contribute to Fedora: <https://fedoramagazine.org/how-to-contribute-to-fedora>

Get involved

CentOS Stream 9

- Being developed right now, in the open
- File and fix bugs: <https://bugzilla.redhat.com>
 - Product: Red Hat Linux Enterprise 9
 - Version: CentOS Stream
- Follow development and send pull requests:
 - <https://gitlab.com/redhat/centos-stream>
 - <https://kojihub.stream.centos.org>
- Download and test daily composes:
 - <https://composes.stream.centos.org/test>

The screenshot displays the GitLab web interface for the 'centos-stream' project. The left sidebar shows navigation options: Subgroup overview, Epics (0), Issues (0), Merge requests (54), Packages & Registries, Analytics, Wiki, and Members. The main content area is titled 'Merge requests' and shows a list of 54 open requests. Each request entry includes a title, the branch path, the creator, the creation time, and the last update time. The requests are sorted by 'Created date' in descending order. The first request is 'Add several patches dealing with the false positive zipbomb detection' by jamartis, created 13 hours ago. Other requests include 'Initial test commit.', 'Use sources with ./configure in to match Fedora', and several 'Add gating definition for RHEL 9' requests. The interface also shows a search bar at the top and a pagination bar at the bottom.

Open	Merged	Closed	All
54	441	80	575

Recent searches	Search or filter results...	Created date	Sort
Add several patches dealing with the false positive zipbomb detection redhat/centos-stream/rpms/unzip!1 · created 13 hours ago by jamartis · updated 13 hours ago			
Initial test commit. redhat/centos-stream/rpms/kmod-kvdo!1 · created 14 hours ago by Andrew Walsh · updated 8 hours ago			
Use sources with ./configure in to match Fedora redhat/centos-stream/rpms/kronosnet!4 · created 15 hours ago by ccaulfie · updated 15 hours ago			
Add gating definition for RHEL 9 redhat/centos-stream/rpms/ansible-freeipa!1 · created 1 day ago by Francisco Triviño García · updated 1 day ago			
Add gating definition for RHEL 9 redhat/centos-stream/rpms/oddjob!1 · created 1 day ago by Francisco Triviño García · updated 1 day ago			
Add gating definition for RHEL 9 redhat/centos-stream/rpms/lpa-healthcheck!5 · created 1 day ago by Francisco Triviño García · updated 1 day ago			
Add gating definition for RHEL 9 redhat/centos-stream/rpms/lpa!3 · created 1 day ago by Francisco Triviño García · updated 1 day ago			
Disable non-core packages redhat/centos-stream/rpms/pki-core!2 · created 1 day ago by Endi S. Dewata · updated 23 hours ago			
Add gating definition for RHEL 9 redhat/centos-stream/rpms/certmonger!2 · created 1 day ago by Francisco Triviño García · updated 1 day ago			
Add gating definition for RHEL 9 redhat/centos-stream/rpms/slapi-nls!1 · created 1 day ago by Francisco Triviño García · updated 1 day ago			
Build jit based on what postgresql server does redhat/centos-stream/rpms/pg_repack!1 · created 1 day ago by Filip Janus · updated 1 day ago			
Build without reflections optional dependency redhat/centos-stream/rpms/jna!2 · created 1 day ago by Ondrej Dubaj · updated 1 day ago			
Add pregenerated man pages redhat/centos-stream/rpms/cyrus-imapd!3 · created 1 day ago by Pavel Zhukov · updated 1 day ago			
don't ship systemd-timesyncd redhat/centos-stream/rpms/systemd!3 · created 2 days ago by David Tardon · updated 2 days ago			
disable 'make check' redhat/centos-stream/rpms/nvml!2 · created 2 days ago by Jmoyer1 · updated 2 days ago			
drop dependency on libfido2 redhat/centos-stream/rpms/systemd!2 · created 2 days ago by David Tardon · updated 1 day ago			
do not use BFQ as default scheduler redhat/centos-stream/rpms/systemd!1 · created 2 days ago by David Tardon · updated 1 day ago			
Switch to bundled gnu lib redhat/centos-stream/rpms/psutils!1 · created 2 days ago by Ondrej Pohorelsky · updated 2 days ago			
Changes crypto-policies from recommends to requires redhat/centos-stream/rpms/libssh!1 · created 3 days ago by Sahana Prasad · updated 3 days ago			
Bug 1935623 Fixing failed addon install because of disabled SHA-1 algorythm redhat/centos-stream/rpms/firefox!6 · created 3 days ago by Jhorak · updated 3 days ago			

< Prev 1 2 3 Next >

Thank you!

Resources

Related talks

- Hyperscale SIG update (CentOS Dojo May 2021): <https://tinyurl.com/47av47ye>
- CentOS Stream on the desktop (CentOS Dojo May 2021): <https://tinyurl.com/36h3zf73>
- Hyperscale SIG introduction (CentOS Dojo FOSDEM 2021): <https://tinyurl.com/2ez2m6hp>
- Speeding up DNF and RPM using Copy on Write (CentOS Dojo FOSDEM 2021): <https://tinyurl.com/6ub6ct49>
- CentOS Stream at Facebook (DevConf.cz 2021): <https://tinyurl.com/jbtw6ska>
- Upgrading CentOS on the Facebook fleet (DevConf.cz 2020): <https://tinyurl.com/3chkmfz5>
- systemd @ Facebook in 2019 (All Systems Go 2019): <https://tinyurl.com/vm2fxyfv>
- Building better FLOSS community relations (DevConf.cz 2017): <https://tinyurl.com/y7gx6nro>
- Really large scale system configuration (Atmosphere 2014): <https://tinyurl.com/bk4pjm3u>

FACEBOOK Infrastructure