

Building the future with CentOS Stream

SCaLE 19x

Davide Cavalca
Production Engineer



Agenda

01 CentOS at Meta


02 Contributing upstream

03 Hyperscale SIG

04 Get involved

CentOS at Meta

CentOS at Meta



Why CentOS?

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

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

Policy deviations

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

Major OS upgrades

- CentOS Linux 5 -> 6 (~2013-2016)
- CentOS Linux 6 -> 7 (2016-2018)
- CentOS Linux 7 -> CentOS Stream 8 (2018-2022)
- CentOS Stream 8 -> 9 (2022-2023)
- Current status
 - Fleet on CentOS Stream 8
 - Completed qualification for CentOS Stream 9
 - Kicked off the mass migration earlier this month

Major OS upgrades

- Reprovisioning for OS upgrades
 - Clean slate
 - Deprecated unwanted features
 - Policy changes coupling
- Leverage the general host maintenance window
- Tooling and automation for rollouts

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

Containers

- Also running CentOS
- Container images built from production repos
- Update cycle somewhat decoupled from hosts
- Update by rolling a new container

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

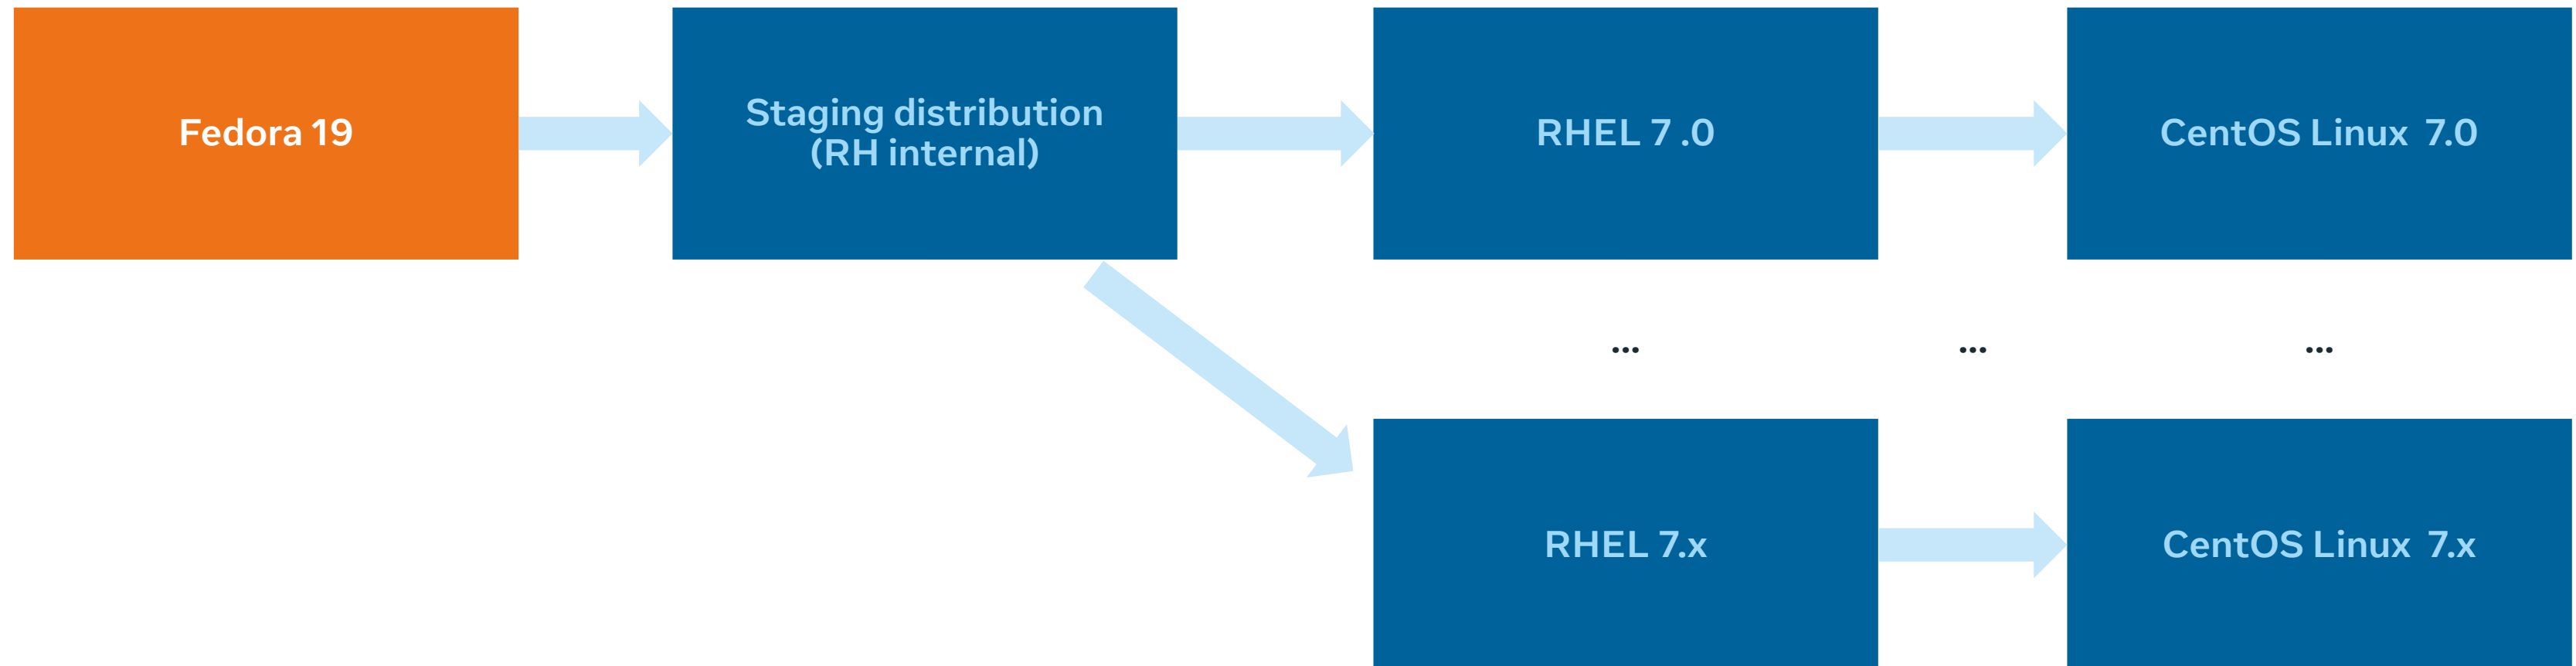
Contributing upstream

How

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

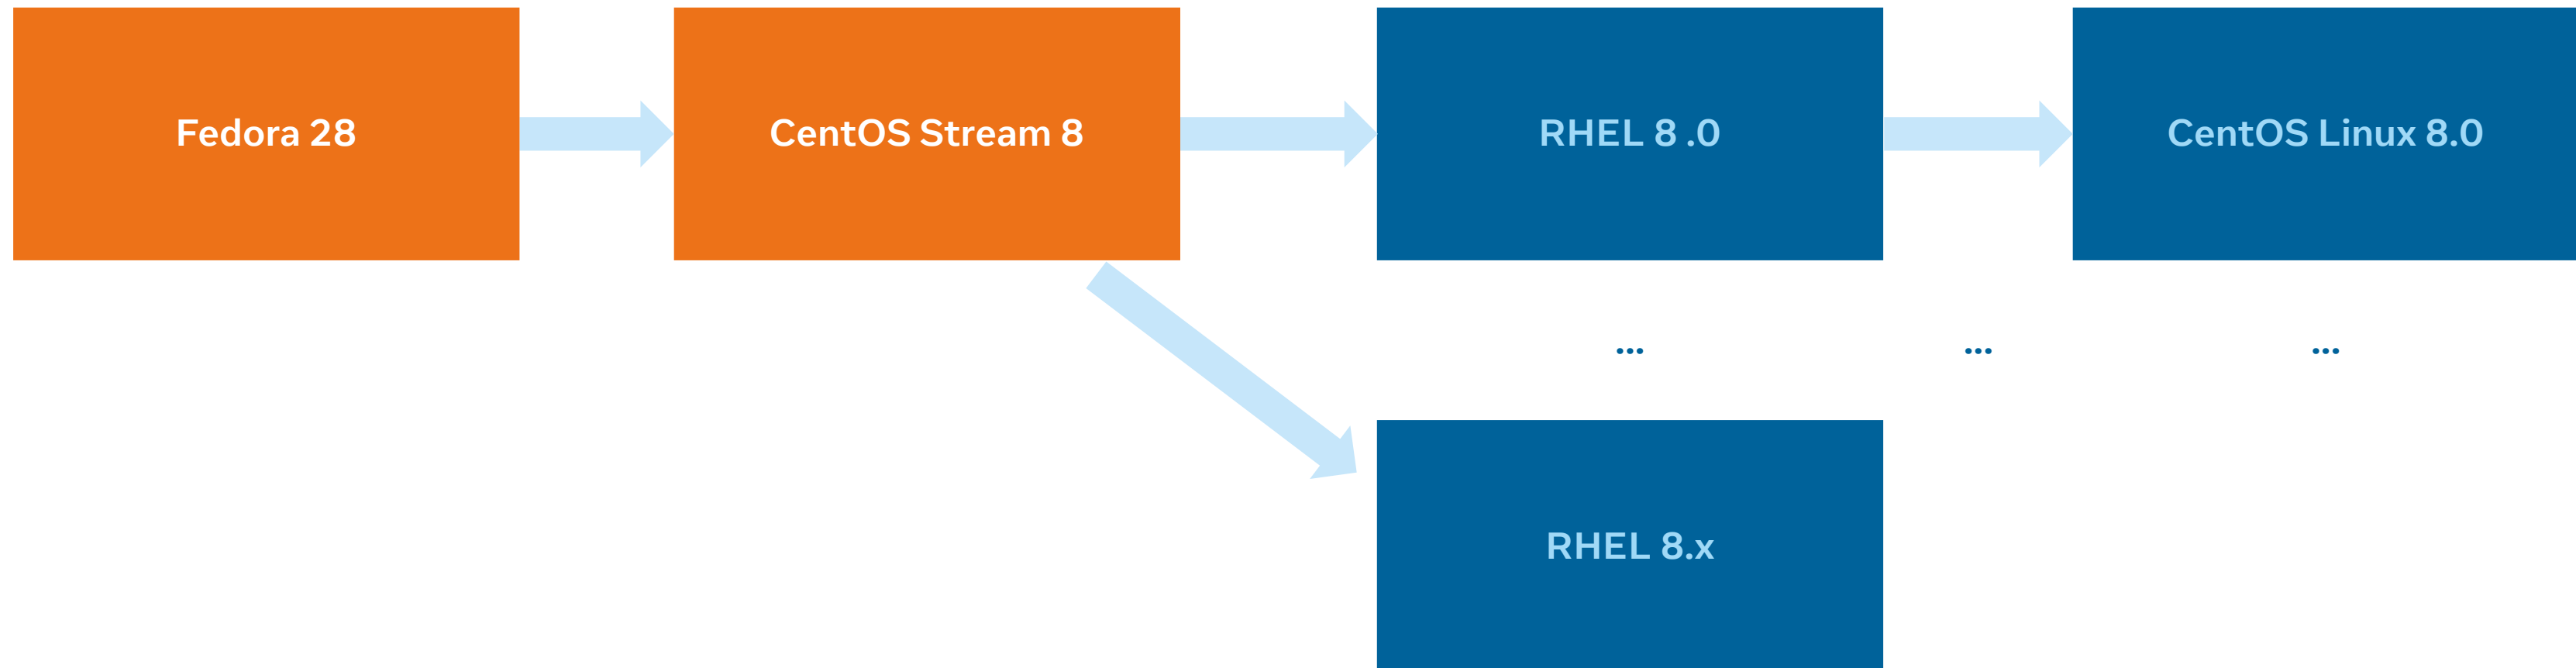
Contributing upstream

CentOS Linux 7




Contributing upstream

CentOS Linux 8 and CentOS Stream 8



Contributing upstream

CentOS Stream 9



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>
- Drive and submit change proposals

Fedora change proposals

Completed

- F33: btrfs by default
- F34: btrfs with zstd compression by default
- F34: systemd-oomd by default
- F35: btrfs by default for Fedora Cloud
- F36: relocate rpmdb to /usr

In progress

- F37: GPT for BIOS installs by default
- F37: Fallback hostname
- -fno-omit-frame-pointer by default
- Linux Firmware Minimization
- DNF RPM Copy-on-Write
- 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>
- ebranch: <https://pagure.io/epel/ebranch>

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 ELN to cover more packages via ELN Extras
 - <https://github.com/fedora-elN>

Contributing upstream

Fedora ELN at Meta

- Meta opensource project builds via PackIt
 - <https://copr.fedorainfracloud.org/groups/g/meta/coprs/>
- ELN Extras workload for packages we care about
 - https://tiny.distro.builders/config-workload--eln_extras_meta.html
- Continuous testing and integration pipeline
 - Covering 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>

Contributing upstream

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/production>

CentOS Stream 9 feature contributions

- systemd-oomd package and default configuration
- Packaging macros for third-party nginx modules
- PipeWire with WirePlumber and JACK compatibility
- Wayland support for the GNOME Classic session
- SDL2 support for GNOME Wayland
- libva and wayland updates
- ethtool 5.16 rebase

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>
- <https://sigs.centos.org/hyperscale>
- #centos-hyperscale on Libera.Chat and Matrix

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`

Faster-moving package backports

- 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, wireshark ...

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>
 - <https://git.centos.org/rpms/systemd/tree/c9s-sig-hyperscale>

systemd

- 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>
 - Leverages the CentOS OpenShift CI environment

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

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`

Kernel

- 5.14 kernel based on the CentOS Stream 9 kernel
- Development tree: <https://pagure.io/centos-sig-hyperscale/linux>
- Build for CentOS Stream 8 and CentOS Stream 9
- Btrfs support
- Work in progress, currently available in the experimental repo
- Kernel userspace: btrfs-progs, compsize, ethtool, kpatch

And more...

- Container image: <https://quay.io/centoshyperscale/centos>
 - Minimal container image based on Hyperscale repos and packages
 - Build from scratch:
<https://pagure.io/centos-sig-hyperscale/containers-releeng>
- Live media spins
 - Live DVD ISO images with Hyperscale repos and packages
 - Leveraging our kernel with btrfs support out of the box
 - <https://sigs.centos.org/hyperscale/spins/workstation/>
- Work in progress: cloud images, btrfs transactional updates, better testing, ...

Get involved

Get involved

Participate in the community

- 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>
- Read and contribute to the blog: <https://blog.centos.org>
- 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>
- Related talks: <https://sigs.centos.org/hyperscale/internal/talks>

Get involved

Join us in Boston in two weeks!

- Boston University, Aug 16-20th 2022
 - Hyperscale meetup
 - CentOS Dojo
 - DevConf.US
- More details: <https://tinyurl.com/34dy2ymz>



Questions?

THANK YOU FOR YOUR TIME

