

Hyperscale SIG update

January 2023

Davide Cavalca
CentOS Connect, FOSDEM 2023

Agenda

- SIG recap
- Deliverables and recent work
- Coming up

Hyperscale SIG

- CentOS Stream focus
- Large-scale infrastructure
- Foster cross-company collaboration on packages and tooling
- Bring in-house development out in the open
- Open to anybody interested in working in this space

SIG health

- Established in Jan 2021 with 6 founding members
- 29 members as of Jan 2023, steadily growing
- `#centos-hyperscale` on Libera.Chat
- `#centos-hyperscale:fedoraproject.org` on Matrix
- Bi-weekly IRC meetings: <https://sigs.centos.org/hyperscale/internal/meetings>
- Monthly hackathons / video hangouts
- Meetups!

Documentation

- Charter: <https://wiki.centos.org/SpecialInterestGroup/Hyperscale>
- User documentation: <https://sigs.centos.org/hyperscale>
- Activity reports: <https://sigs.centos.org/hyperscale/communication/reports/>
- Conference talks: <https://sigs.centos.org/hyperscale/internal/talks>
- Issue tracker: <https://pagure.io/centos-sig-hyperscale/sig>

Hyperscale SIG scope

- Faster moving package backports
- Policy and configuration alternatives
- Large-scale testing enablement
- Kernel
- Live DVD ISOs

Package backports

- Delivered to the Hyperscale main repository
 - `dnf install centos-release-hyperscale`
- Drop in replacements for stock CentOS packages
- File bugs on <https://pagure.io/centos-sig-hyperscale/package-bugs>
- Built against (and requiring) EPEL
- Targeting x86_64 and aarch64
- Recently added: zsh, fish, iperf3, dmidecode, dwarves, wireshark, fio, linuxptp
 - c8s: <https://cbs.centos.org/koji/packages?blocked=0&tagID=2249>
 - c9s: <https://cbs.centos.org/koji/packages?blocked=0&tagID=2378>
- Work in progress: openssh, qemu

Package updates

- Automated tracking for packages that need to be updated
- Files issues whenever a new upstream build supersedes a SIG package
- Leverages MQTT notifications
- Running in the OpenShift CI
- <https://pagure.io/centos-sig-hyperscale/package-updates>
- In the future: automated rebuilds?

systemd

- SIG branch tracking the latest upstream (currently 250, soon 252)
- Builds for CentOS Stream 8 and 9 available
- Defaults to unified hierarchy (cgroup2)
- systemd-oomd (needs PSI)
- systemd-networkd, systemd-resolved
- systemd-journald improvements (including CoW optimizations)
- Experimental SELinux enablement
- <https://github.com/systemd/systemd/blob/main/NEWS>

systemd

- Public repo to track development and stage patches
 - <https://pagure.io/centos-sig-hyperscale/systemd>
- Public repo for releng tools and scripts
 - <https://pagure.io/centos-sig-hyperscale/systemd-releng>
- Contribution guide
 - <https://sigs.centos.org/hyperscale/internal/systemd/>
- Contbuild on CentOS CI
 - Based on the current Hyperscale packaging
 - Daily rebuilds of the upstream git master for CentOS Stream 8 and CentOS Stream 9
 - In the process of being updated to support newer releases
 - In the future: End-to-end VM-based testing

Intel-optimized packages

- Intel-Zlib backport with SSE2 optimized slide hash function
 - <https://github.com/intel/zlib>
 - Major deflate performance improvement on x86_64
- Experimental glibc backport with perf improvements and bugfixes
- ABI and API compatible drop-in replacement
- CentOS Stream 8 work in progress:
<https://cbs.centos.org/koji/packages?blocked=0&tagID=2616>
- CentOS Stream 9 coming soon

CentOS Stream 9

- Available and ready for use:
<https://blog.centos.org/2021/12/introducing-centos-stream-9/>
- Hyperscale is targeting CentOS Stream 8 and 9 concurrently
- Upstream 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
 - libdvdnav-devel addition to CRB
 - ethtool 5.16 rebase

Large-scale testing

- Provide a way to test distro-wide changes in production settings
- DNF/RPM Copy-on-Write
 - <https://fedoraproject.org/wiki/Changes/RPMCoW>
 - Requires patched packaging stack
 - Currently deployed in production at FB
- Released to a dedicated “experimental” repo
 - `dnf install centos-release-hyperscale-experimental`
- Ongoing development
 - <https://hackmd.io/@chantra/BJ1Hedw0F>
 - <https://github.com/rpm-software-management/rpm/discussions/2057>
 - <https://github.com/rpm-software-management/rpm/pull/2378>

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
- Secure boot is still outstanding: <https://pagure.io/centos-infra/issue/307>
- Available in the experimental repo

Kernel collaboration work

- CentOS Stream 9 kernel contributions
 - zstd sync with 5.16 and configuration enablement
 - CONFIG_CRYPTOLIB_BLAKE2B
 - Build fixes
- Contribution guide: <https://sigs.centos.org/hyperscale/internal/kernel/>
- kmod-btrfs for CentOS Stream 9
 - Targeting the stock CentOS Stream 9 / RHEL 9 kernel
 - Published by the CentOS Kmods SIG

Kernel userspace

- Btrfs
 - btrfs-progs backported from Fedora, currently at 5.16.2
 - compsize 1.5
 - Restored btrfs support in the storage stack and installer
 - Anaconda installer development tree: <https://pagure.io/centos-sig-hyperscale/anaconda>
 - Released to dedicated “spin” repo for media
 - Used for Hyperscale spin media
- ethtool 5.16
 - Hyperscale backport for CentOS Stream 8
 - Rebased in CentOS Stream 9
- kpatch 0.9.7
 - Includes `kpatch-build`
 - Backported clang PGO support and other fixes

Container image

- Minimal container image based on Hyperscale repos and packages
- <https://quay.io/centoshyperscale/centos>
- Build from scratch and automatically published from OpenShift CI
 - <https://pagure.io/centos-sig-hyperscale/containers-releng>
- CentOS Stream 8

```
podman run -a stdin,stdout,stderr -t quay.io/centoshyperscale/centos:stream8
```

- CentOS Stream 9

```
podman run -a stdin,stdout,stderr -t quay.io/centoshyperscale/centos:stream9
```

Live media spin

- Two CentOS Stream 8 spins: GNOME and KDE Plasma
- Live DVD ISO images with Hyperscale repos and packages
- Leveraging our kernel with btrfs support out of the box
- Download: <https://sigs.centos.org/hyperscale/spins/workstation/>
- Report issues: <https://pagure.io/centos-sig-hyperscale/spin-bugs>
- CentOS Stream 9 spins coming soon

Coming up

- CBS-based live media image builds
- Updated QEMU package in EPEL
- Btrfs transactional updates
- Hyperscale-enabled cloud images

Resources

- <https://wiki.centos.org/SpecialInterestGroup/Hyperscale>
- <https://sigs.centos.org/hyperscale>
- #centos-hyperscale on Libera.Chat
- #centos-hyperscale:fedoraproject.org on Matrix
- Bi-weekly meetings and monthly VC hangout
 - <https://www.centos.org/community/calendar>
- <https://pagure.io/centos-sig-hyperscale/sig> issue tracker
- centos-devel mailing list

Questions?