What's new and what's next in Fedora CoreOS

Dusty Mabe
Principal Engineer at Red Hat

Mike Nguyen
Senior Engineer at Red Hat

Agenda

- What is Fedora CoreOS?
- What happened last year?
- What's new since last year?
- What's coming soon?

What is Fedora CoreOS?

An emerging Fedora edition

- Came from the **merging** of two communities:
  - CoreOS Inc's Container Linux
  - Project Atomic's Atomic Host

- Incorporates **Container Linux**
  - Philosophy
  - Provisioning Stack
  - Cloud Native Expertise

- Incorporates **Atomic Host**
  - Fedora Foundation
  - Update Stack
  - SELinux Enhanced Security

Philosophy behind Fedora CoreOS

- **Automated provisioning**
  - All nodes start from **~same starting point**
  - Use Ignition to provision a node on **first boot**

- **Immutable infrastructure**
  - **Automate** deployment and system configuration
  - Update configs and **re-provision** to apply changes

- Additional software runs in **containers**
  - Makes host updates more **reliable**

- **Automatic updates**
  - No interaction for administrators

None of this happens without... Passing Tests!

✓	☼	kola-aws	13 hr - #325	4 days 0 hr - #312	1 hr 11 min
✓	☼	kola-azure	13 hr - #80	12 days - #57	45 min
✓	☼	kola-gcp	13 hr - #169	3 days 14 hr - #164	15 min
…	☼	kola-kubernetes	N/A	N/A	N/A
✓	☁	kola-openstack	13 hr - #340	1 day 10 hr - #337	45 min

Multiple Update Streams

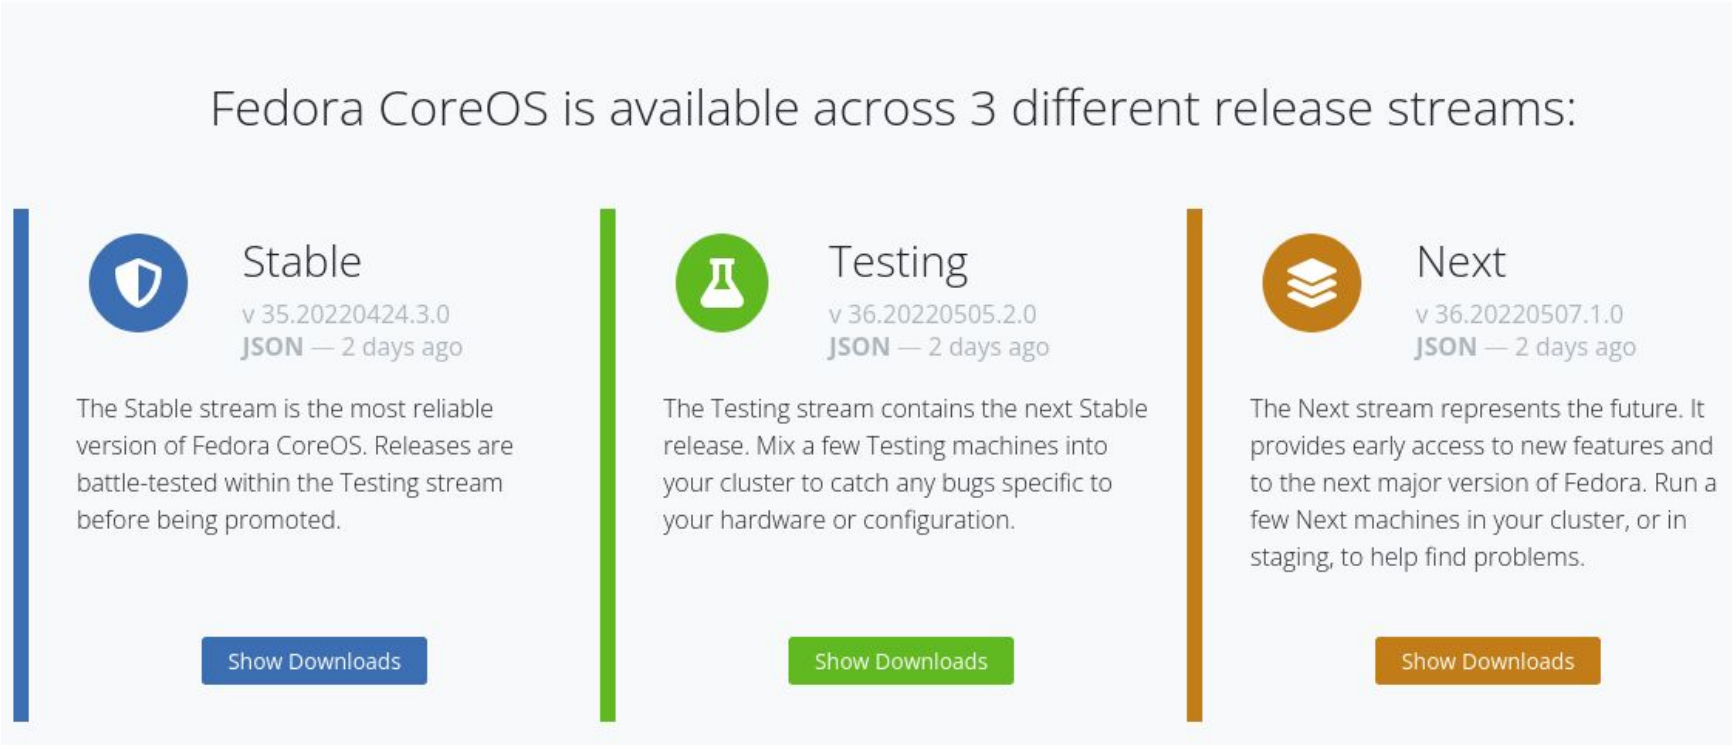

Supported platforms and architectures

- Available for a plethora of **cloud/virt platforms**
  - Alibaba, AWS, Azure, DigitalOcean, Exoscale, GCP, IBM Cloud, OpenStack, Vultr, VMware, QEMU/KVM
  - Directly launchable on AWS & GCP

- Several options for **Bare Metal**
  - Live ISO (automated or interactive installations)
  - PXE (network) boot
  - Raw and 4K native disk images

- Multiple Architectures (**x86_64**, **aarch64**, **s390x**)
  - (**ppc64le** support coming soon)

What happened last year?

cgroups v2 by default

- Switched to **v2 by default**
- Did this later than the rest of Fedora
- Some container software wasn't ready
  - Most notably Kubernetes

https://docs.fedoraproject.org/en-US/fedora-coreos/kernel-args/#_removing_existing_kernel_arguments

Reliable live changes to the system

- New options to change the system content **live** in a **safe**, and **atomic** way
- rpm-ostree install --apply-live strace
  - **Install** a package into a new (offline) deployment
  - Atomically switch the running system to this deployment to **apply the changes live** (still RO)
- No longer need a reboot

rpm-ostree v2021.1 & rpm-ostree v2021.3 & https://coreos.github.io/rpm-ostree/apply-live/

Kernel arguments in Ignition

- Add, remove, replace kernel arguments **via Ignition**
- Applied on **first boot**, will trigger a reboot

```
# Disabling CPU
# vulnerability mitigations
variant: fcos
version: 1.4.0
kernel_arguments:
  should_exist:
    - mitigations=off
  should_not_exist:
    - mitigations=auto,nosmt
```

```
# Staying on cgroups v1
variant: fcos
version: 1.4.0
kernel_arguments:
  should_exist:
    - systemd.unified_cgroup_hierarchy=0
```

https://docs.fedoraproject.org/en-US/fedora-coreos/kernel-args/#_modifying_kernel_arguments_via_ignition

/boot is now read-only

- Manually modifying content in /boot is **discouraged**
- Change **kernel arguments** with:
  - rpm-ostree kargs
- Change **boot order** with:
  - rpm-ostree rollback / update / deploy

https://docs.fedoraproject.org/en-US/fedora-coreos/storage/#_mounted_filesystems

Encrypted storage via LUKS in Ignition

- Unlock via a keyfile, TPM2 or a Tang server (via Clevis)
- Includes support for the **root partition**
  - Requires unlocking via a TPM2 or a Tang server

```
# LUKS for / using TPM2
variant: fcos
version: 1.4.0
boot_device:
  luks:
    tpm2: true
```

```
# LUKS for another device
variant: fcos
version: 1.4.0
storage:
  luks:
    - name: data
      device: /dev/vdb
      clevis:
        tpm2: true
  filesystems:
    - path: /var/lib/data
      device: /dev/mapper/data
      format: xfs
      label: DATA
      with_mount_unit: true
```

https://docs.fedoraproject.org/en-US/fedora-coreos/storage/#_encrypted_storage_luks

RAID support in Ignition

- Setup any RAID level (0, 1, 5, etc.) on first boot **via Ignition**
- Mirrors EFI System Partition (ESP) & BIOS bootloader
- Side effect: ESP no longer mounted (empty /boot/efi)

```
# Mirror boot disk with RAID1
variant: fcos
version: 1.4.0
boot_device:
  mirror:
    devices:
      - /dev/sda
      - /dev/sdb
```

```
# Move / to RAID0
variant: fcos
version: 1.4.0
storage:
  raid:
    - name: myroot
      level: raid0
      devices:
        - /dev/disk/by-id/virtio-disk1
        - /dev/disk/by-id/virtio-disk2
  filesystems:
    - device: /dev/md/myroot
      format: xfs
      wipe_filesystem: true
      label: root
```

https://docs.fedoraproject.org/en-US/fedora-coreos/storage/#_reconfiguring_the_root_filesystem

What’s new in Fedora CoreOS?

(since August 2021)

Added Platforms

- Added aarch64 - including AWS images
- Added s390x - including s390x IBMCloud images
- Added support for Nutanix
- Became the base for `podman machine`

DNF Count Me support

- Enables **privacy preserving** and reliable system counting
- Only reports a **large approximation** of the age of a system
- Only reaches out to **official** Fedora repositories servers
- **No other information** sent or stored

https://fedoramagazine.org/getting-better-at-counting-rpm-ostree-based-systems/
https://github.com/coreos/fedora-coreos-tracker/issues/717

CountMe Stats - All Nodes

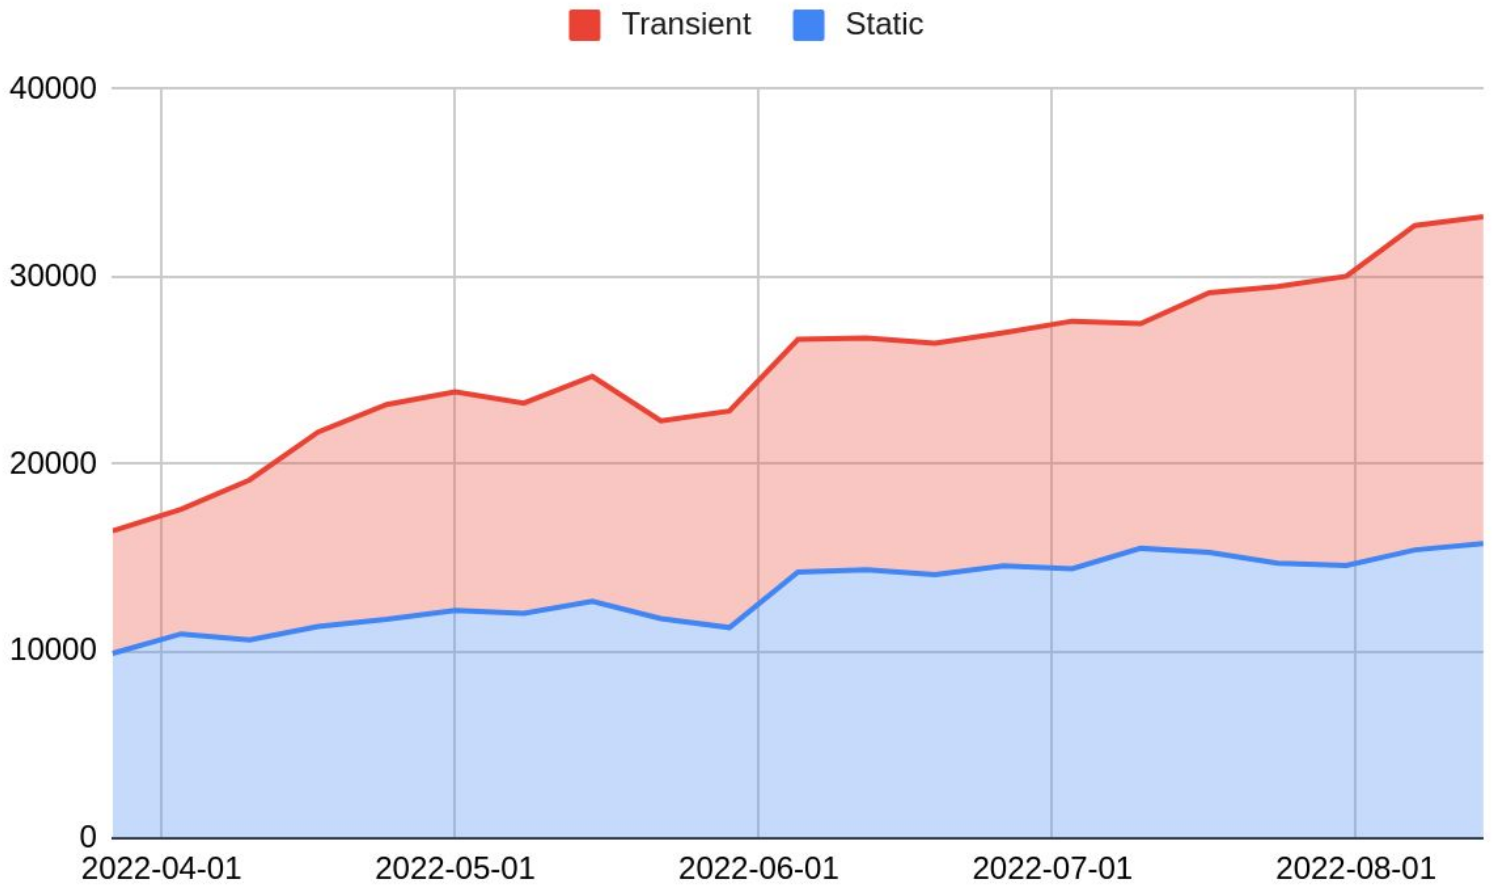

CountMe Stats - Fedora Release

CountMe Stats - Architecture

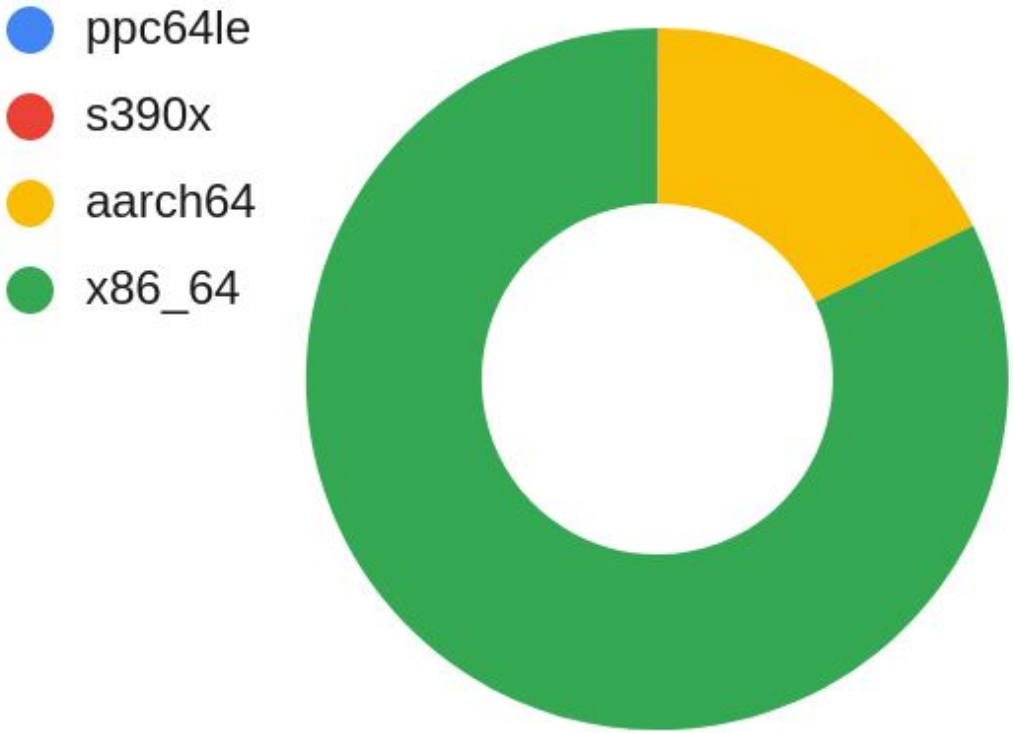

iptables using nftables by default

- iptables still using legacy backend instead of nftables one
- **Unintended** consequence of alternatives(8) 's behaviour
  - Configuration stored in a mix of /var and /etc
  - **Incompatibility** with rpm-ostree strict split between configuration and data
- Easy **workaround** available
- Full fix requires **adjustments** to alternatives(8) or an alternative(!)

https://docs.fedoraproject.org/en-US/fedora-coreos/alternatives/
https://github.com/coreos/fedora-coreos-tracker/issues/676
https://github.com/coreos/fedora-coreos-tracker/issues/677

ostree commits in container images

- New commands to export an ostree commit to a **container image**
- Enables **rebasing** to the content of a container image:
- Enables **running** an ostree commit as a container for testing and debugging:

https://lists.fedoraproject.org/archives/list/devel@lists.fedoraproject.org/thread/B23FZILDI3J73OMION2IDEYMLKNKN5YE/

Building and Testing against Rawhide

- We are now building and testing a **rawhide** stream
  - Suite of automated tests now complement rawhide!
  - Helps identify unexpected breakage from new features.
  - Now participate closer upstream with developers and get general problems fixed.

And.. A lot of boring stuff

- Investments in
  - Automated Testing
  - Automated builds

What’s coming soon in Fedora CoreOS?

FCOS as a top level Fedora Edition

- Closer Proximity to Fedora Releases
- Working closer with Fedora QA
- Participating more in the Fedora Change Process
- Continuing to deliver Fedora CoreOS consistently

More enablement - More Platforms

- Azure ARM instances
- GCP ARM instances
- Azure Community Galleries
- Power PC (ppc64le) architecture support
- Adding support for Kubevirt
- Secure Execution support for S390x

CoreOS Layering / OSTree Native Containers

- Fedora CoreOS OSTree is additionally offered as a container
- Customize Fedora CoreOS by performing a container build
  - FROM: quay.io/fedora/fedora-coreos
- Makes individual derivation and distribution easier
  - Dockerfile & Container registry

CoreOS Layering Example

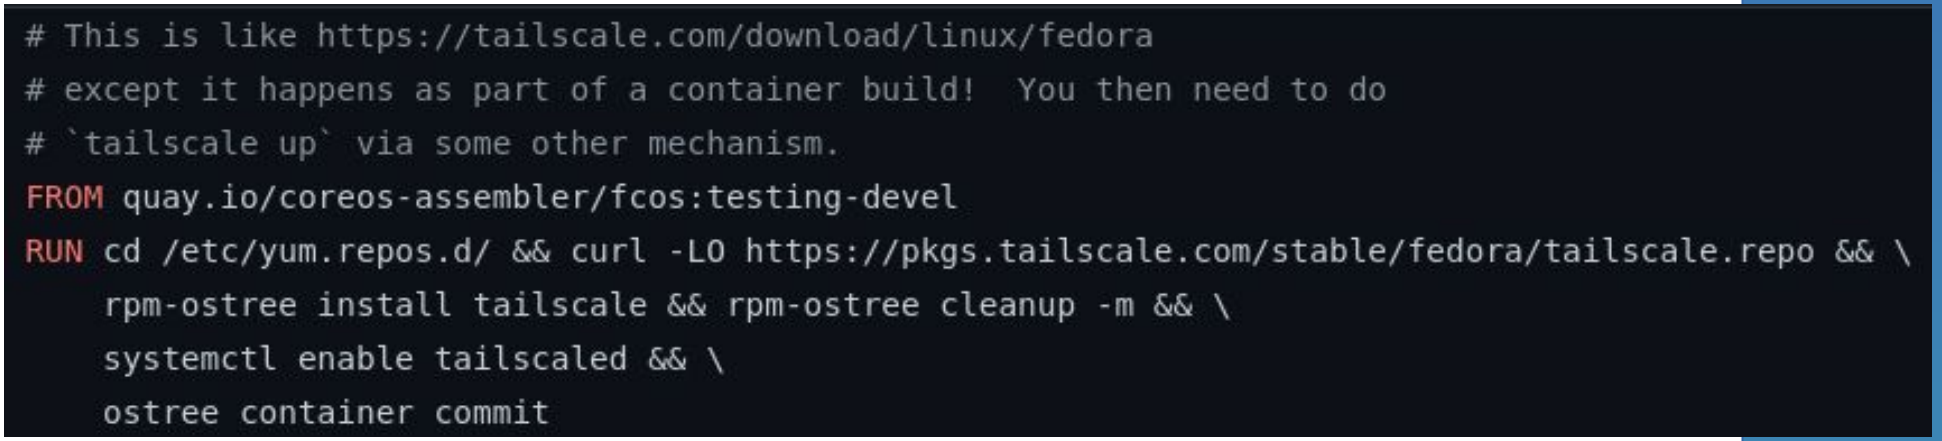

```
# This is like https://tailscale.com/download/linux/fedora
# except it happens as part of a container build!  You then need to do
# `tailscale up` via some other mechanism.
FROM quay.io/coreos-assembler/fcos:testing-devel
RUN cd /etc/yum.repos.d/ && curl -LO https://pkgs.tailscale.com/stable/fedora/tailscale.repo && \
    rpm-ostree install tailscale && rpm-ostree cleanup -m && \
    systemctl enable tailscaled && \
    ostree container commit
```

- https://github.com/coreos/coreos-layering-examples

CoreOS Layering Example

- podman build -t myfcos:latest .
- podman push myfcos:latest quay.io/$USER/myfcos:latest
- rpm-ostree rebase --experimental \
  ostree-unverified-registry:quay.io/$USER/myfcos:latest

Questions

Get involved!

- Web: https://getfedora.org/coreos
- Issues: https://github.com/coreos/fedora-coreos-tracker/issues
- Forum: https://discussion.fedoraproject.org/tag/coreos
- Mailing list: coreos@lists.fedoraproject.org
- IRC: Libera.chat #fedora-coreos
- Other talks to get started:
  - Fedora CoreOS Introduction (Jul 13, 2020)
  - Getting Started with Fedora CoreOS (Mar 17, 2021)