

INTRODUCING BTRFS

The butter of Fedora

NEAL GOMPA

Fedora Project Contributor

 ngompa@fedoraproject.org

 [@Det_Conan_Kudo](https://twitter.com/Det_Conan_Kudo)

 [@Conan_Kudo](https://github.com/Conan_Kudo)



WHO AM I?

- Open Source Advocate
- Contributor and package maintainer in Fedora, openSUSE, Mageia, and OpenMandriva
- Member of Fedora Engineering Steering Committee (FESCo)
- Member of openSUSE Board
- DevOps Engineer at Datto, Inc.
- Developer and maintainer of package build and release pipeline using OBS
- Packager and Backporter

Neal Gomba

✉ ngomba@fedoraproject.org



SO... BTRFS?



BTRFS

From the [Btrfs wiki](#):

Btrfs is a copy on write (CoW) filesystem for Linux aimed at implementing advanced features while focusing on fault tolerance, repair and easy administration.

COPY ON WRITE

The label Copy on Write (CoW) refers to a type of filesystem optimization strategy where each modification to the filesystem is written in a new location while the original remains preserved.

By doing this, it's possible to preserve each instance of the filesystem and move back and forth through the instances.

It is NOT a replacement for proper backups, but it does provide some safety that isn't possible in traditional filesystems.

VOLUME LIMITS

64-bit filesystem, max volume size is 16 EiB

- Approximately 18,446,744 terabytes!
- At 100GB per 4K (QHD) full-length film, it'd take 184,467,441 copies to fill the whole volume at max size!
- This is more data than what is even possible (or even desired!) to record today on any single disk or disk array.

FEATURES

- Space efficient storage/packing of small files
- Space efficient indexing of directories
- Subvolumes & quota support for subvolumes
- Read-only and writable snapshots
- Sending/receiving volume data with efficient deltas
- SSD awareness and SSD-specific optimizations
- Integrated disk management & multiple disk support
 - RAID 0, 1, 5, 6, 10 support
 - Dynamic resizing (shrink/grow) arrays/volumes after initial array creation
- Transparent on-disk compression
- Seeding from other filesystems
- [And much more...](#)

SUBVOLUMES AND SNAPSHOTS

Subvolumes are subsections of a volume that can be independently managed. This is useful if you want to have different snapshotting schedules for portions of your volume.

Snapshots are instances of (sub)volumes that are preserved. With the appropriate tools and configuration, snapshots can be used as a means to provide “Time Machine” style data recovery or even to save a system from a bad software install/upgrade.

DEVELOPER COMMUNITY

It is principally developed by:



Western Digital[®]

ORACLE[®]

Neal Gomba

✉ ngomba@fedoraproject.org

fedora^f

USER COMMUNITY

It is used in production by:



WHY BTRFS

As a filesystem that is developed within the mainline kernel, it takes advantage of facilities provided in the kernel to be more efficient at doing operations on devices.

Linux distributions also have support for Btrfs out of the box, and can be used with minimal effort.

BTRFS AND FEDORA



CURRENT STATE AND FUTURE PLANS



Fedora 34

- Anaconda has been configured to install non-server variants with Btrfs
- Disk images of desktop variants provide Btrfs-based images
- /boot is not on Btrfs currently by default
 - Installation is possible with /boot as Btrfs subvolume or separate Btrfs volume
- Zstd compression by default
- Disk encryption uses LUKS
 - LUKS with Btrfs means only full disk encryption is possible

Fedora 35+

- /boot on Btrfs by default
- Online/Live native disk encryption
 - Pending upstream work
- Full support for Btrfs for osbuild
 - Subvolume creation is missing
- Factory reset/recovery for OEM setup
- Simpler setup for full system snapshotting and boot-to-snapshot
 - Pending coordination with bootloader team and snapper developers
- Server/desktop with atomic updates and transferrable snapshots
 - Based on openSUSE's approach with libdnf plugin and microdnf/PK

Neal Gomba

✉ ngompa@fedoraproject.org

FUTURE HOPES AND DREAMS

Linux has the potential to offer a superior storage experience for desktops and servers using Btrfs. The unified storage management interface and comprehensive feature set for all use-cases enables many opportunities further up the stack:

Perhaps one day, we can see these features leveraged by desktop environments and third-party software to solve user problems and make the functionality of the Linux desktop surpass that of proprietary alternatives.

RESOURCES

- Btrfs wiki: <https://btrfs.wiki.kernel.org/>
- openSUSE Btrfs documentation: <https://en.opensuse.org/SDB:BTRFS>
- Btrfs + Snapper (F31 Edition):
<https://dustymabe.com/2019/12/29/fedora-btrfs-snapper-the-fedora-31-edition/>
- Nest with Fedora Btrfs talk: <https://www.youtube.com/watch?v=iHjhouSxIrc>
- Btrfs at Facebook: <https://lwn.net/Articles/824855/>
 - YouTube video: <https://www.youtube.com/watch?v=U7gXR2L05IU>
- Fedora Btrfs Pagure project: <https://pagure.io/fedora-btrfs/project>

QUESTIONS?

