Btrfs, Two Years In

Neal Gompa (Conan Kudo [ニール・ゴンパ])

Who am I?

- Professional technologist
- Contributor and package maintainer in the Fedora Project
- Contributor and package maintainer in Mageia Linux
- Contributor to RPM, DNF, and various related projects
- Diligent follower of the telecommunications industry
- DevOps Engineer at Datto, Inc.

Contact Points:

- Twitter: @Det_Conan_Kudo
- Google+: +NealGompa

Recap: What is Btrfs?

From the Btrfs wiki:

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

Recap: Features of Btrfs

- Space efficient storage/packing of small files
- Space efficient indexing of directories
- Subvolumes & quota support for subvolumes
- Read-only and writable snapshots
- Sending/receiving snapshot data
- 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...

Recap: Who makes Btrfs?

It is principally developed by:

Recap: Who uses Btrfs?

It is used in production by:

My use of Btrfs

- Nearly every computer I've set up since April 2015 has been configured with varying Btrfs configurations
- Desktops/Workstations:
  - Subvolumes for root filesystem and home folder
  - Fedora systems with /boot on ext4 (due to lack of support for subvolume boot)
  - openSUSE systems with /boot on btrfs as subvolume for snapshotting
  - VM disk image directory (/var/lib/libvirt/images) is an XFS filesystem
- Servers
  - Unified root filesystem as btrfs and snapshotted regularly, RAID 1
  - Specialized data partitions with Btrfs on RAID 10
  - CentOS/Fedora systems have separate /boot on xfs (for already noted reason)
  - Database and VM disk image directories are broken out as XFS filesystems

Noted caveats

Some general CoW-centric caveats:

- Applications and services that are very heavy on I/O can have issues
  - Examples: Databases, virtual machines
- Always leave ~10% of “space” free, or otherwise write operations will fail
  - “Copy on Write” means that at least for a given moment, it is possible for a file to take 2x the space.

Some Btrfs-specific caveats:

- RAID 5 / RAID 6 is not completely sane yet, so should be avoided for the time being.
  - The Btrfs developers have indicated that this is something that they expect to resolve this year.
- Use at least Linux kernel 4.4 to leverage significant performance improvements
  - Noted exception: Red Hat Enterprise Linux/CentOS 7.3 and newer have kernel 4.4+ Btrfs code

The End

Questions?