

# Encrypted btrfs subvolumes: new features

# Agenda

Intro: btrfs, status

Technical aspects: fscrypt, motivation, extent-based encryption; status, learnings, solutions

Future goals: more features from LUKS and bcachefs

Practical usage: soon

# Introduction to btrfs

# Intro to advanced filesystems

- Not a judgement on quality, just a convenient alias for a particular set of features.
- Reflinks, subvolumes, snapshots, checksums
- Btrfs, XFS, Bcachefs

# Motivation: btrfs

- Btrfs has long wanted to have encryption, but doesn't want to give up checksumming or reflinking.
- By having per-subvolume encryption, individual user homedirs can have unique keys.

# Technical challenges for btrfs

# What is fscrypt?

- Kernel library providing a standard encryption interface across filesystems using it.
- Used on Android
- Ext4, f2fs, ceph, ubifs as yet
- One master key per directory tree
- No mixing keys within one tree
- Can delete files without their key
  
- Only encrypts filenames and data (vs LUKS/dm-crypt, which encrypts everything)
- Crypto either with crypto api or blk-crypto
  - With blk-crypto, filesystem never sees encrypted data
- struct inode embeds struct fscrypt\_inode\_info, storing on disk as struct fscrypt\_context:
  - Encryption is based on file + file offset
  - Key either master key+nonce applied to plaintext, or a derived key from master key + nonce

# Difficulties for advanced filesystems

- **No mixing keys within one tree**
  - Breaks nested subvolumes with different master keys
- Crypto either with crypto api or blk-crypto
  - **With blk-crypto, filesystem never sees encrypted data**
    - Unsafe to store checksums of plaintext
- **struct inode embeds struct fscrypt\_inode\_info, storing on disk as struct fscrypt\_context**
  - **Encryption is based on inode + file offset**
    - One piece of data can be reflinked into two inodes at different offsets. How to make both inodes decrypt it successfully? Awkward...



# Extent-based encryption

- Still has a struct `fsencrypt_inode_info` / `struct fsencrypt_context` for inodes.
- **struct `fsencrypt_extent_context` per extent**
  - **Encryption is based on extent + extent offset**
  - No issue reflinking an extent into two inodes anymore
  - Stores key, so in theory every extent can have a different key
- Takes more metadata space usually

Current state

# History

- Design 1 in Oct '21 by Omar Sandoval
  - Per-extent context contained nonce only
  - Encryption using master key directly only
  - Patches Jun-Oct '22
  - Risks of master key reuse for too much data
  - Crypto api only
  - Checksum encrypted data
- Design 2 in Nov '22 by me
  - Per-extent context reusing 'normal' per-inode context struct
  - Patches Jan-Aug '23
  - Insufficiently elegant
  - Blk-crypto only
  - Checksummed unencrypted data

# Current state

- Design 3 in Sep '23 by Josef Bacik
  - Per-extent context with nonce and key (must match inode key for now)
  - Encryption restricted to derived key from inode context + extent nonce
  - V5 in flight
  - Doesn't support nested subvols with different keys or full range of key options, but enough information is in the context to do so
  - Still needs review, hasn't gotten into the kernel

# Addresses previous difficulties

- **Still doesn't allow changing keys within one tree**
  - Nested subvolumes still don't work, but enough info is stored to allow changing key between inodes.
- Extent-based only with blk-crypto
  - **Adds a callback to blk-crypto to allow checksumming encrypted data**
- **Encryption is based on inode + extent + extent offset**
  - **Addresses reflinking between inodes with the same key, can be extended to allow reflinking between inodes with different keys**

Future goals

# Bcachefs has different features

- Doesn't use fscrypt
- Only one encryption key per filesystem
- Everything is encrypted: no access to anything, even for deletion, when the key isn't loaded
- Authenticated encryption instead of encryption + checksums of encrypted data
- Less options for encryption algorithm

# LUKS (dm-crypt + dm-integrity) has different features

- Only one encryption key per filesystem
- Everything is encrypted: no access to anything, even for deletion, when the key isn't loaded
- Authenticated encryption instead of encryption + checksums of encrypted data
- Encryption key changes
  - Useful for repudiation or changing to a newer encryption algorithm
- Encrypts everything



# Key change motivation

- There's a Fedora proposal to use btrfs encryption one day
  - initial unencrypted or encrypted image installed on disk by OEM
  - Company or user sets new key on /, installs own packages, sets up homedir template
  - User sets new key for homedir
- Meta once and may again want to install an unencrypted package in subvolume, run in container with per-subvolume key for anything written by package.
- Also useful for companies that require changing passwords every so often.

How will it be  
usable?

Almost ready  
to use once  
it's in the  
kernel

- Obviously any new feature may have bugs, but early adopters' testing will be invaluable
- Uses the same kernel keyring facility as LUKS
  - systemd-cryptsetup-generator already knows how to ask for LUKS password and send it to the kernel, just need to use it for encrypted btrfs.
- Systemd-homedir hopefully will one day be able to have a per-use key for a per-user subvolume, mounted on login.

# FACEBOOK Infrastructure