dm-user

Linux Plumbers
August 24th 2020

dm-user

- Block device in userspace
- Ship in Android S (in OTA flow), then upstream
  - Most of the work for S is in userspace
- Smells like FUSE, but with a block device
  - Userspace daemon talks to miscdev
- drivers/md/dm-user.c
  - Originally a dm-snap-user, so just for DM snapshots
- Simple implementation now, plans for more features
  - and fix the deadlocks...
- Target Linus' tree
  - One other possible Android use case already (dm-bow)
  - Also an internal Google use case
  - More here?

Virtual A/B Compression

- Virtual A/B updates use device mapper snapshots
  - Android has its own user tools, but doesn’t touch this part of the kernel
- Snapshots merged when update is deemed successful
  - Updates require much less downtime
  - Failures can be rolled back
- Generic COW format is 10x larger than Android-specific OTA format
  - Byte based, with circular dependencies
  - bsdiff and compression
  - Android-specific tuning (verity, reaches inside APKs)
- dm-user allows Android-specific COW format
  - Currently simple, but will likely continue to evolve

Current ABI

```
struct dm_user_message {
        u64 type, flags;
        u64 sequence_number;
        u64 data_length;
        u8 data[];
};
```

- One open() per block device
- Two streams of messages
  - read() gets a new request
  - write() responds to a request
- Sequence number so userspace can reorder responses
  - Entire response buffered at once
- Data in-line, buffers allocated by userspace
- Everything is blocking
  - Large BIOs fragmented by kernel

Issues and Plans

- Deadlocks
- Userspace can only reply to one message at a time
- Daemon can’t die
- Every BIO round trips to userspace
- Single-queue only