F2FS

Jin-Soo Kim
(*jinsoo.kim@snu.ac.kr*)

Systems Software &
Architecture Lab.

Seoul National University

Fall 2022

(C. Lee et al., FAST 2015)

Mobile Workloads

- Facebook mobile app issues 150% more random writes than sequential writes [Kim et al., FAST '12]

- Facebook and Twitter: Over 80% of total I/Os are random and more than 70% of the random writes are triggered with fsync [Jeong et al., ATC '13]

- Dominant use of SQLite

F2FS

- Based on the Log-structured File System (LFS)
- Flash-friendly on-disk layout
- Cost-effective index structure
- Multi-head logging
- Adaptive logging
- Fsync acceleration with roll-forward recovery

- Publicly available, included in Linux mainline kernel since Linux 3.8
- Being used in commercial Android smartphones

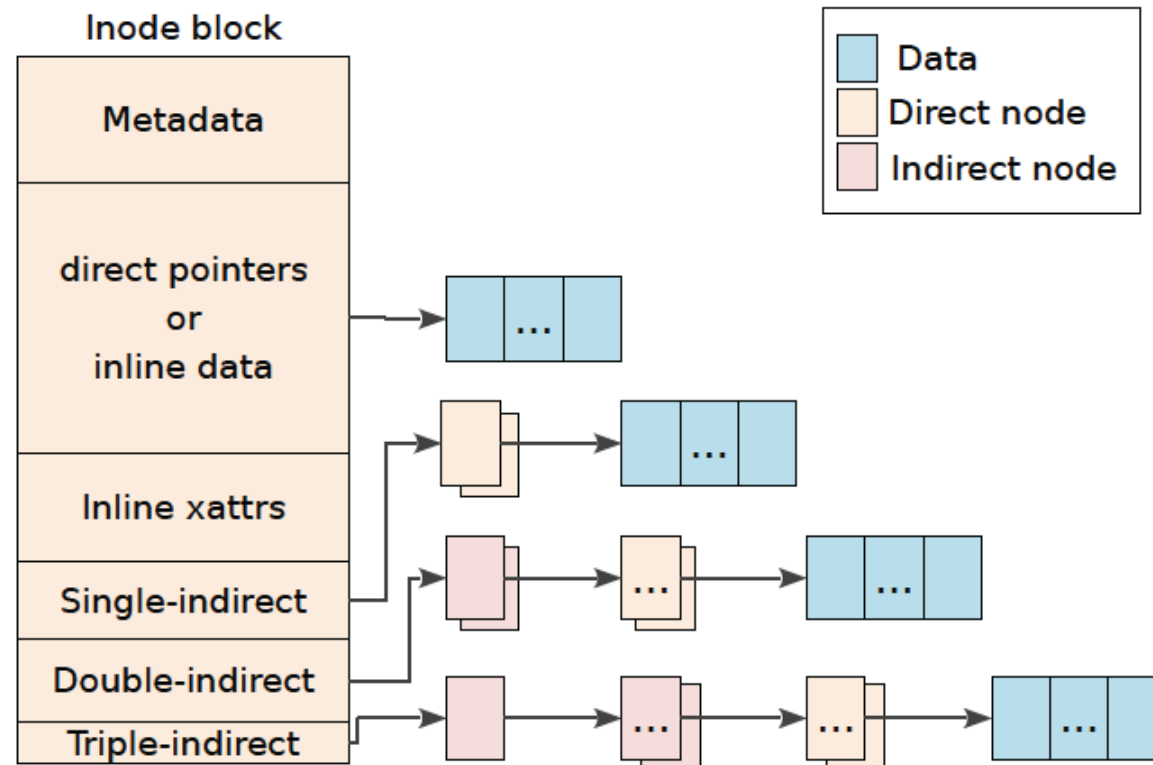[저자직강] F2FS Features in Android (구글 김재극박사)

[고수의 세미나] F2FS, How it works?

Flash-friendly On-Disk Layout

- **Flash awareness**
  - All the metadata are located together for locality
  - Start address of main area is aligned to the zone size
  - Cleaning is done in a unit of section (FTL's GC unit)

- **Cleaning cost reduction**
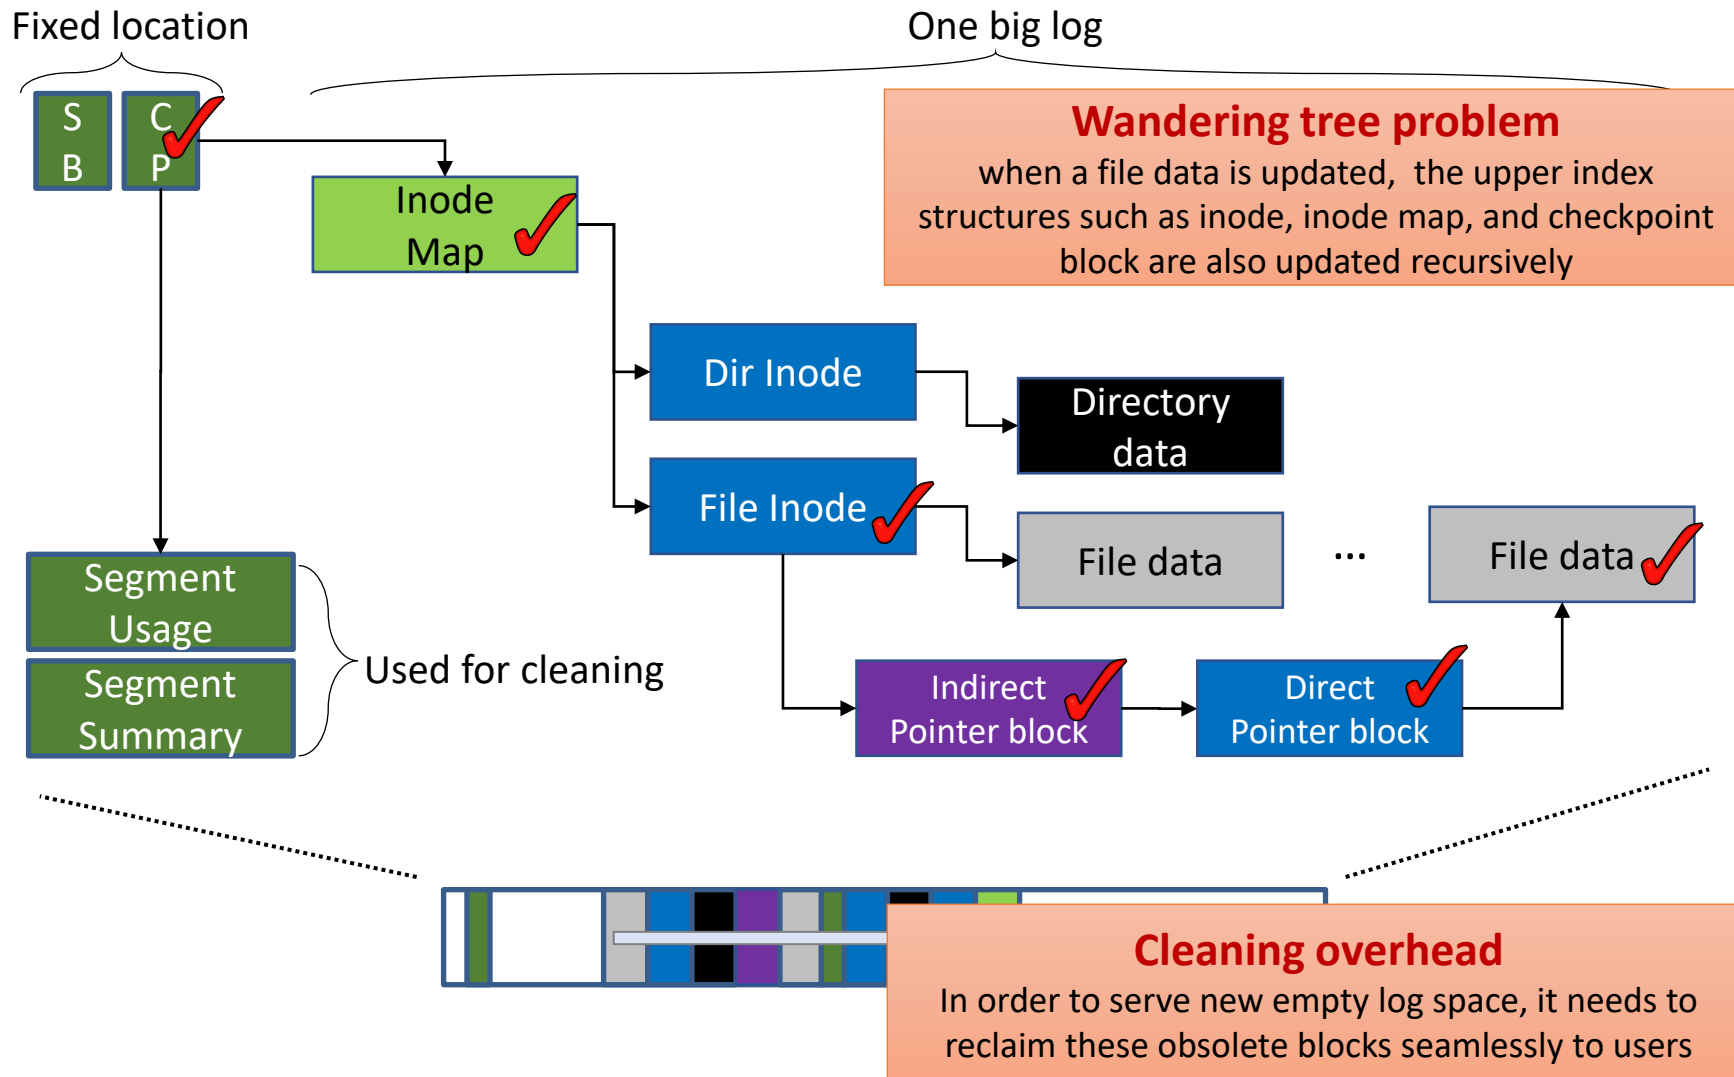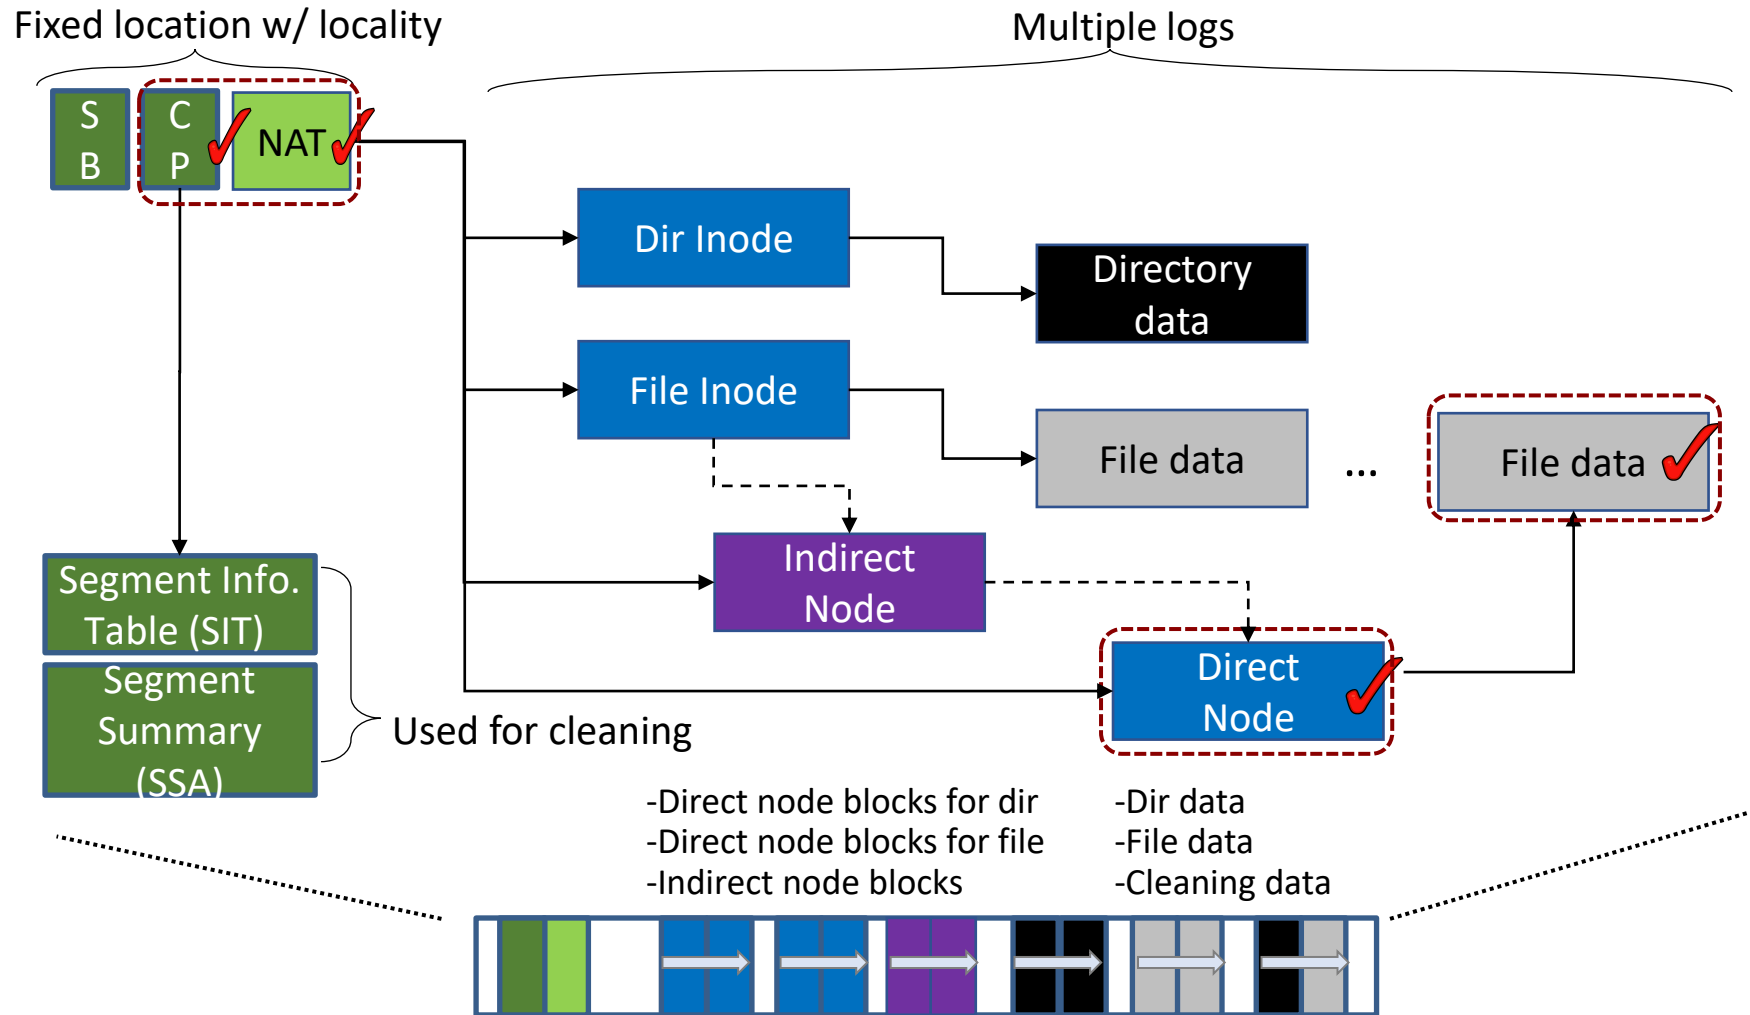 - Multi-head logging for hot/cold data separation
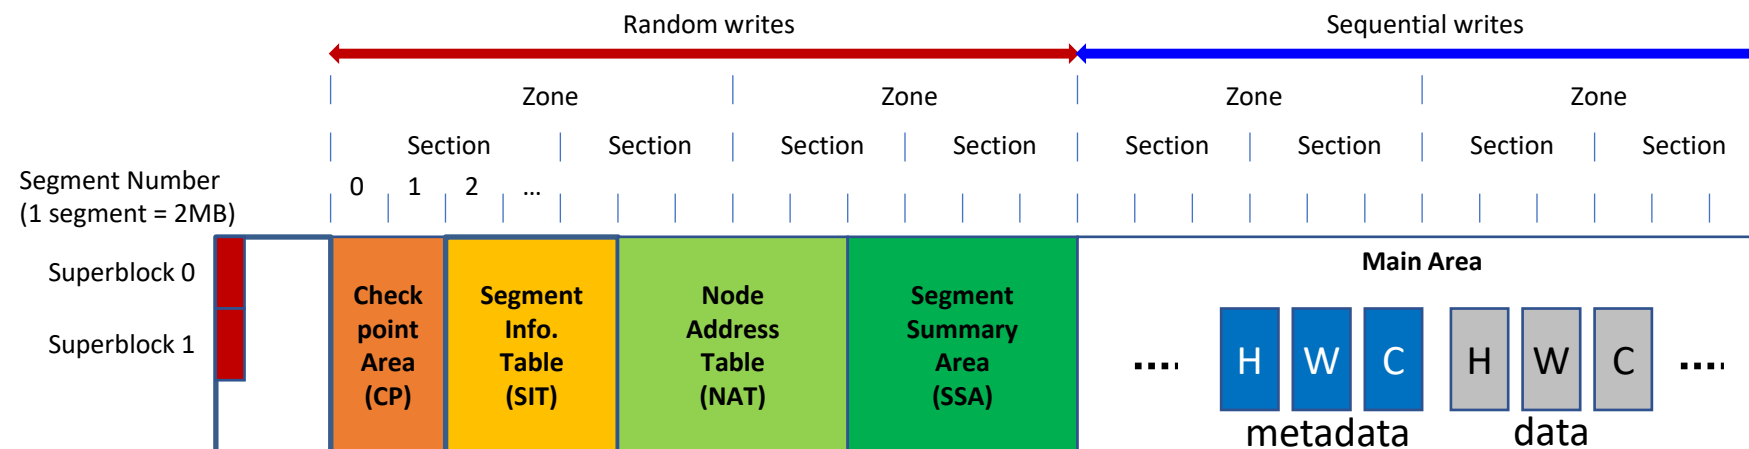
Inode

- Inline data up to 3,692 bytes

- Inline extended attributes up to 200 bytes

LFS Index Structure

Fixed location   One big log

S B

C P ✓

Inode Map ✓

Wandering tree problem
when a file data is updated, the upper index structures such as inode, inode map, and checkpoint block are also updated recursively

Dir Inode

Directory data

File Inode ✓

File data ... File data ✓

Segment Usage

Segment Summary

Used for cleaning

Indirect Pointer block ✓

Direct Pointer block ✓

Cleaning overhead
In order to serve new empty log space, it needs to reclaim these obsolete blocks seamlessly to users

F2FS Index Structure

Multi-head Logging

- **Data temperature classification**
  - Node > Data
  - Direct node > Indirect node
  - Directory > User file

- **Separation of multi-head logs in NAND flash**
  - Zone-aware log allocation for set-associative mapping FTL
  - Multi-stream interface

| Type	Temp.	Objects
Node	Hot	Direct node blocks for directories
	Warm	Direct node blocks for regular files
	Cold	Indirect node blocks
Data	Hot	Directory entry blocks
	Warm	Data blocks made by users
	Cold	Data blocks moved by cleaning; Cold data blocks specified by users; Multimedia file data

Adaptive Logging

- To reduce cleaning cost at highly aged conditions, F2FS changes write policy dynamically

- Append logging (logging to clean segments)
  - Need cleaning operations if there is no free segment
  - Cleaning causes mostly random read and sequential writes

- Threaded logging (logging to dirty segments)
  - Reuse invalid blocks in dirty segments
  - No cleaning needed
  - Cause random writes

Threaded logging writes data into invalid blocks in segment.

segment

Sudden Power Off Recovery

■ **Checkpoint and rollback**

- Maintains shadow copy of checkpoint, NAT, SIT blocks

- Recovers the latest checkpoint

- Keeps NAT/SIT journal in checkpoint to avoid NAT, SIT writes
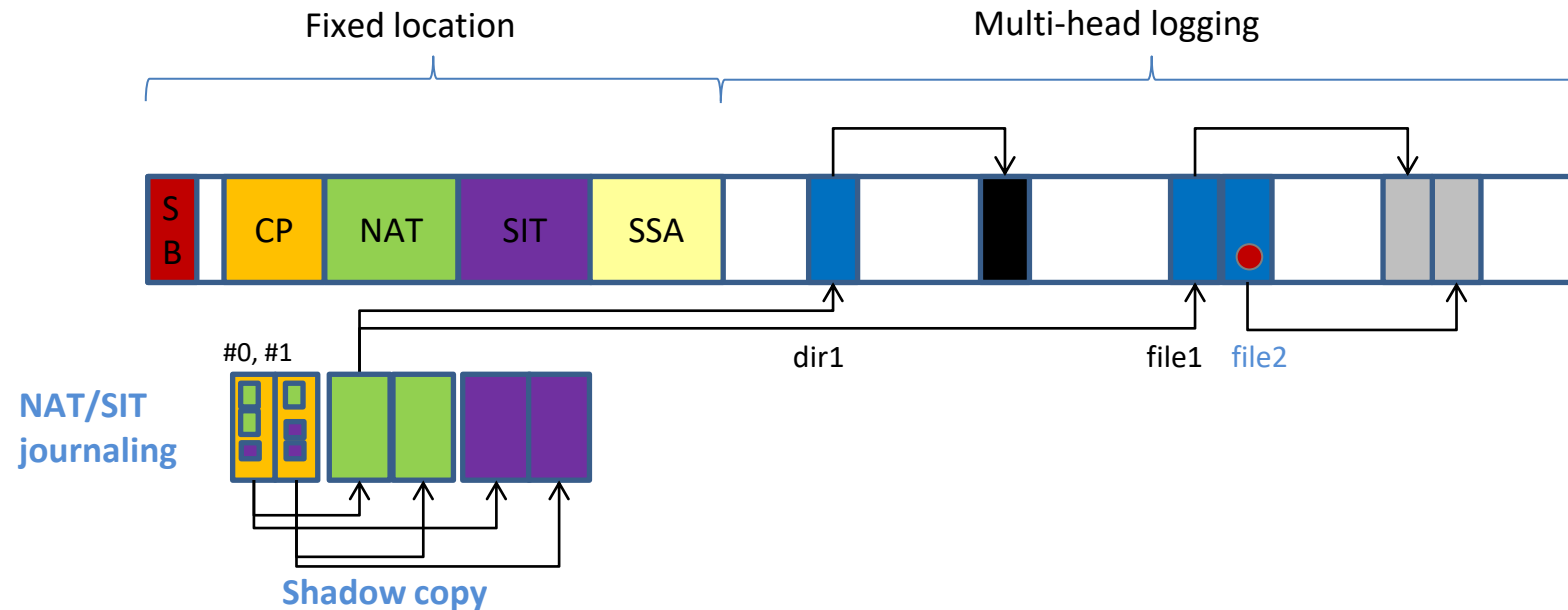
Fsync Acceleration

- **Fsync handling**
  - On fsync, checkpoint is not necessary
  - Direct node blocks are written with fsync mark

- **Roll-forward recovery procedure**
  - Search marked direct node blocks
  - Per marked node block, identify old and new data blocks by checking the difference between the current and previous node block
  - Update SIT; invalidate old data blocks
  - Replay new data block writes; update NAT, SIT accordingly
  - Create checkpoint

Mobile Benchmark

- In F2FS, more than 90% of writes are sequential

- F2FS reduces write amount per fsync by using roll-forward recovery

- Btrfs and Nilfs2 performed poor than Ext4

  - Btrfs: heavy indexing overheads
 - Nilfs2: periodic data flush

(a) iozone

(b) SQLite

Server Benchmark

- Performance gain of F2FS over Ext4 is more on SATA SSD than PCIe SSD

- Discard size matters in SATA SSD due to interface overhead
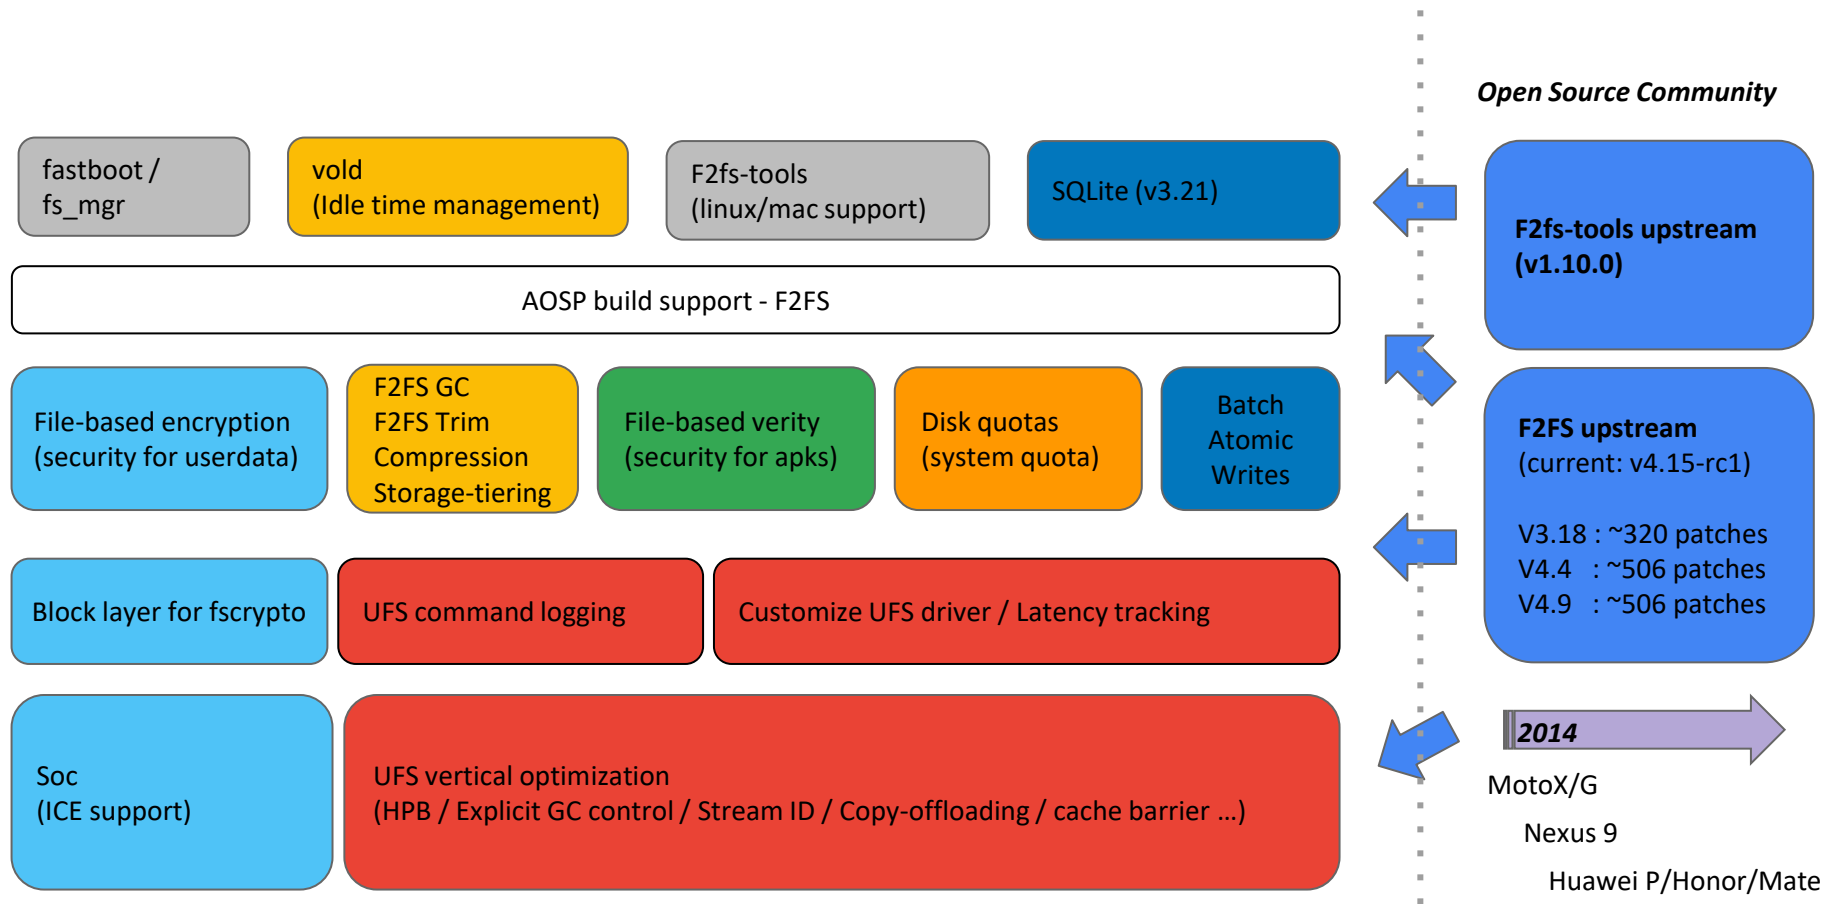

(a) SATA SSD

(b) PCIe SSD

Storage Access Pattern: F2FS ("fileserver")

Storage Access Pattern: Ext4 ("fileserver")

Features for Android

fastboot / fs_mgr	vold (Idle time management)	F2fs-tools (linux/mac support)	SQLite (v3.21)

AOSP build support - F2FS

File-based encryption (security for userdata)	F2FS GC F2FS Trim Compression Storage-tiering	File-based verity (security for apks)	Disk quotas (system quota)	Batch Atomic Writes

Block layer for fscrypto	UFS command logging	Customize UFS driver / Latency tracking

Soc (ICE support)	UFS vertical optimization (HPB / Explicit GC control / Stream ID / Copy-offloading / cache barrier …)

Open Source Community

F2fs-tools upstream (v1.10.0)

F2FS upstream (current: v4.15-rc1)

V3.18 : ~320 patches
V4.4   : ~506 patches
V4.9   : ~506 patches

2014

MotoX/G

Nexus 9

Huawei P/Honor/Mate

F2FS on Android

SAMMOBILE

NEWS   SAMSUNG   REVIEWS   PREMIUM   FIRMWARE

ONE UI 2.5   ONE UI 3.0   SECURITY UPDATE   ANDROID 11   GALAXY NOTE 20 ULTRA

Last updated: August 12th, 2019 at 15:38 UTC+01:00

Galaxy Note 10 uses F2FS, not EXT4 file system: What's the difference?

This year, a couple of Samsung smartphones will make use of a file system other than EXT4, which has been traditionally used on the Android platform almost exclusively throughout its existence. The Galaxy Note 10 series, however, switches gears and adopts the F2FS file system together with UFS 3.0 storage. Today we'll explain some of the benefits to this, and what Samsung fans can expect in the future.