

Preferred I/O for High-Capacity SSDs



the Future of Memory and Storage

Daniel Gomez

Staff Engineer, SSDR

Samsung Semiconductor

August 6, 2025

Index

01

PostgreSQL Data Journey

- A 4K World

02

PostgreSQL Data Journey In HC-SSDs

- A 4K World?

03

PostgreSQL Data Journey In HC-SSDs

- Alignment

04


PostgreSQL Data Journey In HC-SSDs With LBS

- LBS World

01 | PostgreSQL Data Journey

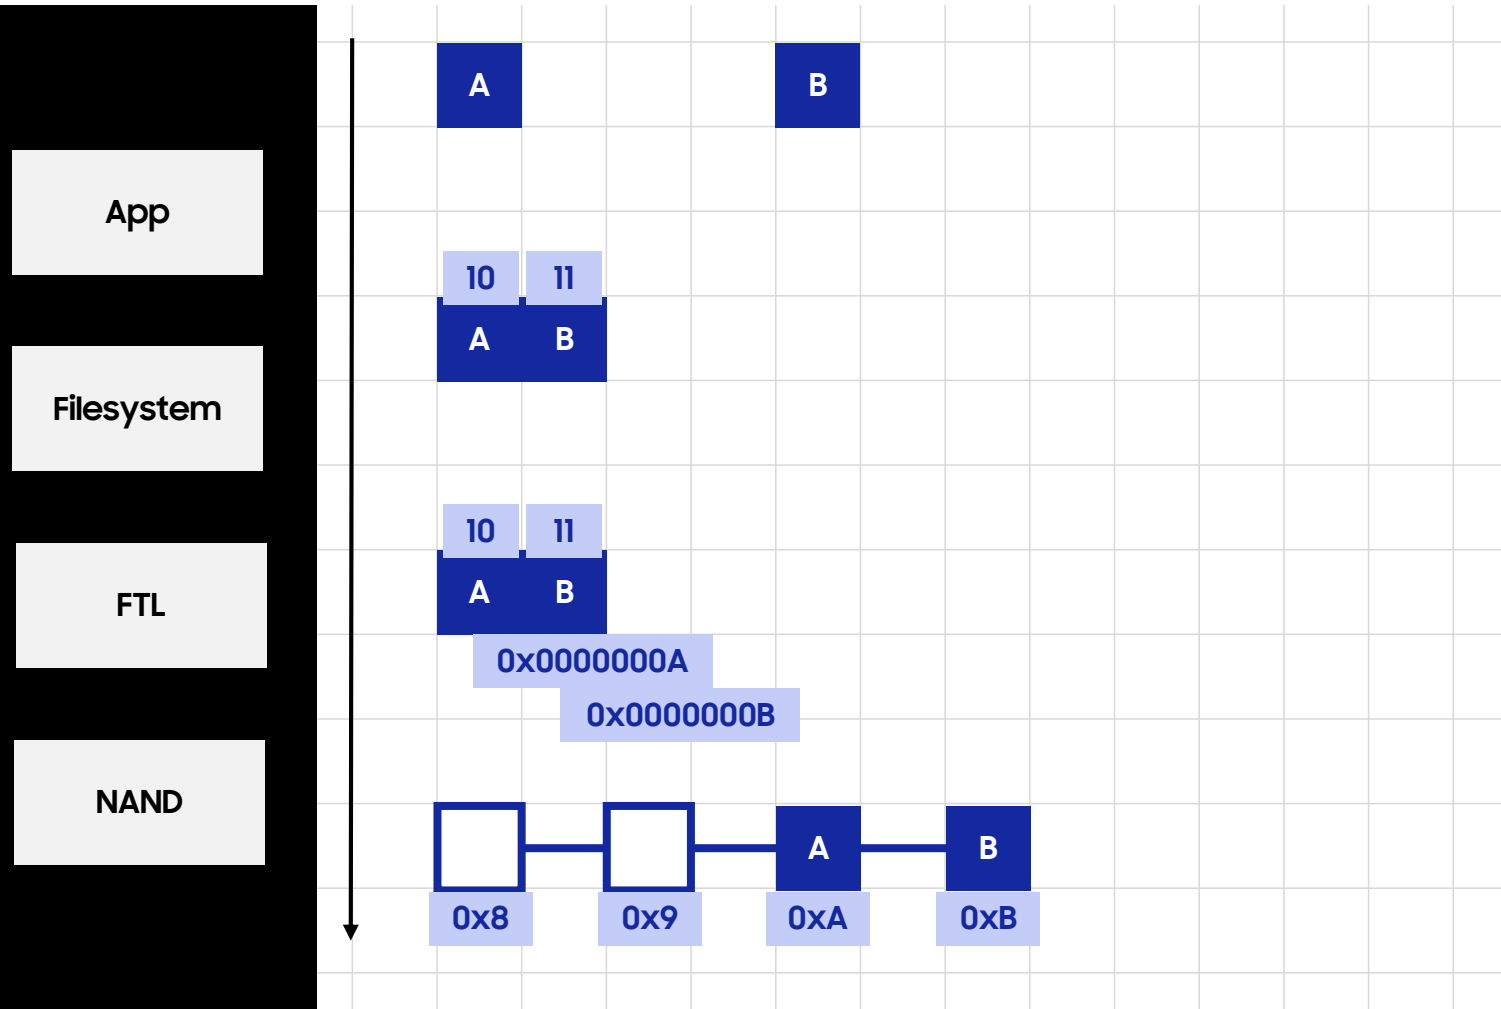
A 4K World

SAMSUNG



01 | PostgreSQL Data Journey


A 4K World



- PostgreSQL's Page Size is 8K
- LBA Format 4K
- Filesystem block size 4K
 - Block in flight: 8K

01 | PostgreSQL Data Journey

A 4K World

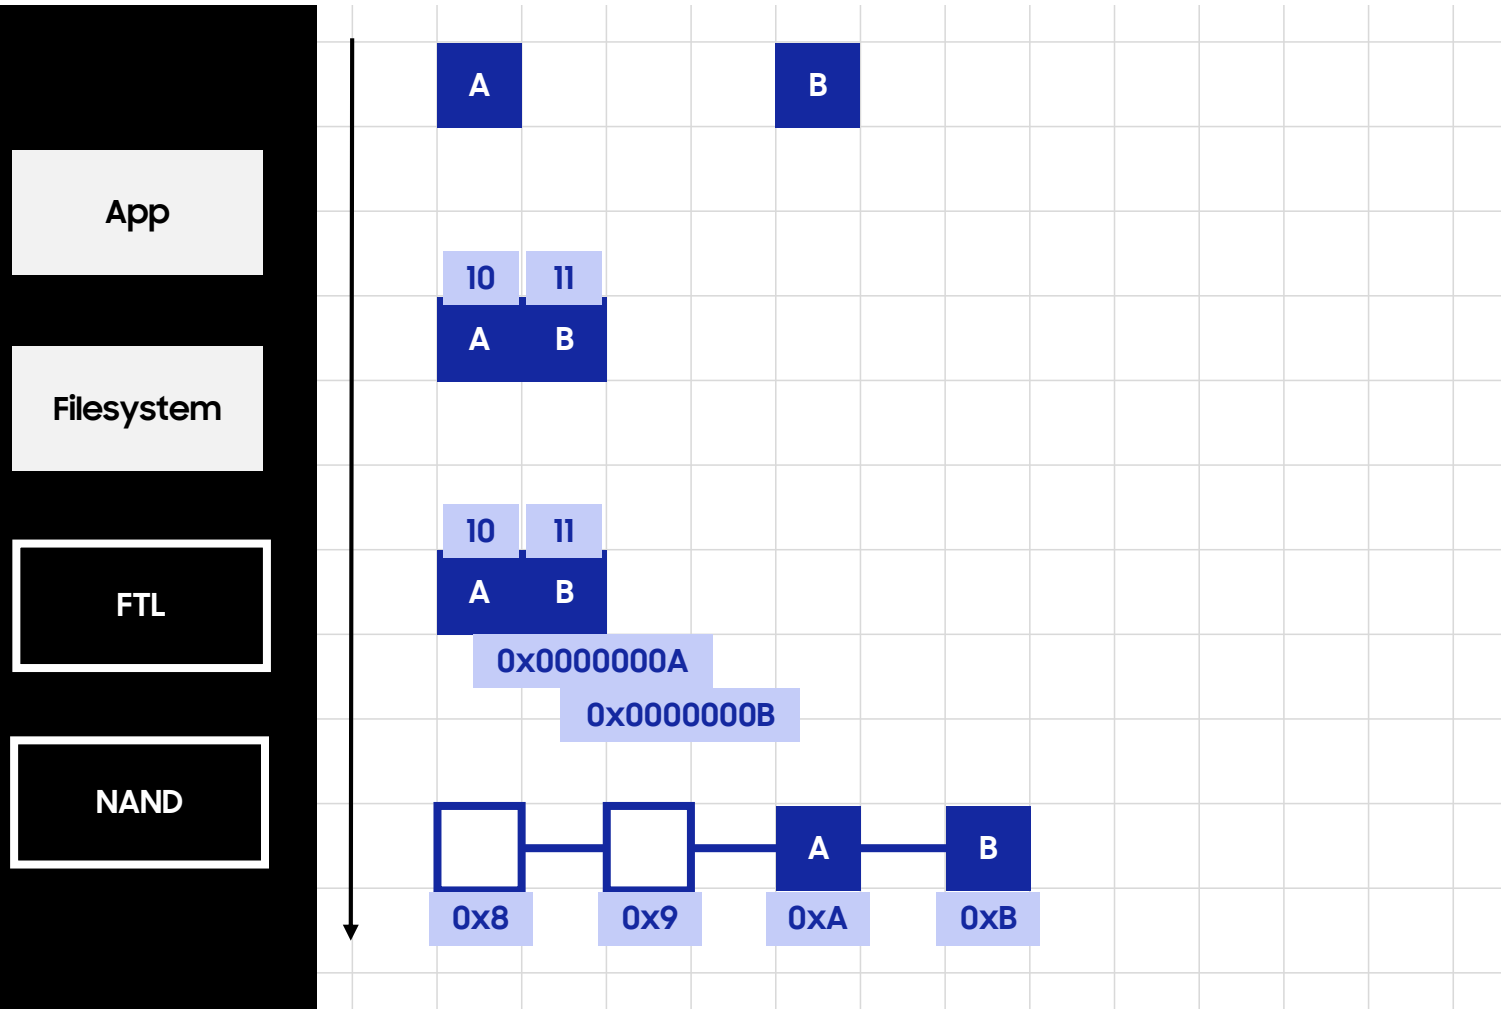


- PostgreSQL's Page Size is 8K
- LBA Format 4K
- Filesystem block size 4K
 - Block in flight: 8K

▪ 17.3 trillion → 16 TiB at 4K

01 | PostgreSQL Data Journey

A 4K World



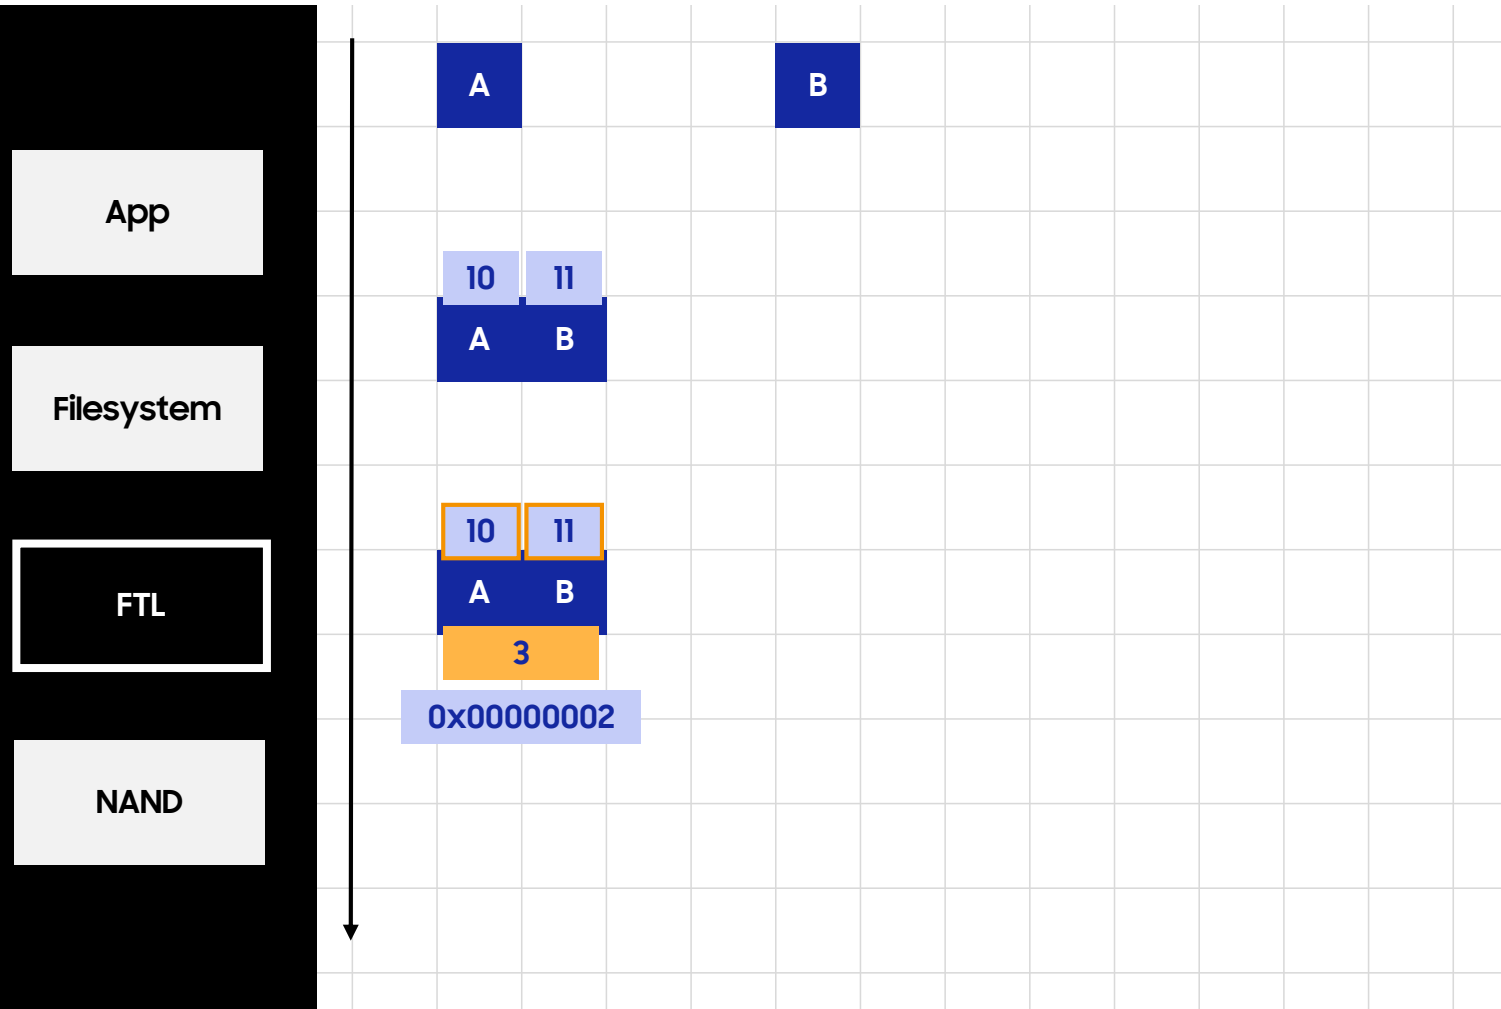
- PostgreSQL's Page Size is 8K
- LBA Format 4K
- Filesystem block size 4K
 - Block in flight: 8K
- 17.3 trillion → 16 TiB at 4K
- 17.3 trillion → 64 TiB at 16K
- Indirection Unit (IU)

High-Capacity SSDs Data Journey

02 | PostgreSQL Data Journey in HC-SSDs



A 4K World?

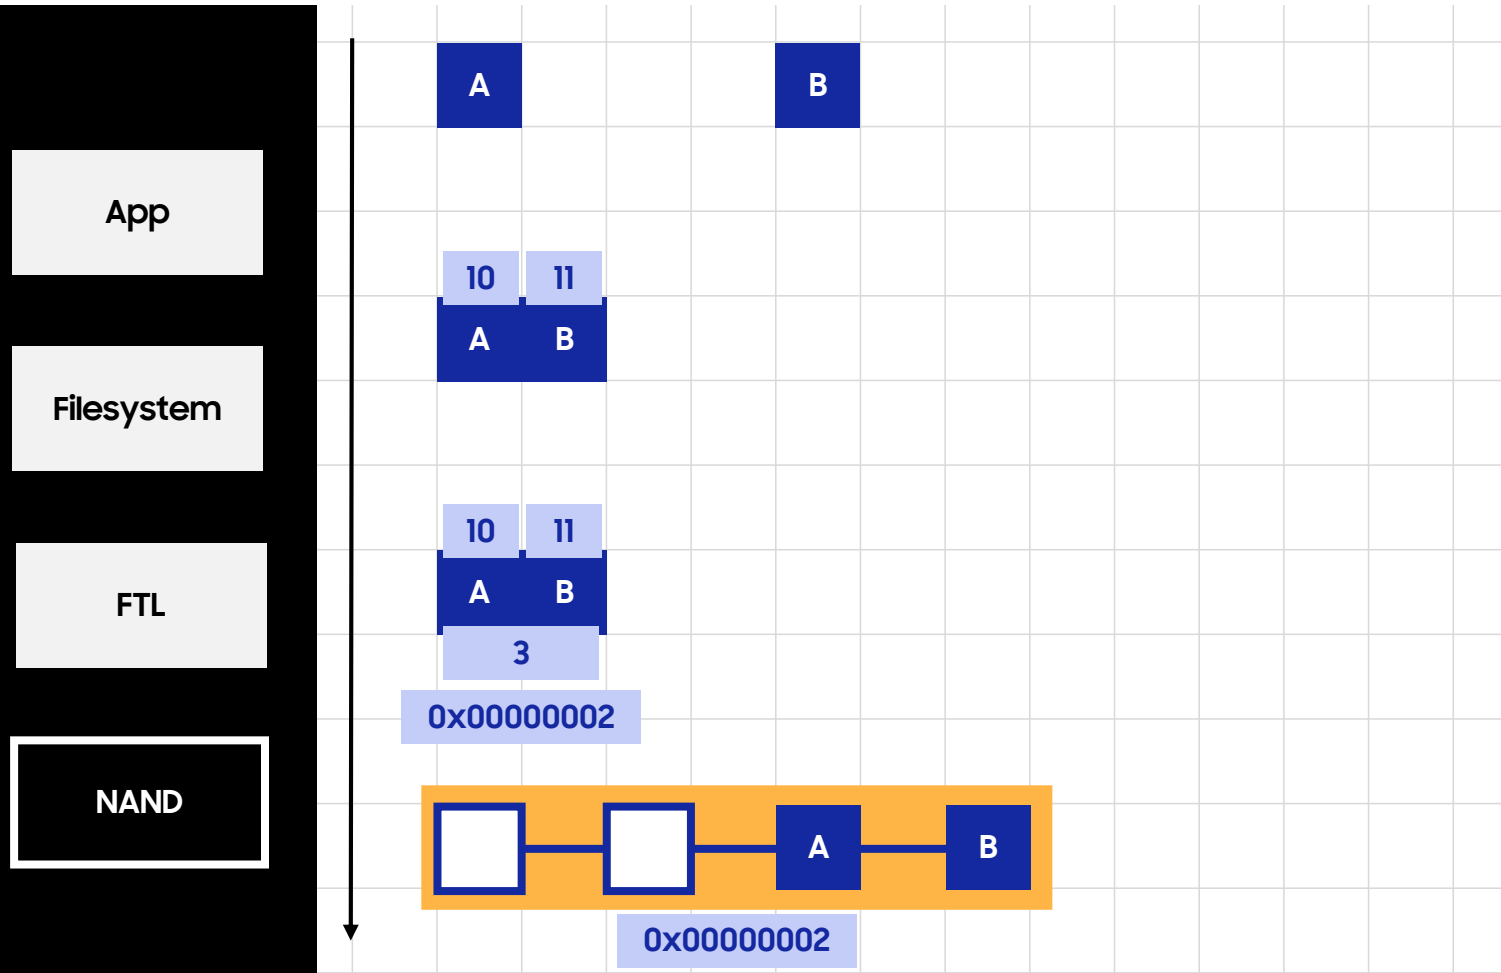


- ~~4K world~~
- PostgreSQL's Page Size is 8K
- LBA Format 4K
- NPWG = IU = 16K
 - NPWA

02 | PostgreSQL Data Journey in HC-SSDs



A 4K World?



- 4K world
- PostgreSQL's Page Size is 8K
- LBA Format 4K
- NPWG = IU = 16K

▪ RMW

▪ Device WAF ↑ ↑

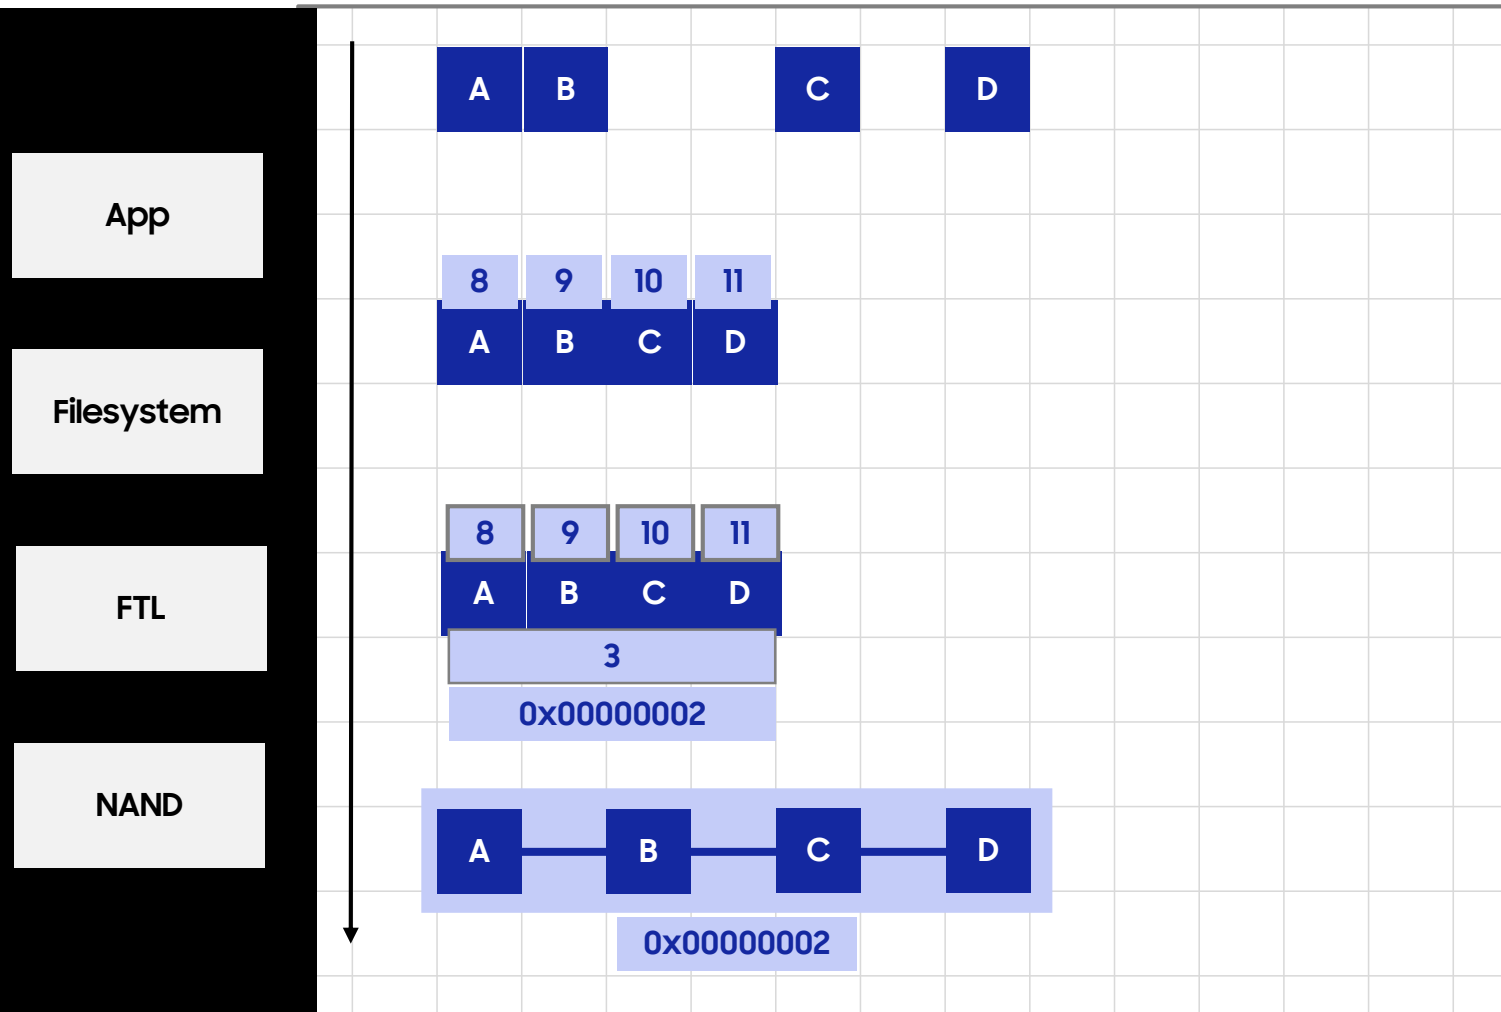
▪ Performance ↓ ↓

**What if I send
16K?**

03 | PostgreSQL Data Journey in HC-SSDs



Alignment



▪ 4K world

▪ PostgreSQL's Page Size is 16K

▪ LBA Format 4K

▪ NPWG = IU = 16K

▪ RMW?

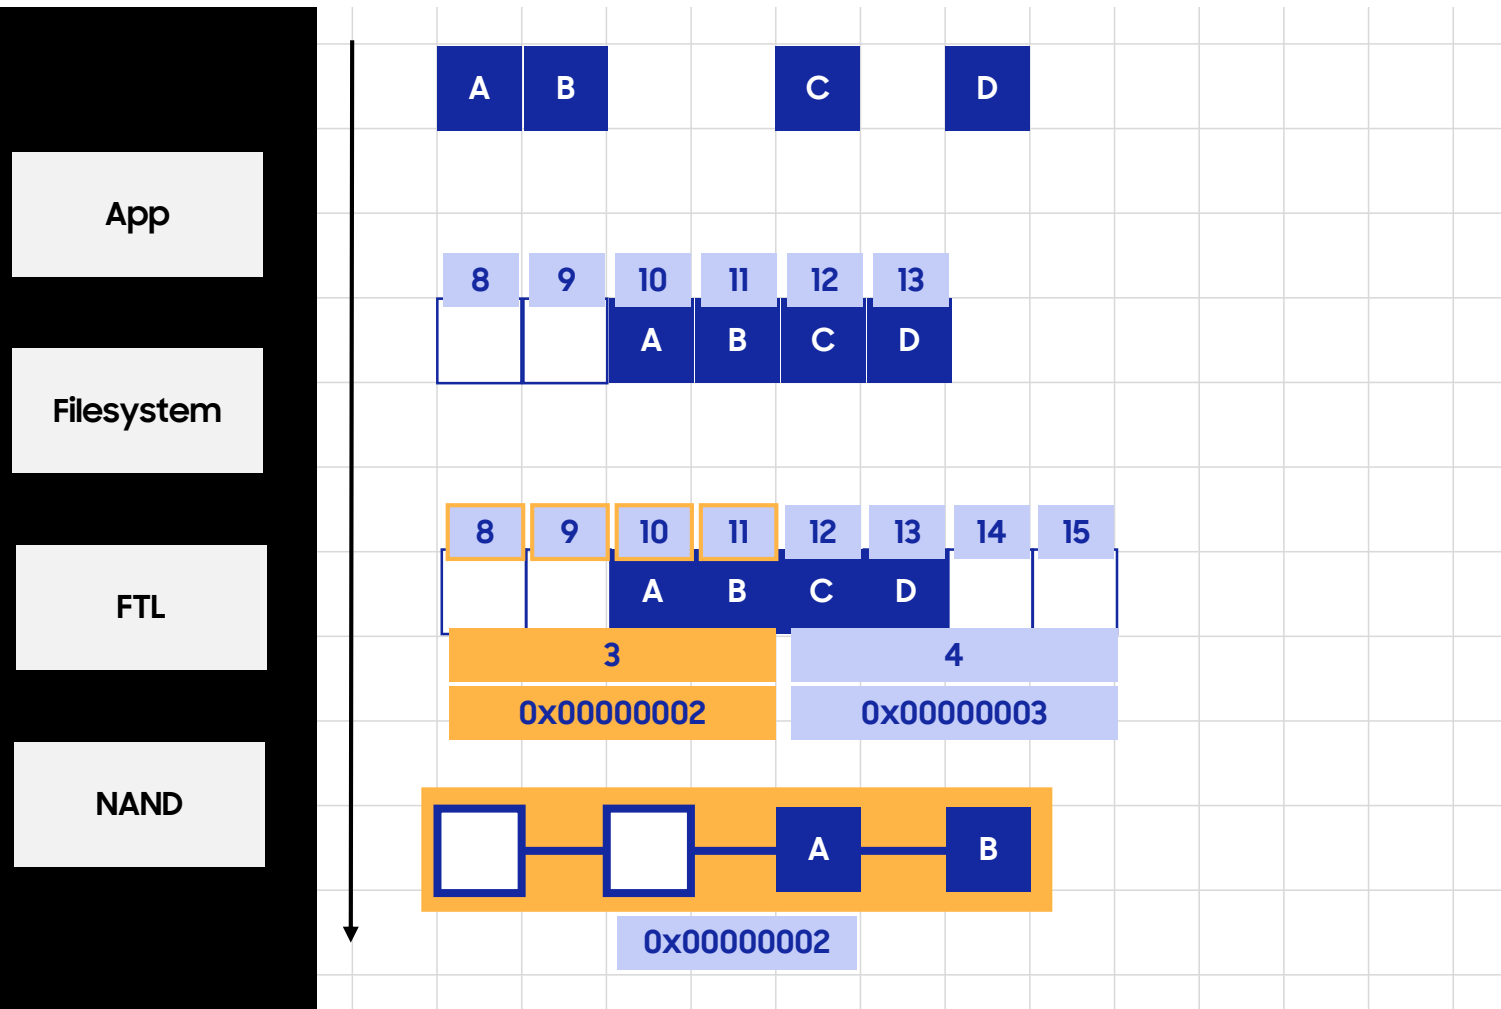
▪ Device WAF?

▪ Performance?

03 | PostgreSQL Data Journey in HC-SSDs



Alignment



▪ 4K world

▪ PostgreSQL's Page Size is 16K

▪ LBA Format 4K

▪ NPWG = IU = 16K

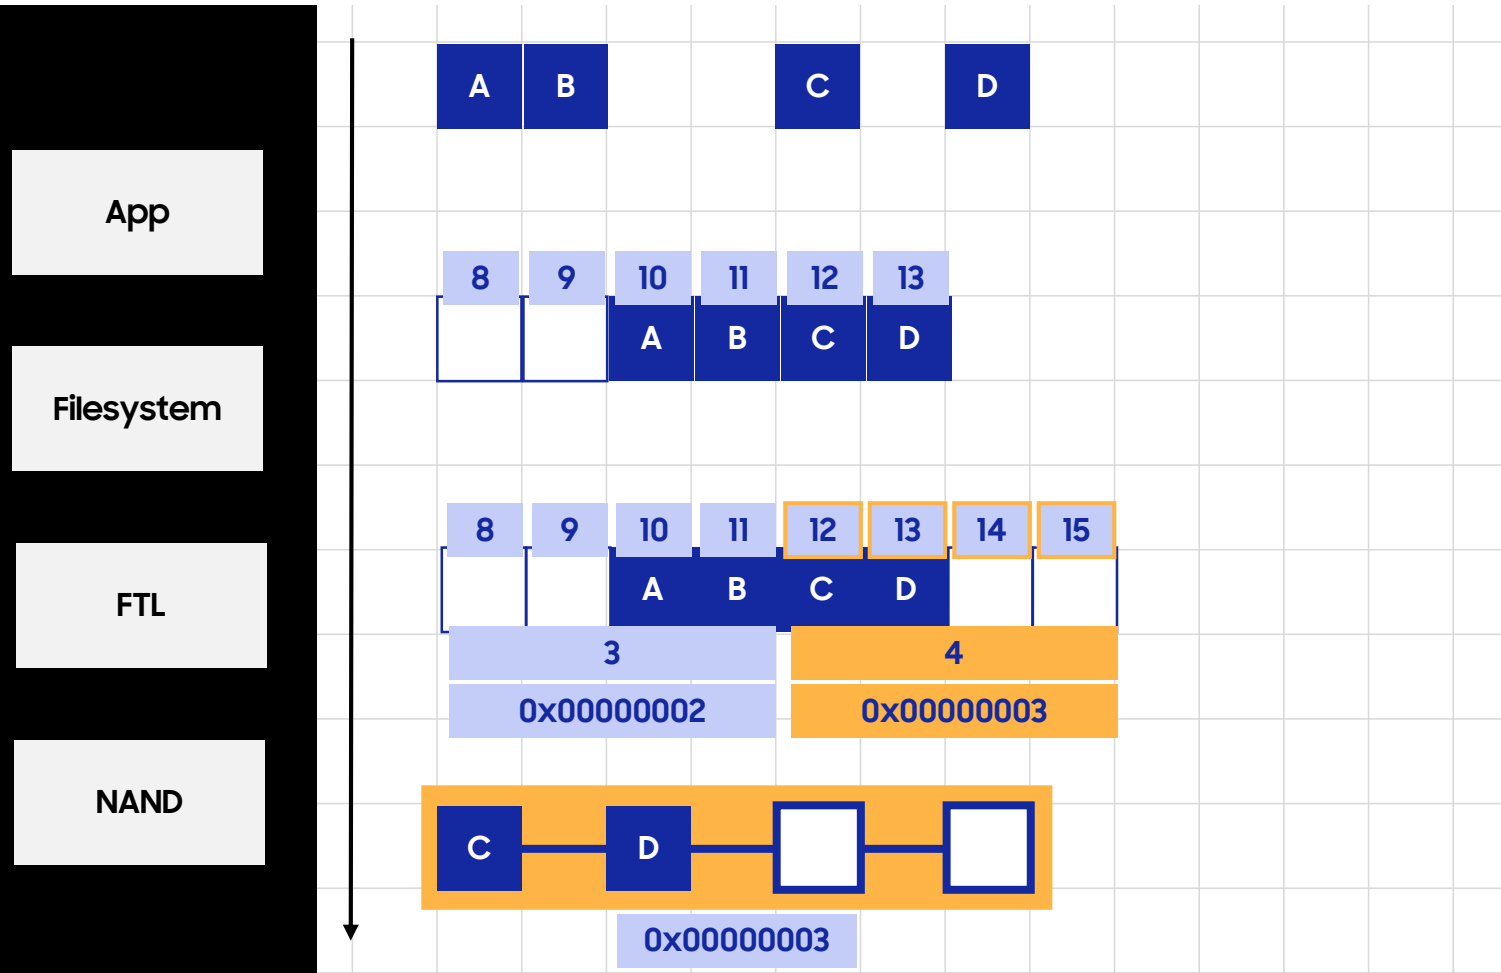
▪ RMW

▪ Device WAF ↑ ↑

▪ Performance ↓ ↓

03 | PostgreSQL Data Journey in HC-SSDs

Alignment



▪ 4K world

▪ PostgreSQL's Page Size is 16K

▪ LBA Format 4K

▪ NPWG = IU = 16K

▪ RMW

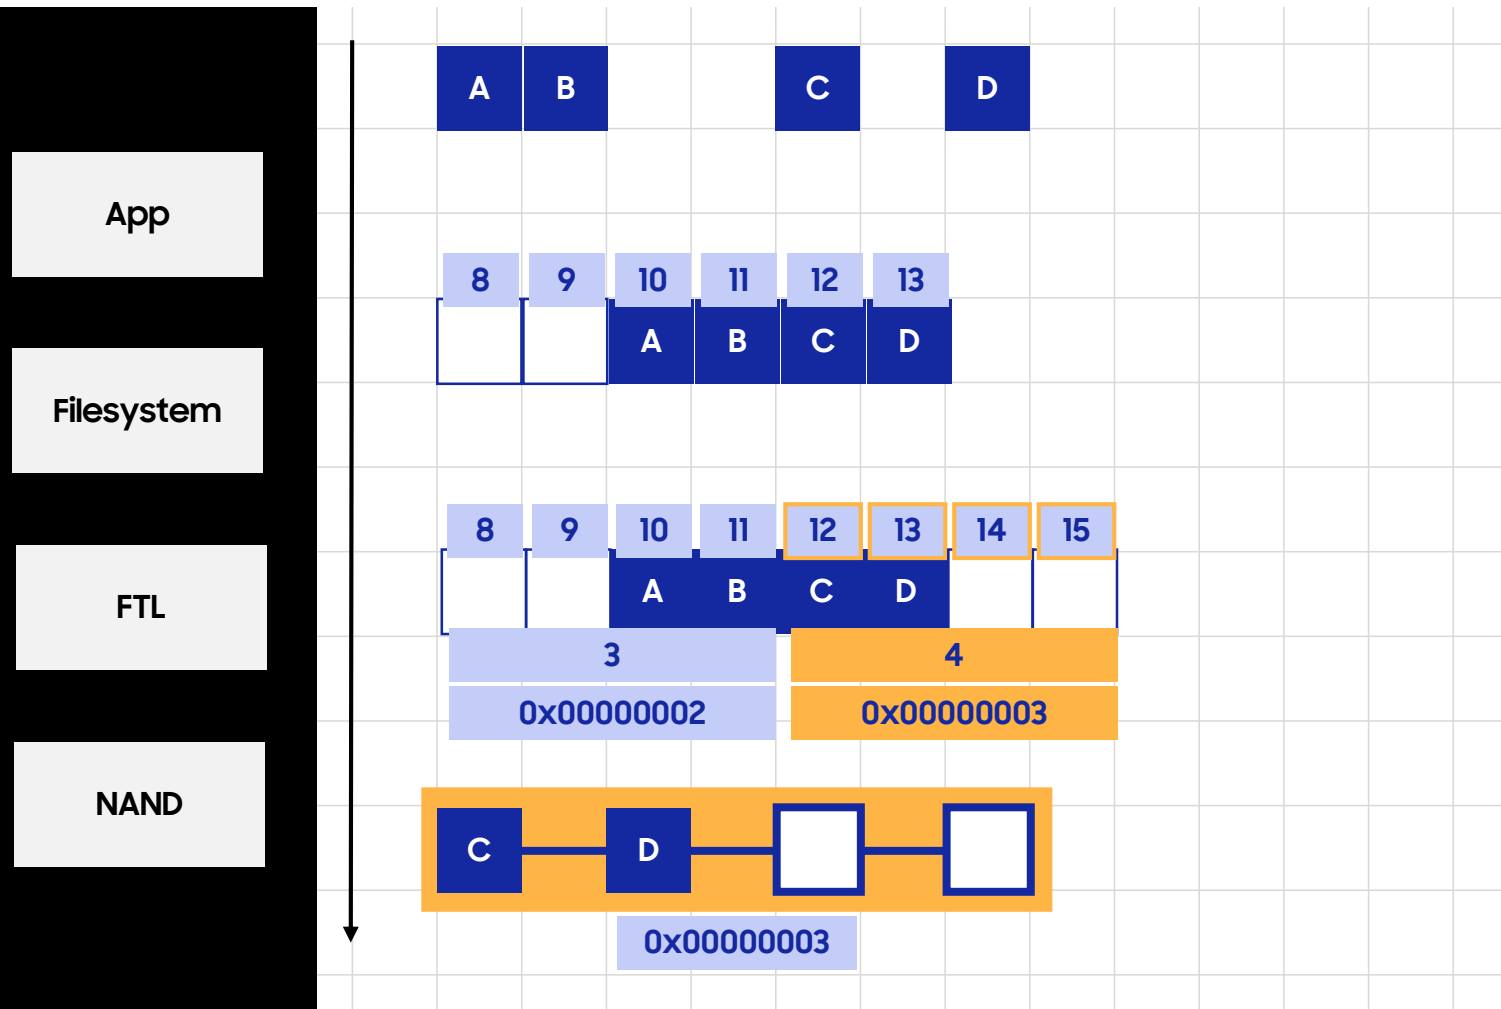
▪ Device WAF ↑ ↑

▪ Performance ↓ ↓

03 | PostgreSQL Data Journey in HC-SSDs



Alignment



- 4K world

- PostgreSQL's Page Size is 16K

- LBA Format 4K

- NPWG = IU = 16K

- RMW

- Device WAF

- Performance

- Alignment

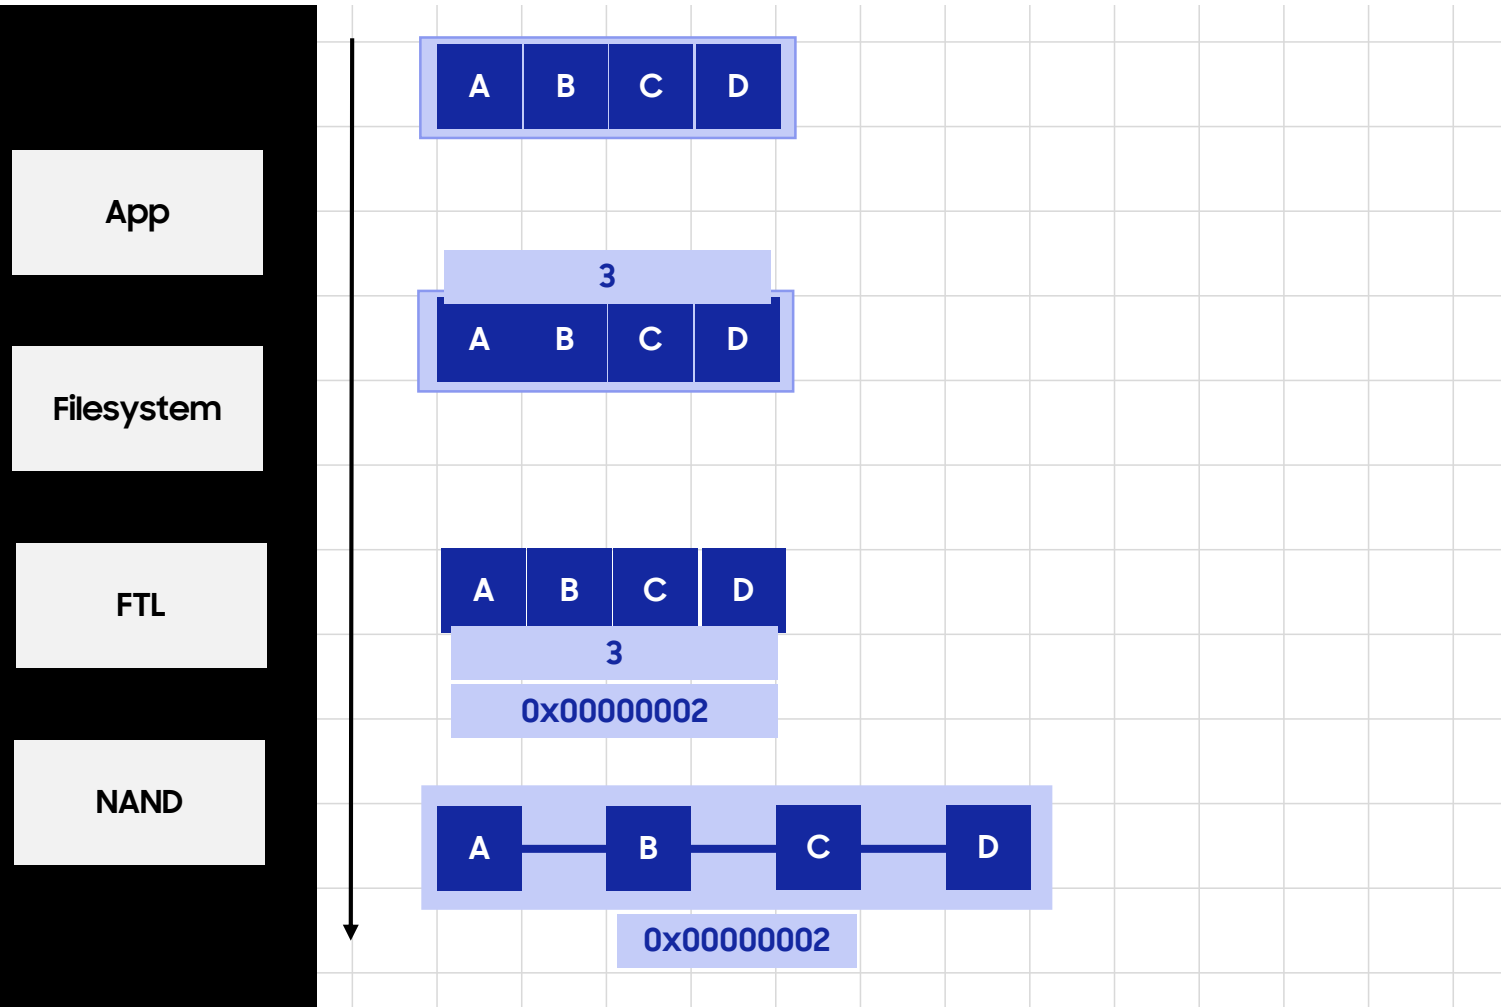
LBS

Large Block Sizes

04 | PostgreSQL Data Journey in HC-SSDs



LBS World



▪ ~~4K world~~

▪ PostgreSQL's Page Size is 16K

▪ LBA Format 4K

▪ NPWG = IU = 16K

▪ ~~RMW~~

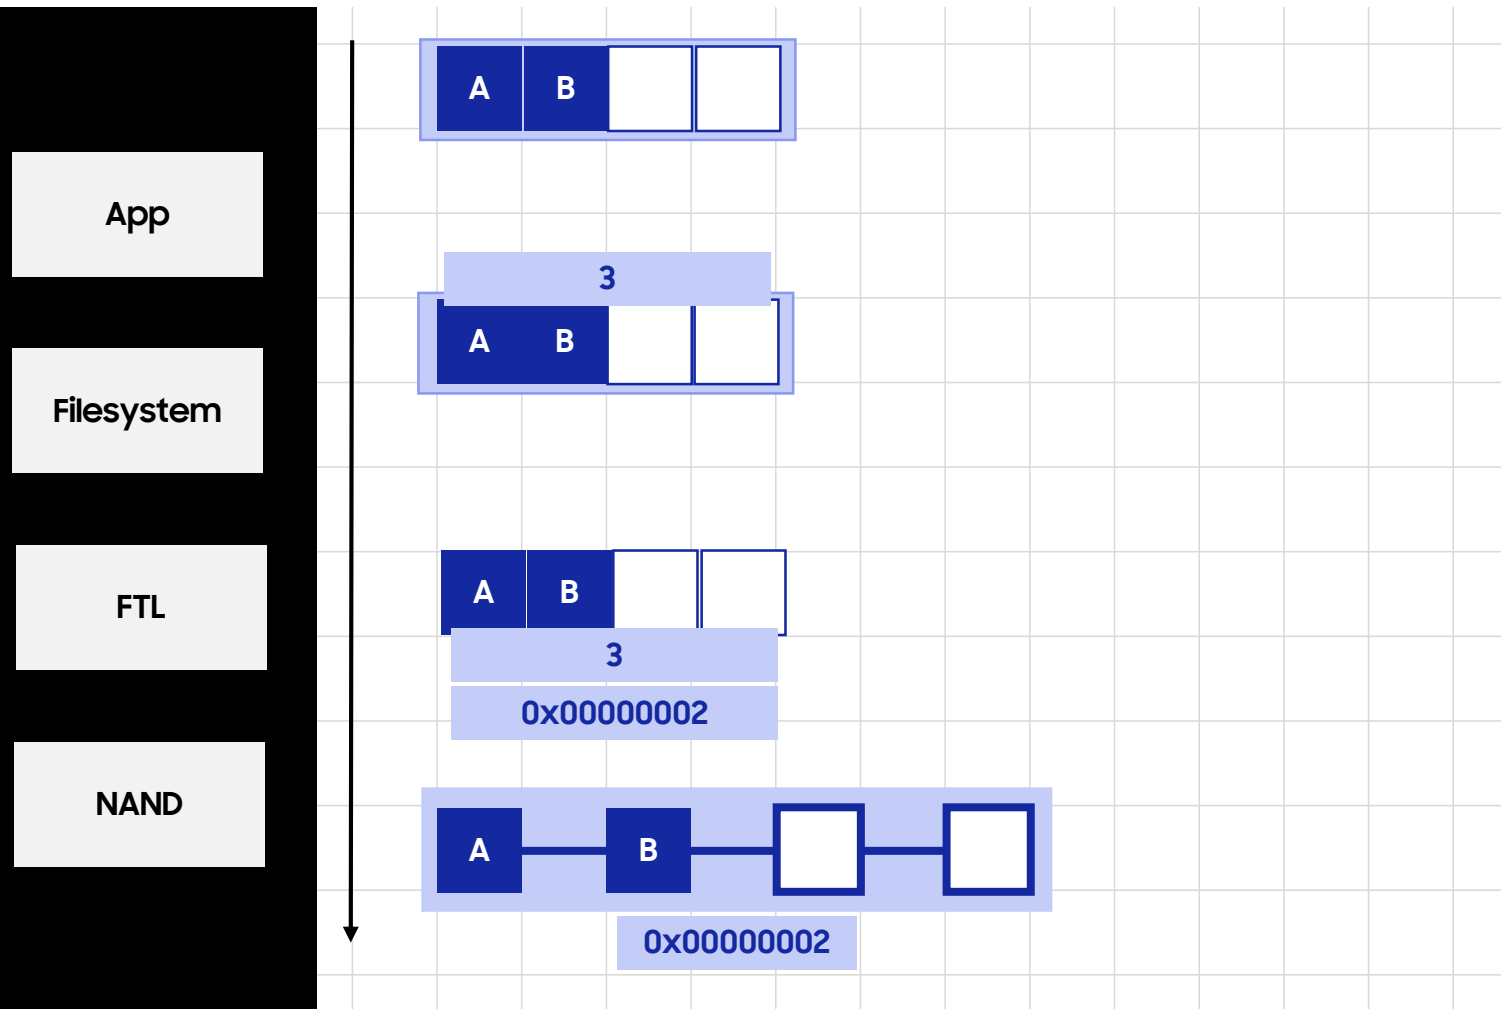
▪ ~~Device WAF~~

▪ Performance ↑↑

▪ Alignment

04 | PostgreSQL Data Journey in HC-SSDs

LBS World



▪ ~~4K world~~

▪ PostgreSQL's Page Size is 8K

▪ LBA Format 4K

▪ NPWG = IU = 16K

▪ ~~RMW~~

▪ ~~Device WAF~~

▪ Performance ↑↑

▪ Alignment

▪ Application WAF ↑↑

05 | Summary



- Indirection Unit (IU)
 - 512-byte (90s) → 4K (98) → >4K (mid-2020s)
- NPWG (Namespace Preferred Write Granularity)
 - NPWA (Namespace Preferred Write Alignment)
- LBS it's ready today
 - Since Linux v6.12 (December 2024)
 - 3-4 year team effort
 - Community effort of +17 years
- LBS Gains
 - Deterministic alignment
 - Increases performance
 - Removes application induced device WAF

**High Capacity SSDs
with 16K IU
up to 256 TiB**

Thank you