



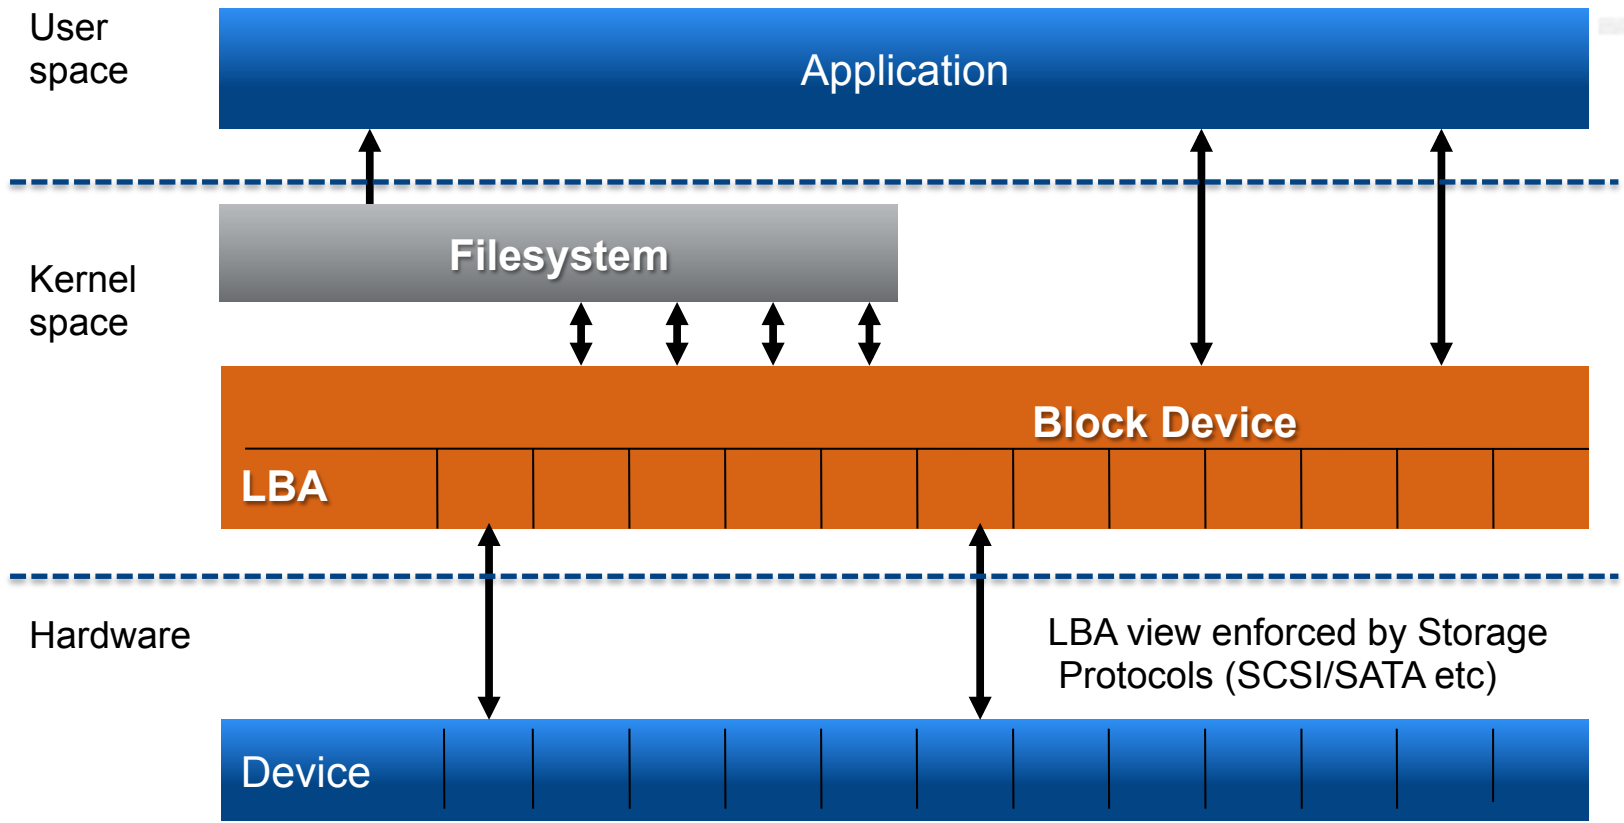
# **NATIVE FLASH SUPPORT FOR APPLICATIONS**

Nisha Talagala



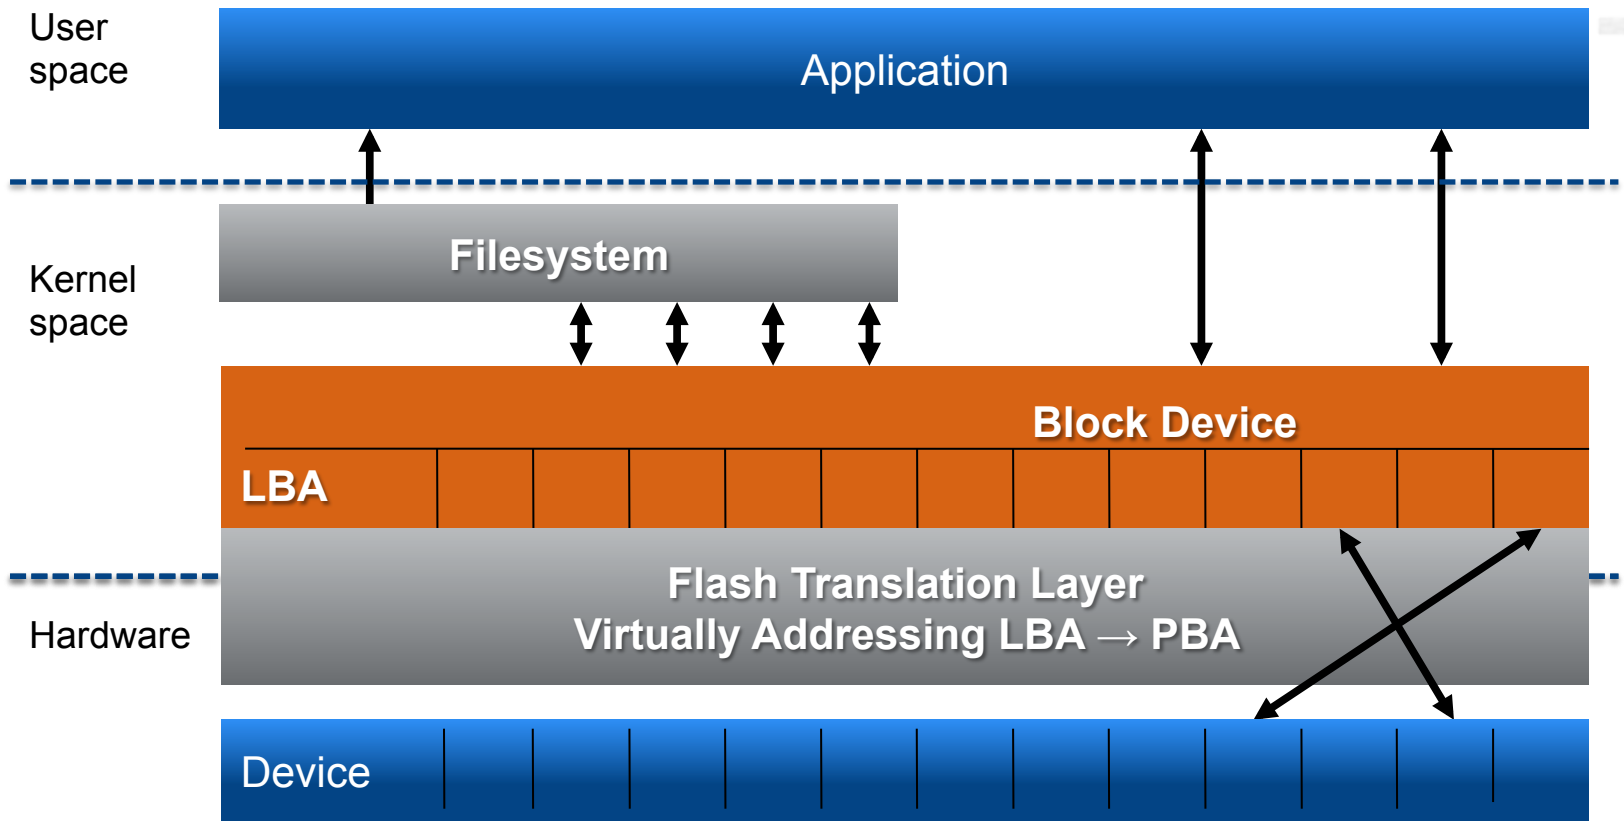
# TRADITIONAL STORAGE STACK

FUSION-io



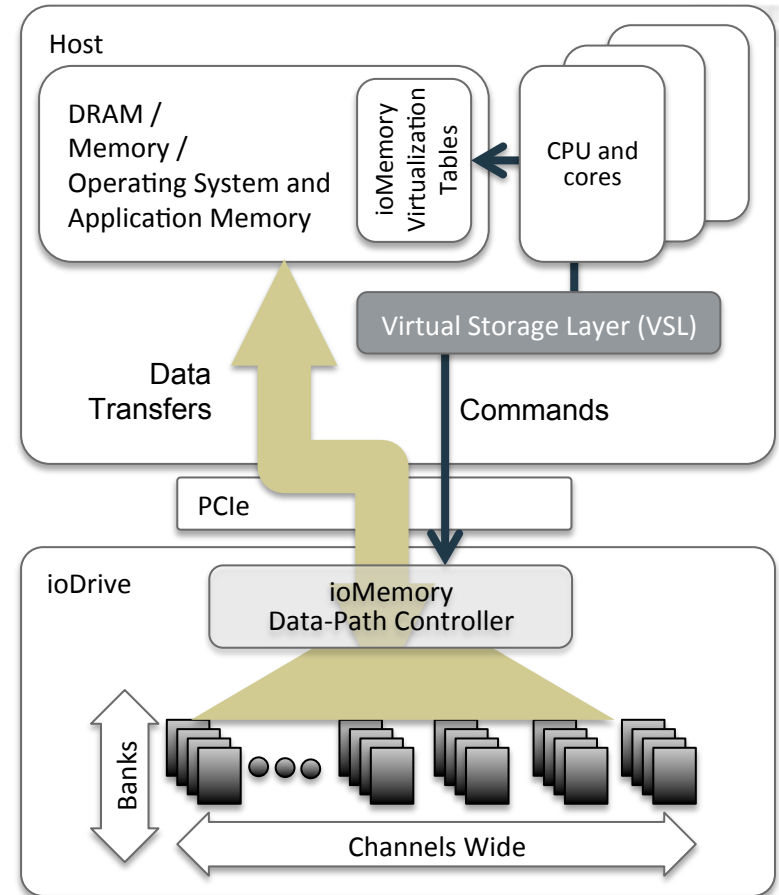


# FLASH IN TRADITIONAL STORAGE STACKS



# VIRTUAL STORAGE LAYER

- ▶ Cut-thru architecture – avoids traditional storage protocols
- ▶ Scales with multi-core
- ▶ HW/SW functional boundary defined as optimal for flash
- ▶ Traditional block access methods for compatibility
- ▶ New access methods, functionality and primitives natively supported by ioMemory





# FLASH MEMORY EVOLUTION

FUSION-io®

FUSION-io®

Native Access



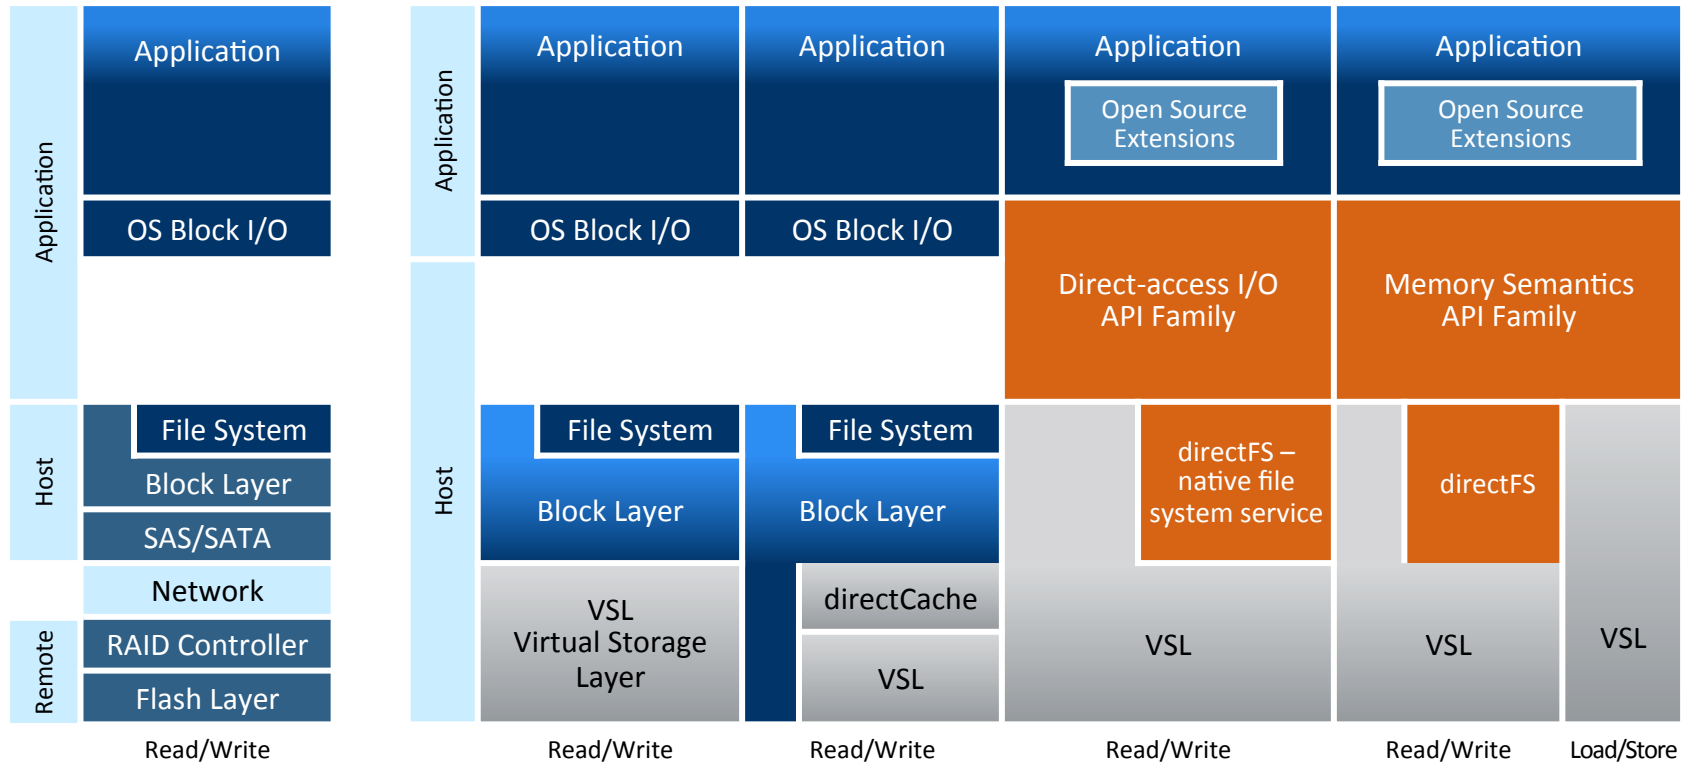
Legacy SSDs

ioMemory as ioDrive

ioMemory as Transparent Cache

ioMemory with direct access I/O

ioMemory with memory semantics





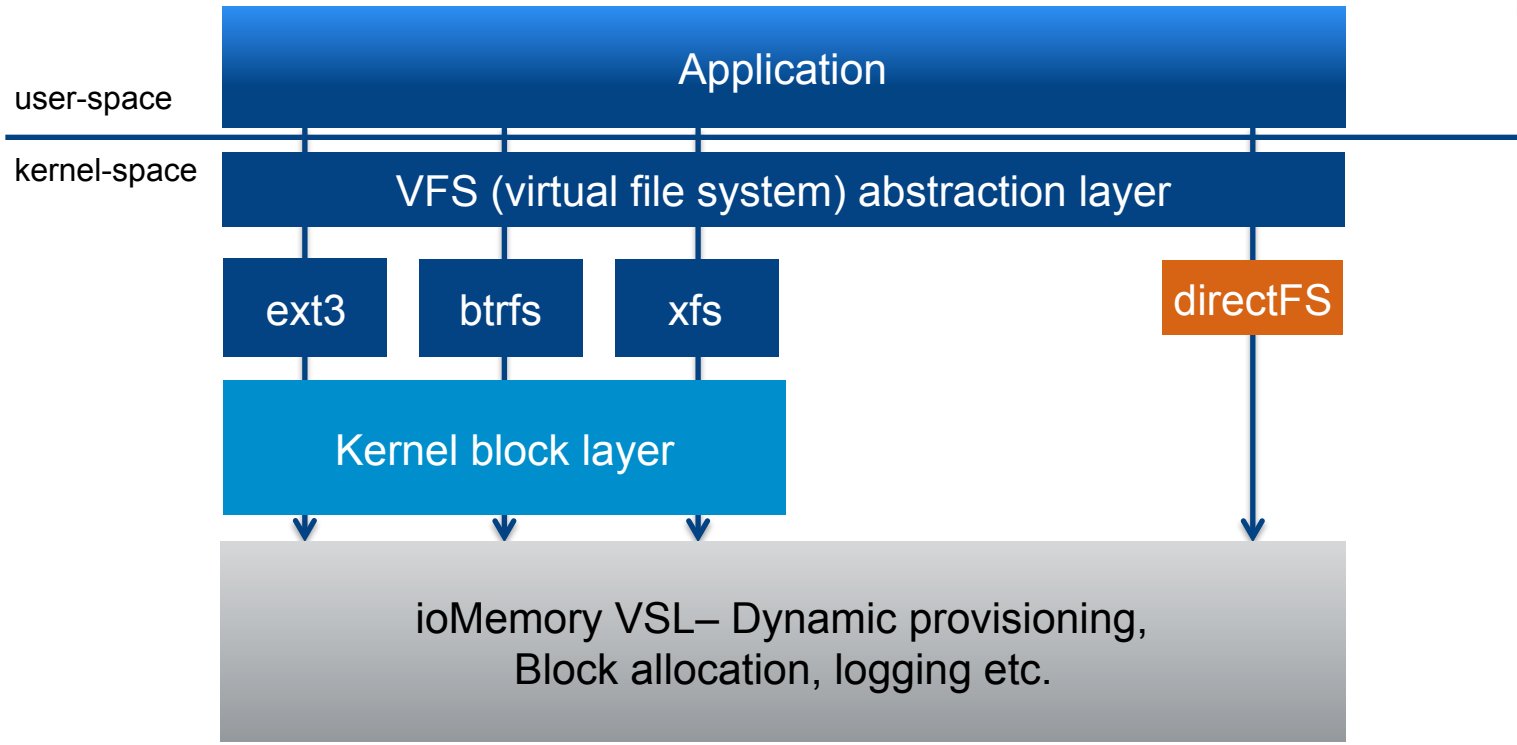
# EXPLOITING NATIVE CHARACTERISTICS OF IOMEMORY

1. Native log-append writes  
incorporates **copy-on-write** basics
2. Native block mapping and allocation  
incorporate **file system** basics
3. Native large virtual address space  
incorporates **sparse semantics**
4. Native storage methods  
incorporate **key-value store** basics



# DIRECTFS – NATIVE FILE NAMESPACE FOR FLASH

FUSION-io





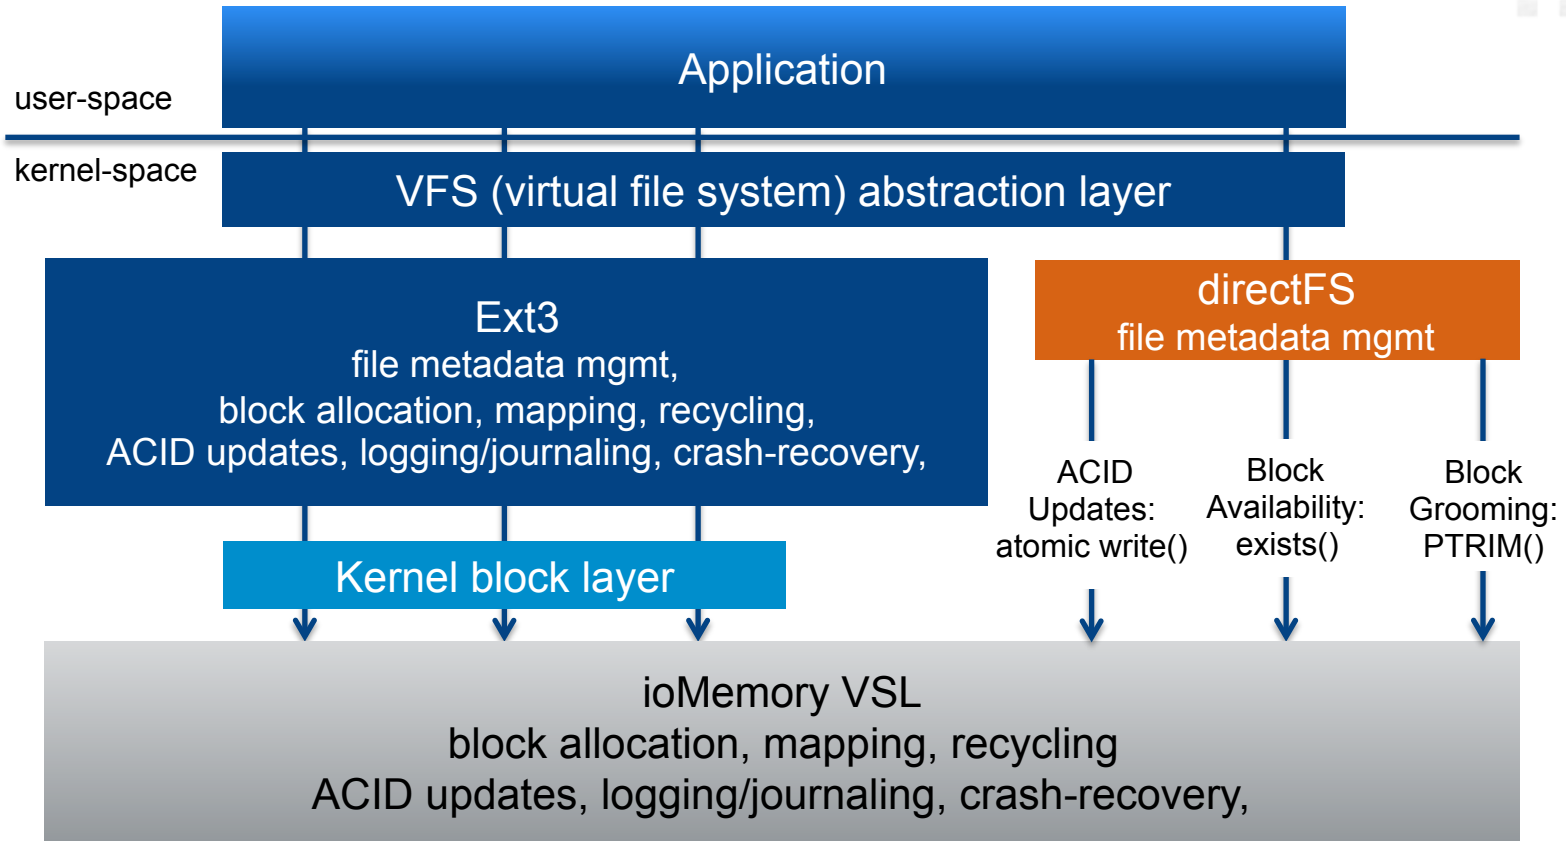
# DIRECTFS

- ▶ Appears as any other file system in Linux
- ▶ Applications can use directFS filesystem unmodified with performance benefits
- ▶ Focuses only on file namespace
- ▶ Employ virtualized flash storage layer's logic for:
  - Large virtualized addressed space
  - Direct flash access
  - Crash recovery mechanisms
- ▶ Exposes VSL Primitives through file namespace
  - Applications can use SDK through directFS or directly to VSL





# DIRECTFS – ELIMINATING DUPLICATE LOGIC





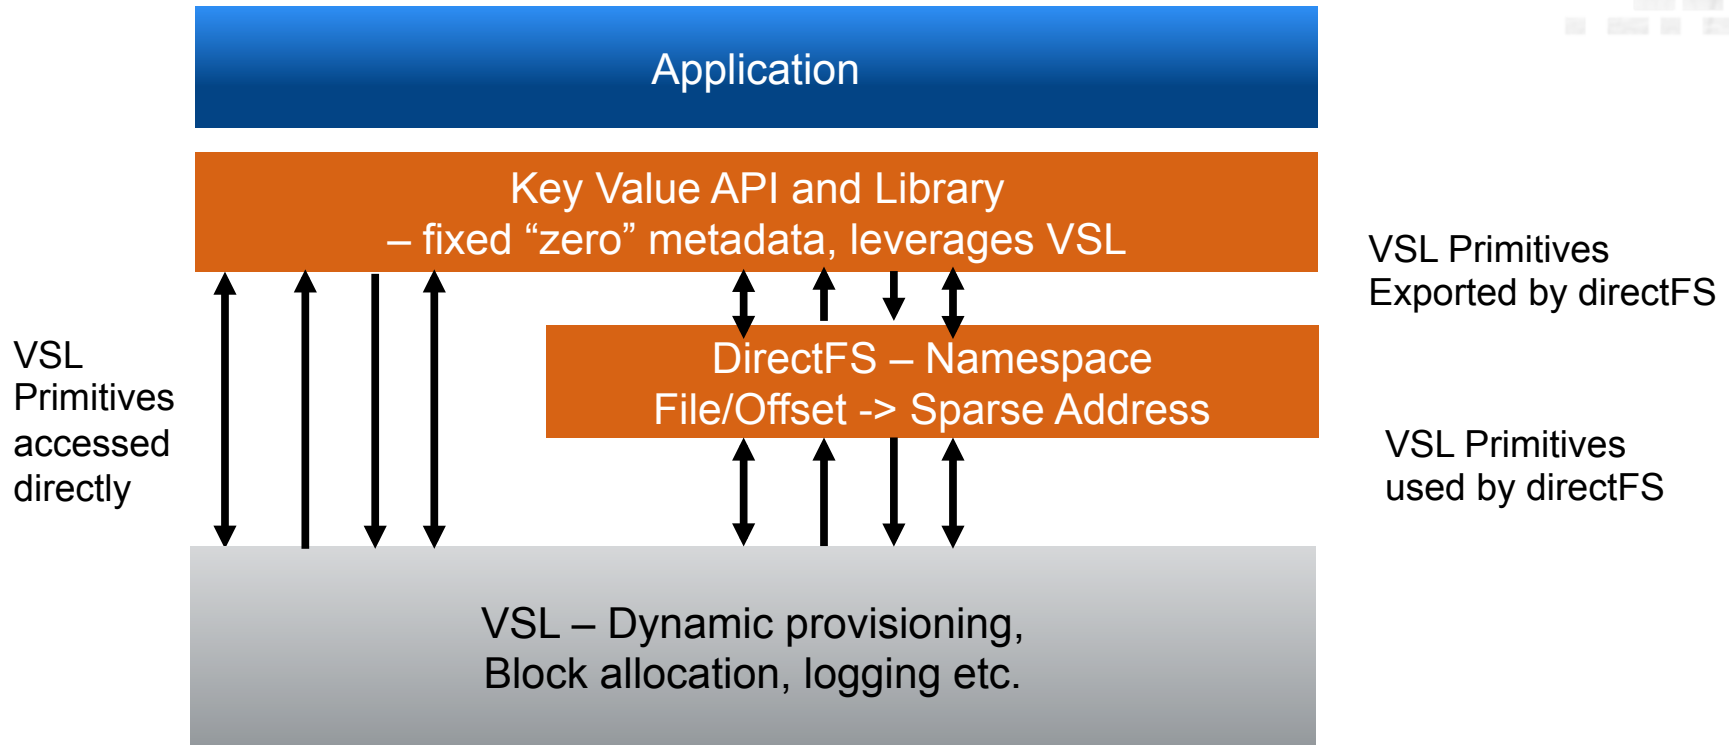
# DIRECTFS SIMPLICITY – LINES OF CODE

FUSION-io

File System	Lines of Code
DFS	6879
ReiserFS	19996
Ext4	25837
Btrfs	51925
XFS	63230



# DIRECTFS – NATIVE FILE NAMESPACE FOR SDK API LIBRARIES

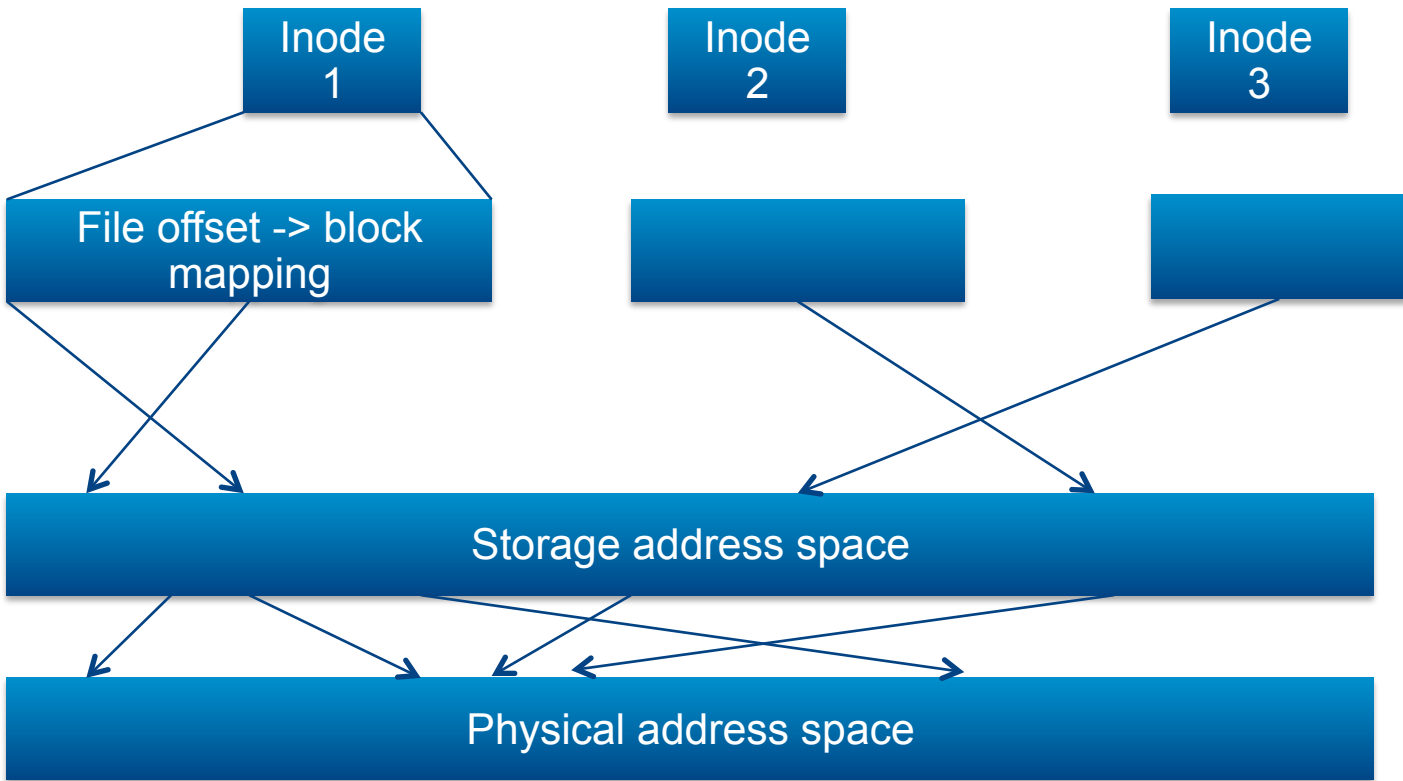


DirectFS exports VSL Primitives



# DIRECTFS – SIMPLIFIED BLOCK MAPPING THROUGH NATIVE SPARSE VIRTUAL ADDRESSING

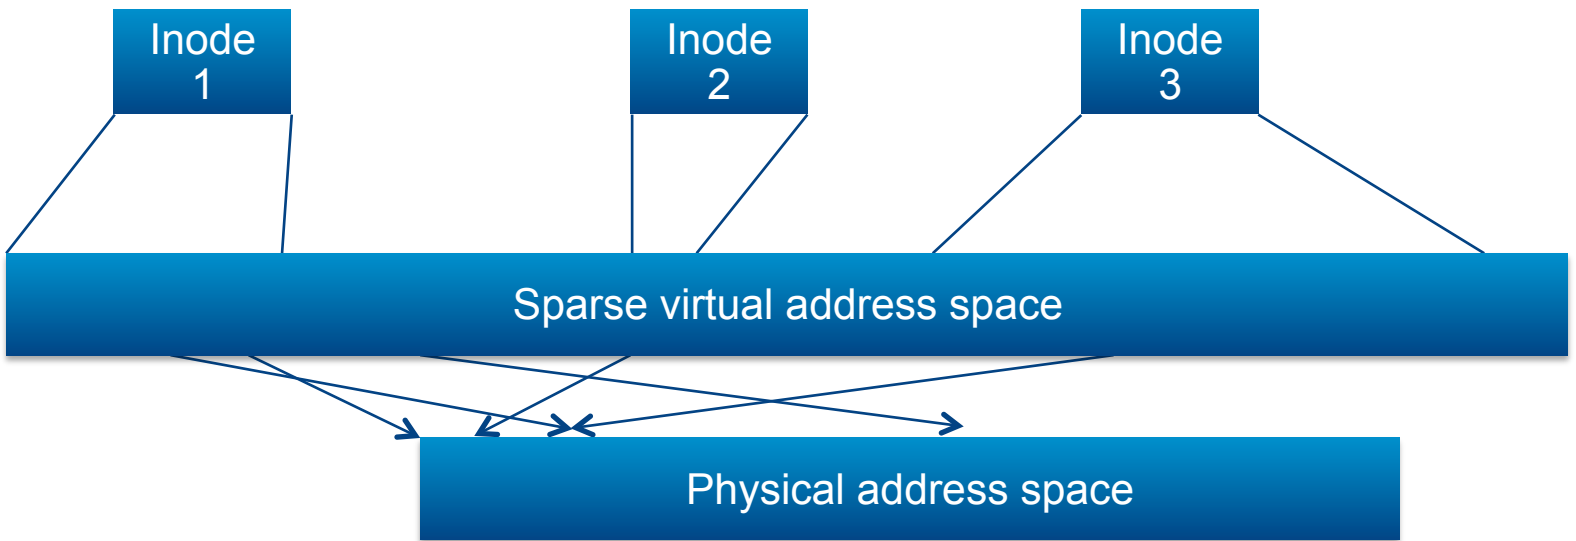
Conventional file systems on block storage





# DIRECTFS – SIMPLIFIED BLOCK MAPPING THROUGH NATIVE SPARSE VIRTUAL ADDRESSING

- ▶ DirectFS – files mapped directly to sparse address space
- ▶ Removes mapping layers
- ▶ Allows file aware NVM optimizations and features





# DIRECTFS – EXAMPLE: FALLOCATE

- ▶ `Fallocate()` – preallocate large files
  
- ▶ In conventional file systems
  - Write zeros for a large file to the physical media
  
- ▶ `directFS`
  - Pre-allocate virtual address space
  - Minimal writing to device – only to commit virtual address allocation
  - No need to write physical zeros to the entire file



# DIRECTFS – CRASH SAFE THROUGH ATOMIC DATA AND METADATA UPDATE OPERATIONS

- ▶ Conventional file systems
  - Journal for power cut safety
  - Allows replay of transactions after crash
  
- ▶ directFS
  - Leverages atomics features of VSL
  - Can implement crash safety without journaling



# DIRECTFS – INITIAL PROTOTYPE PERFORMANCE FOR UNMODIFIED APPS

FUSION-io

Application	Ext3	DFS Prototype	Speedup
Quick Sort	1268	822	1.54
N-Gram	4718	1912	2.47
KNNImpute	303	248	1.22
VM Update	685	640	1.07
TPC-H	5059	4154	1.22





**THANK YOU**