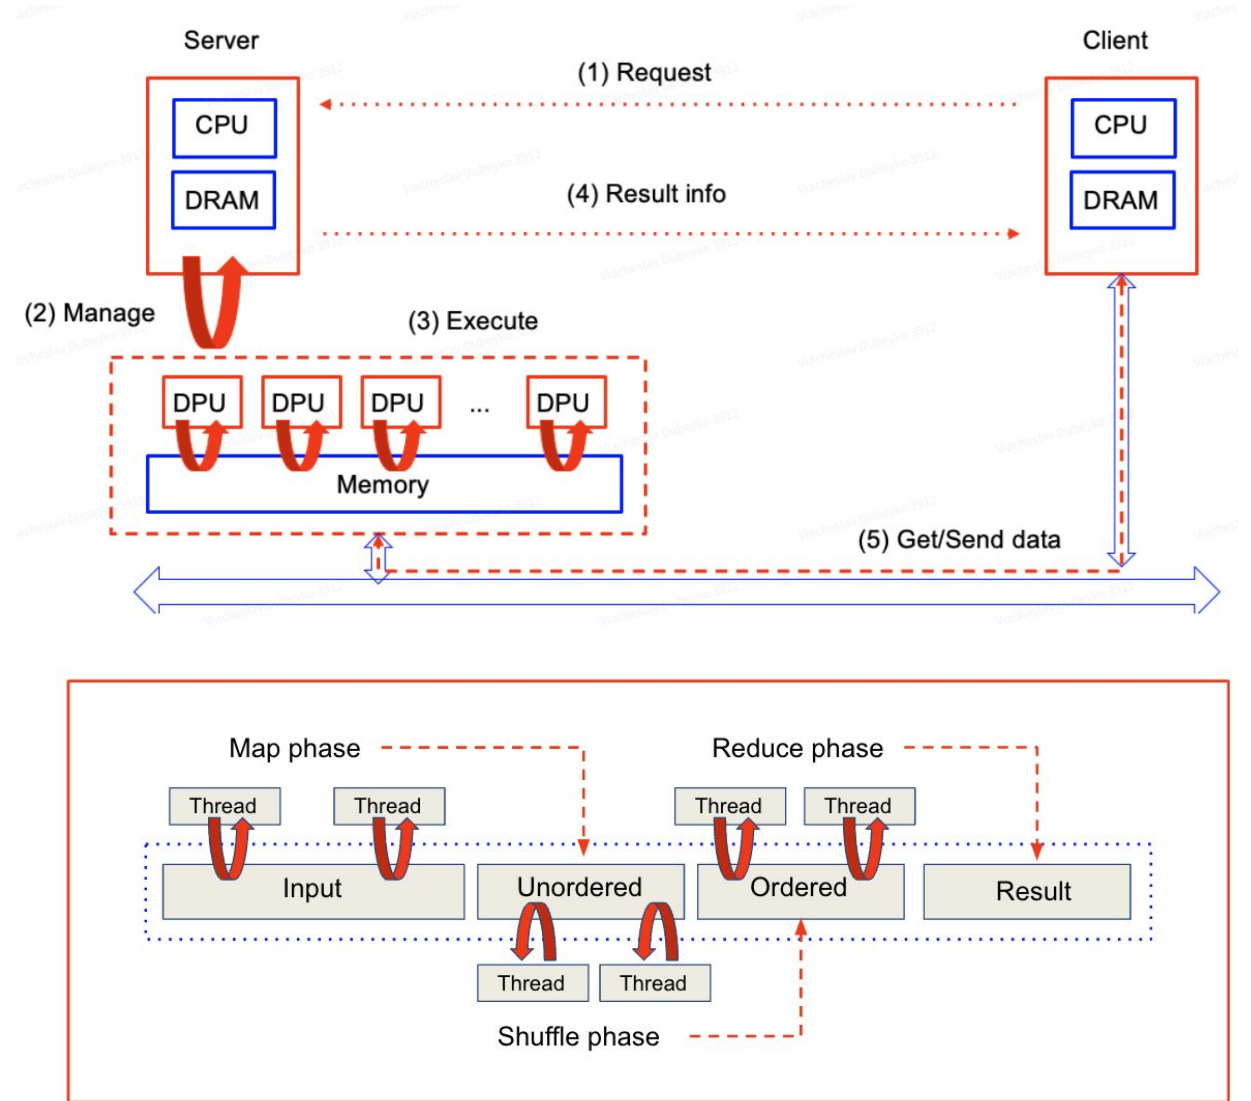
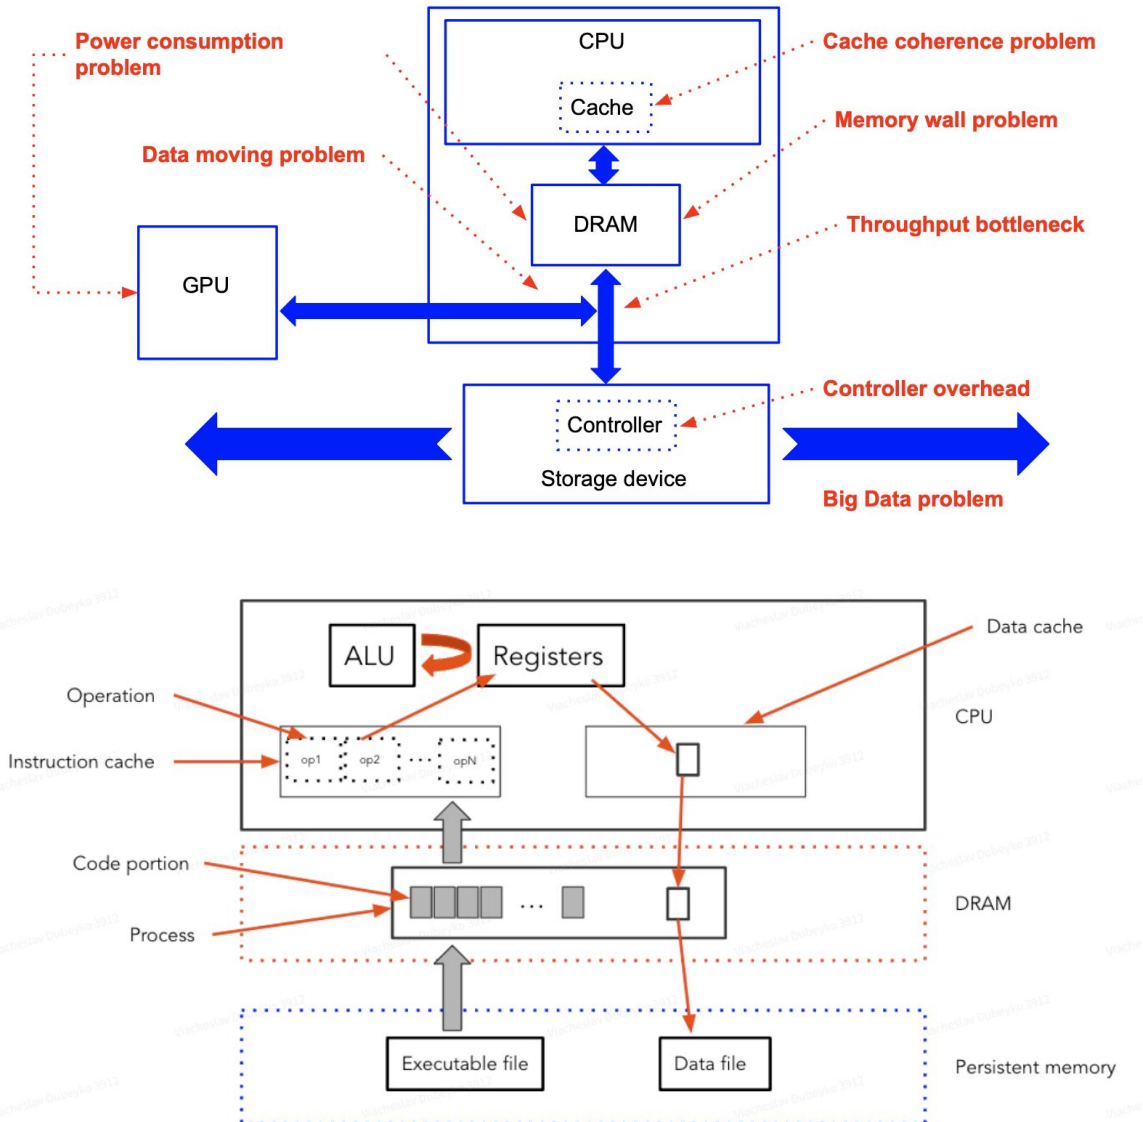




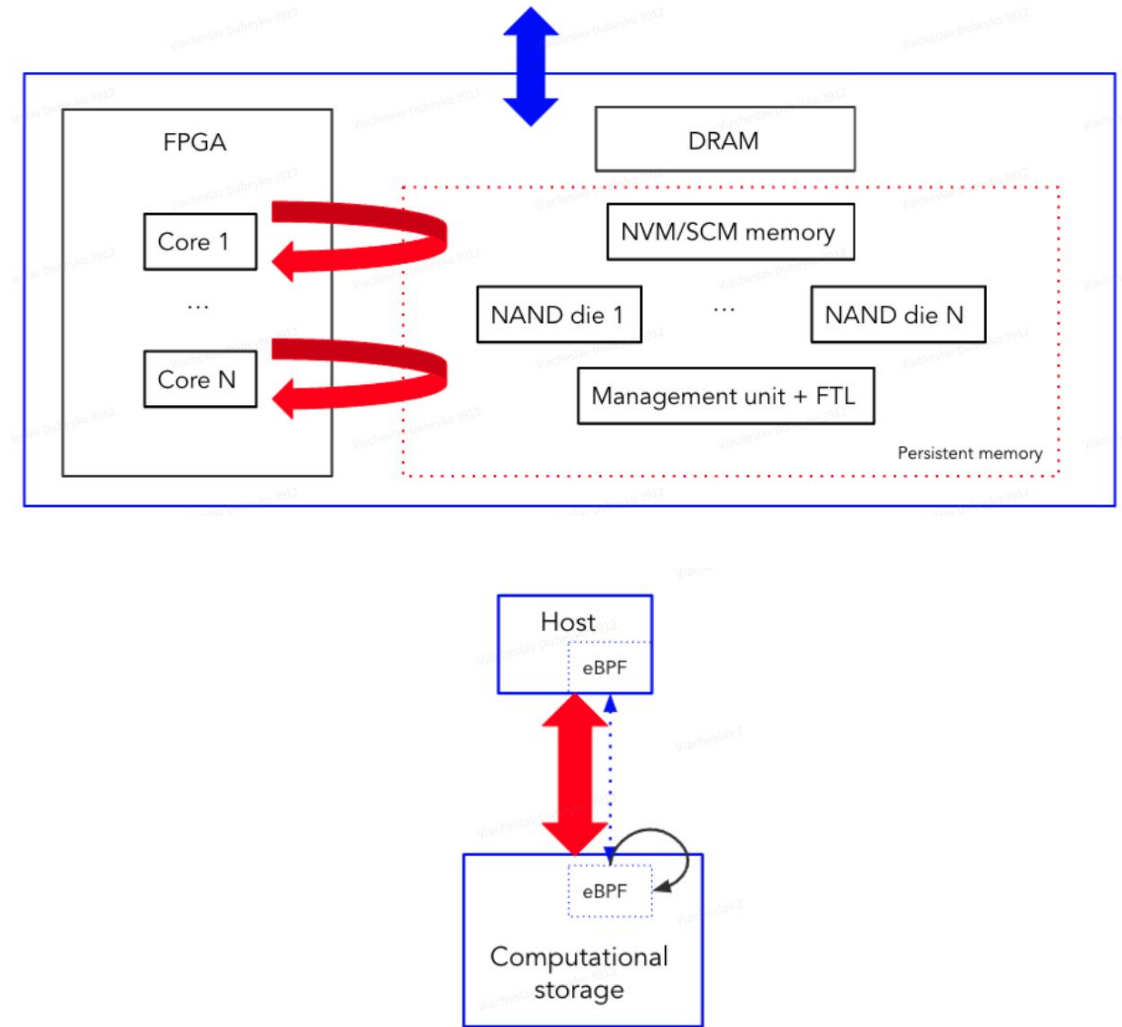
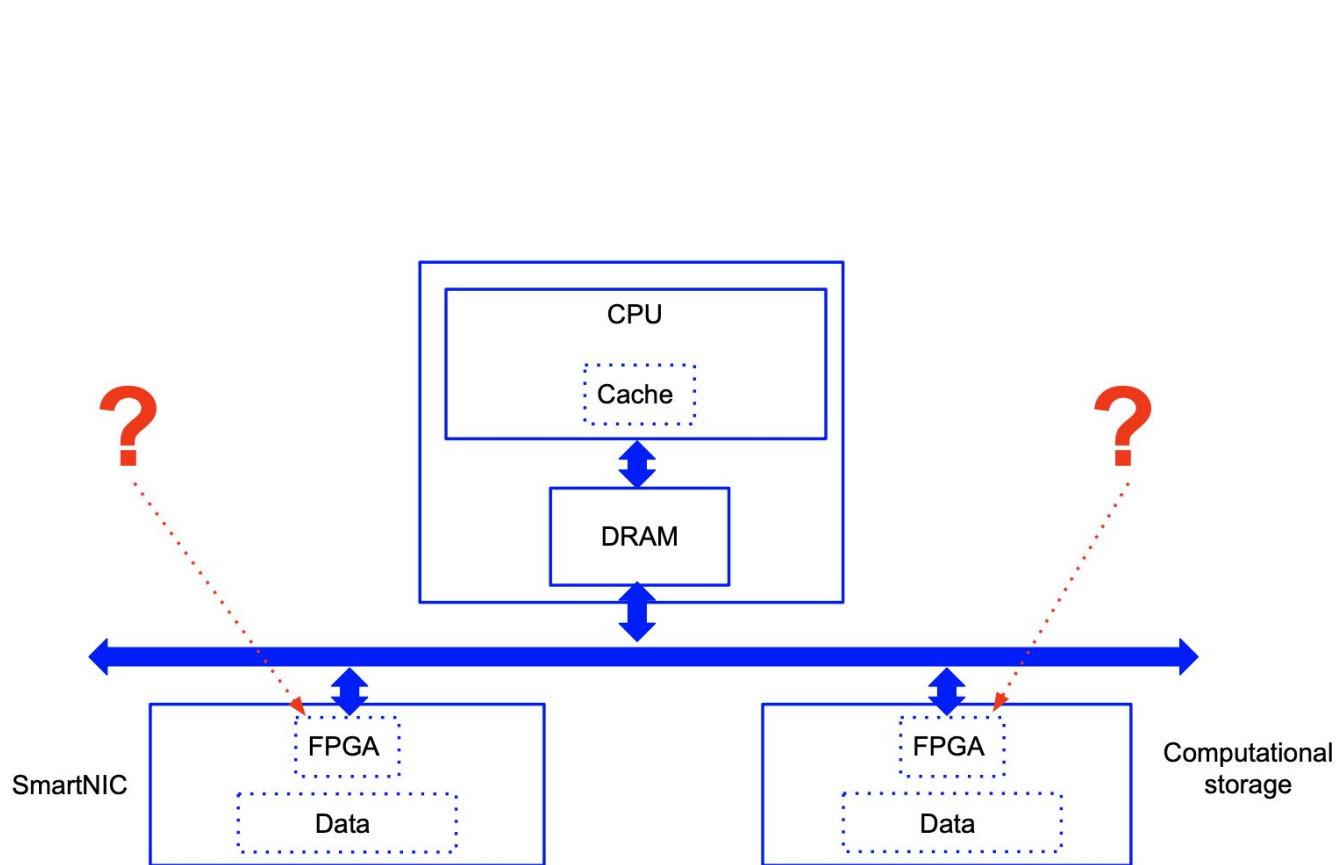
# Computational storage and file system collaboration

Viacheslav Dubeyko <[slava@dubeyko.com](mailto:slava@dubeyko.com)>

# Why do we need data-centric computing?

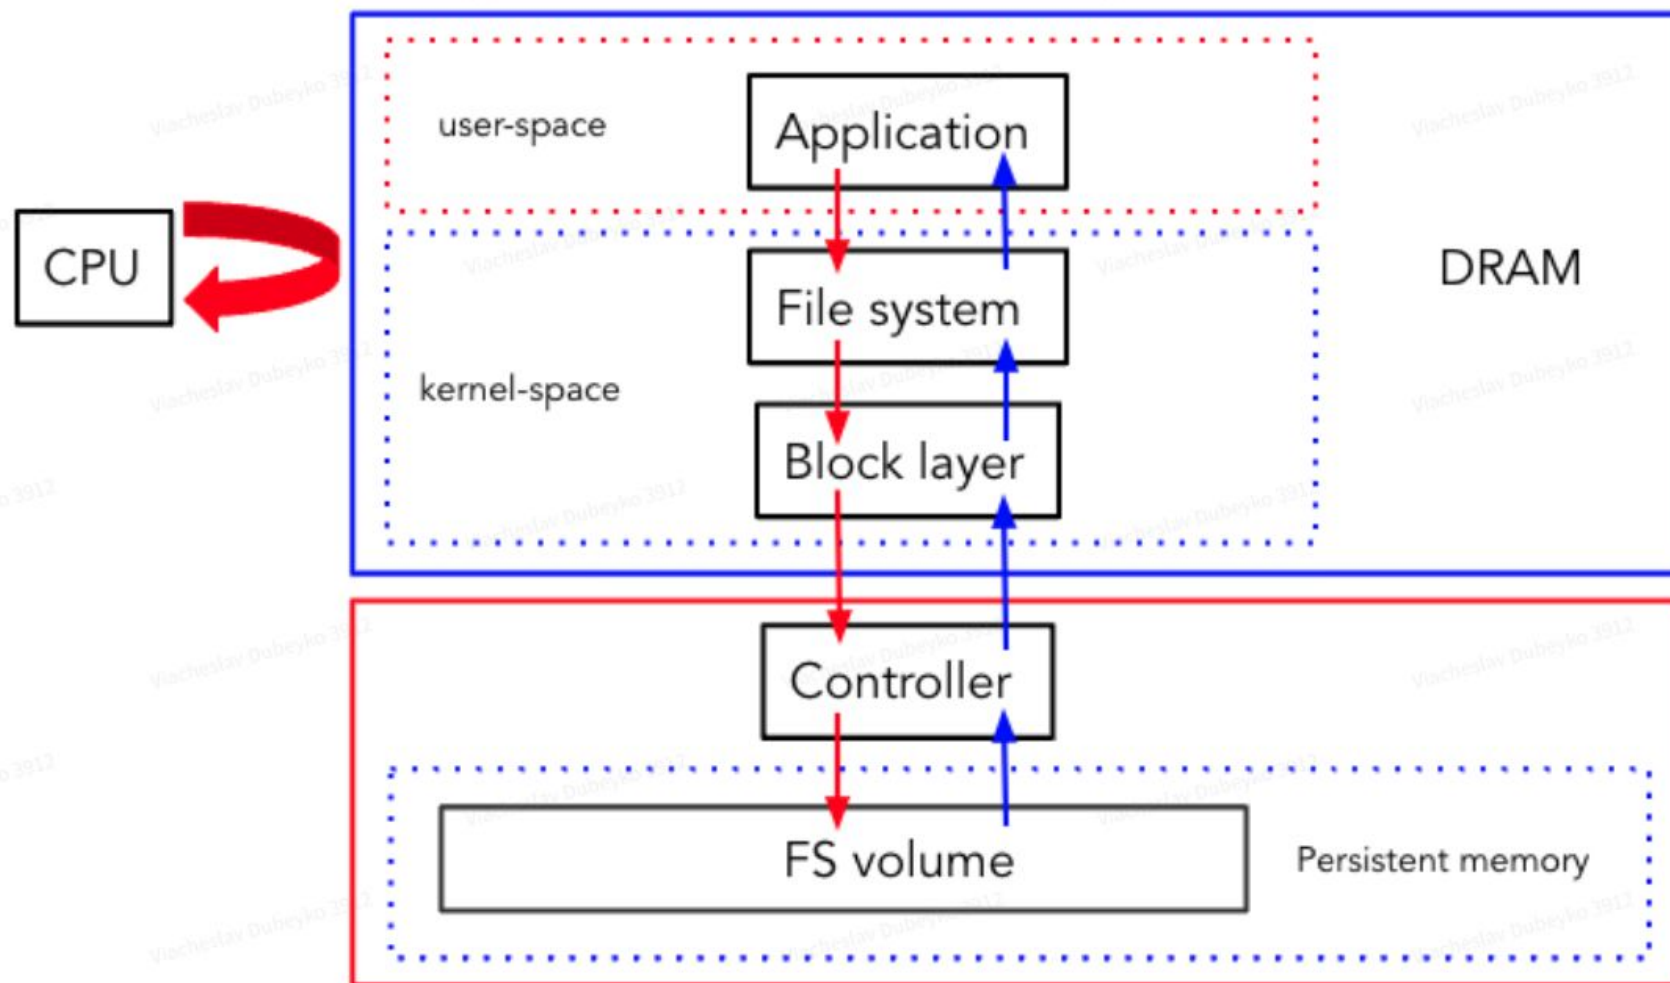


# What is computational storage?

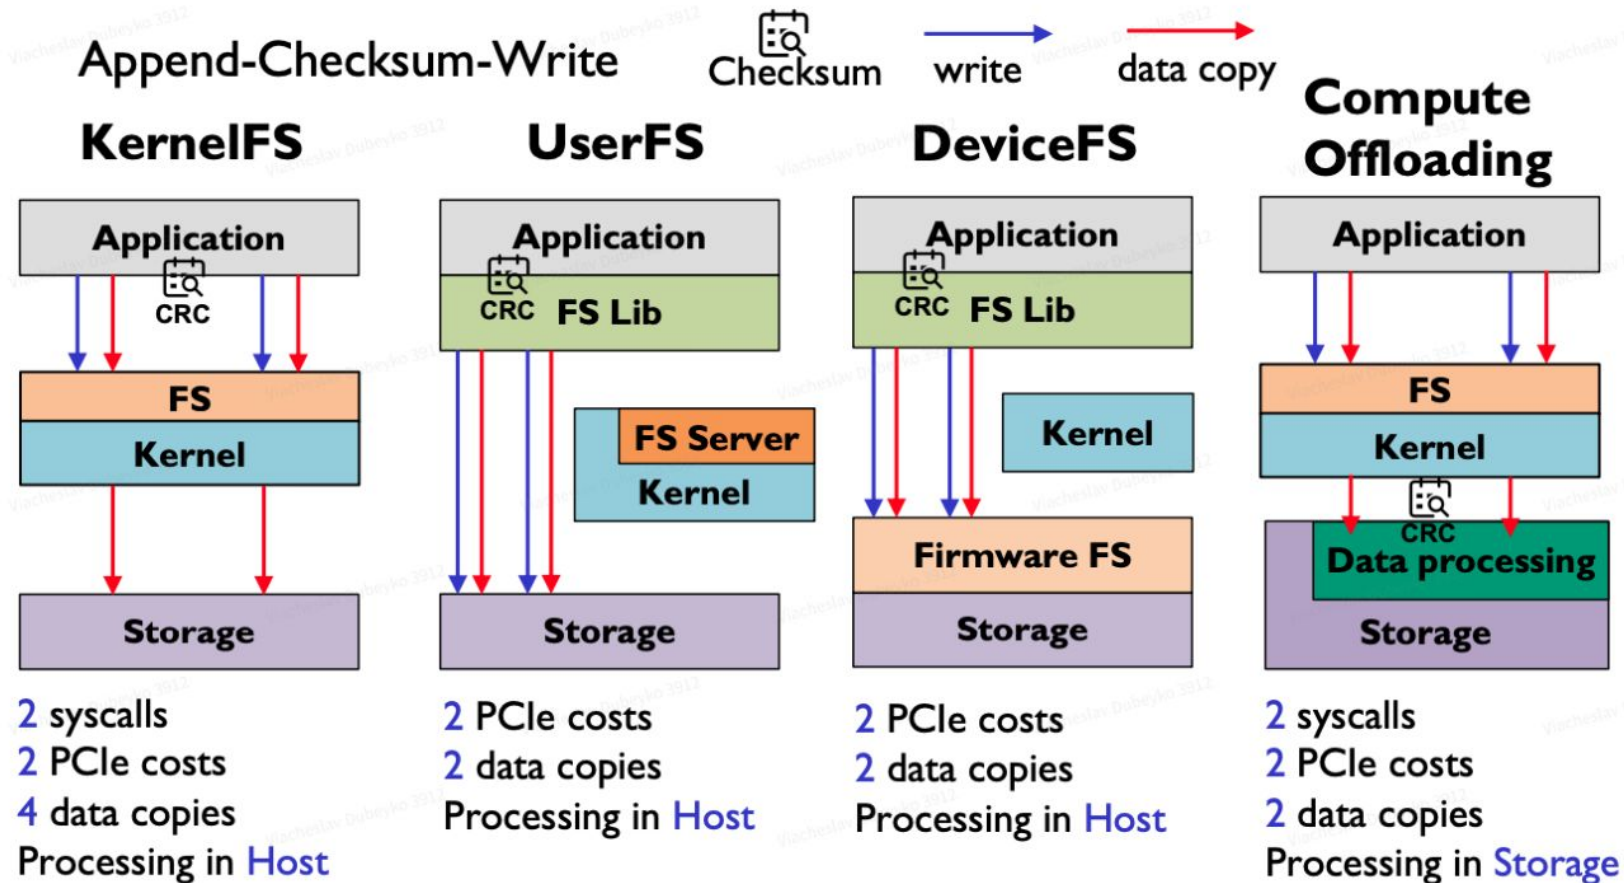




# What's about file system?



# Can file system survive in data-centric world?




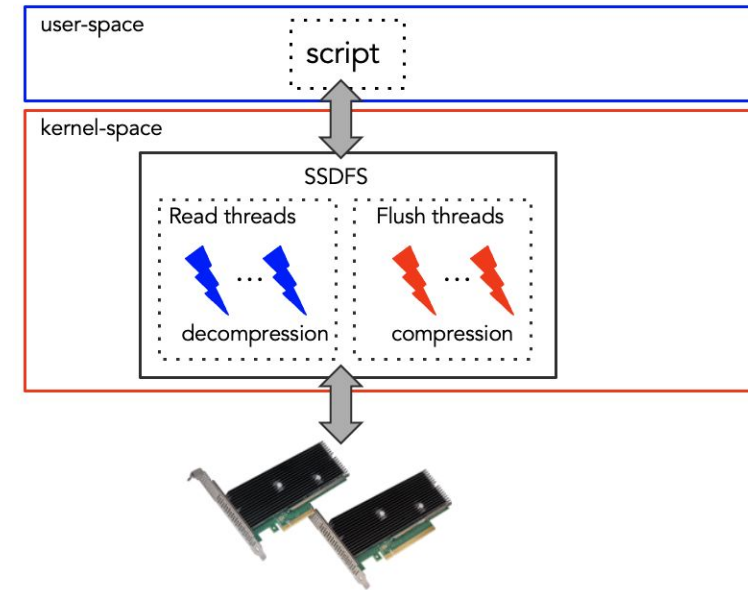
Jian Zhang, Yujie Ren, and Sudarsun Kannan, "FusionFS: Fusing I/O Operations using CISCops in Firmware File Systems", 20th USENIX Conference on File and Storage Technologies (FAST 22) - <https://www.usenix.org/conference/fast22/presentation/zhang-jian>

# Compression/decompression in file system(s)


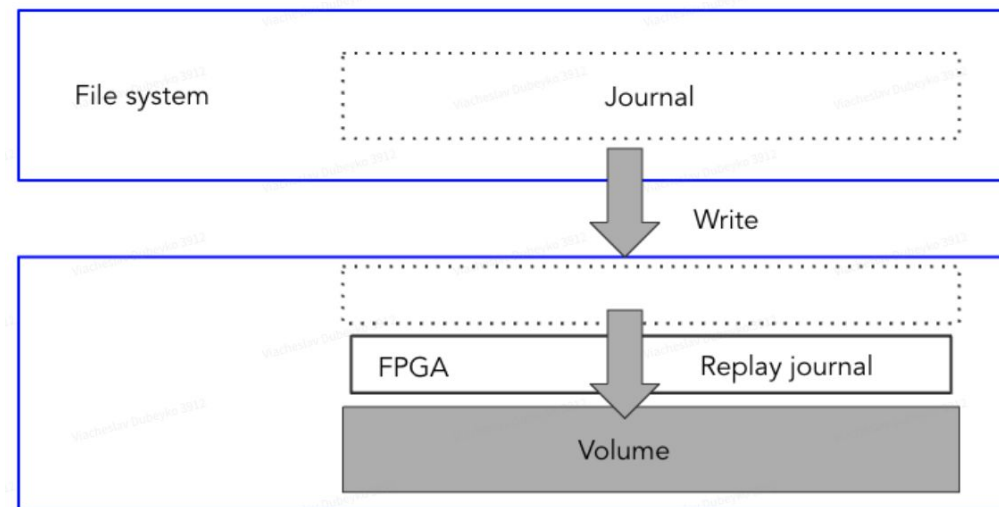
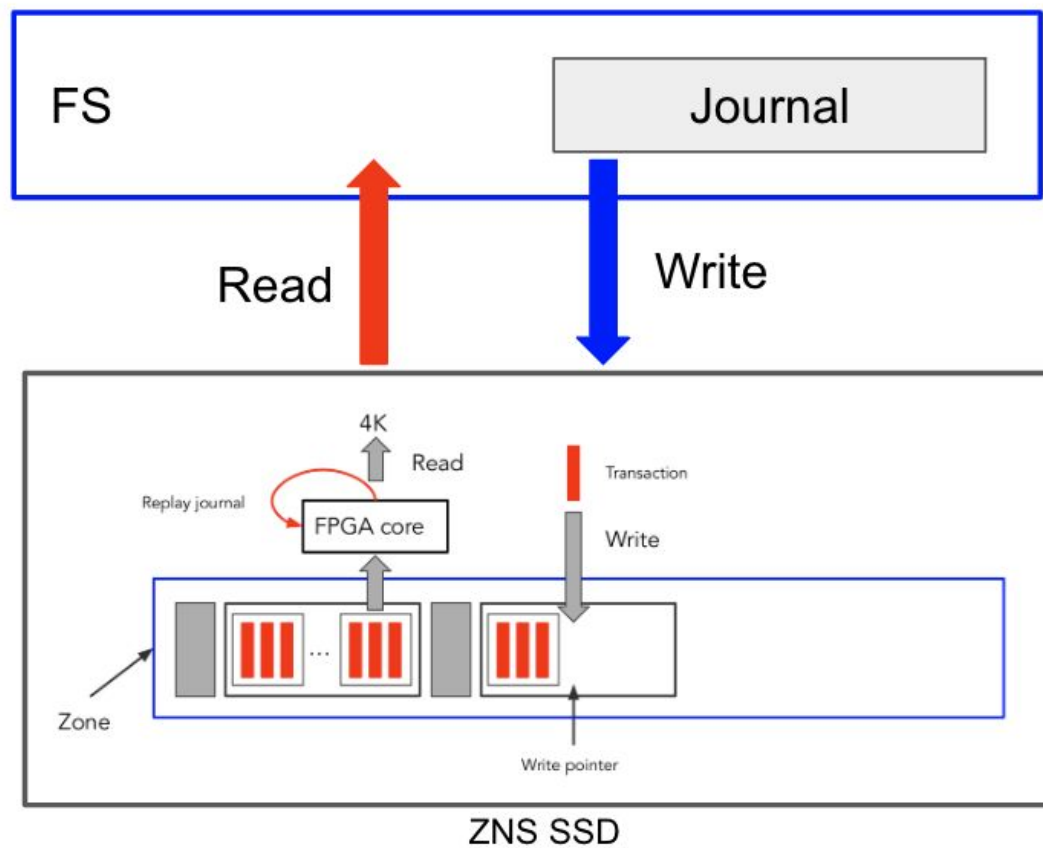


Intel® QuickAssist Adapters 8960/8970 deliver turn-key standard PCI Express (PCIe) access to hardware acceleration for compute intensive applications.

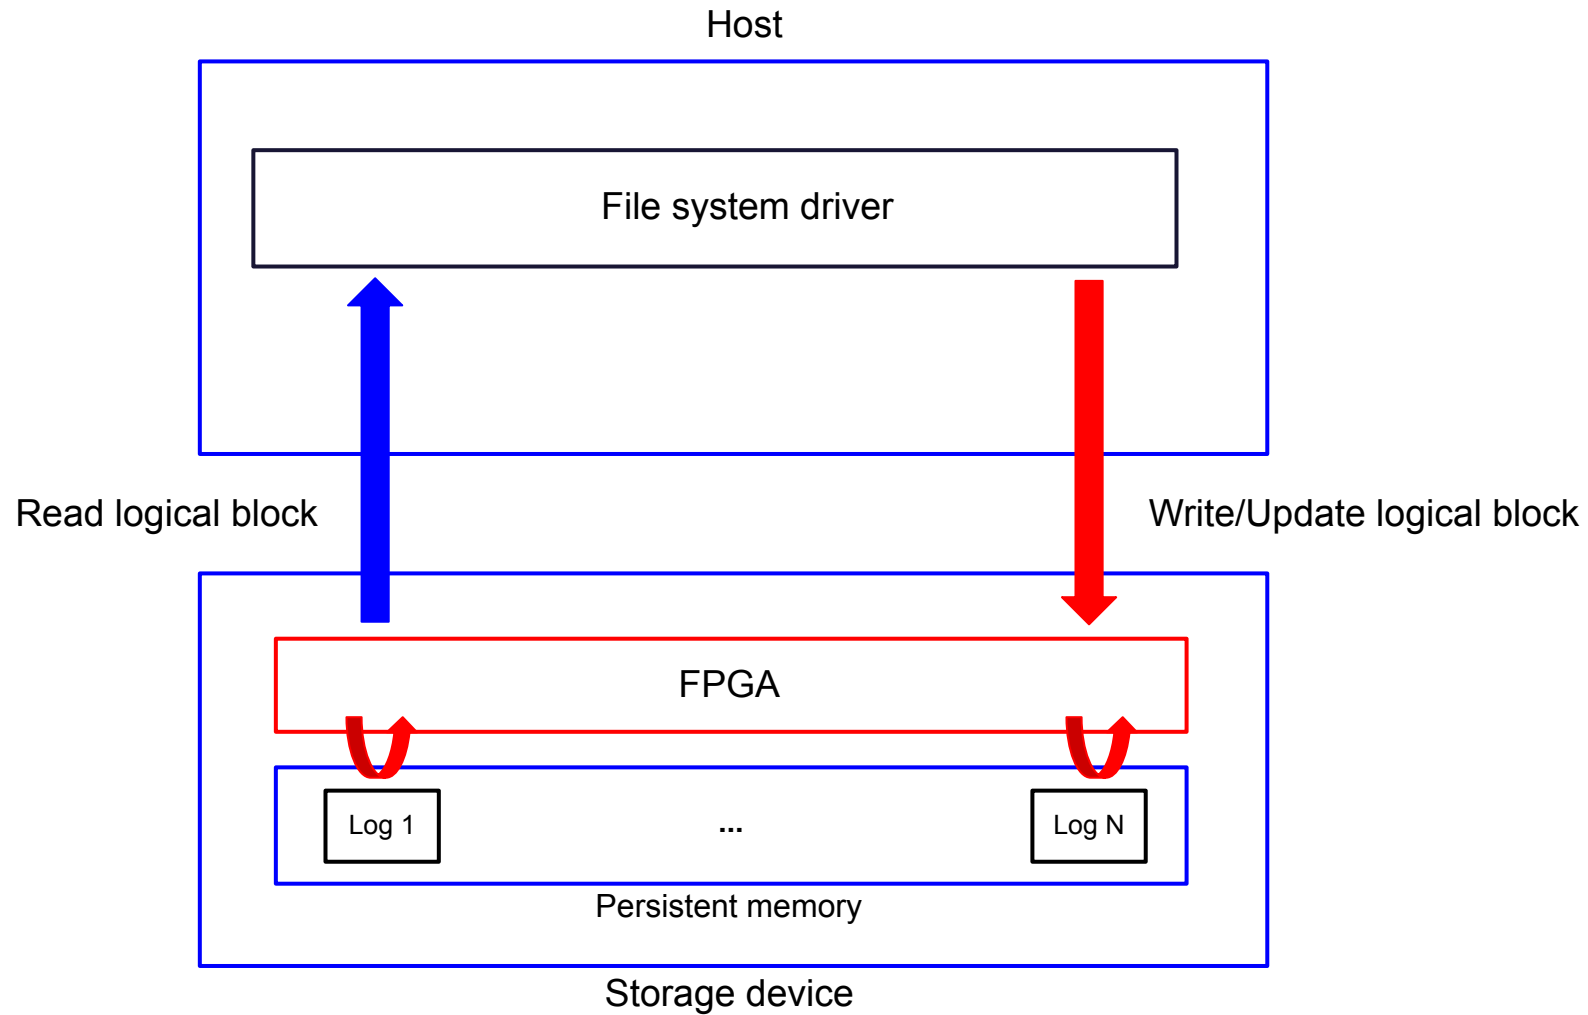
Performance Specifications		
Category	8960	8970
Performance Symmetric Ciphers AES128-CBC AES-XCBC	51Gbps @4KB Packet	103Gbps @4KB Packet
RSA2K Key Decrypts	100K Ops/s	100K Ops/s
Verified Compression Level 1 Dynamic Deflate	37Gbps @64KB	66Gbps @64KB
Decompression Level 1 Dynamic Deflate	54Gbps @64KB	160Gbps @64KB



# ZNS SSD + computational storage?

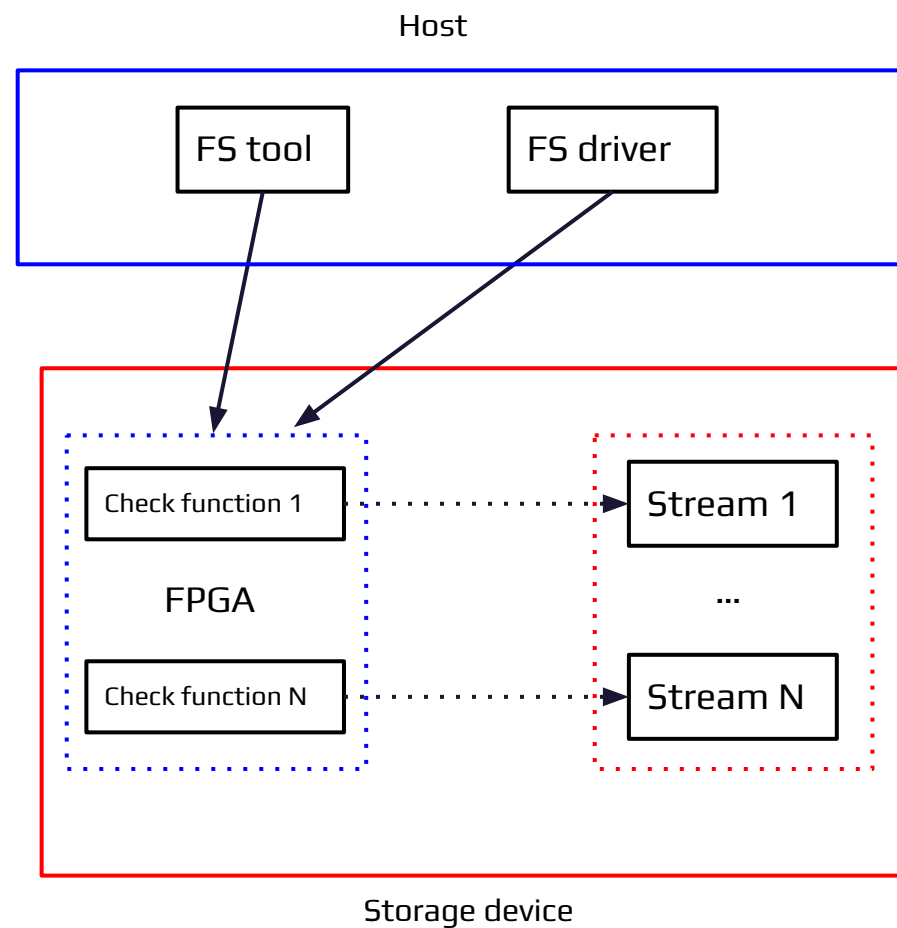
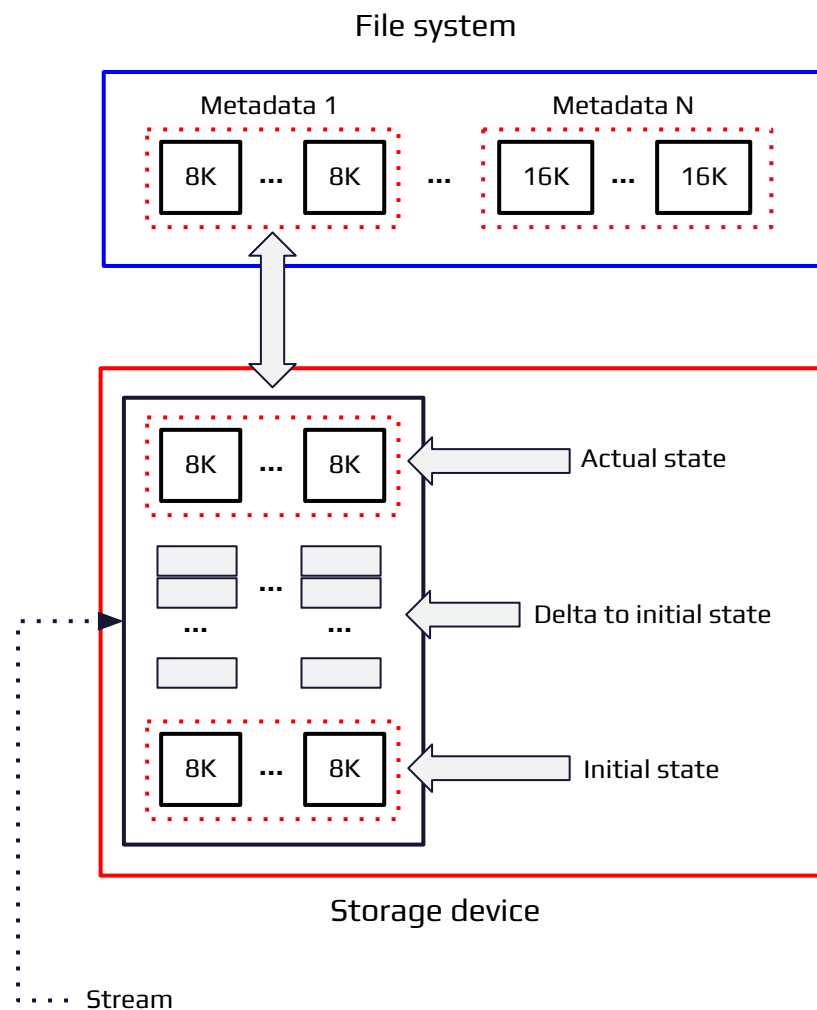


# Log-based and log-structured file systems

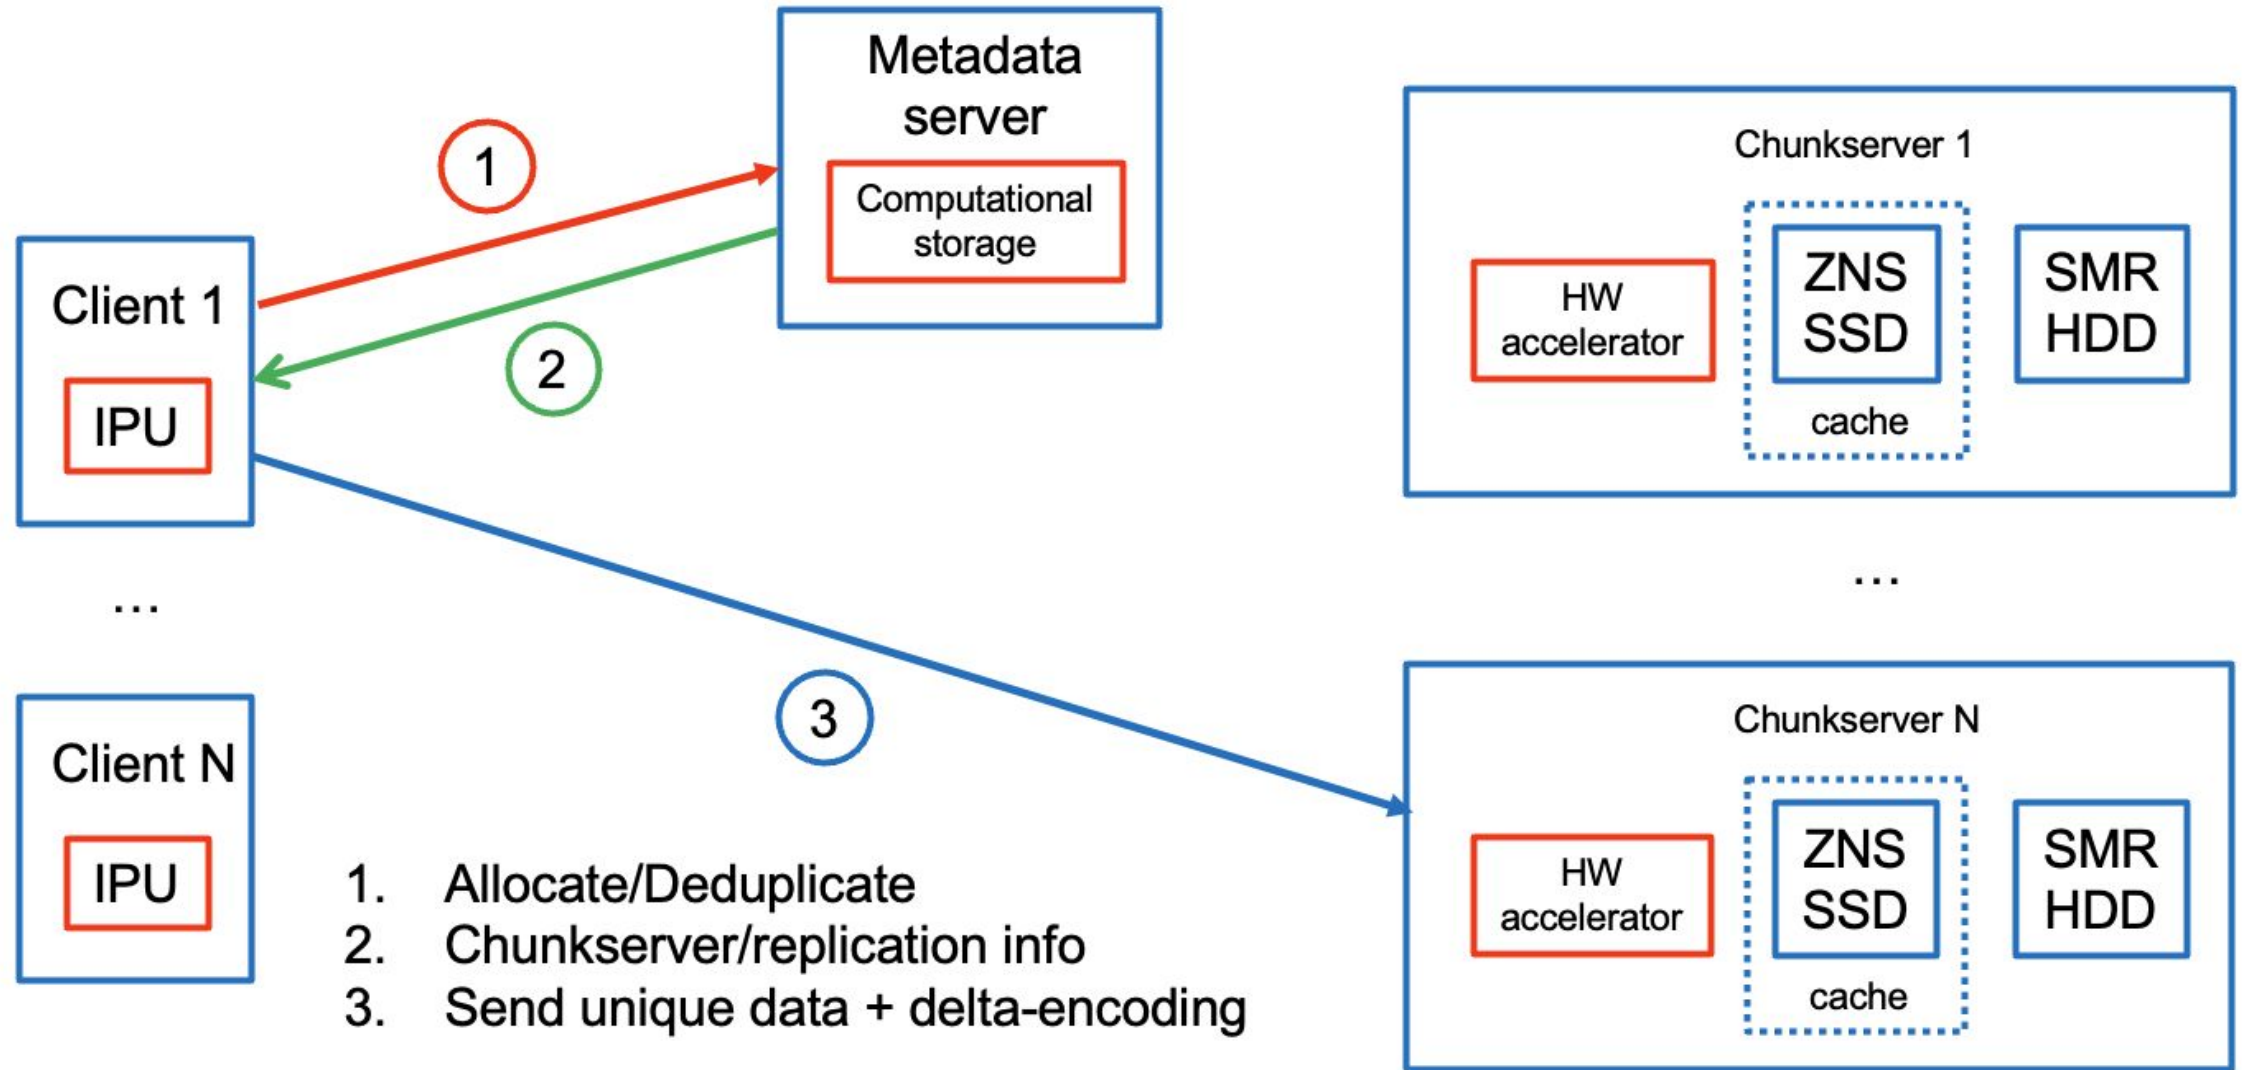




# FSCK-based and snapshot-based corruption recovering

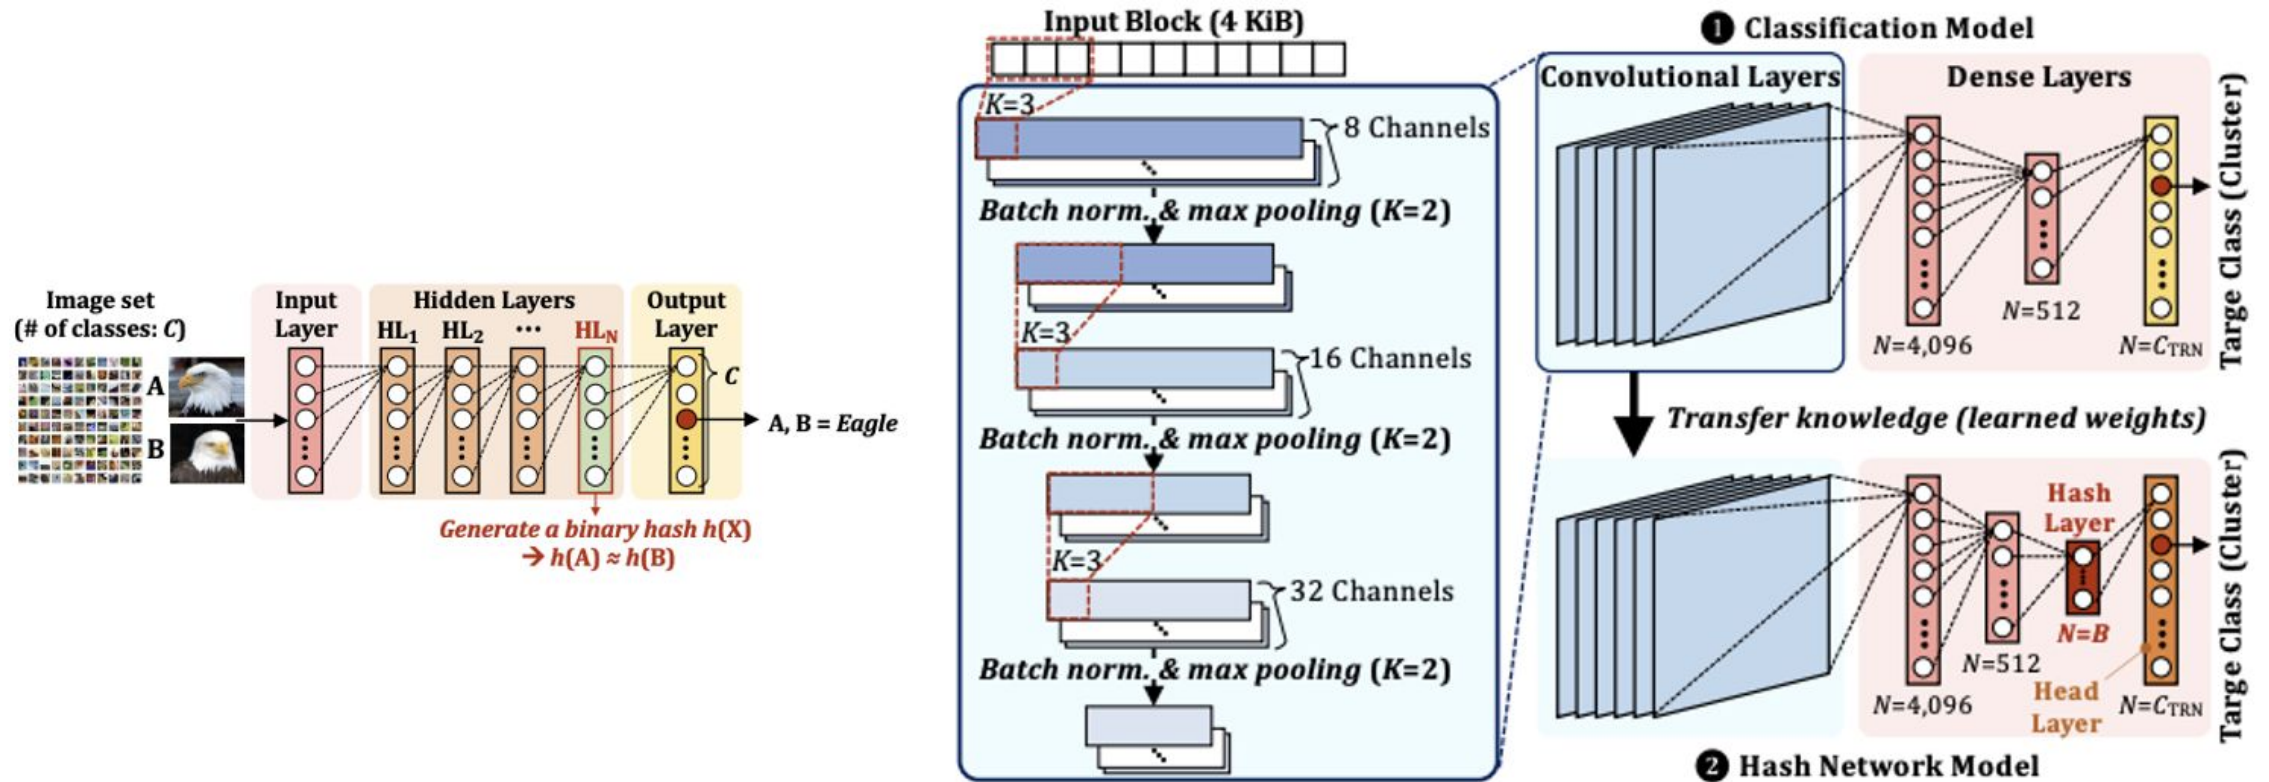


# Delta-encoding and deduplication



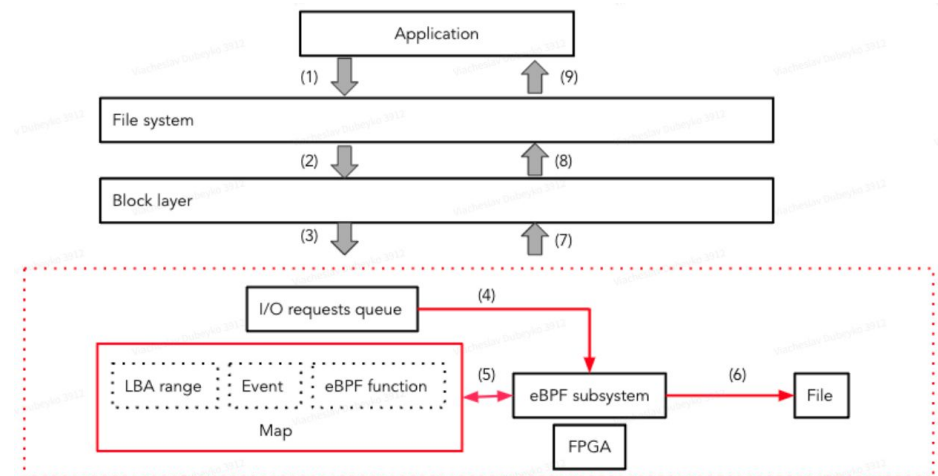
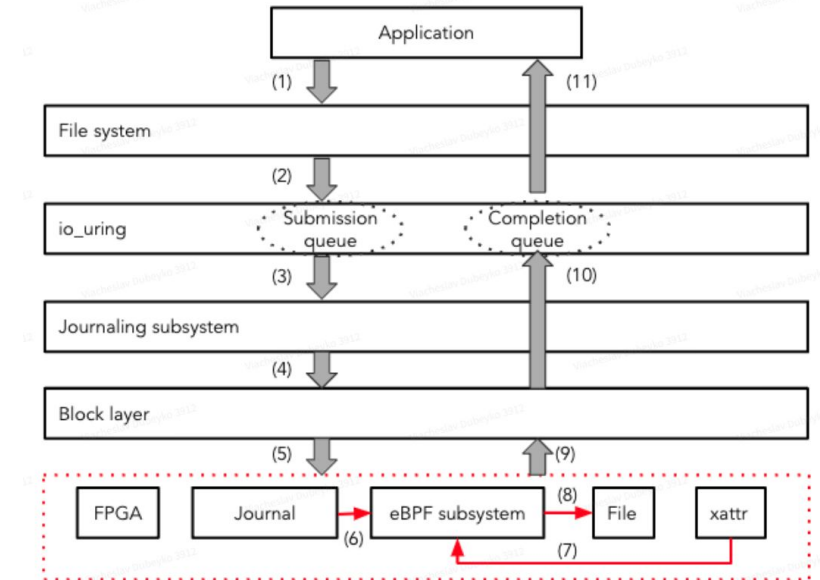
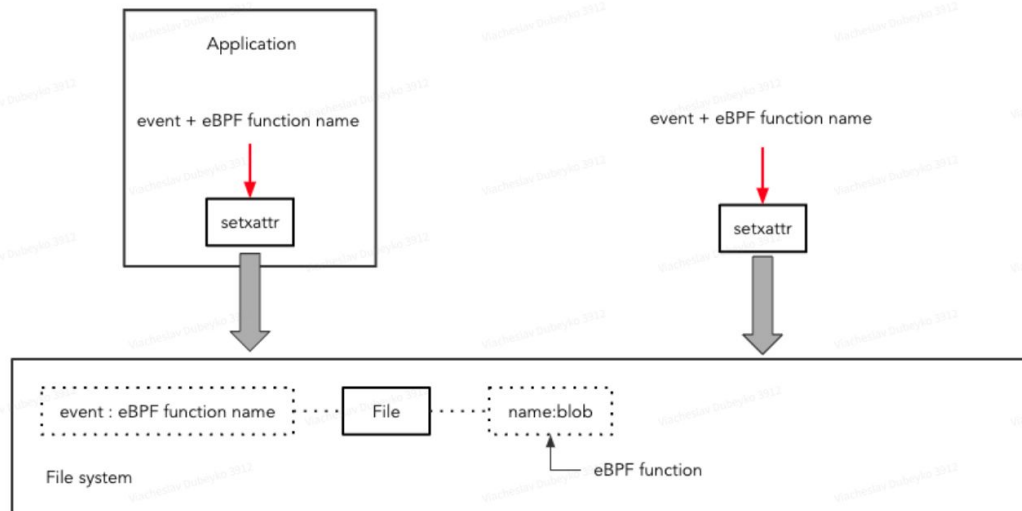
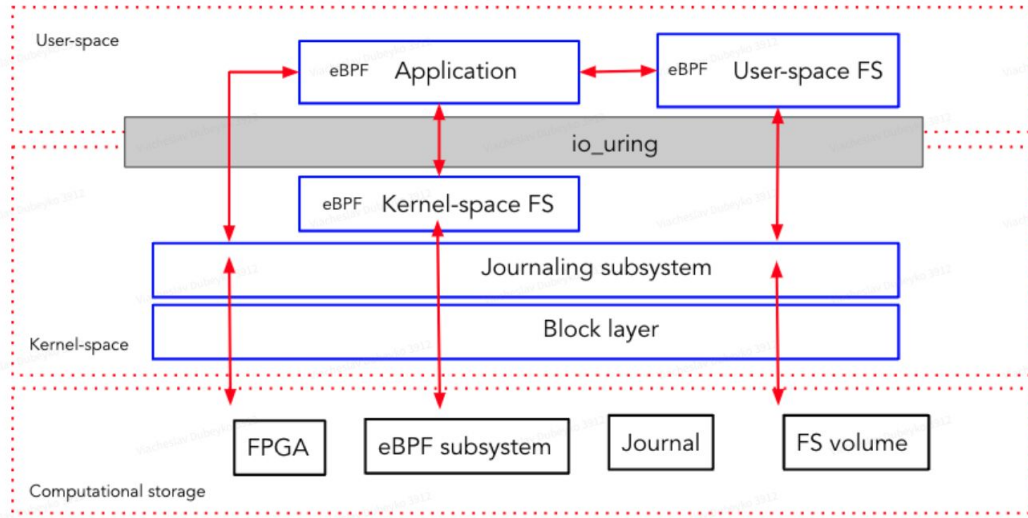
# ML approaches in file systems

## DeepSketch: Machine Learning-Based Reference Search Technique



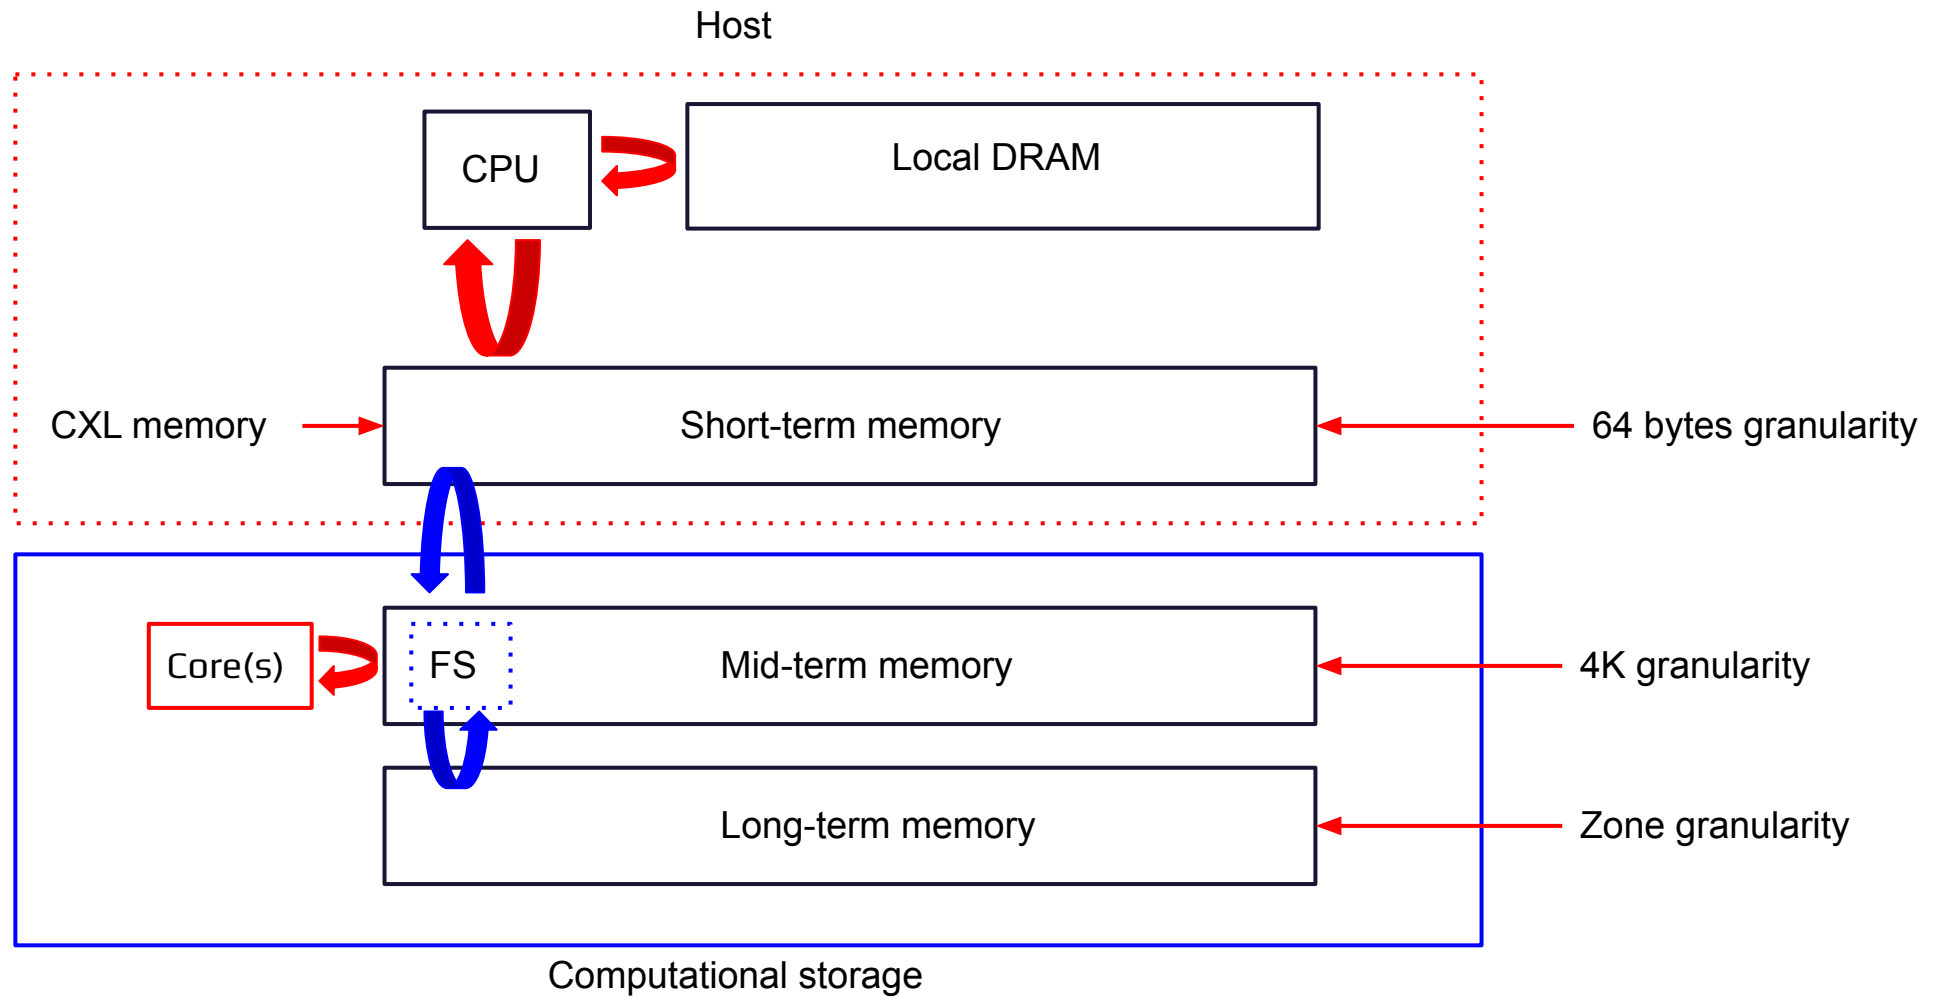
Jisung Park and Jeonggyun Kim and Yeseong Kim and Sungjin Lee and Onur Mutlu. DeepSketch: A New Machine Learning-Based Reference Search Technique for Post-Deduplication Delta Compression. 20th USENIX Conference on File and Storage Technologies (FAST 22), 2022.

# File system based computational storage architecture

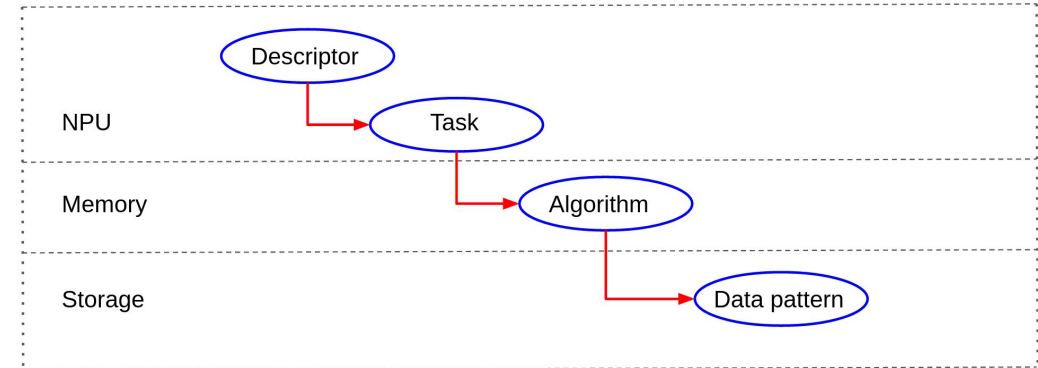
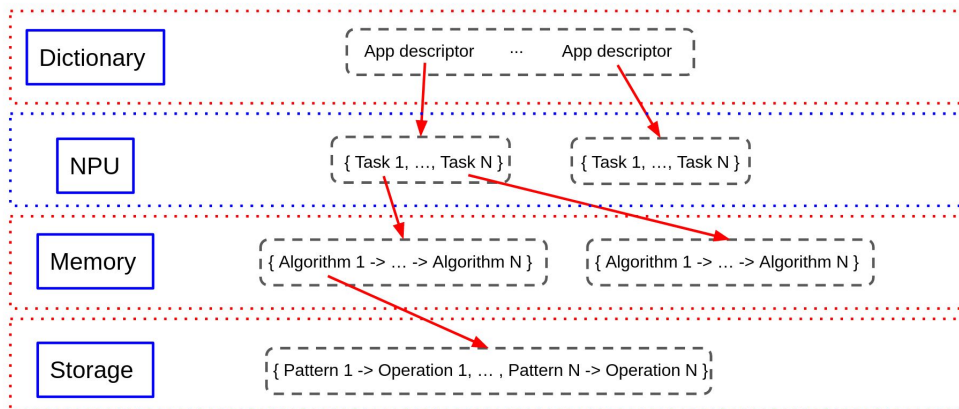





# Data-centric computing stack



# Cognitive computing + neuromorphic computing



A night landscape featuring a starry sky with the Milky Way, snow-capped mountains, and a calm lake reflecting a bright light source on the left.

THANK YOU

QUESTIONS???