

# Hardware Acceleration for RAID5/6, Deduplication & Security for parallel workloads

Vikas Aggarwal

Storage Developer Conference Bangalore 2017



# Team Members

- Venkatesh Bolla
- Achyutha Krishna
- Anil Reddy

**Responsibilities:** Storage software acceleration on OCTEON TX.

# Outline

- RAID 5/6
- Deduplication
- Security
- Open Source Storage Applications
- OCTEON TX Acceleration blocks
- Accelerated Data Flow
- Observations

# Offloading CPU intensive operations

- RAID6
  - XOR
  - Galois Field Multiplication
- Deduplication
  - Fingerprint Generation
  - Fingerprint Lookup
- Security
  - Encryption/Decryption of Data Blocks

# Open Source Storage Software Offload Case Studies

RAID

DEDUPLICATION

SECURITY

MD RAID

DM Dedup

ecryptfs

DM RAID

OpenZFS

DM Crypt

btrfs

btrfs

OpenDedup

# Linux MD RAID

- Implements the following:
  - RAID 0, 1, 5, 6
- MDRAID can offload following to hardware using `async_tx` Linux offload infrastructure:
  - `memcpy`
  - XOR
  - Galois Field Arithmetic
- Current Offload benefits the following RAID variants:
  - RAID 5, 6

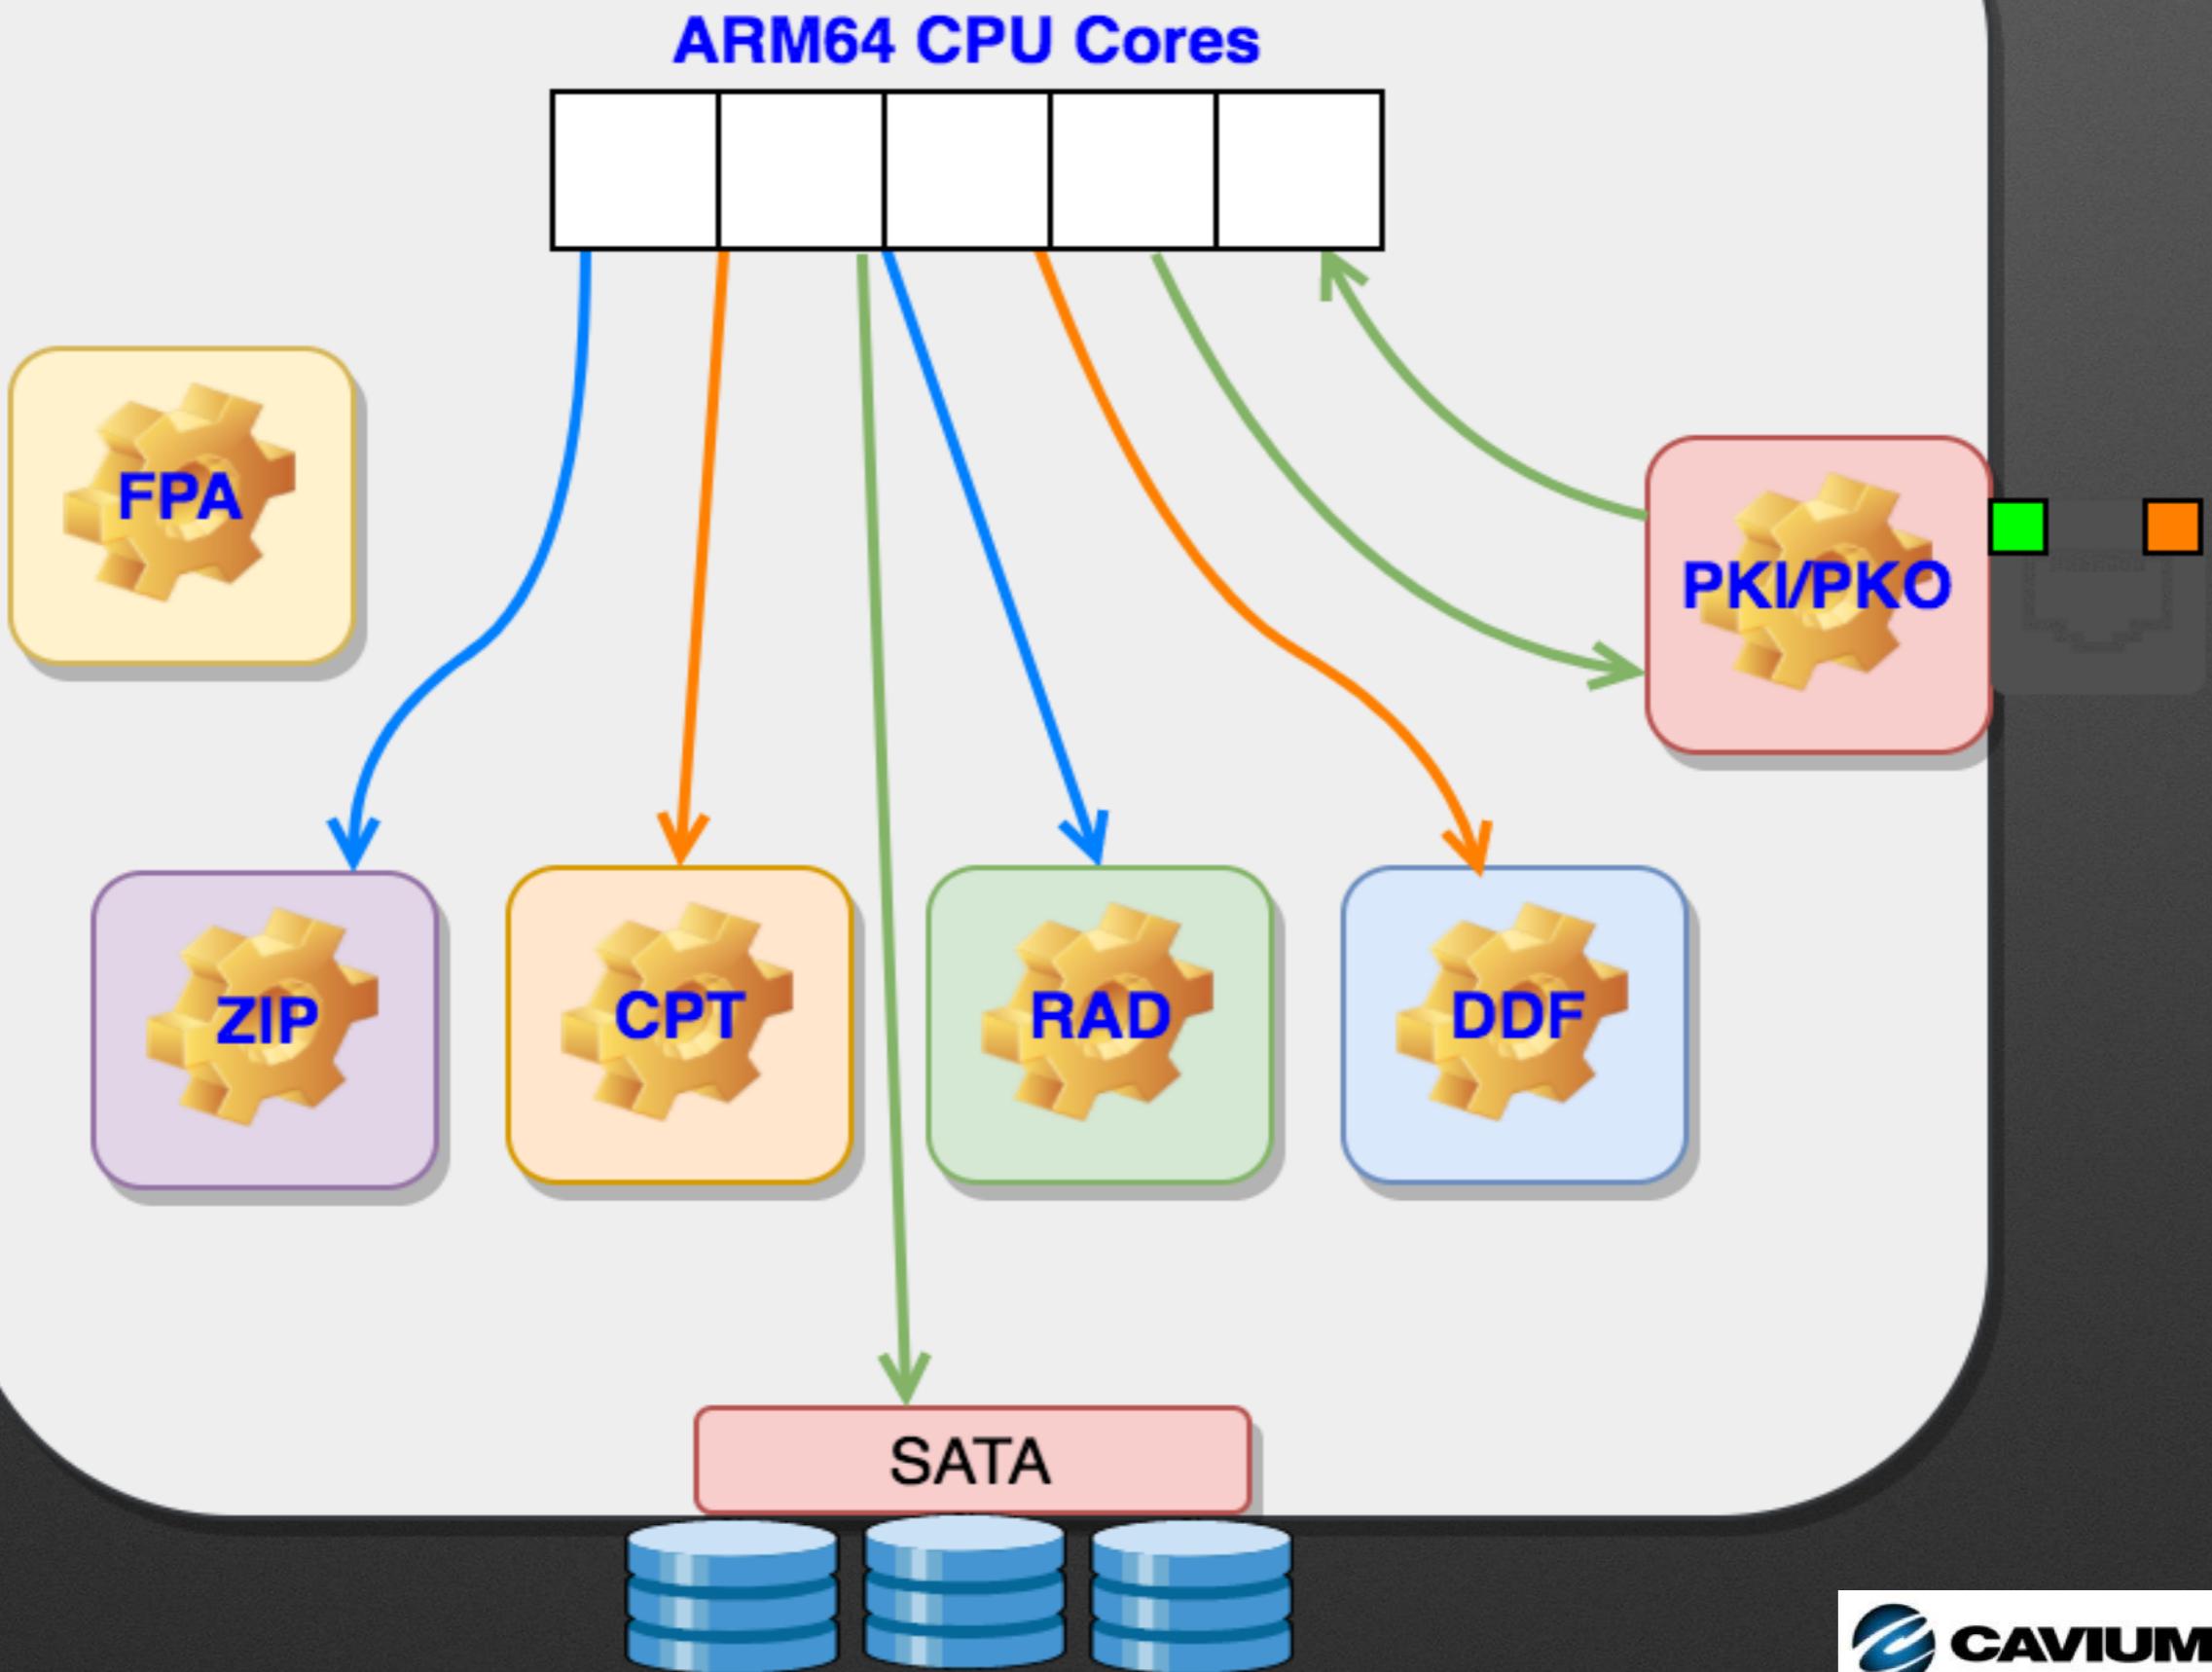
# Linux DM-dedup

- Implements the following:
  - 4KB block fingerprint
  - Fingerprint to PBN(Physical block number) lookup
  - LBN to PBN(Physical block number) lookup
- DM Dedup(modified) can offload following to hardware:
  - Fingerprint (Digest using MD5, SHA1, SHA2)
  - Fingerprint and LBN lookup.
- Current Offload benefits the following:
  - Lookups.

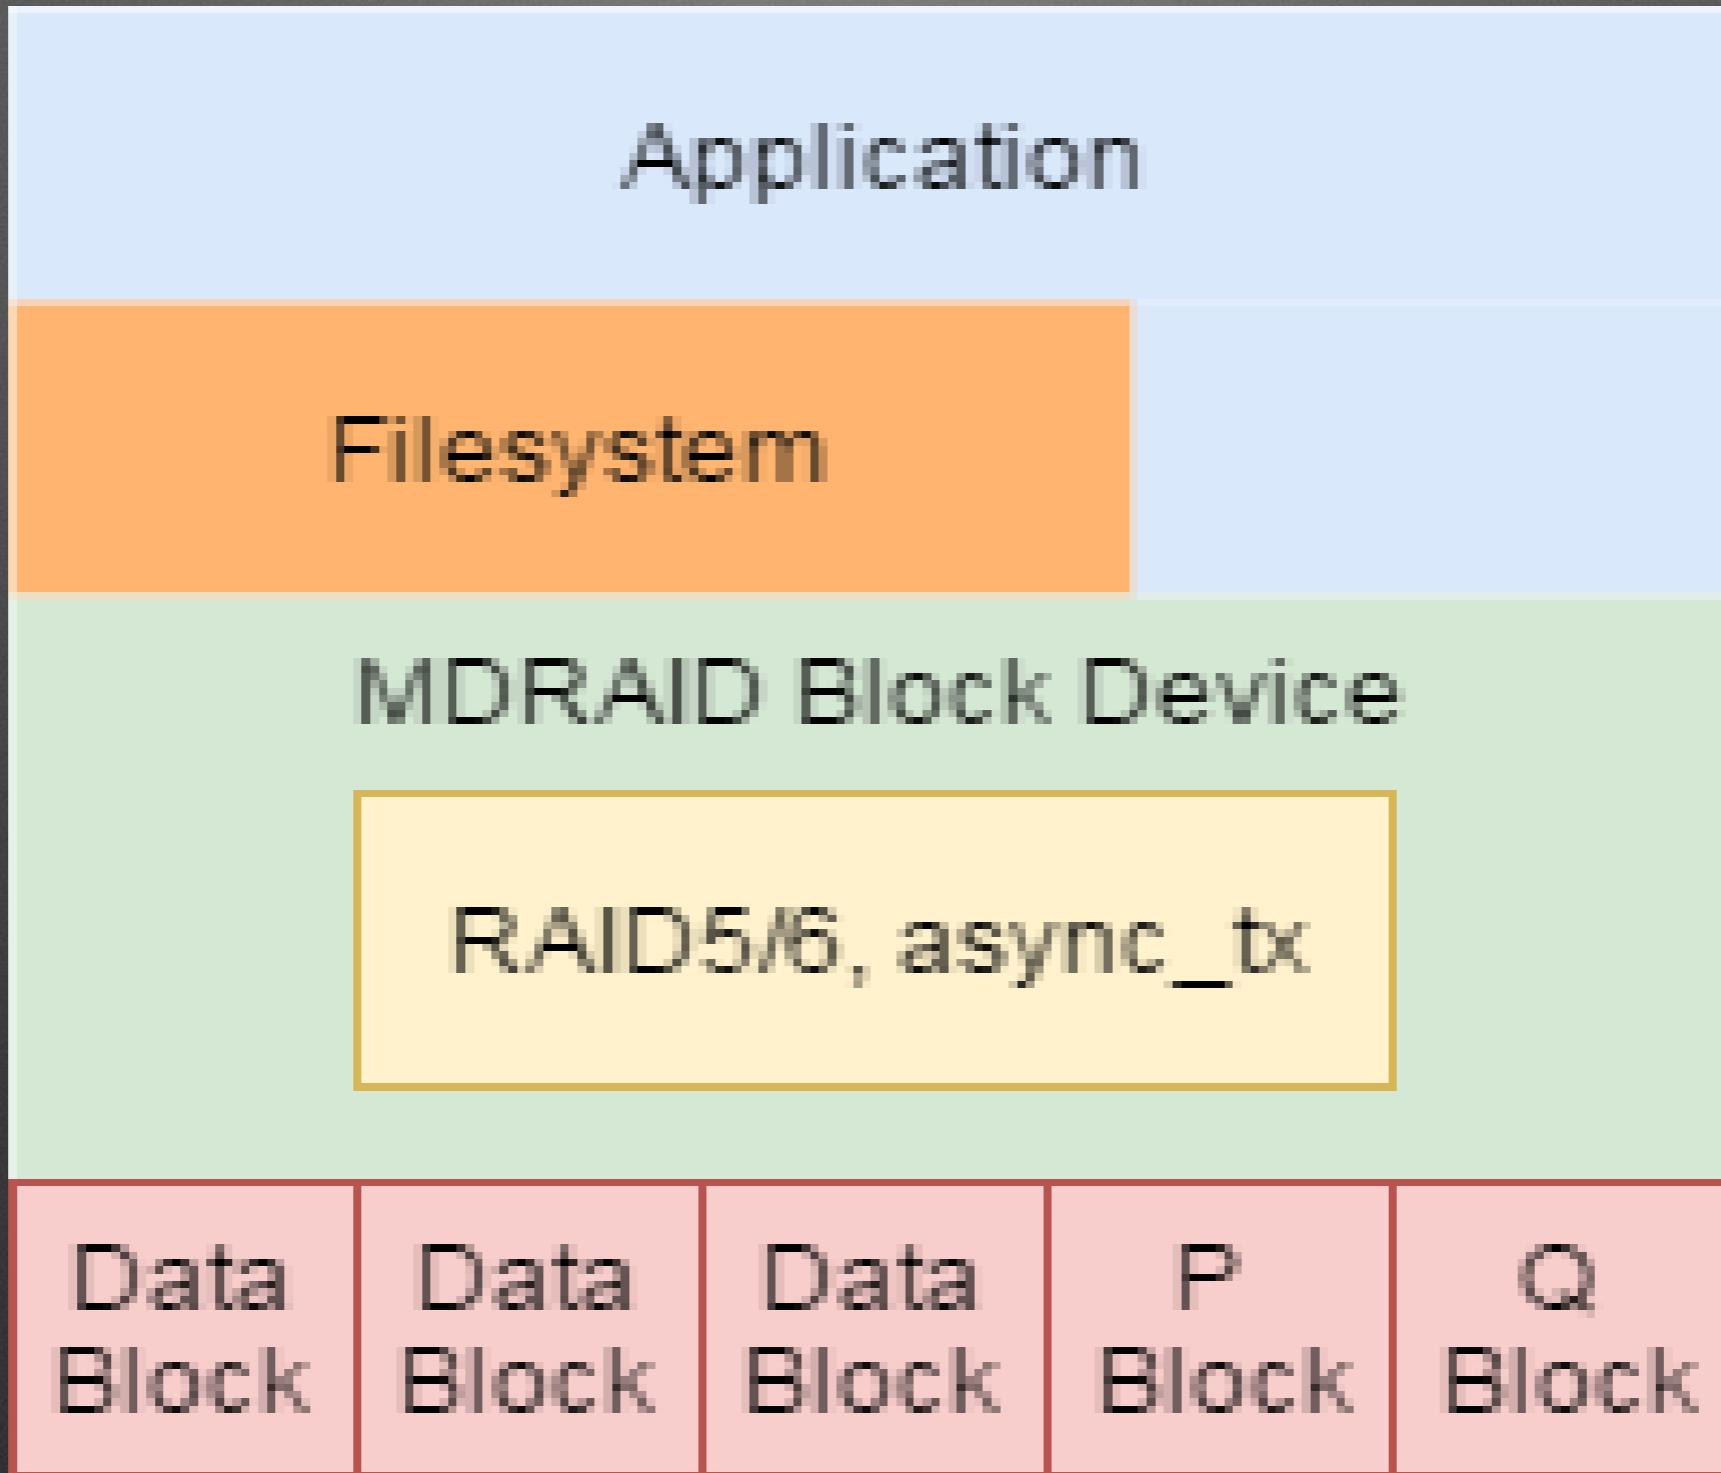
# Linux eCryptfs

- In-kernel standalone implementation.
- Security gets inherited into incremental backups.
- Cryptographic metadata is stored along with encrypted file.
- Supports Linux cryptographic ciphers.
- Utilizes Linux crypto framework
- eCryptfs can offload following to hardware
  - AES CBC
  - DES3 CBC
  - AES XTS
- Offload benefits: Encryption, Decryption.

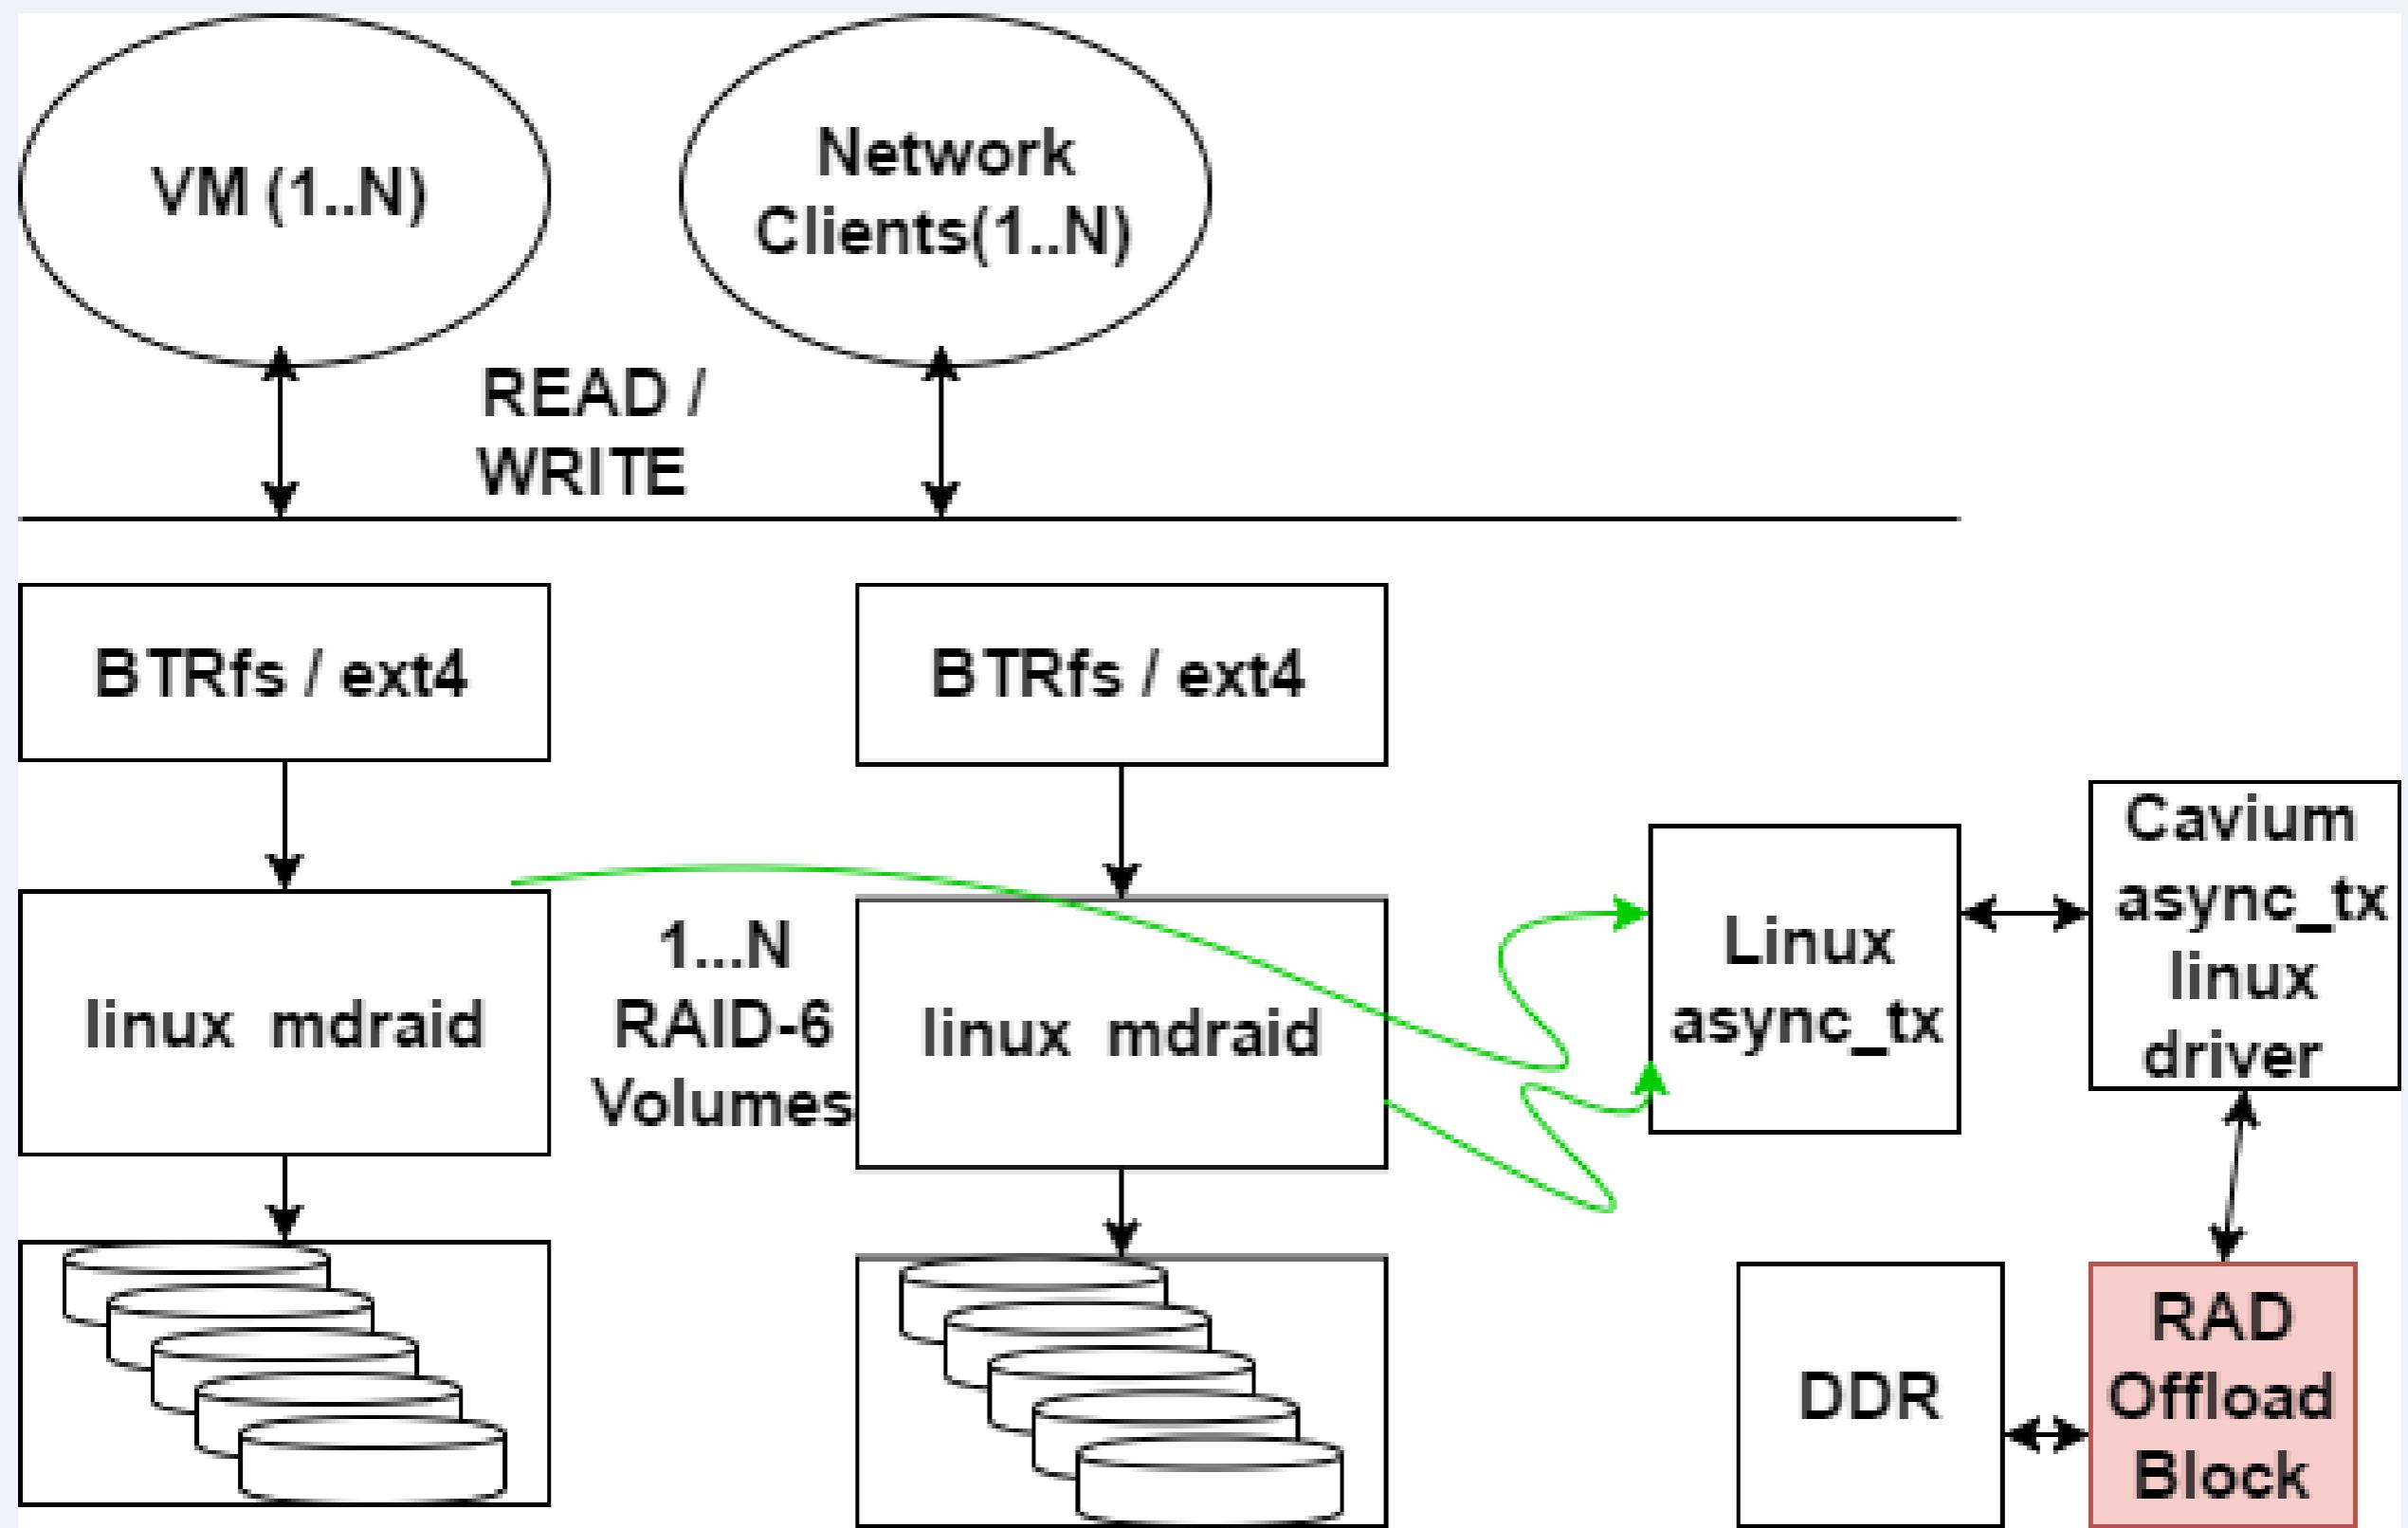
# OCTEON TX as a NAS/SAN SOC platform

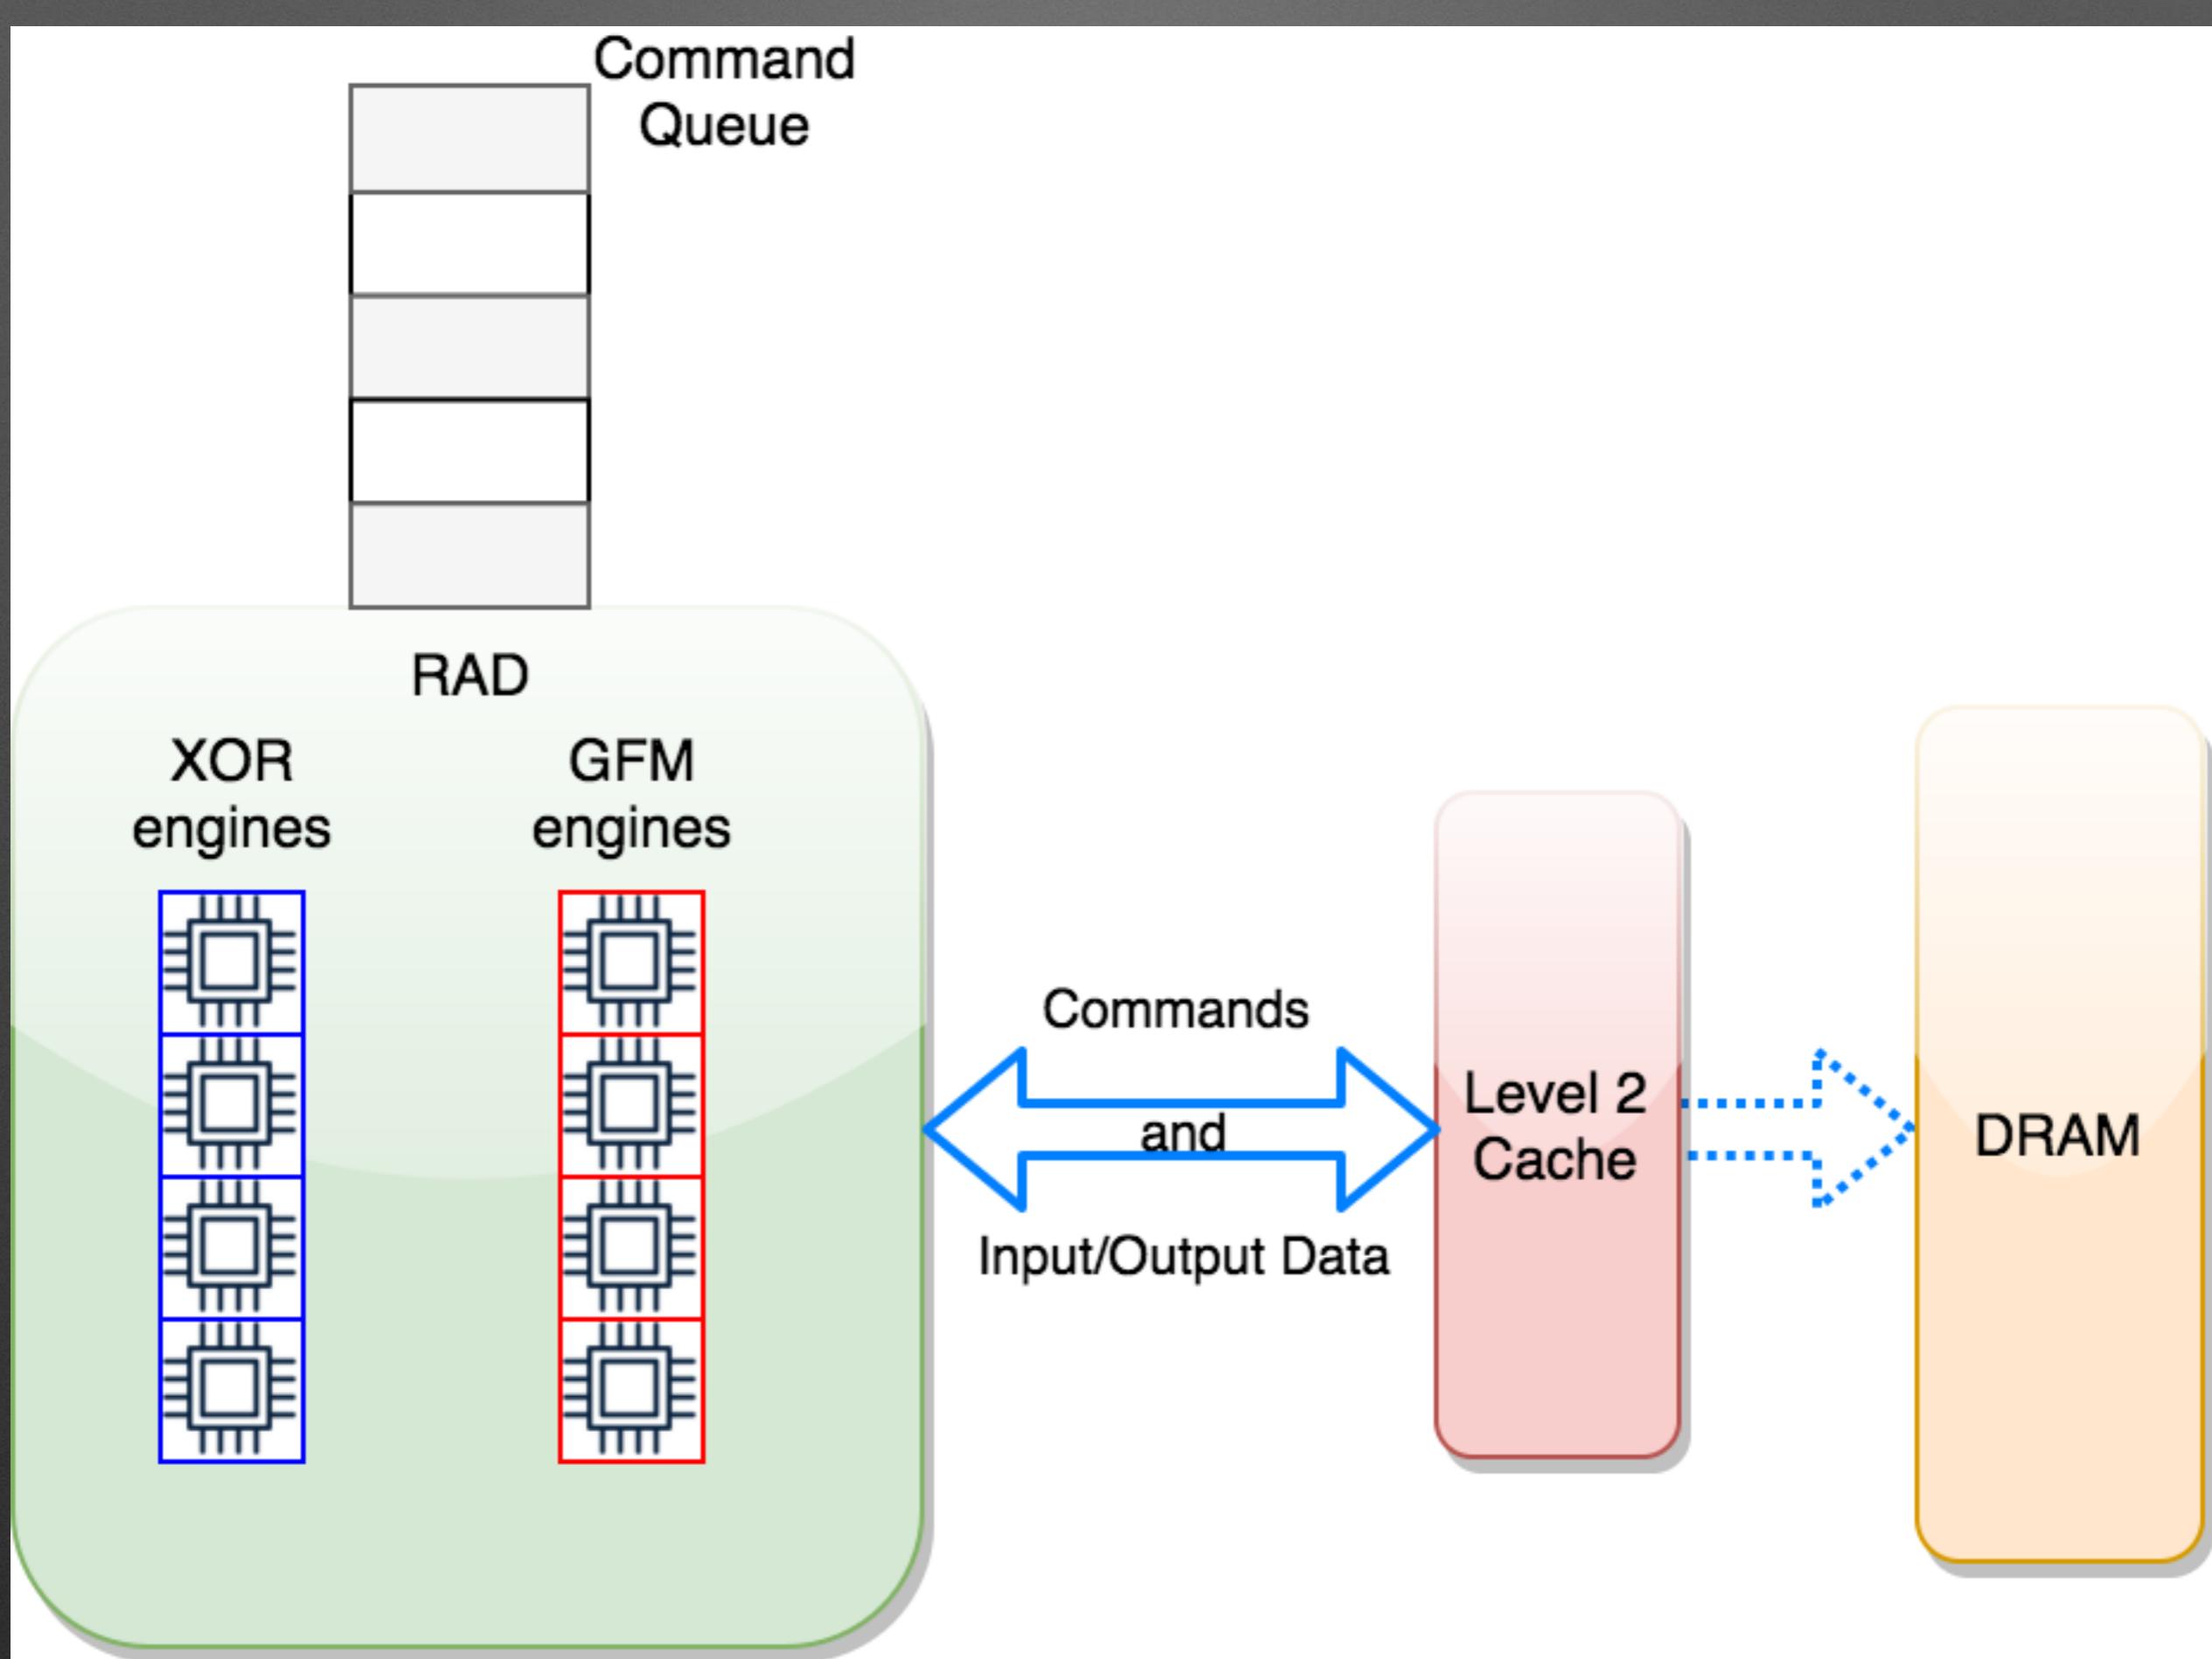


# RAID Architecture

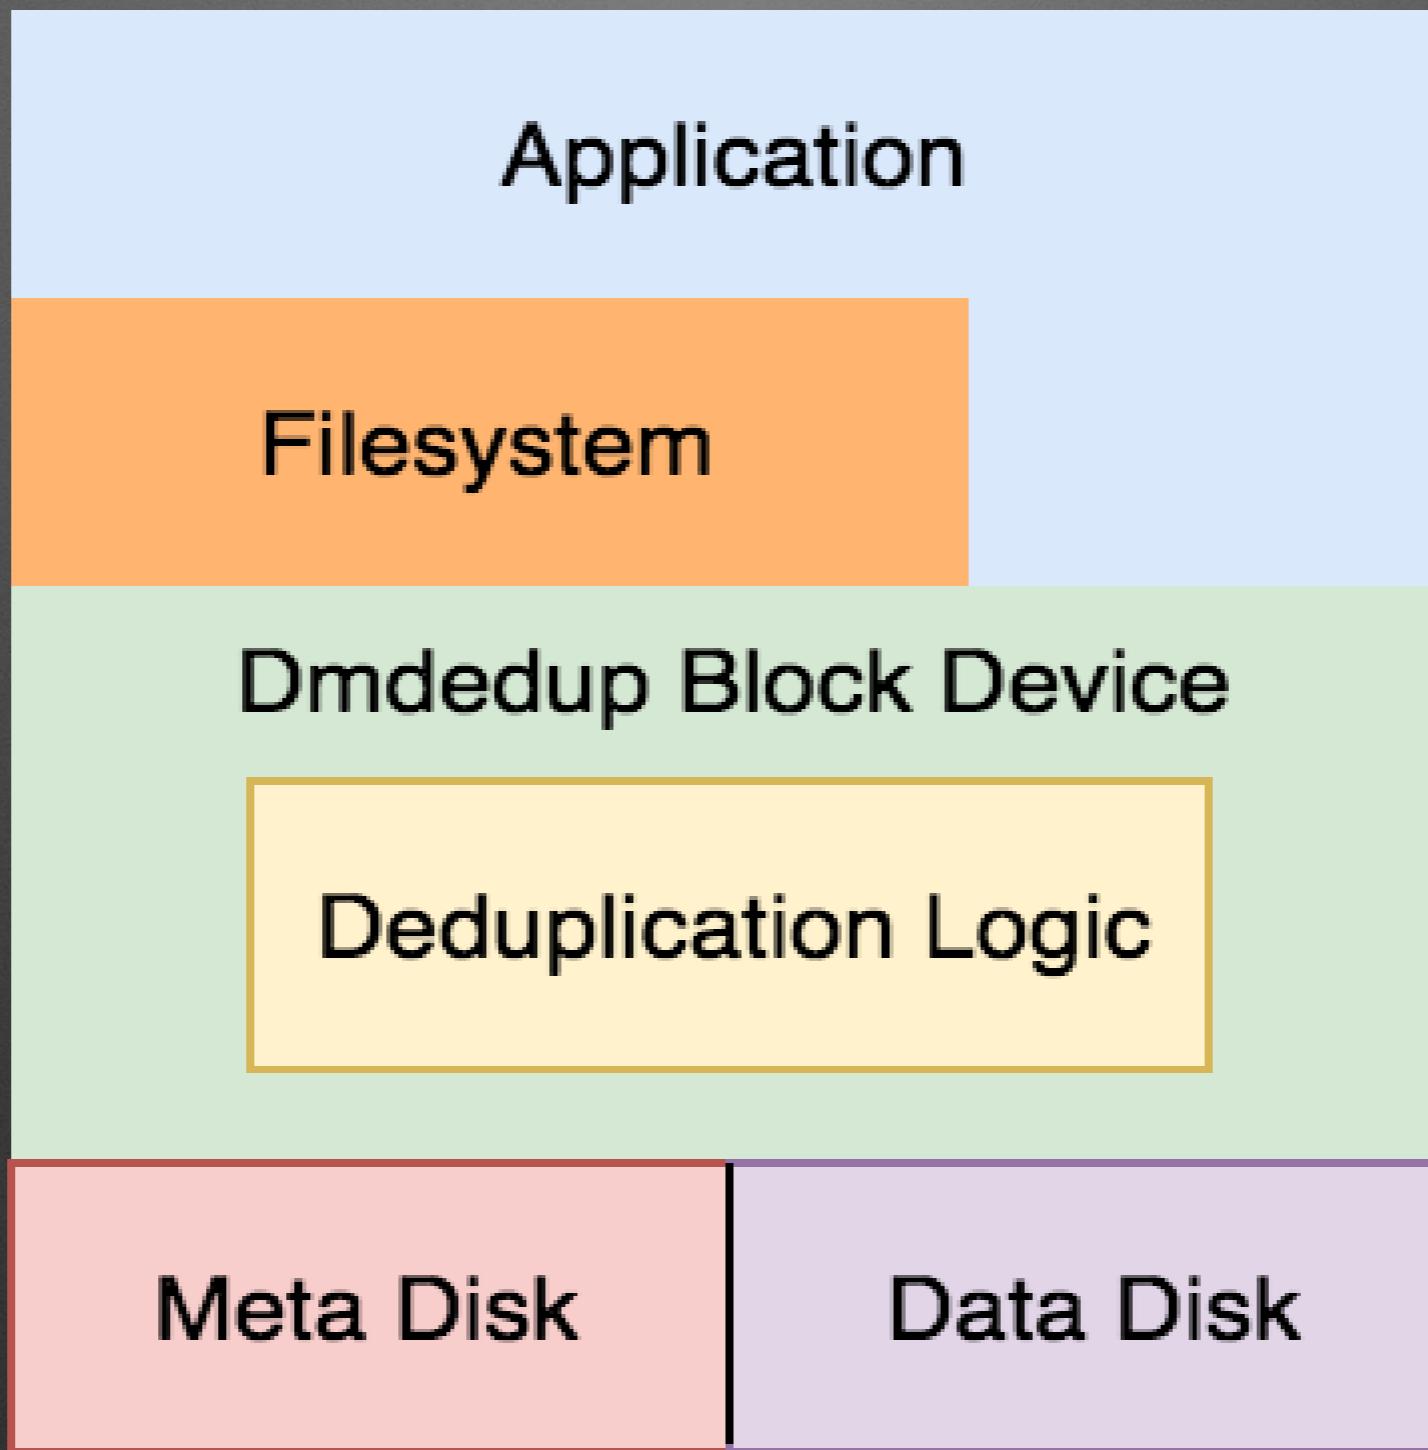


# Linux MDRAID Offload

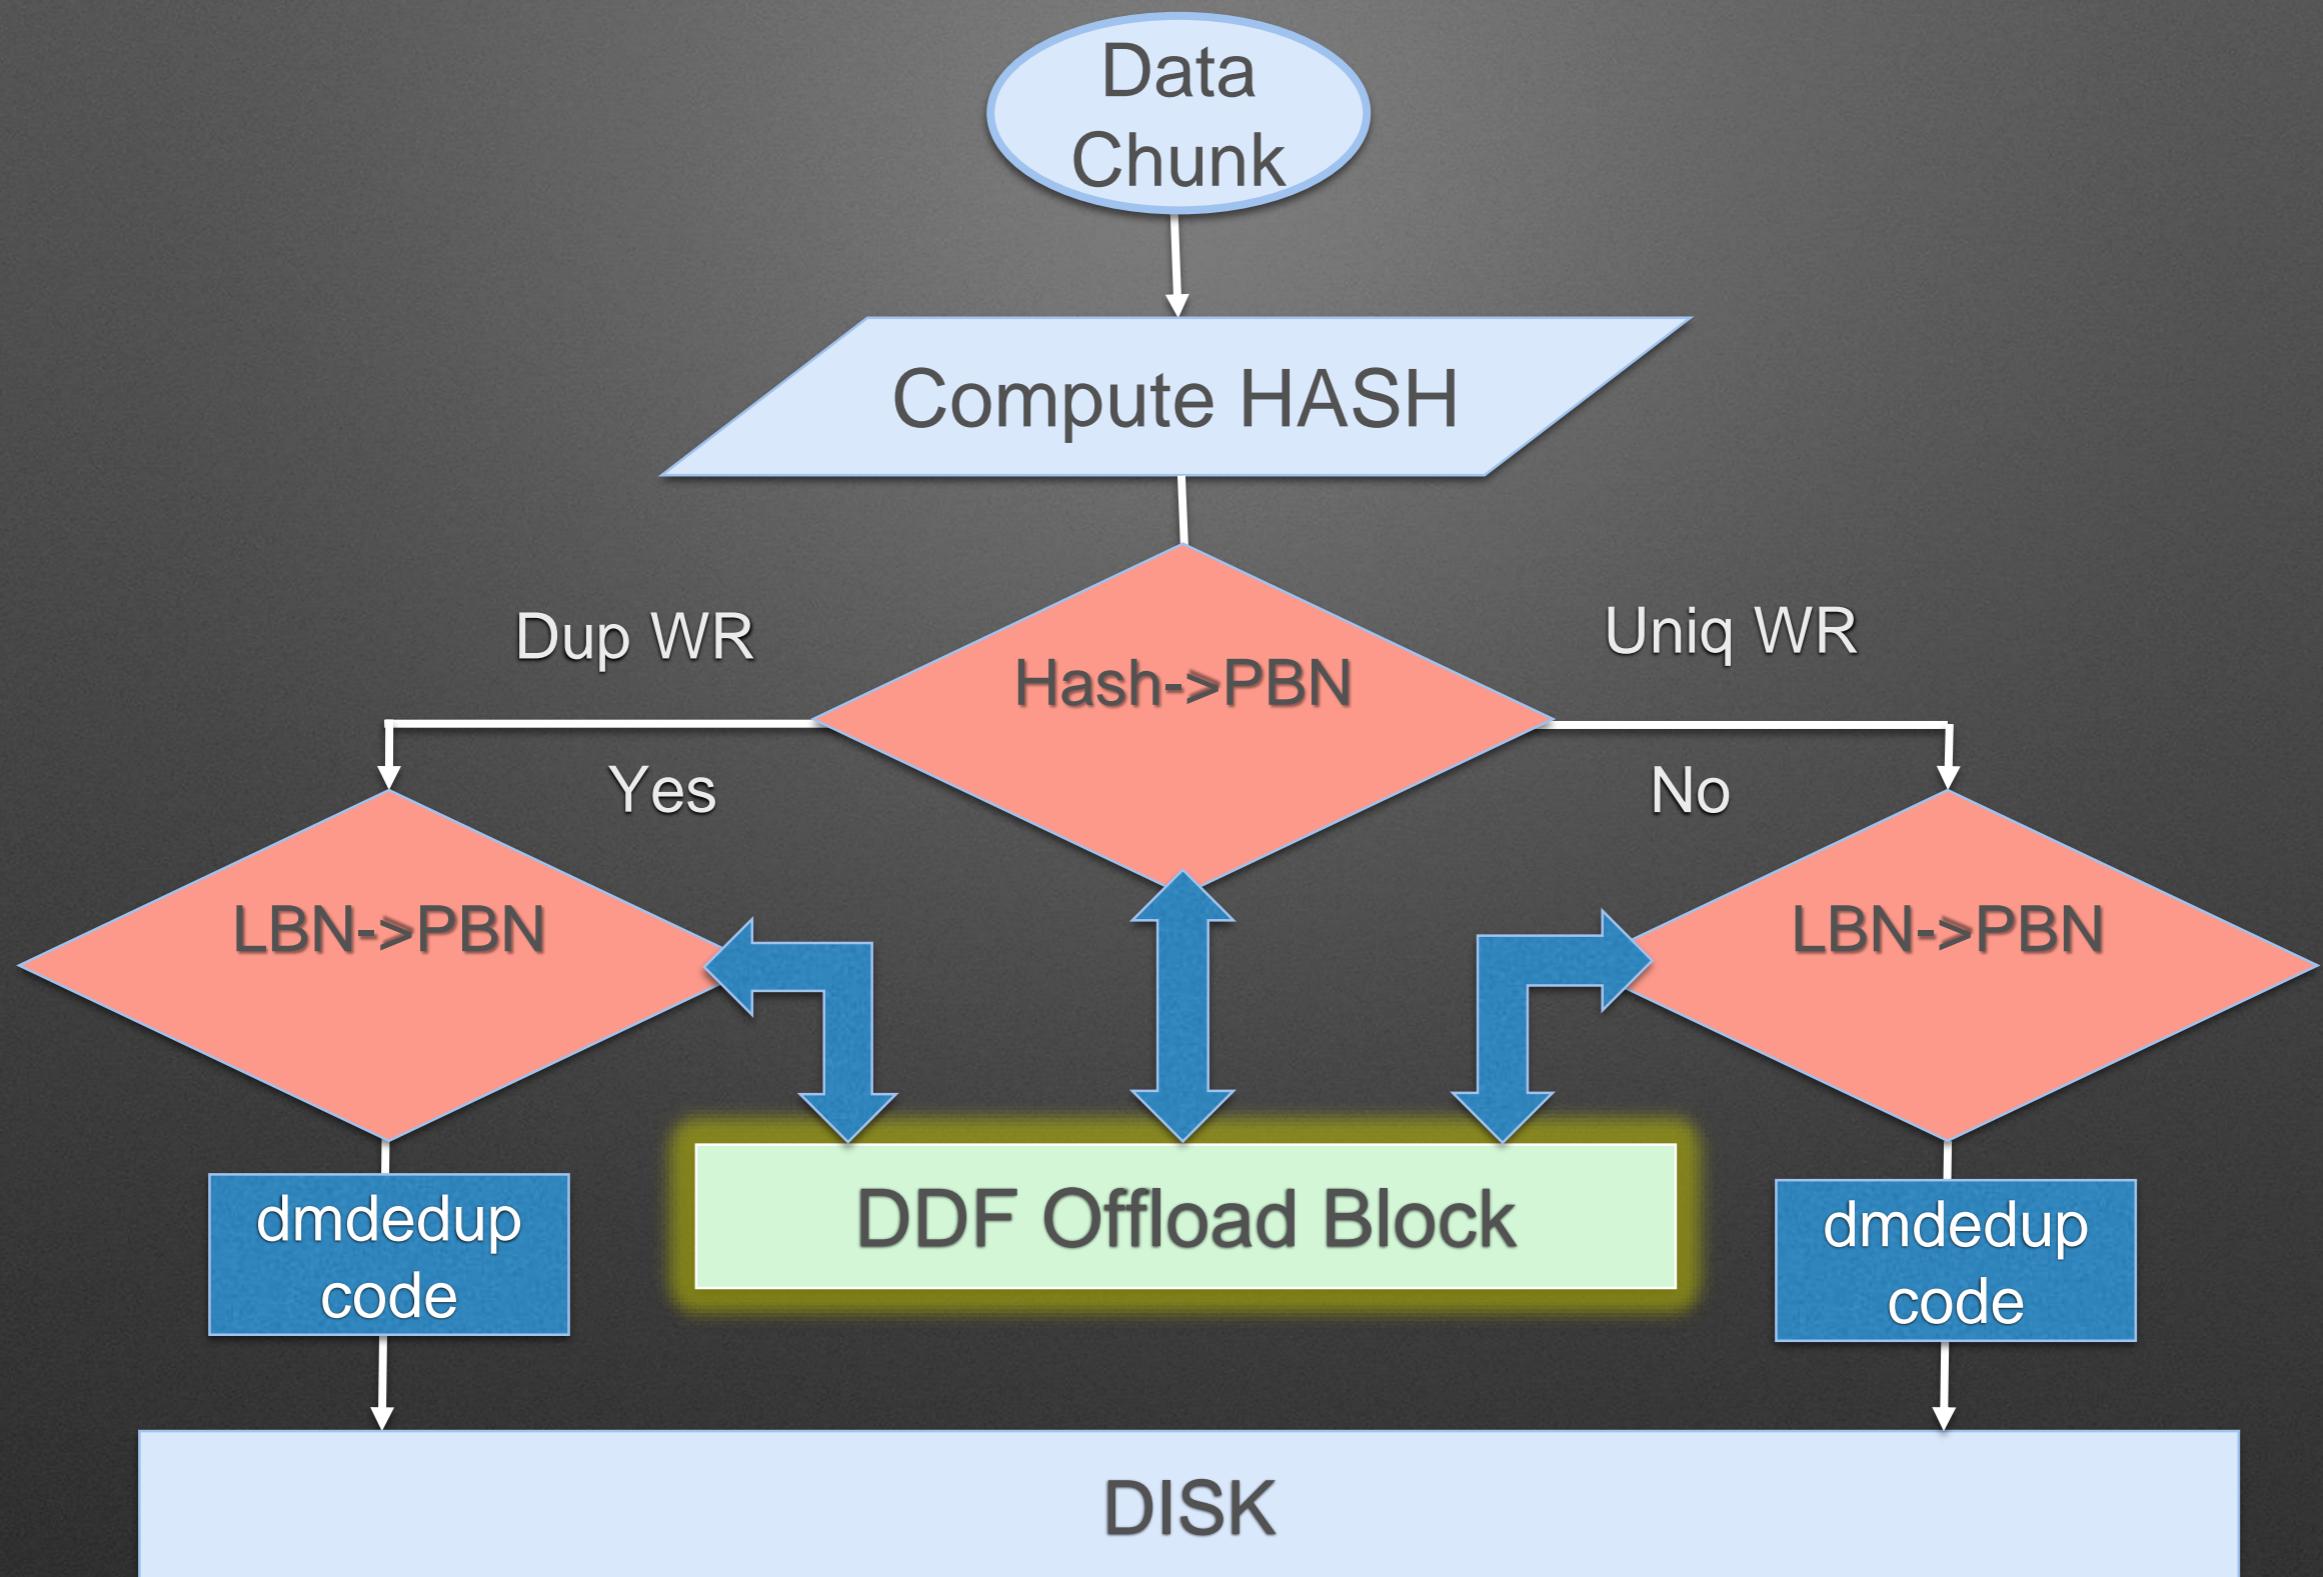




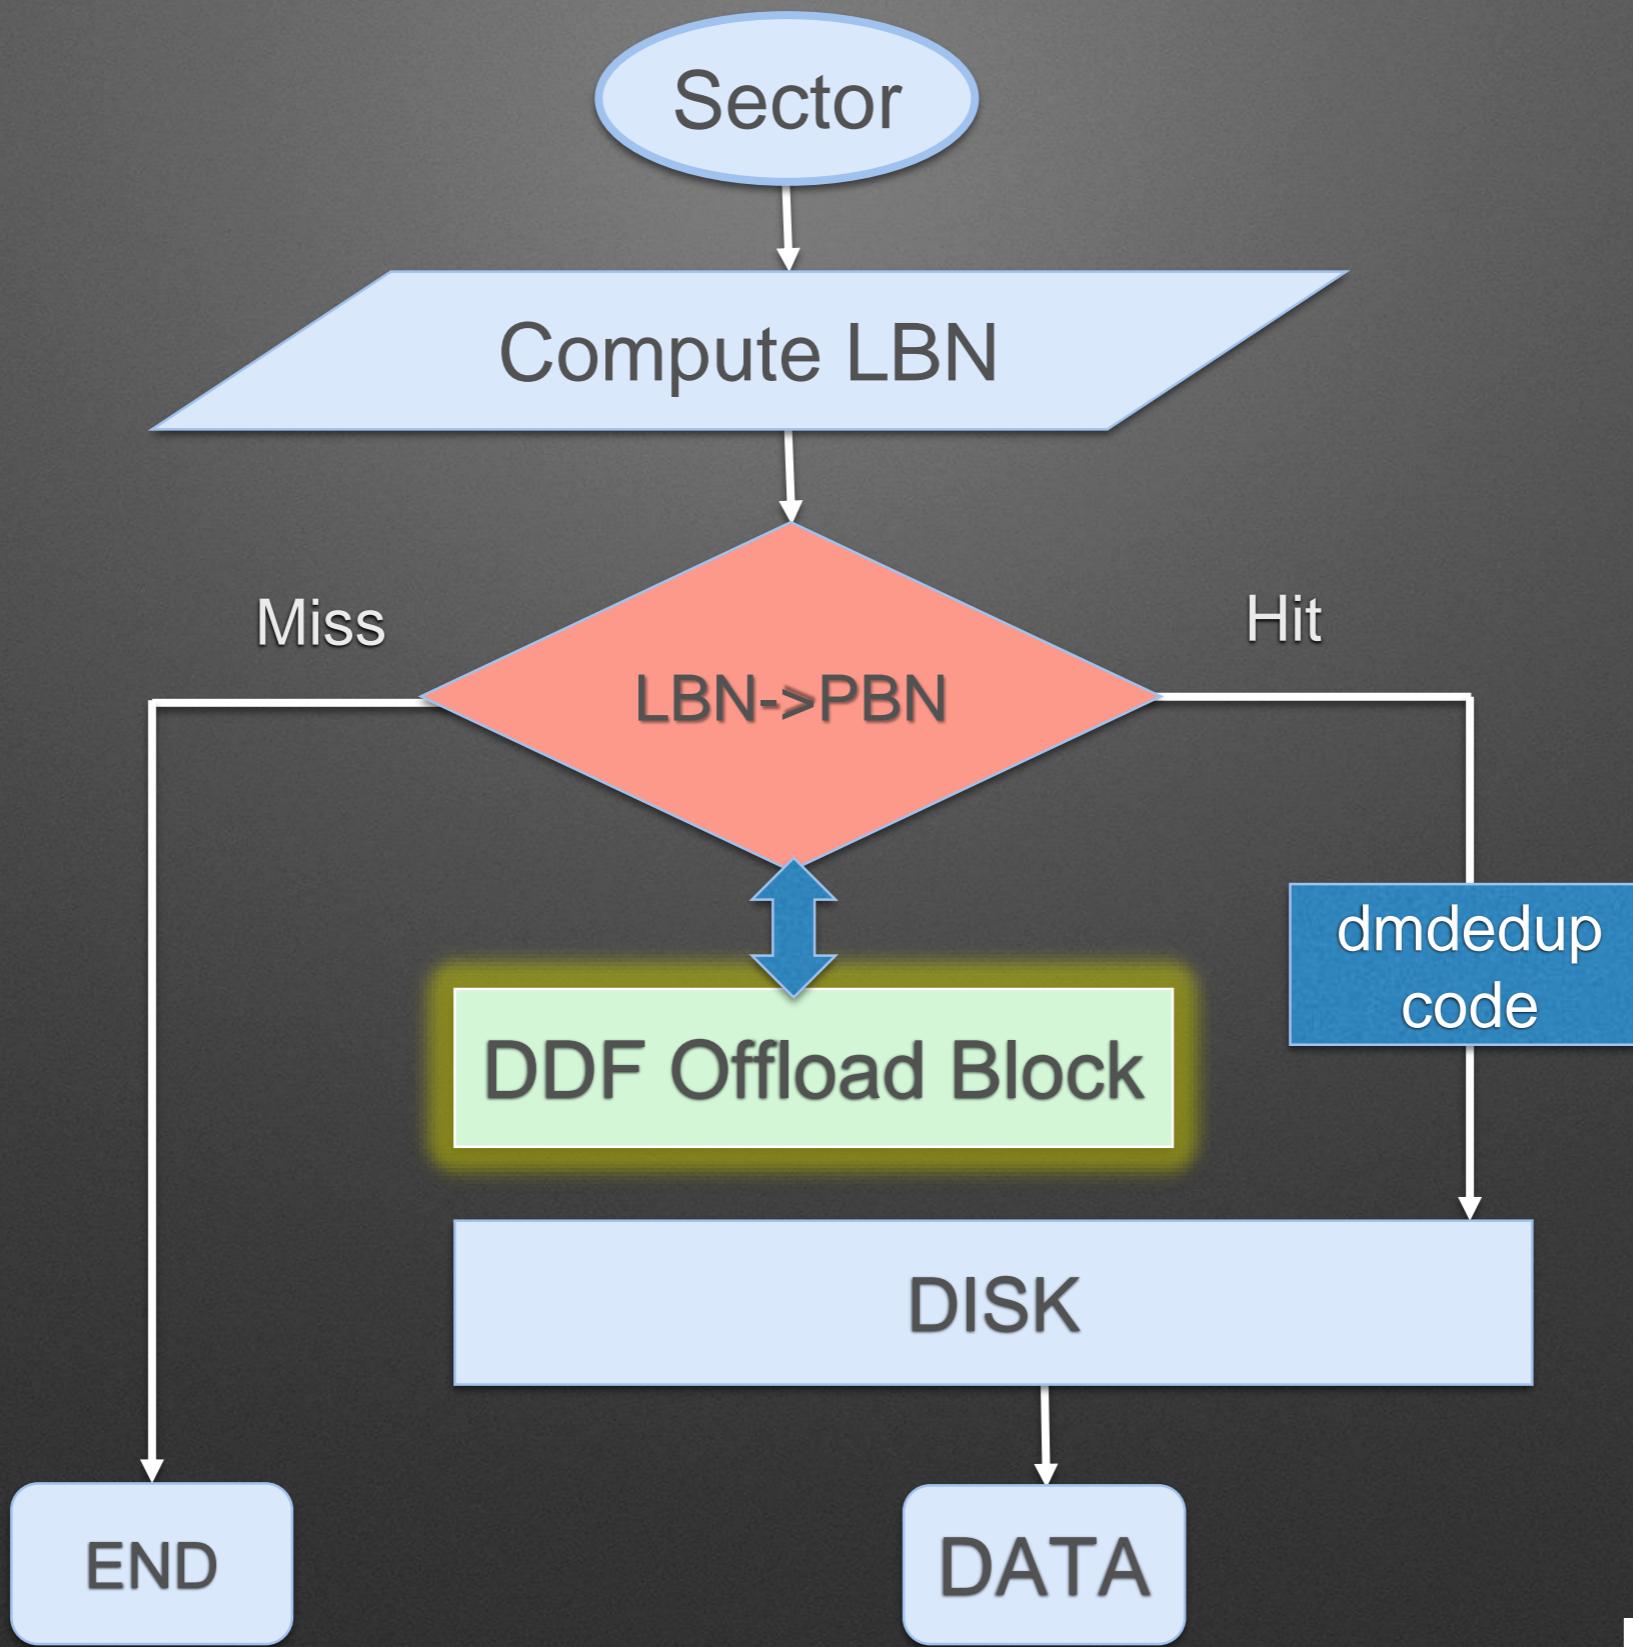
# Dmdedup Architecture

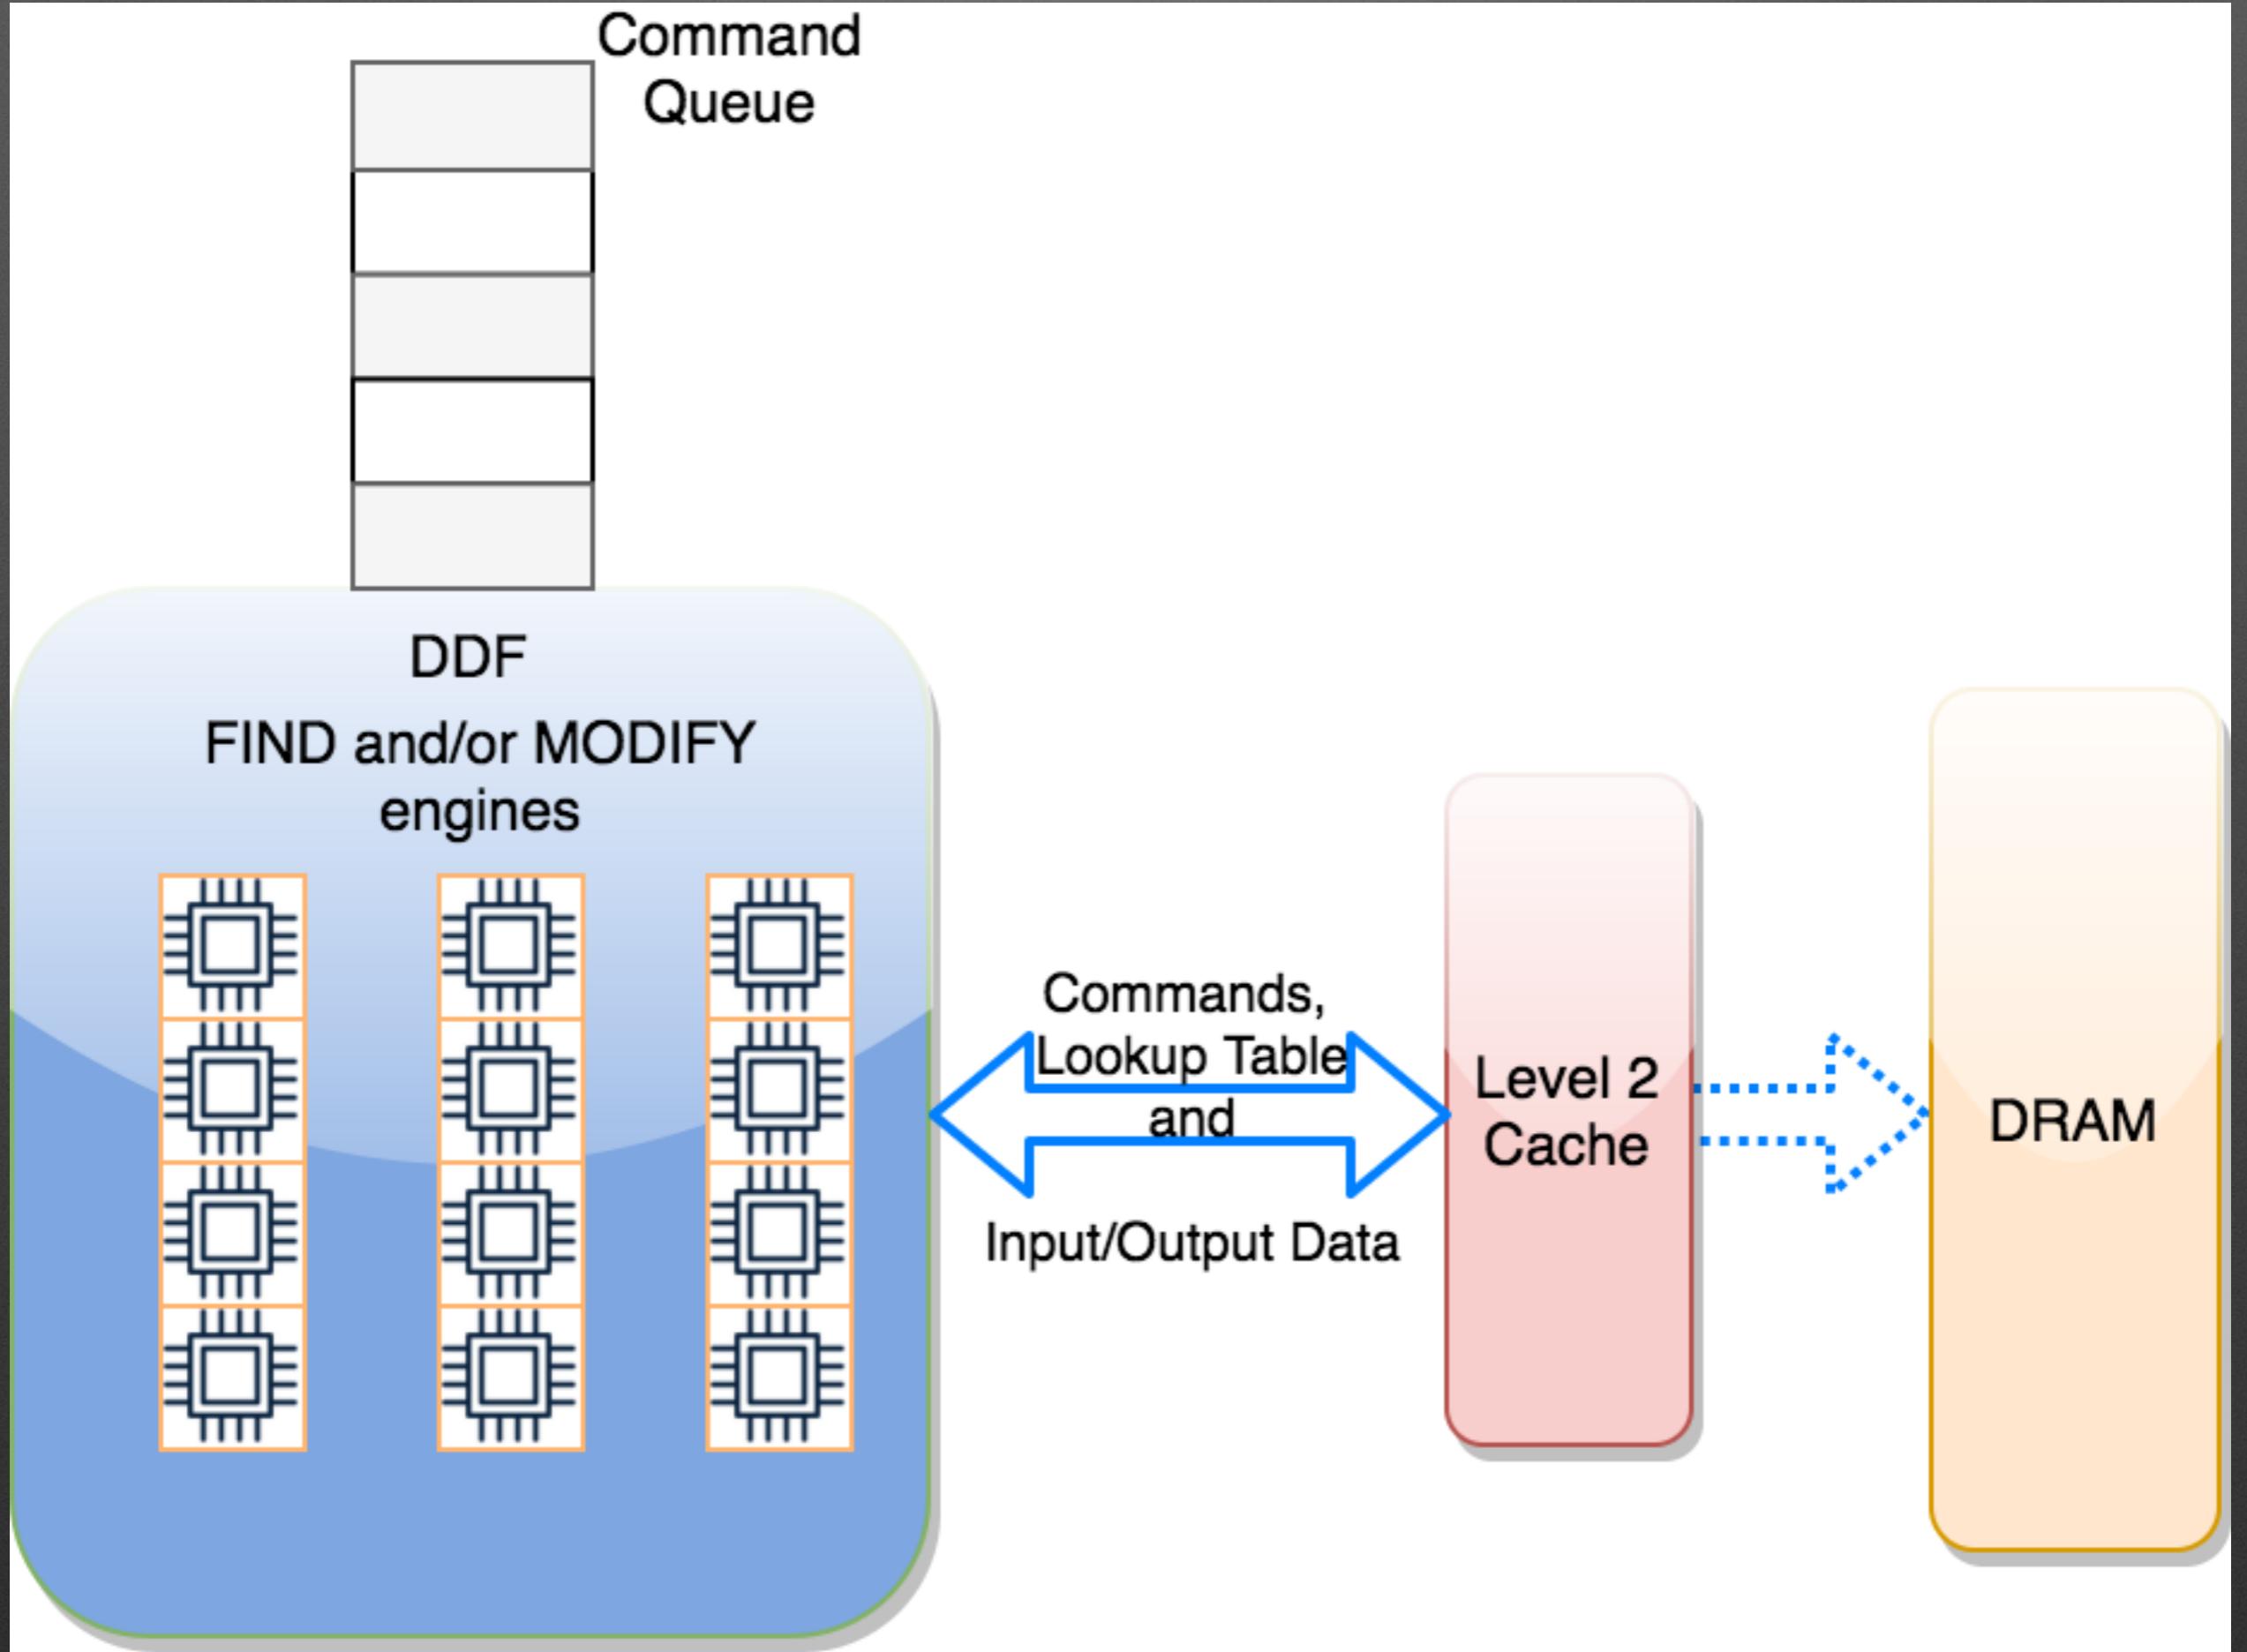


# DmDedup WRITE Offload

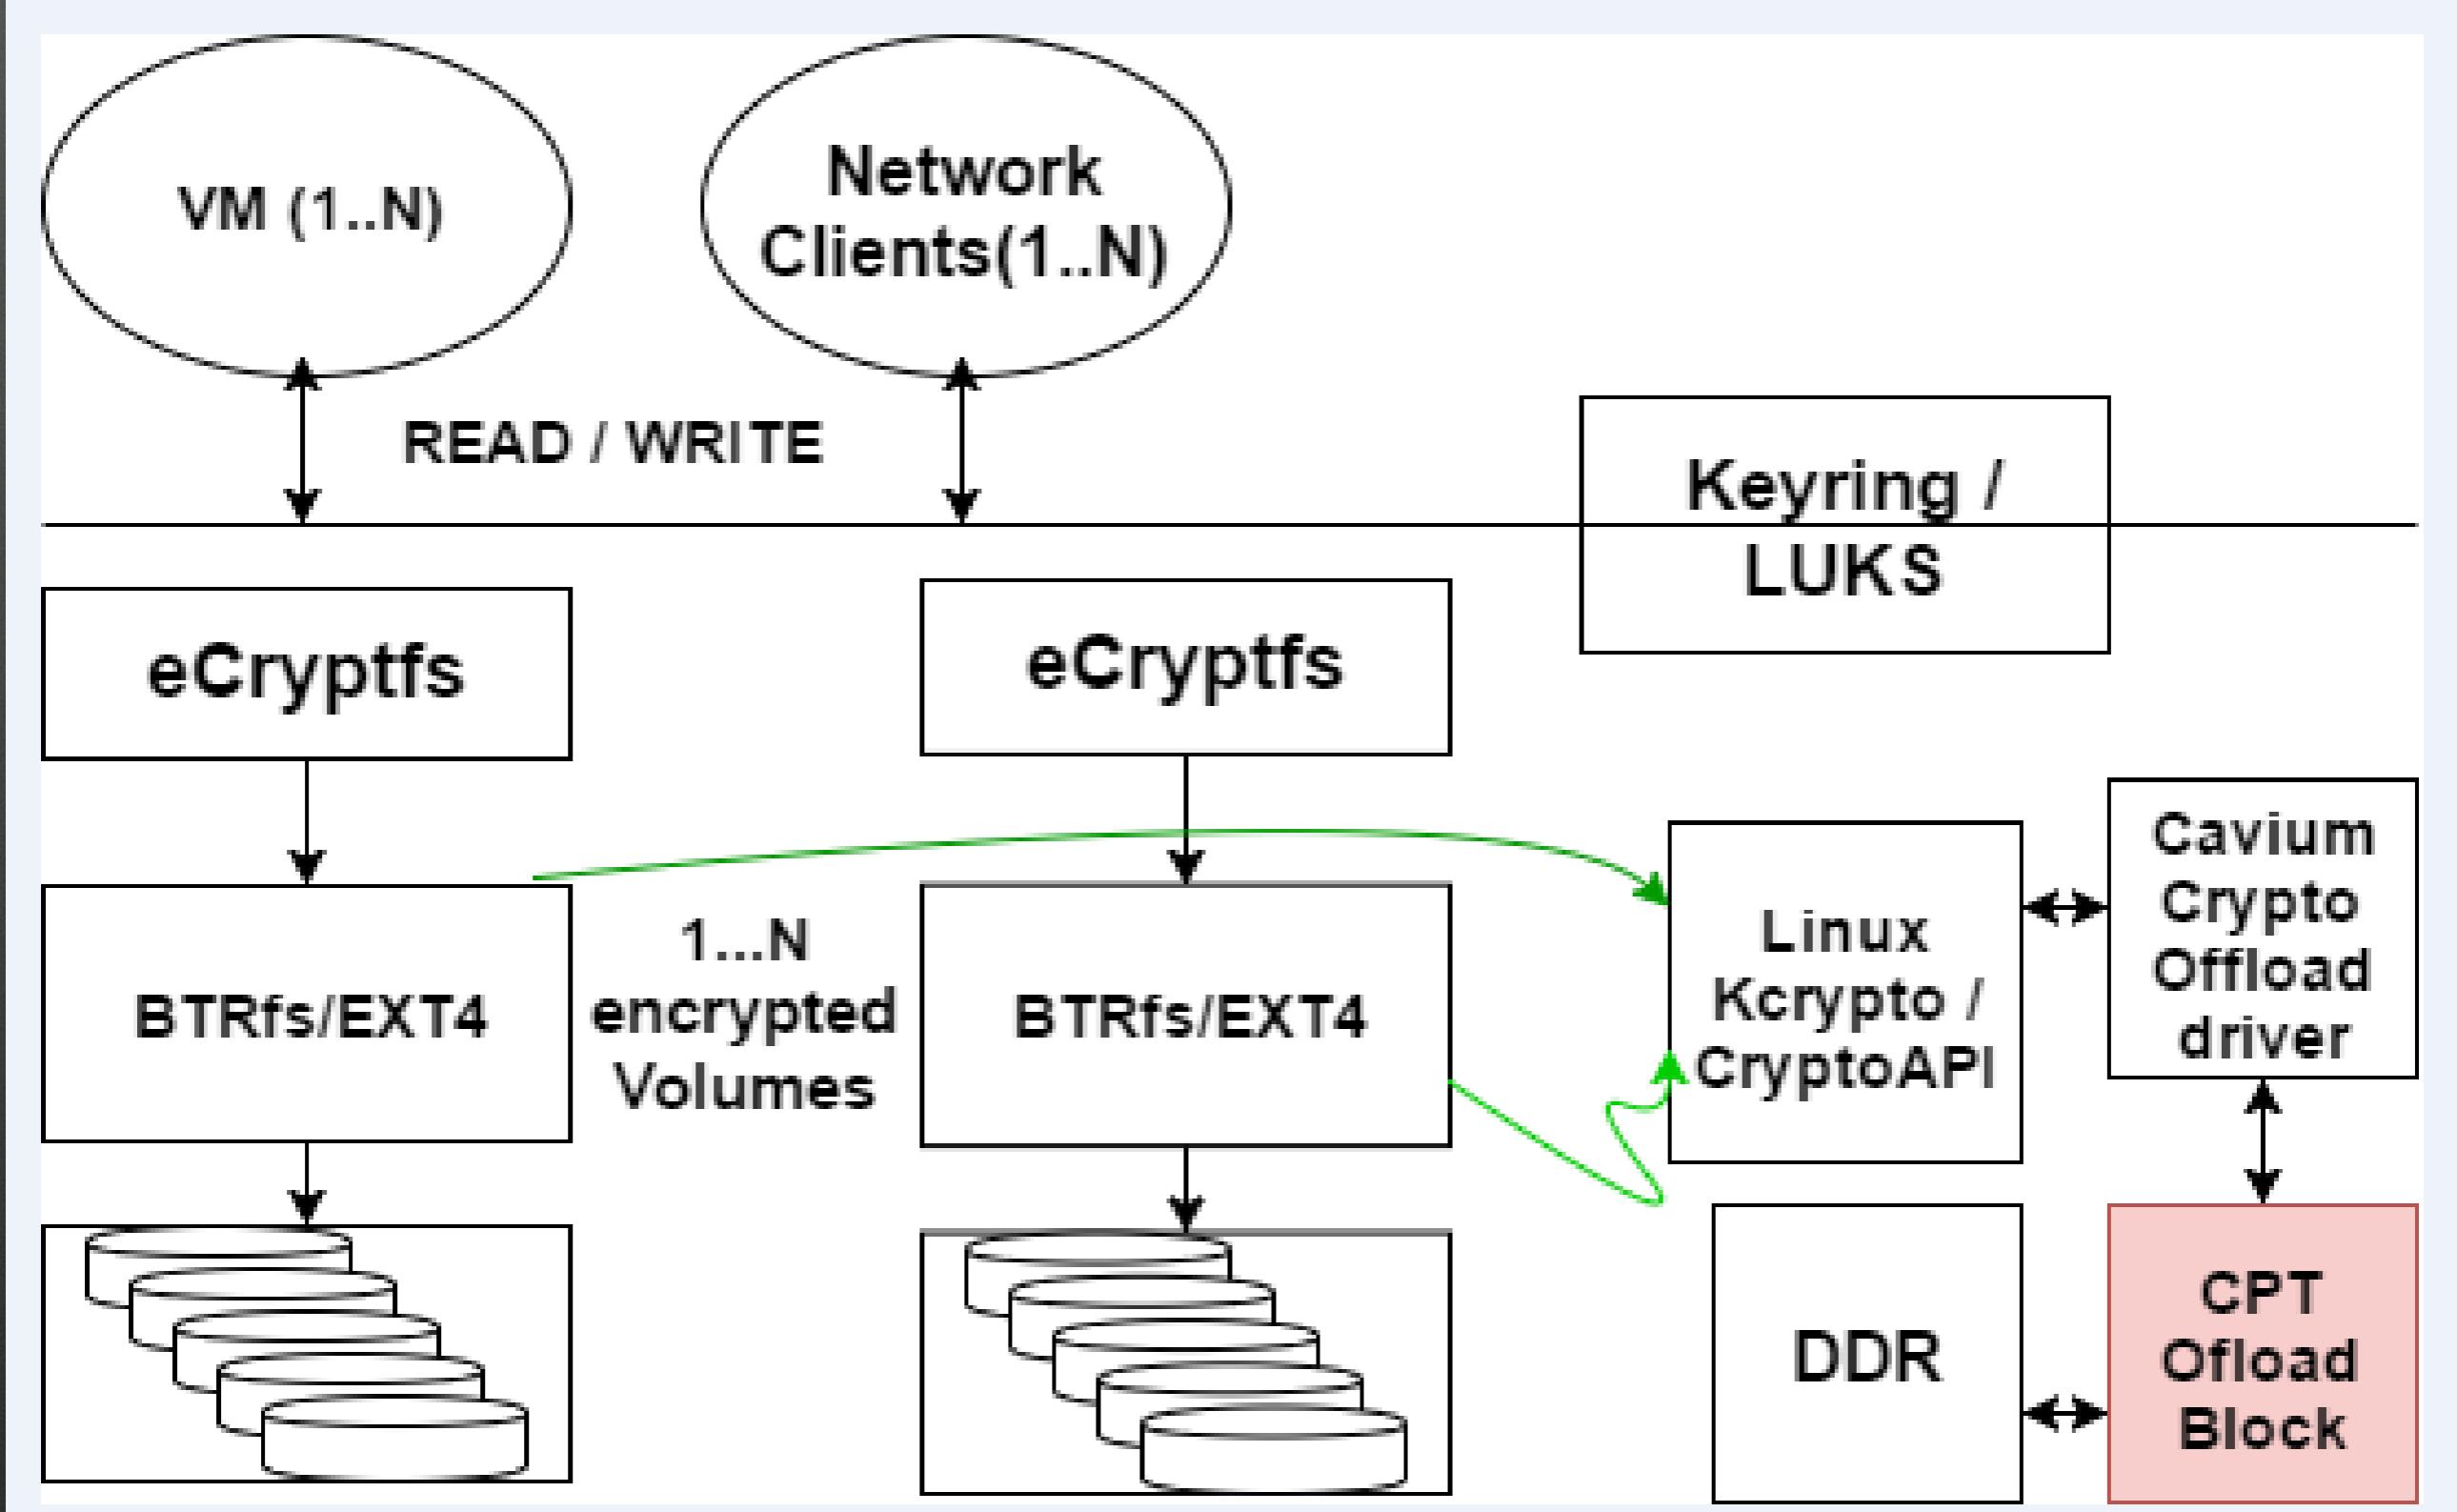


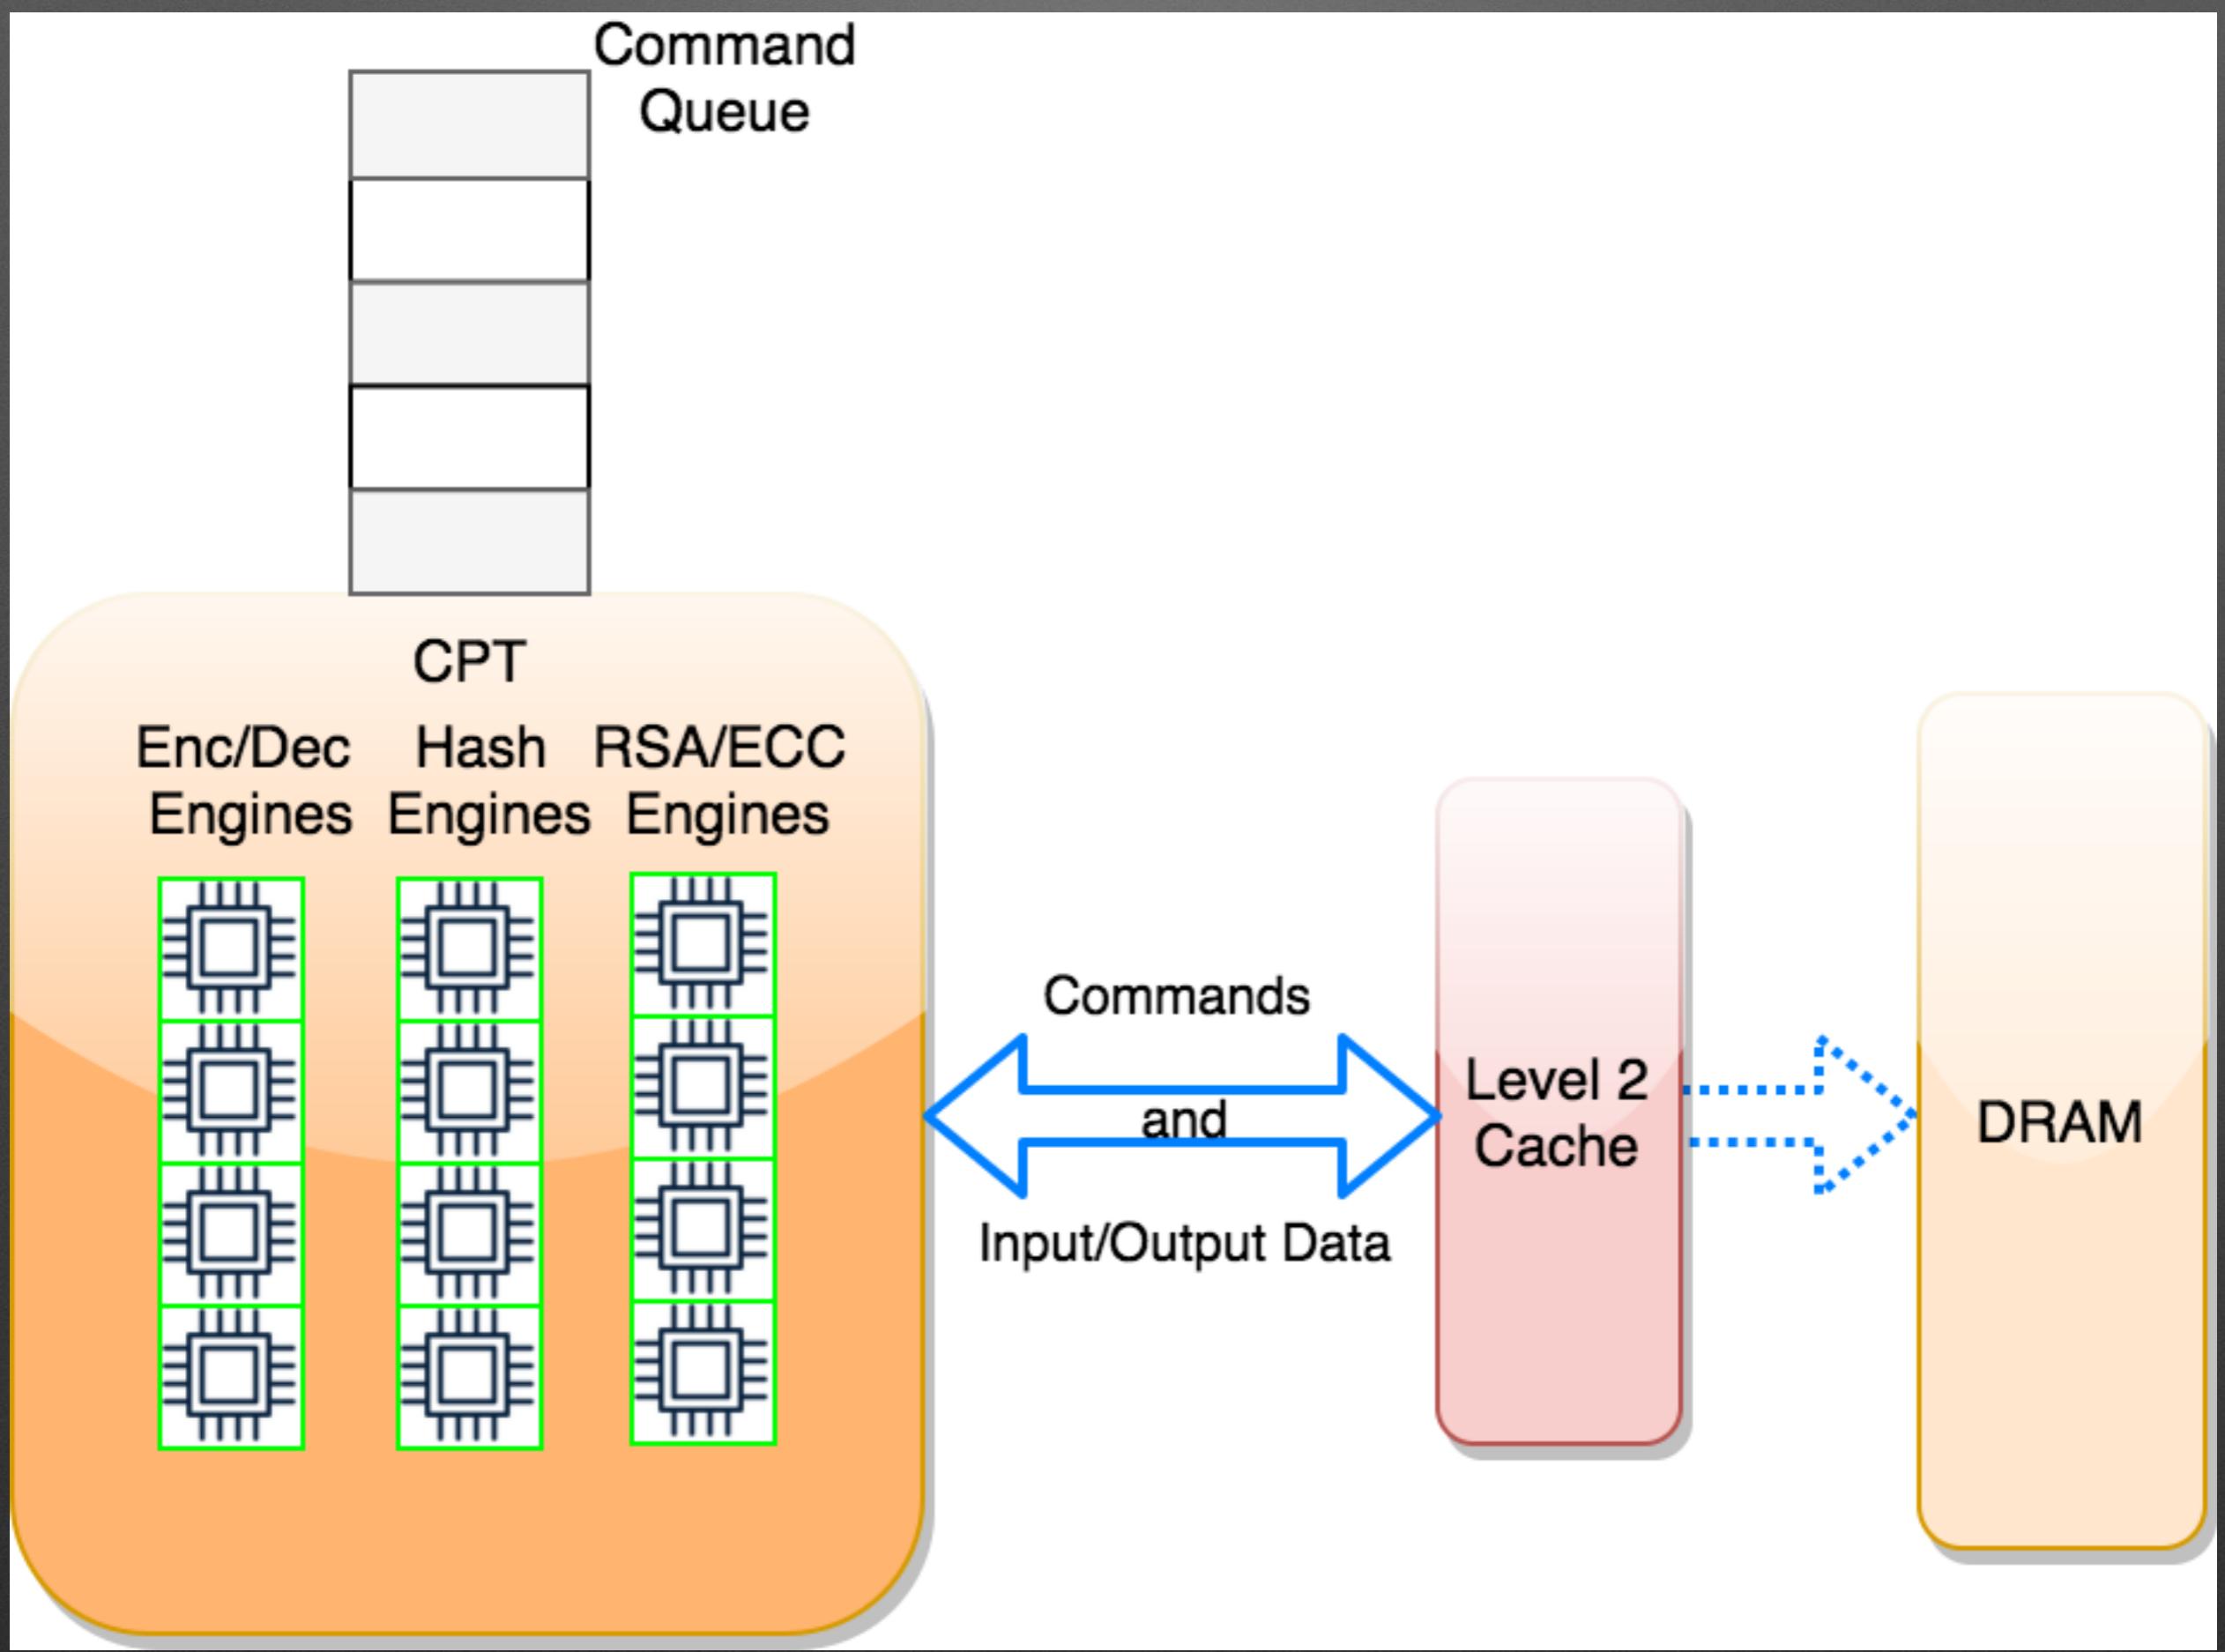
# DmDedup READ Offload





# Storage Security Offload





# Accelerations Integrated

Application1

BTRFS+DMDEDUP+RAID6

DDF

RAD

Application2

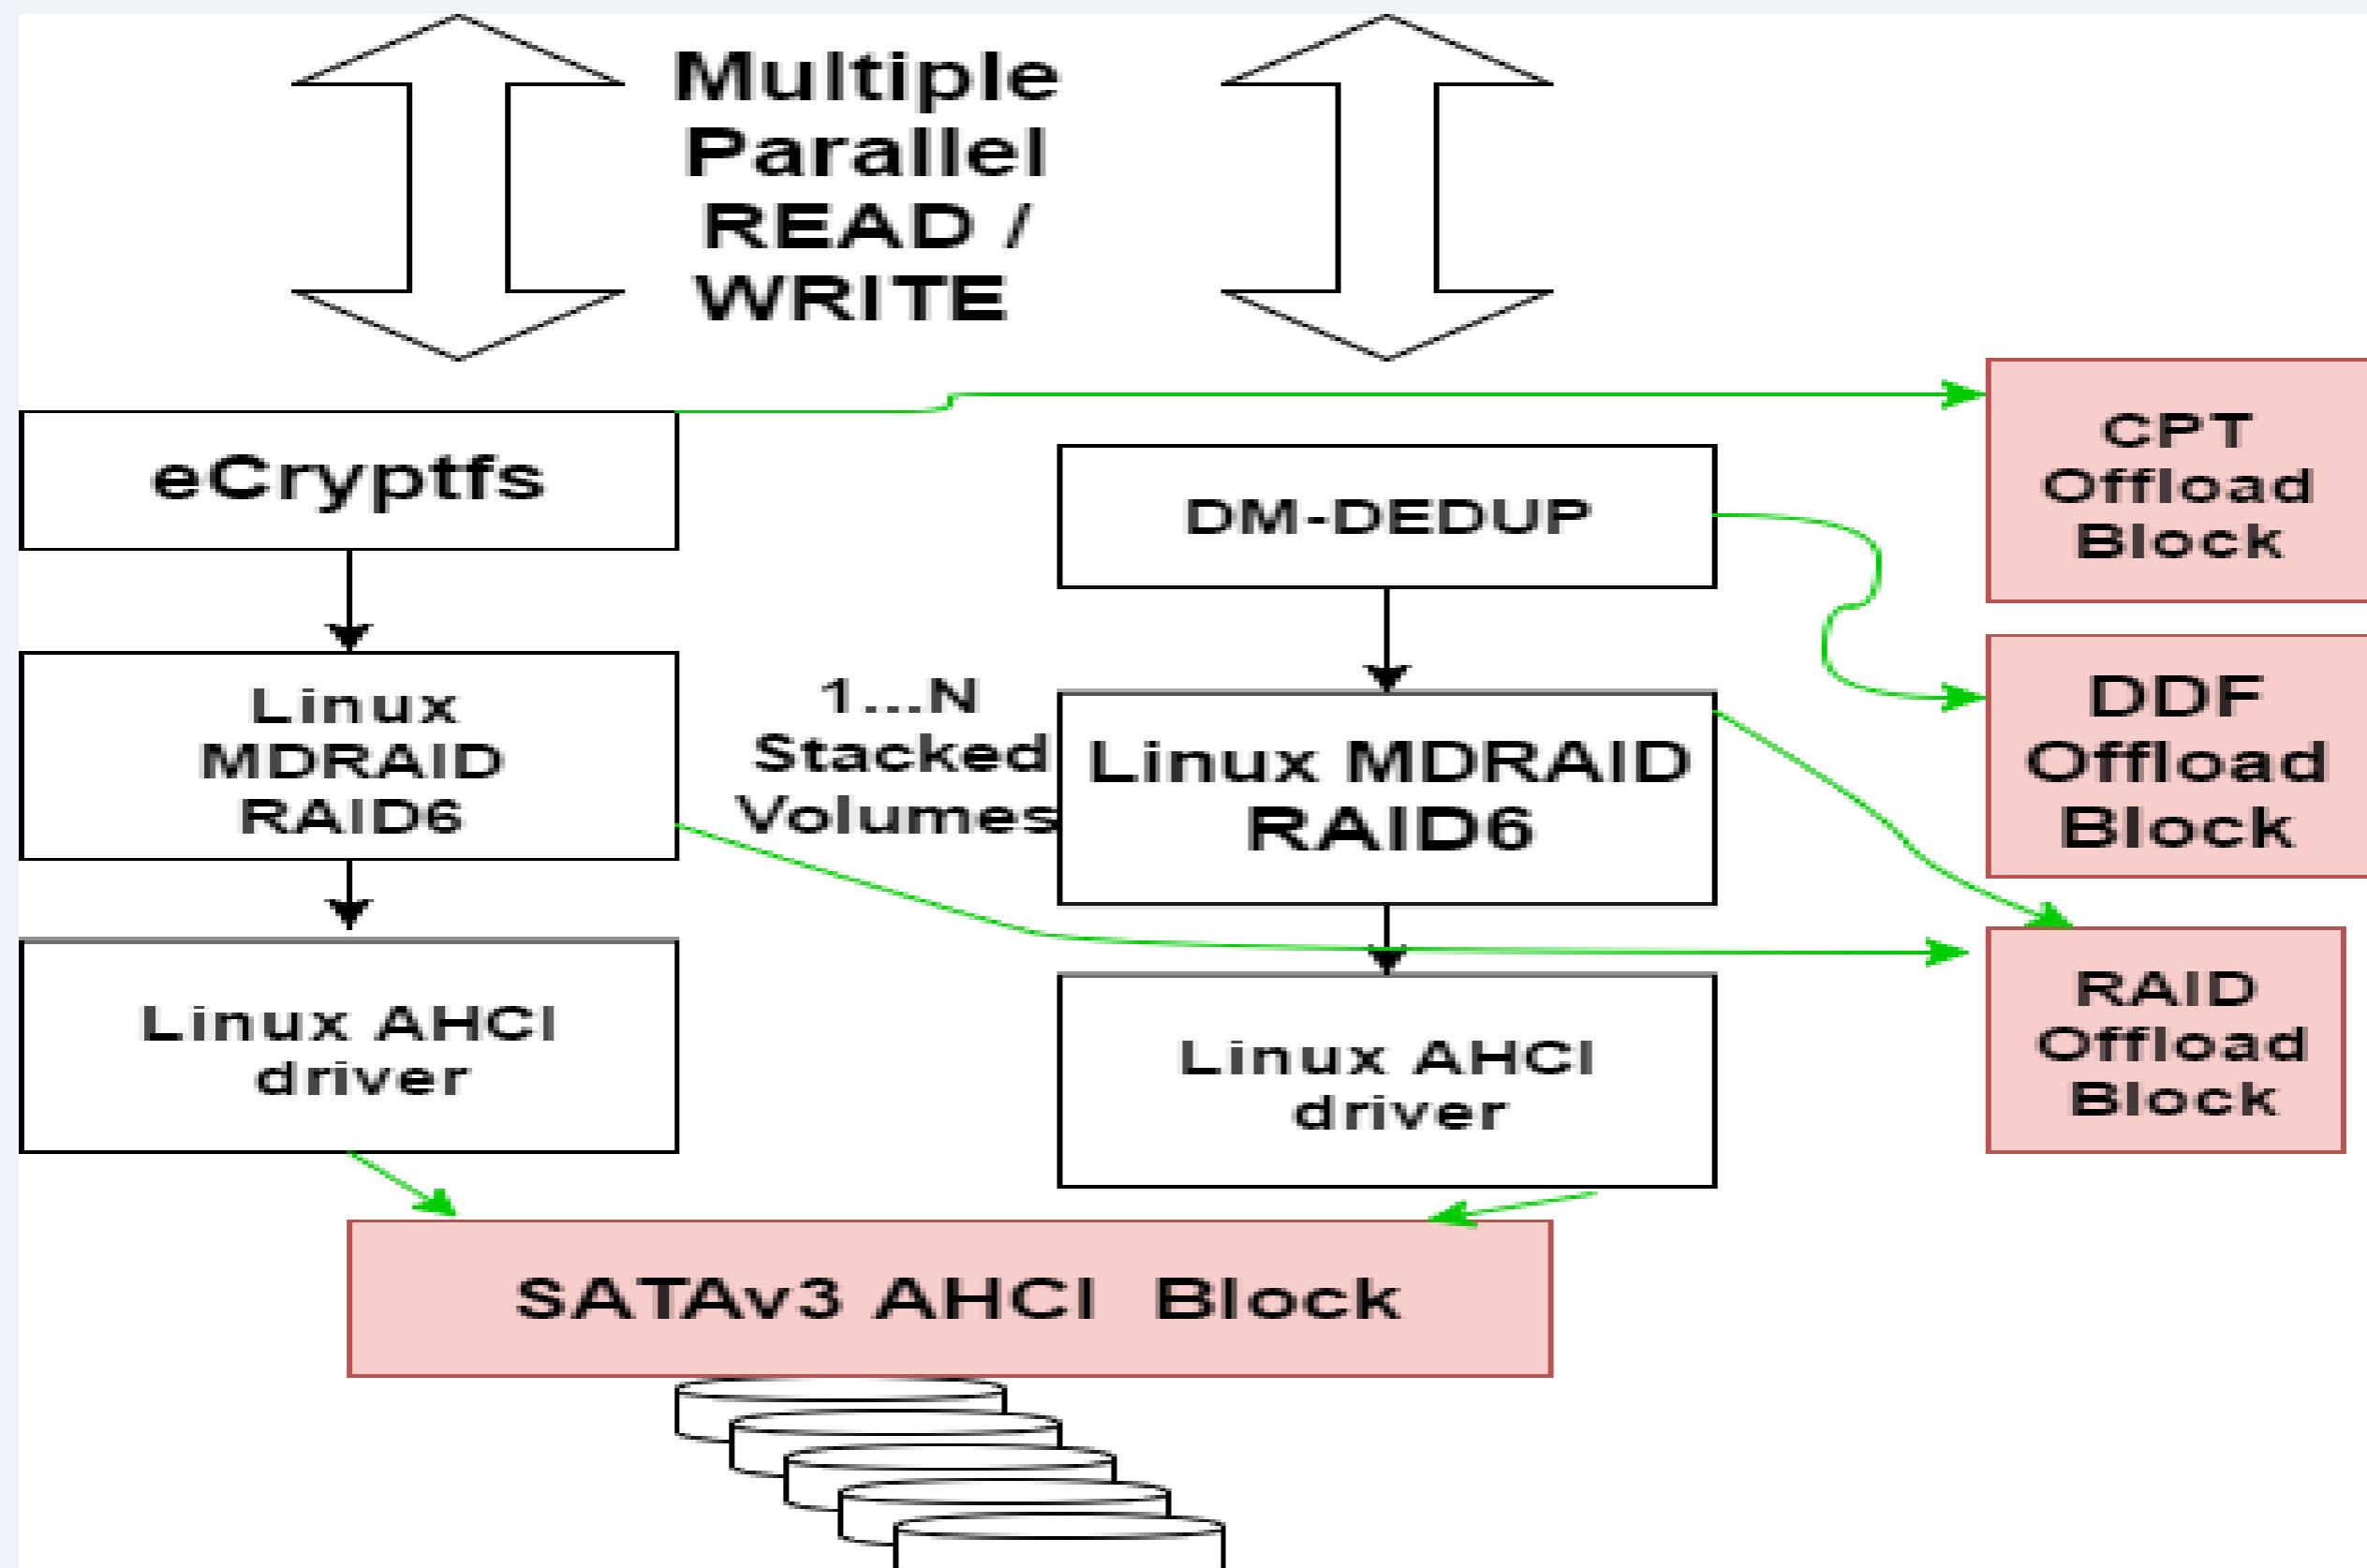
ECRYPTFS+EXT4+RAID6

CPT

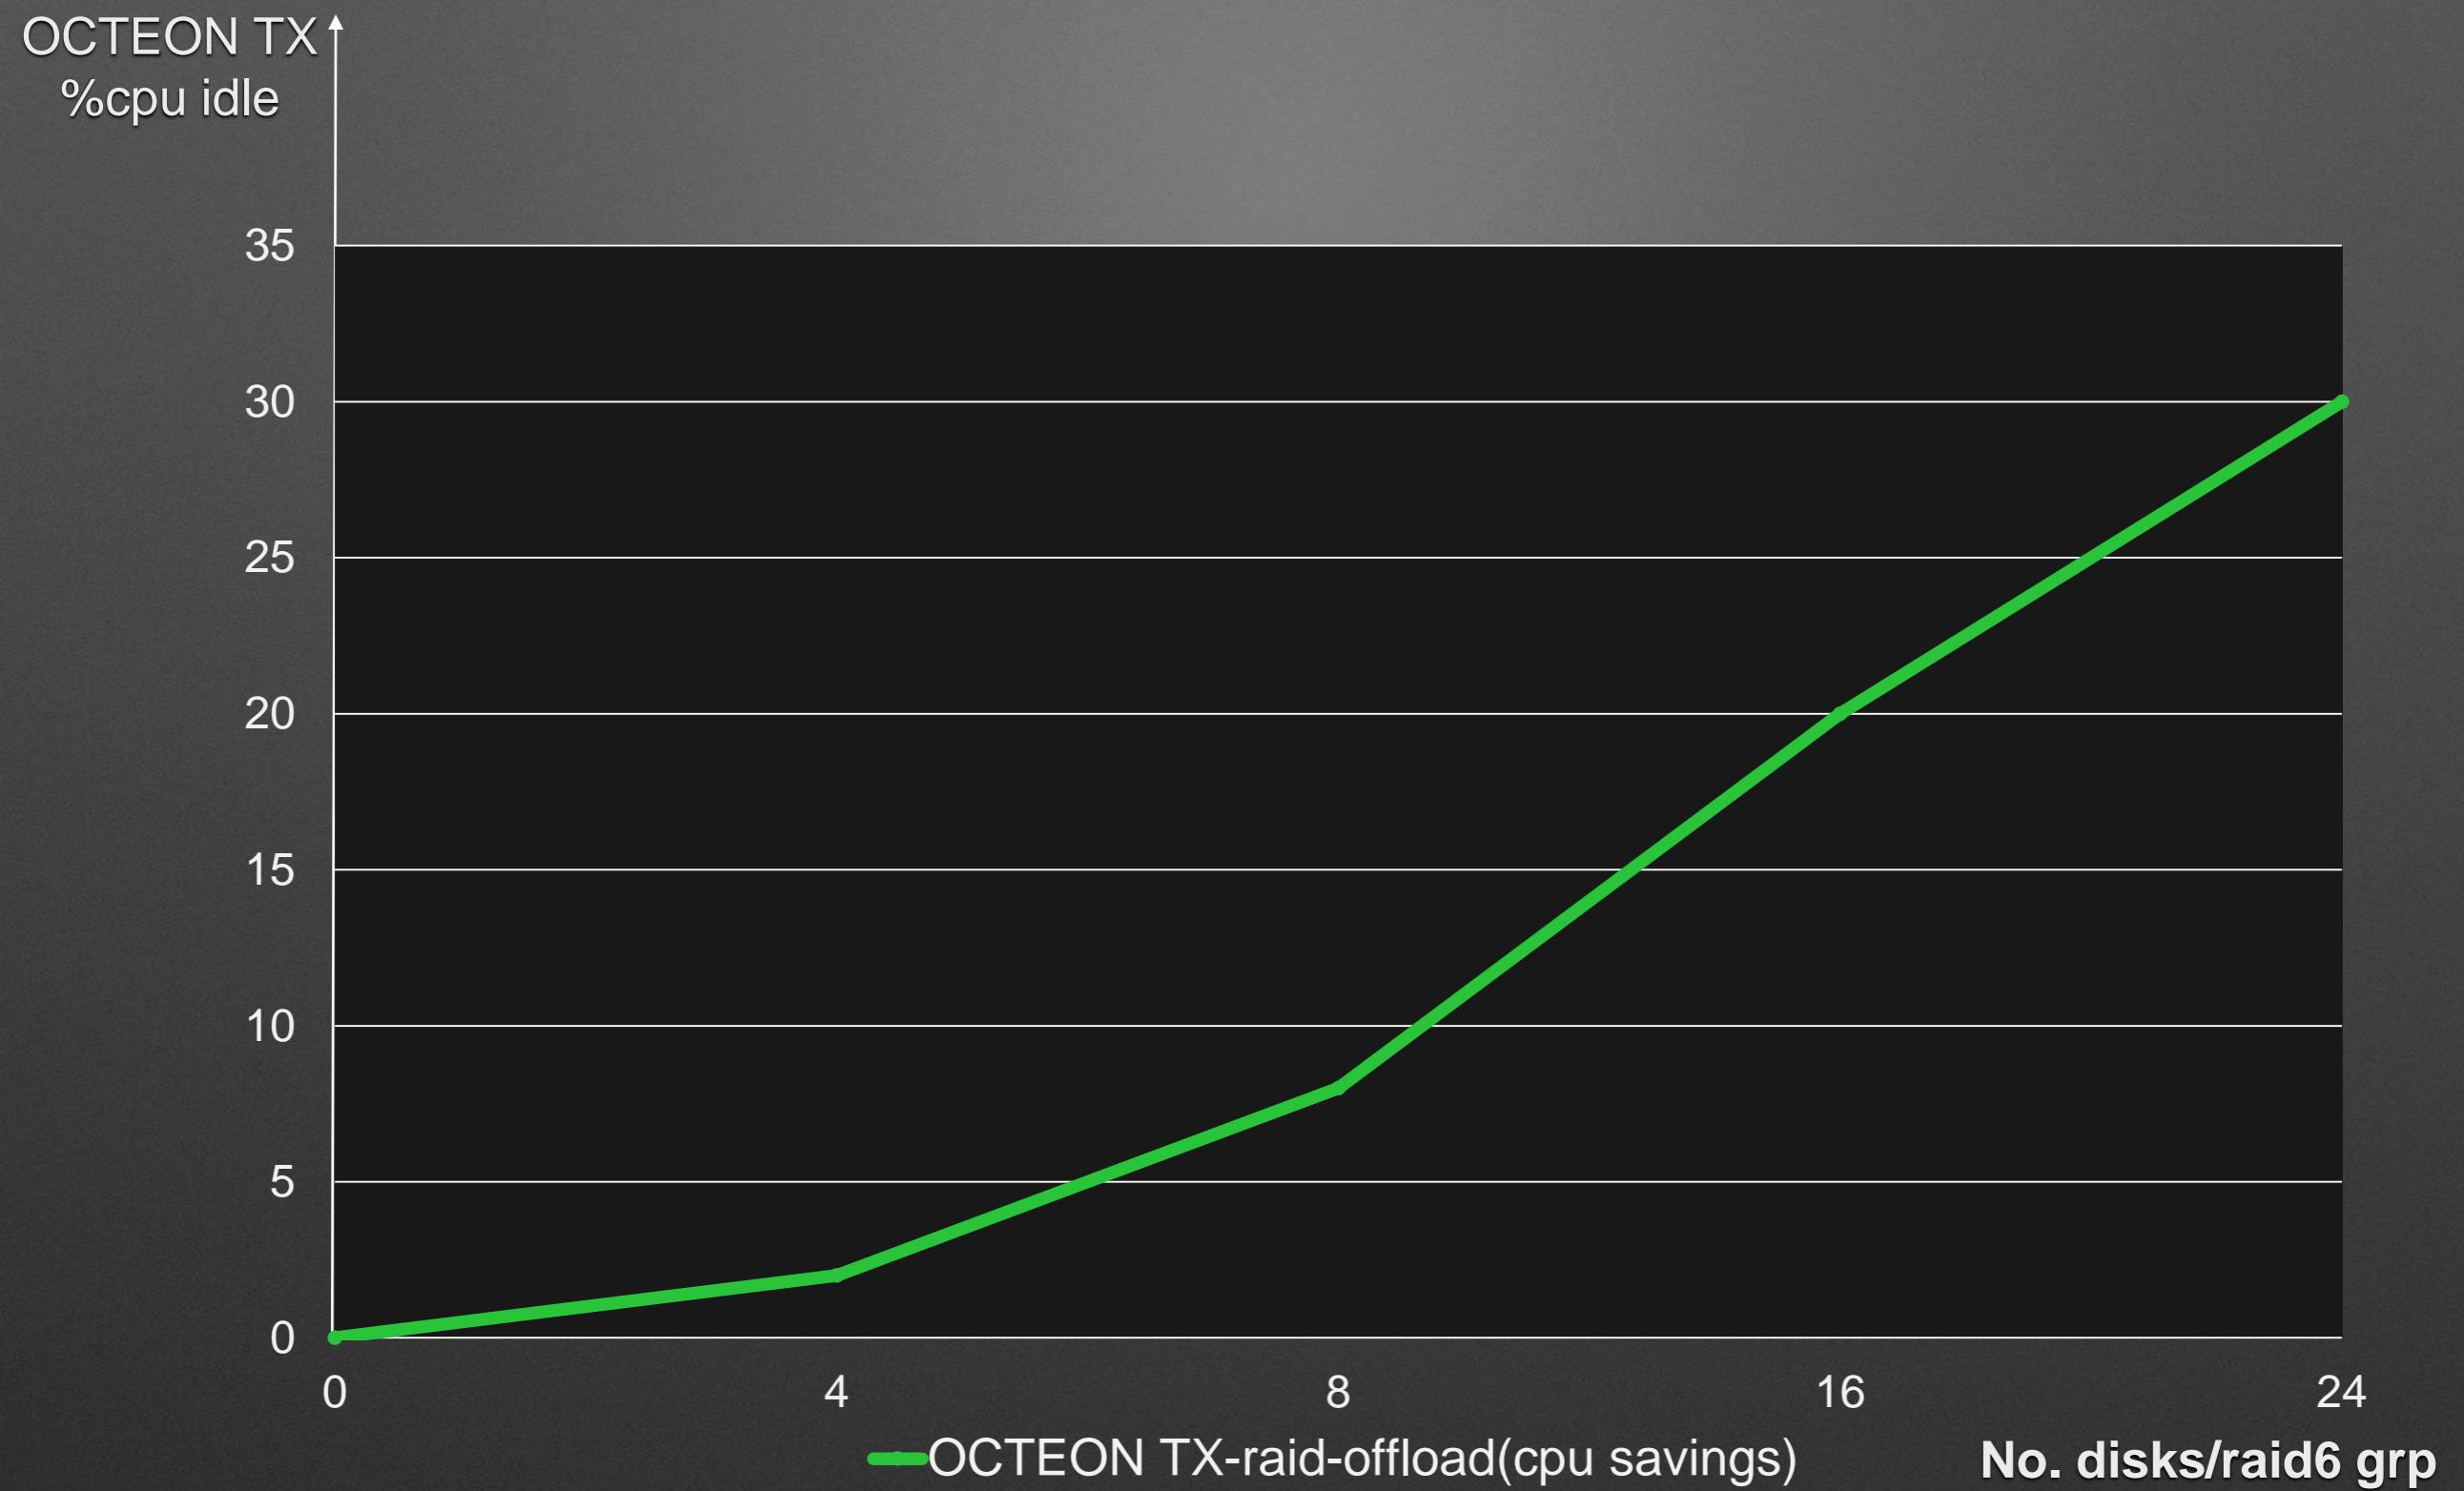
RAD

D D D D P Q D D D D P Q

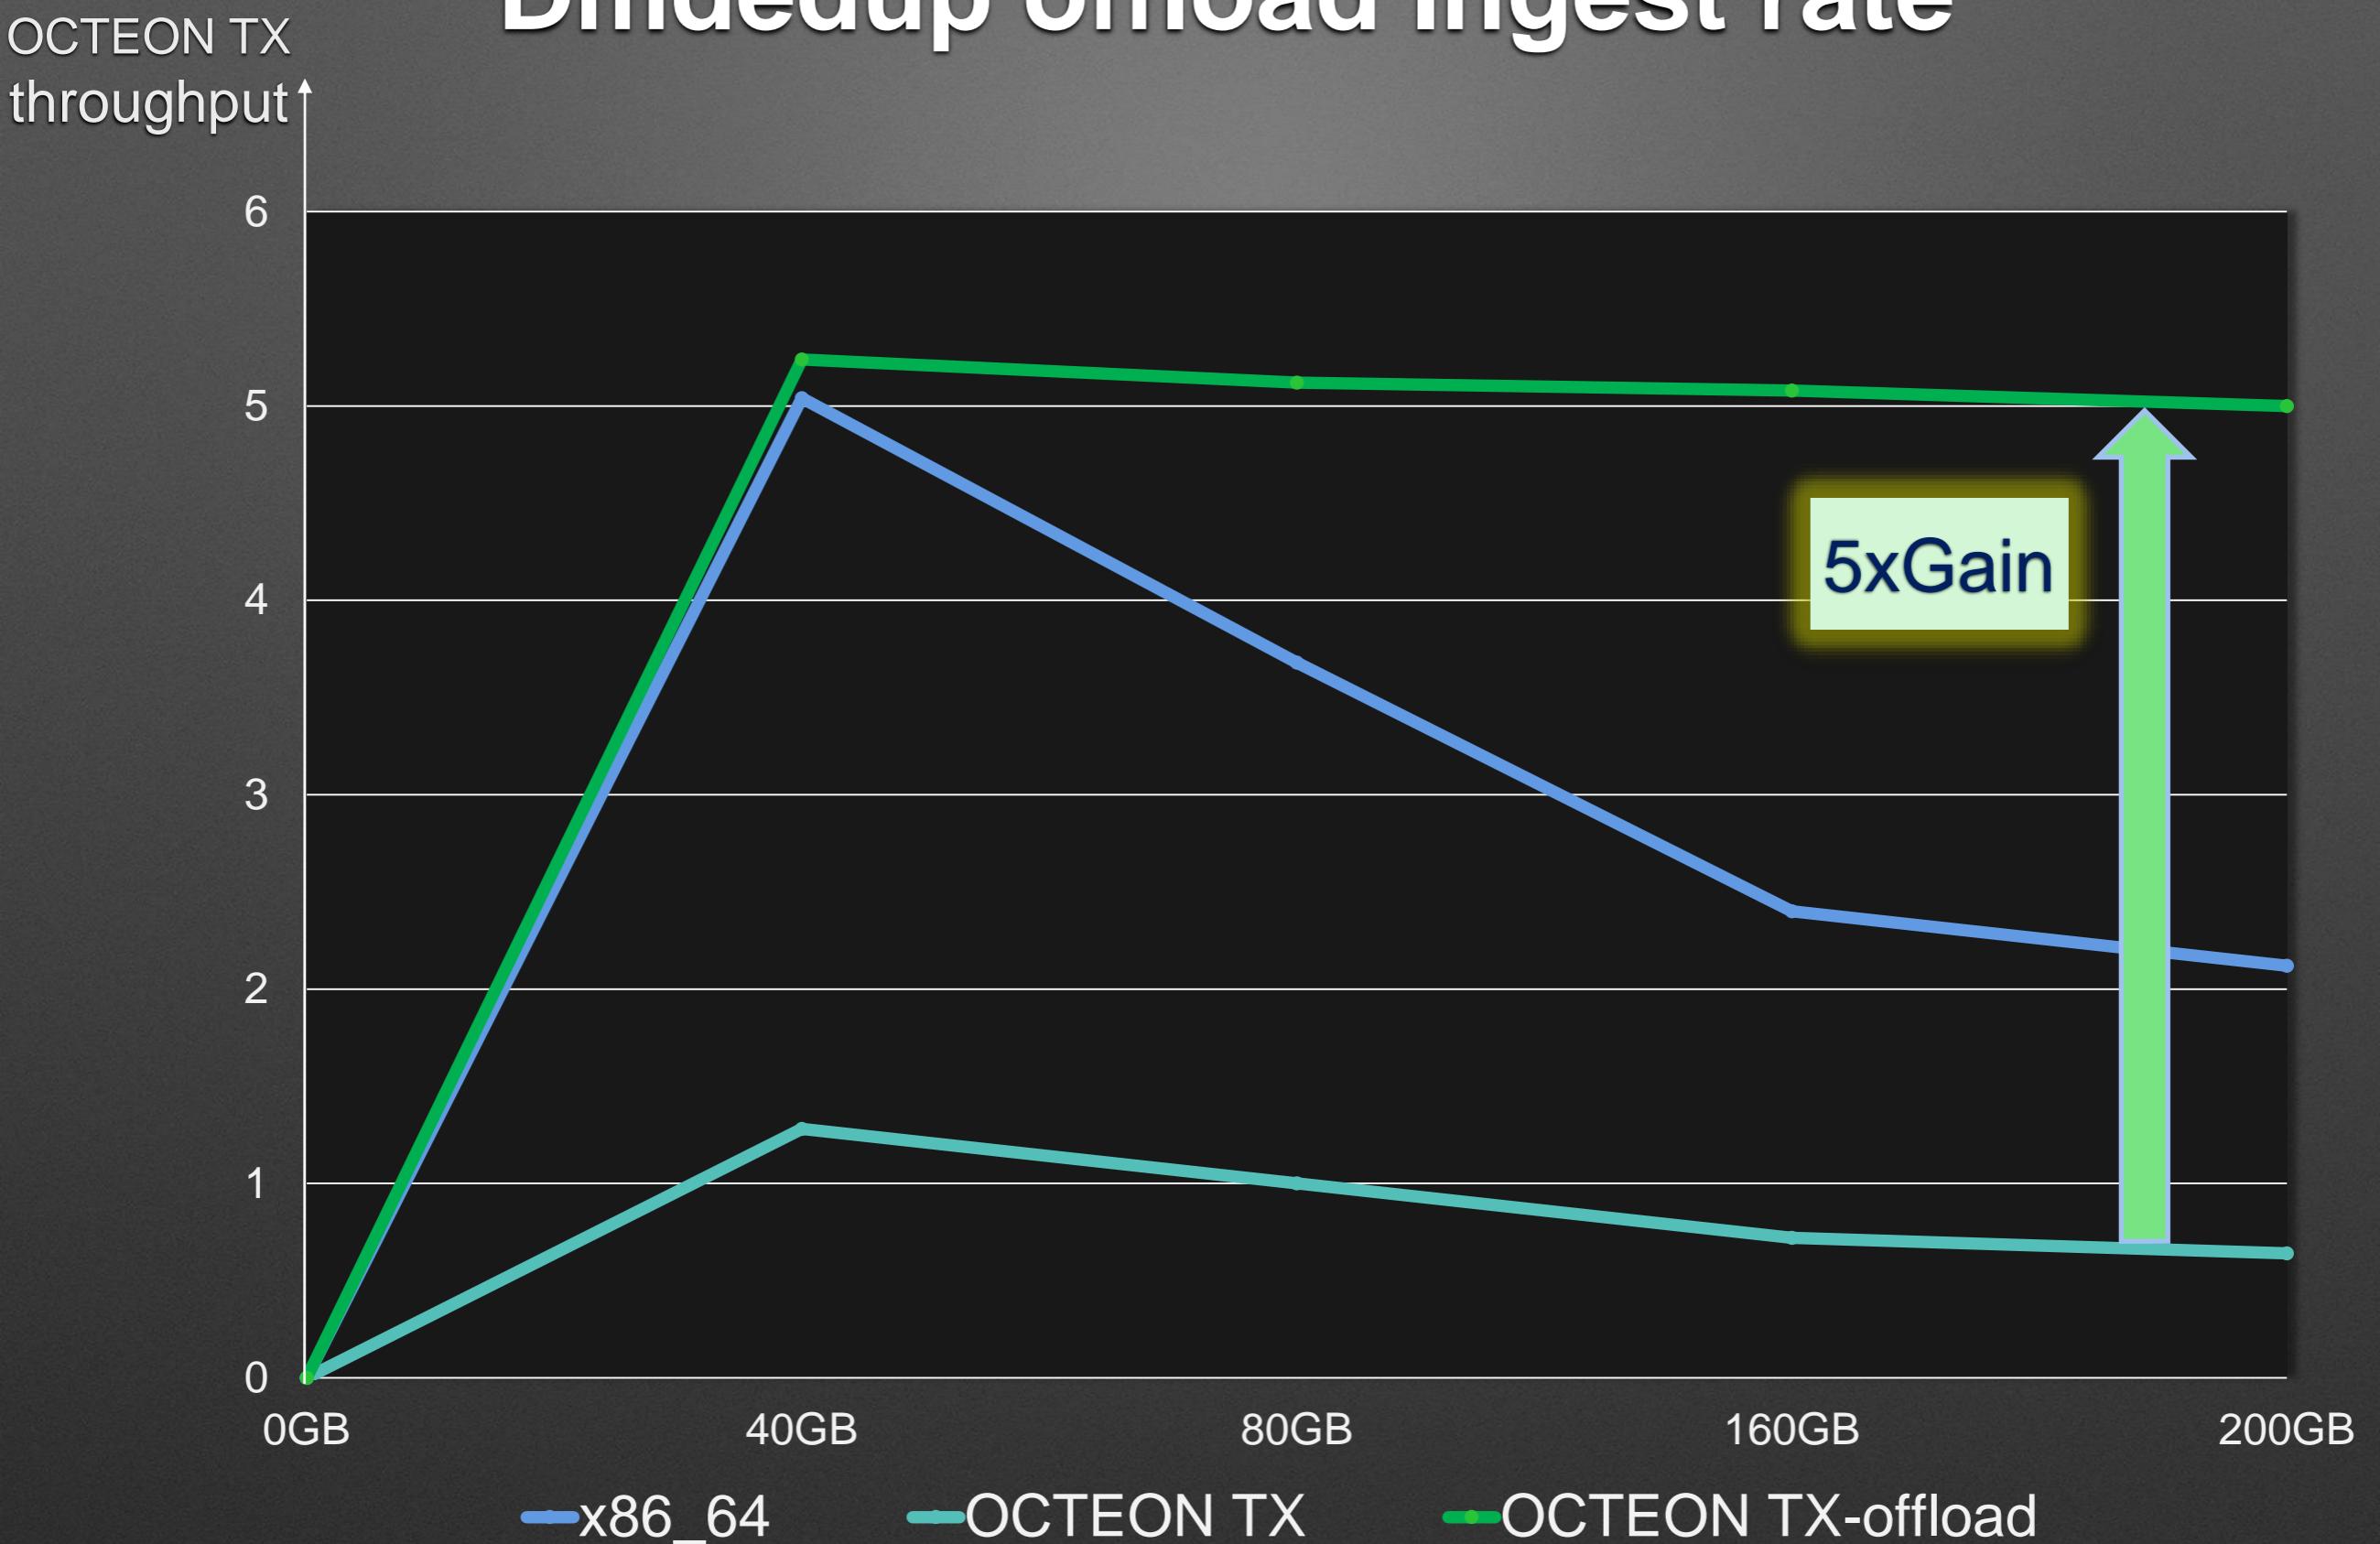
# Accelerations Integrated



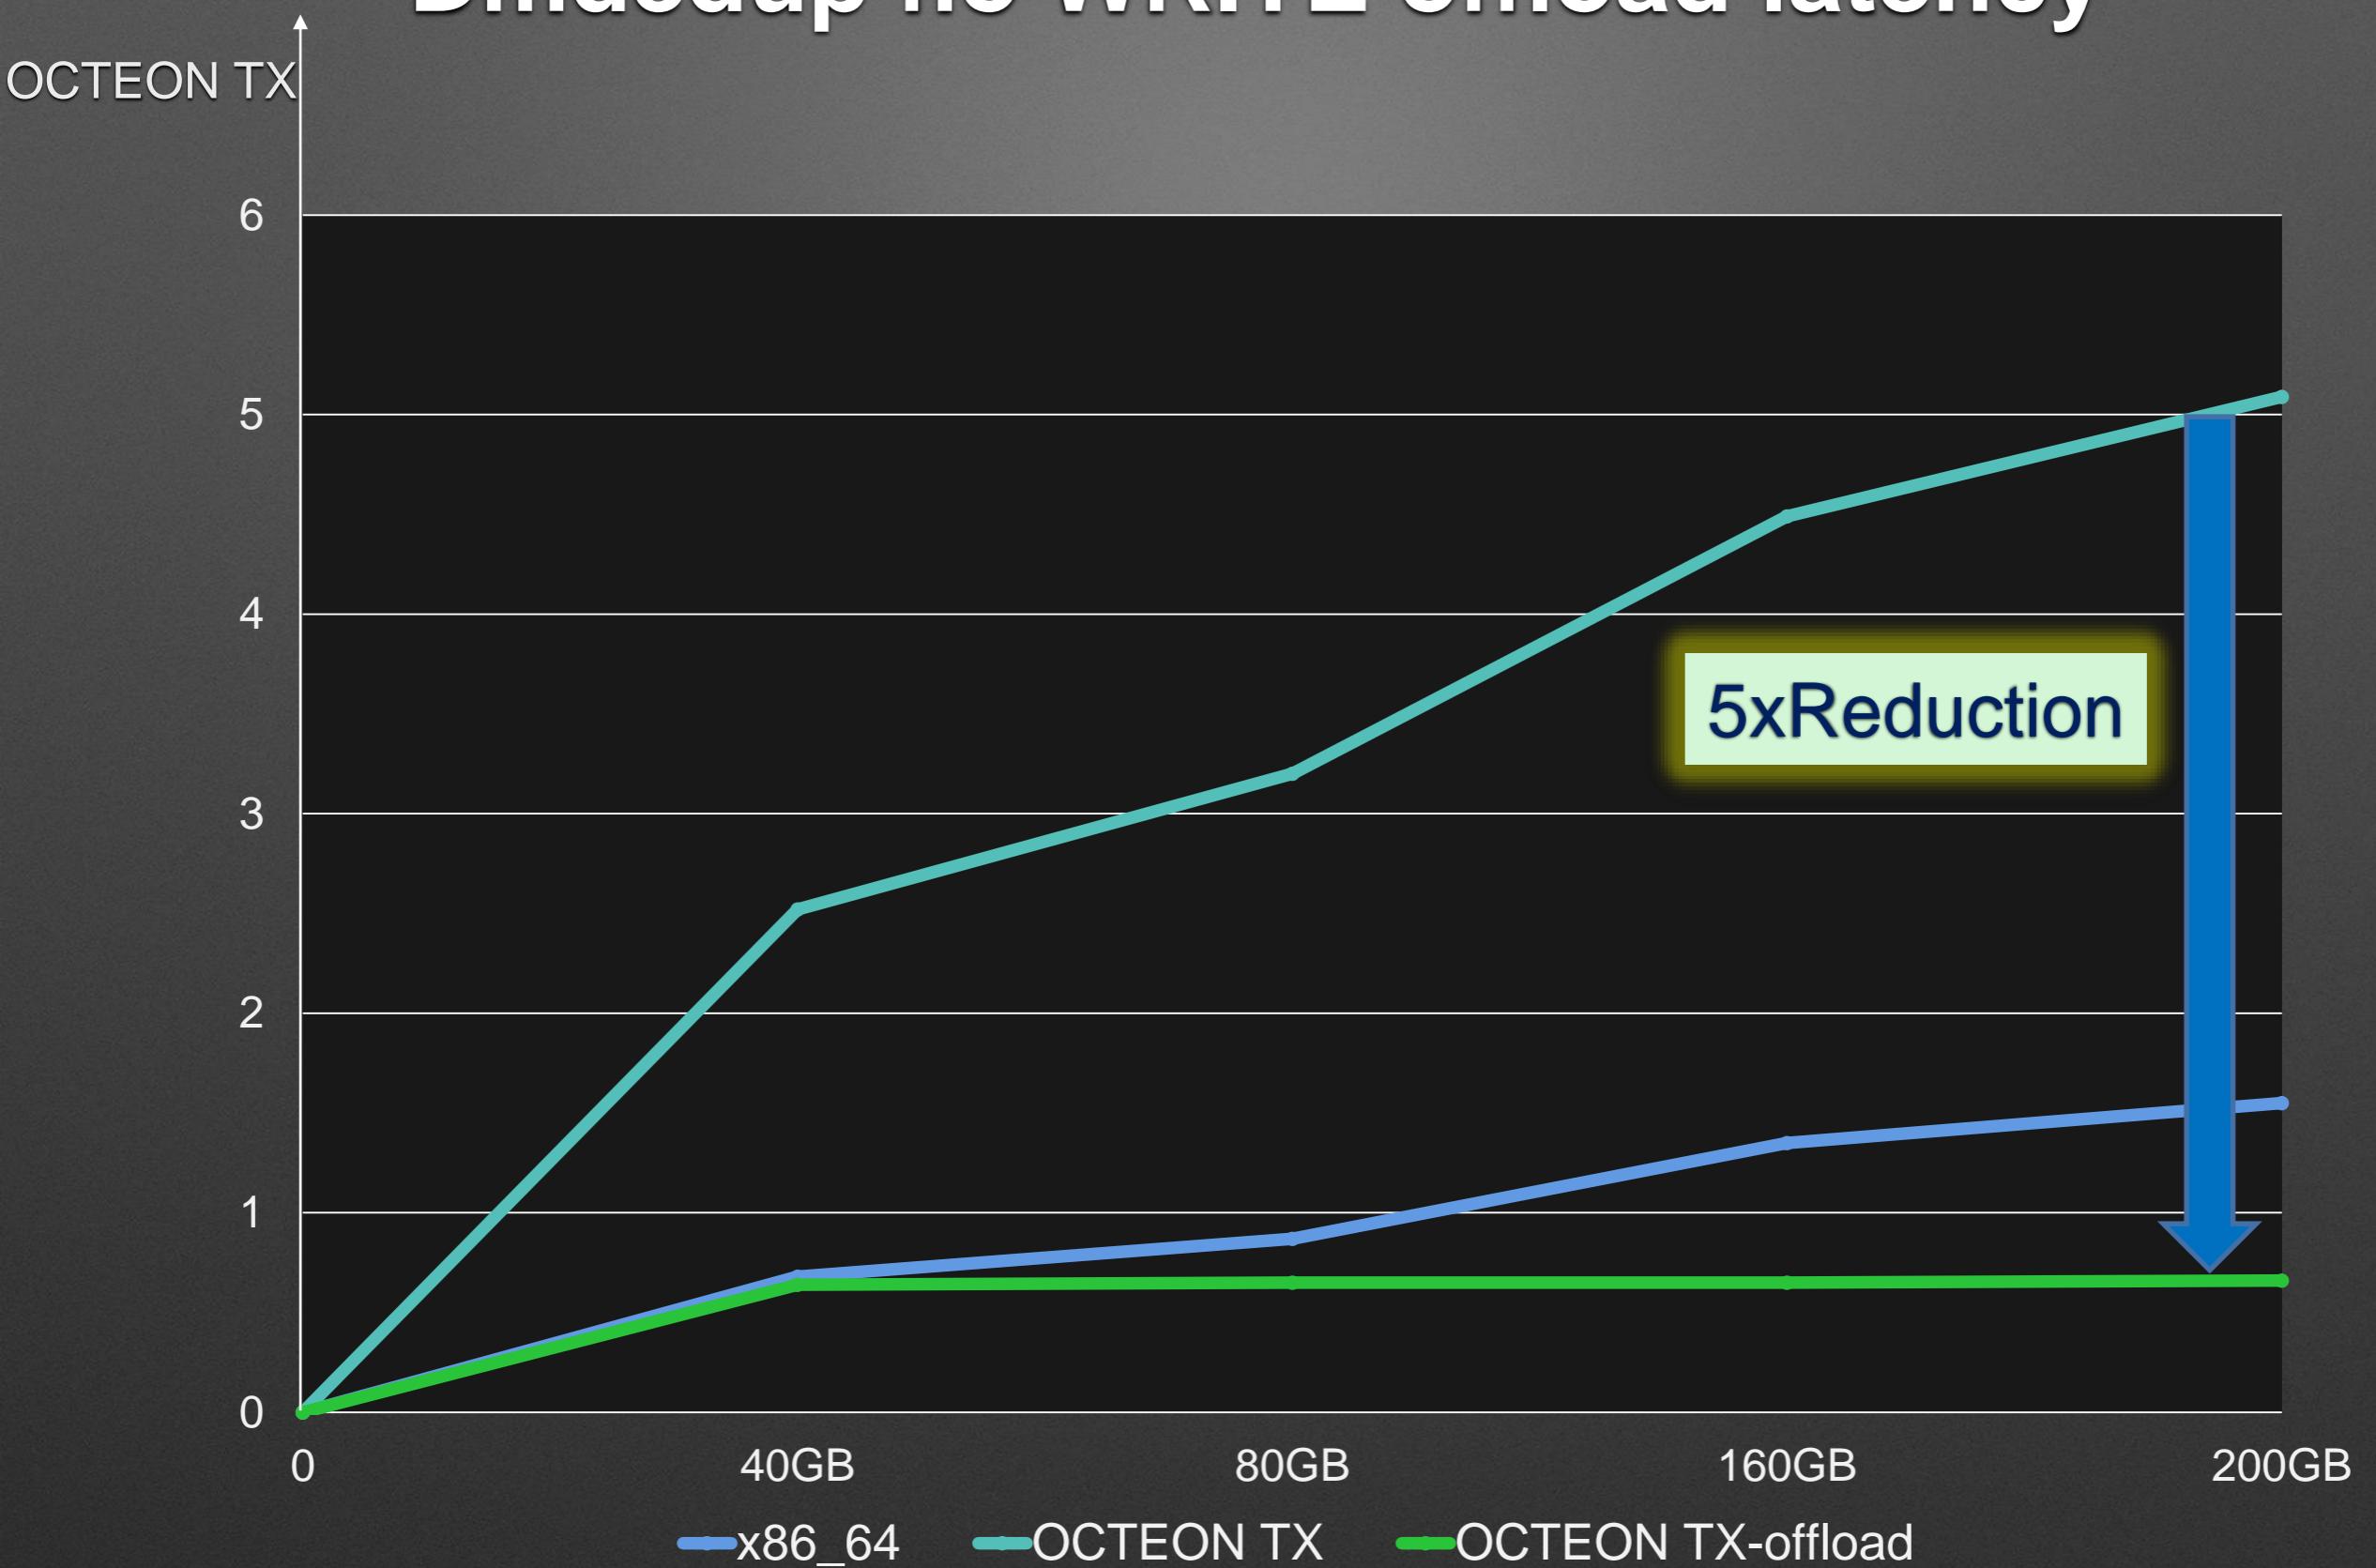
# RAID offload relative CPU savings



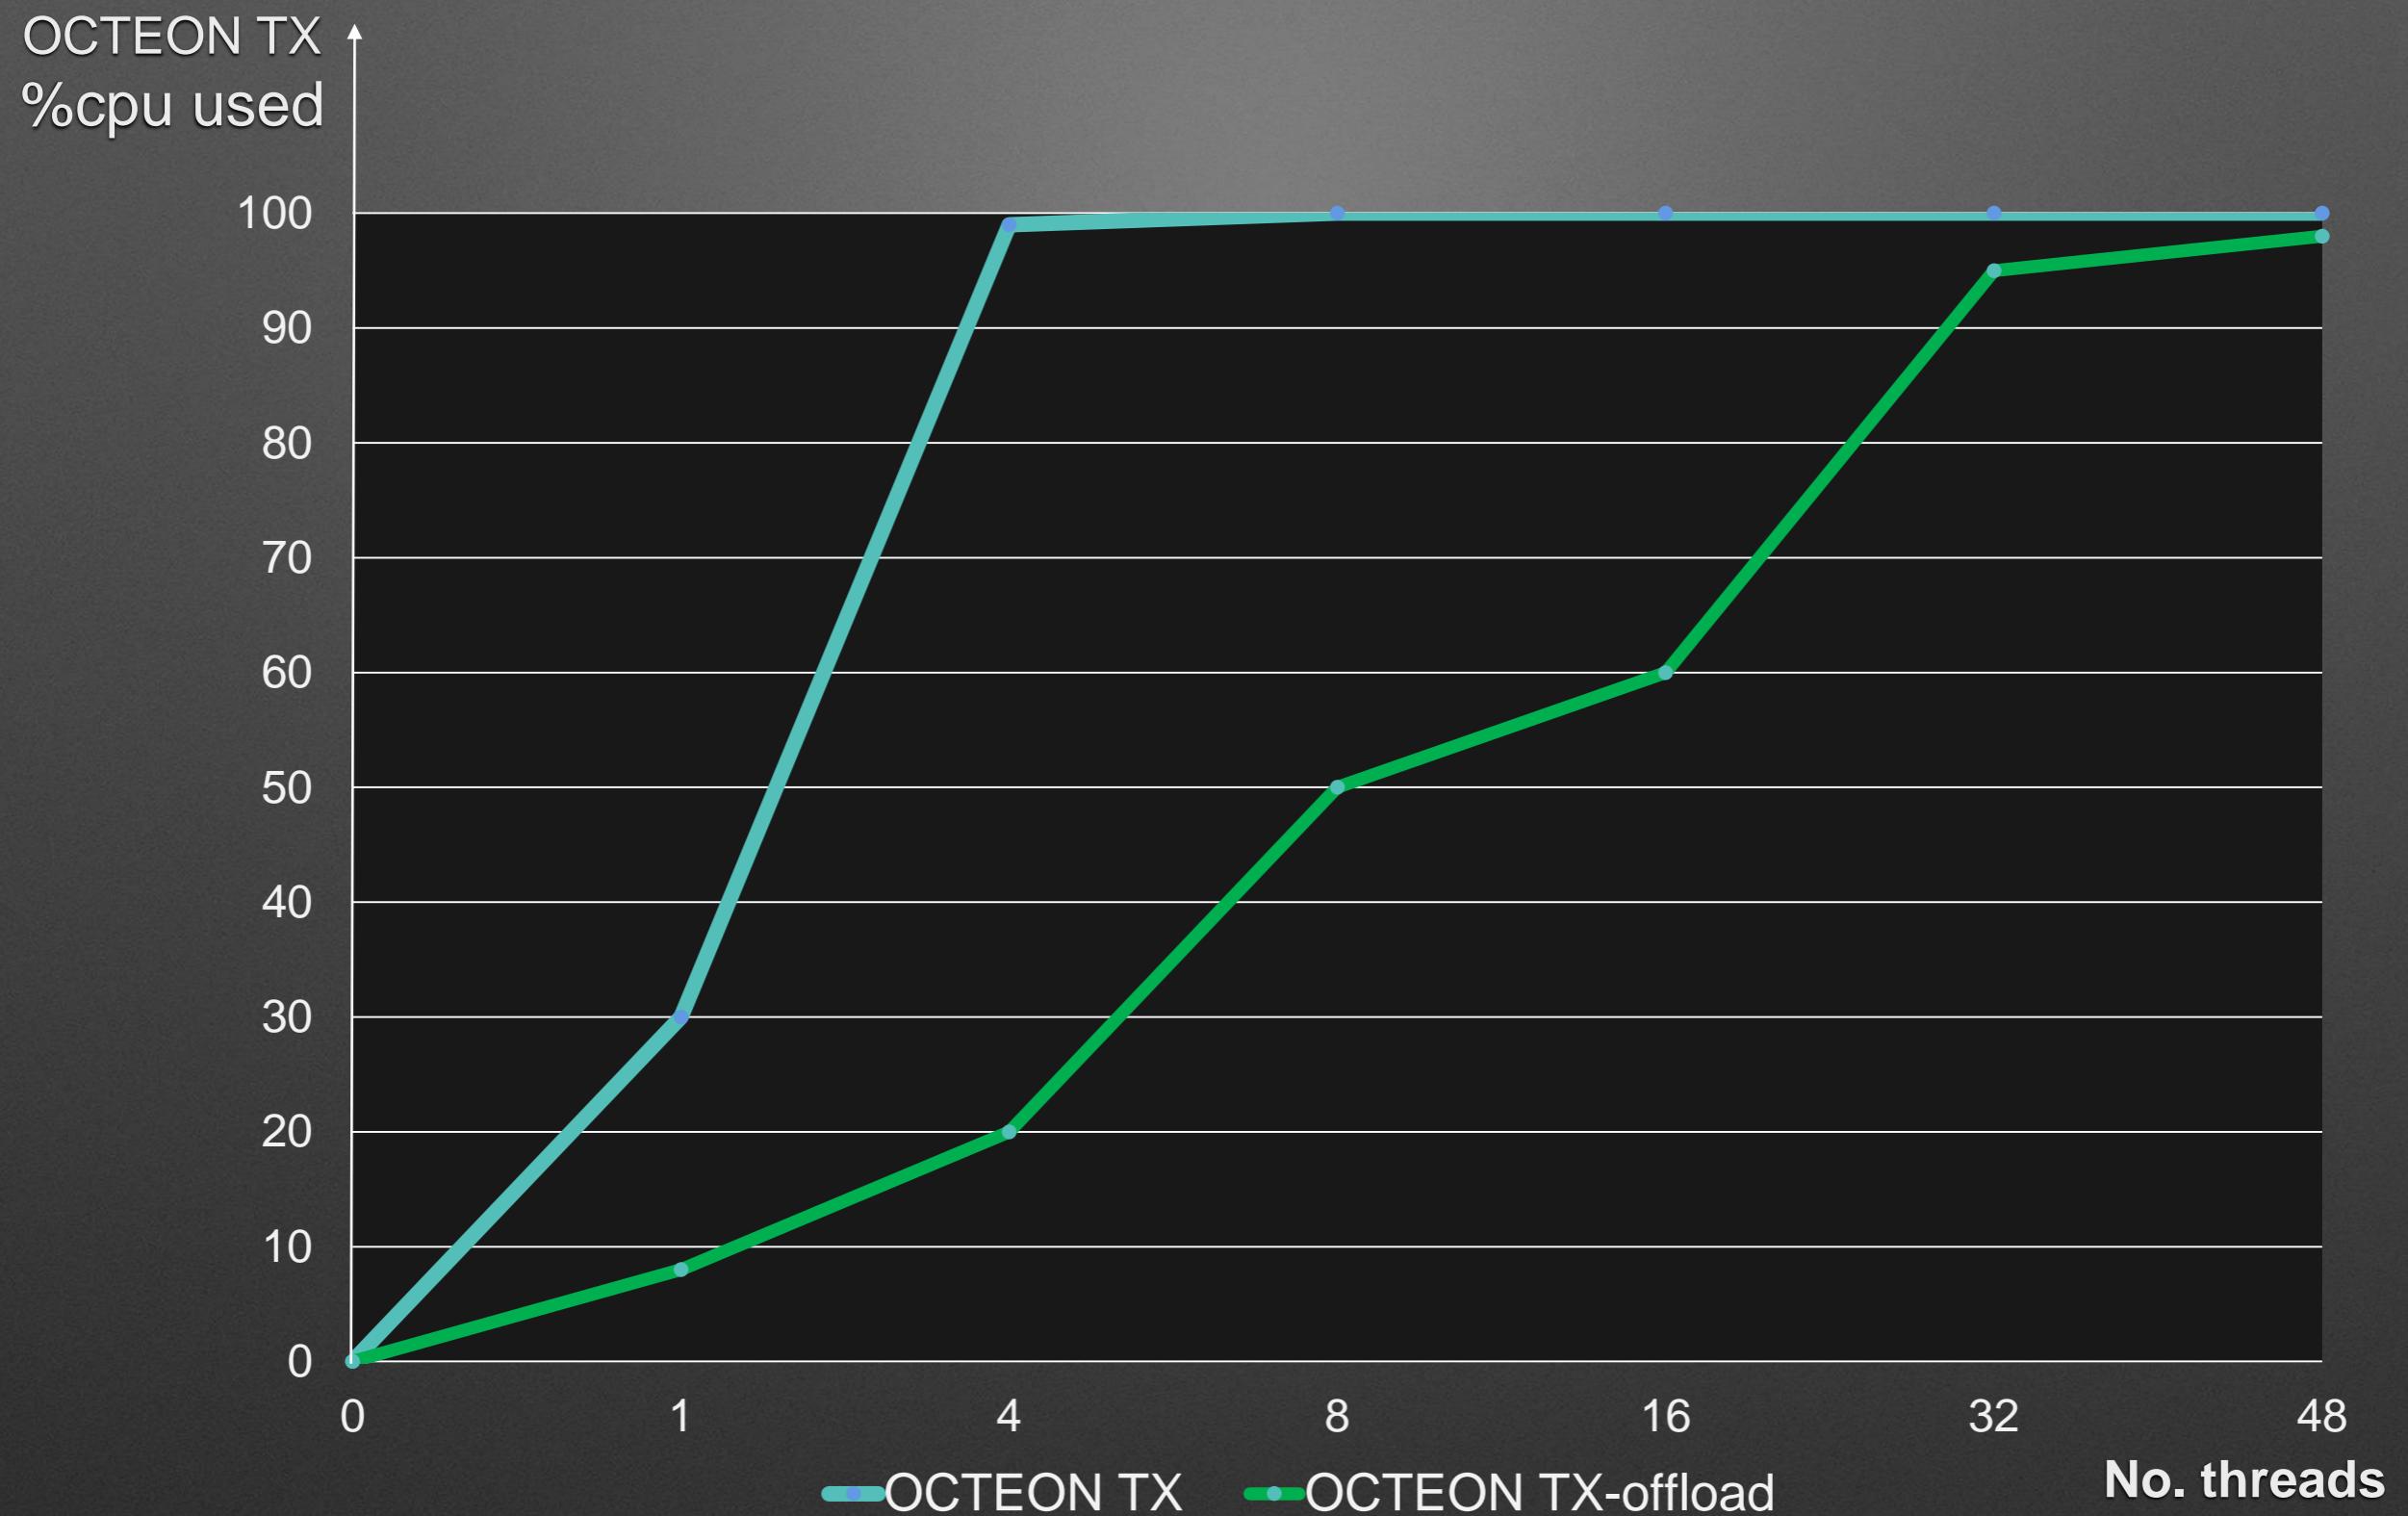
# Dmdedup offload Ingest rate



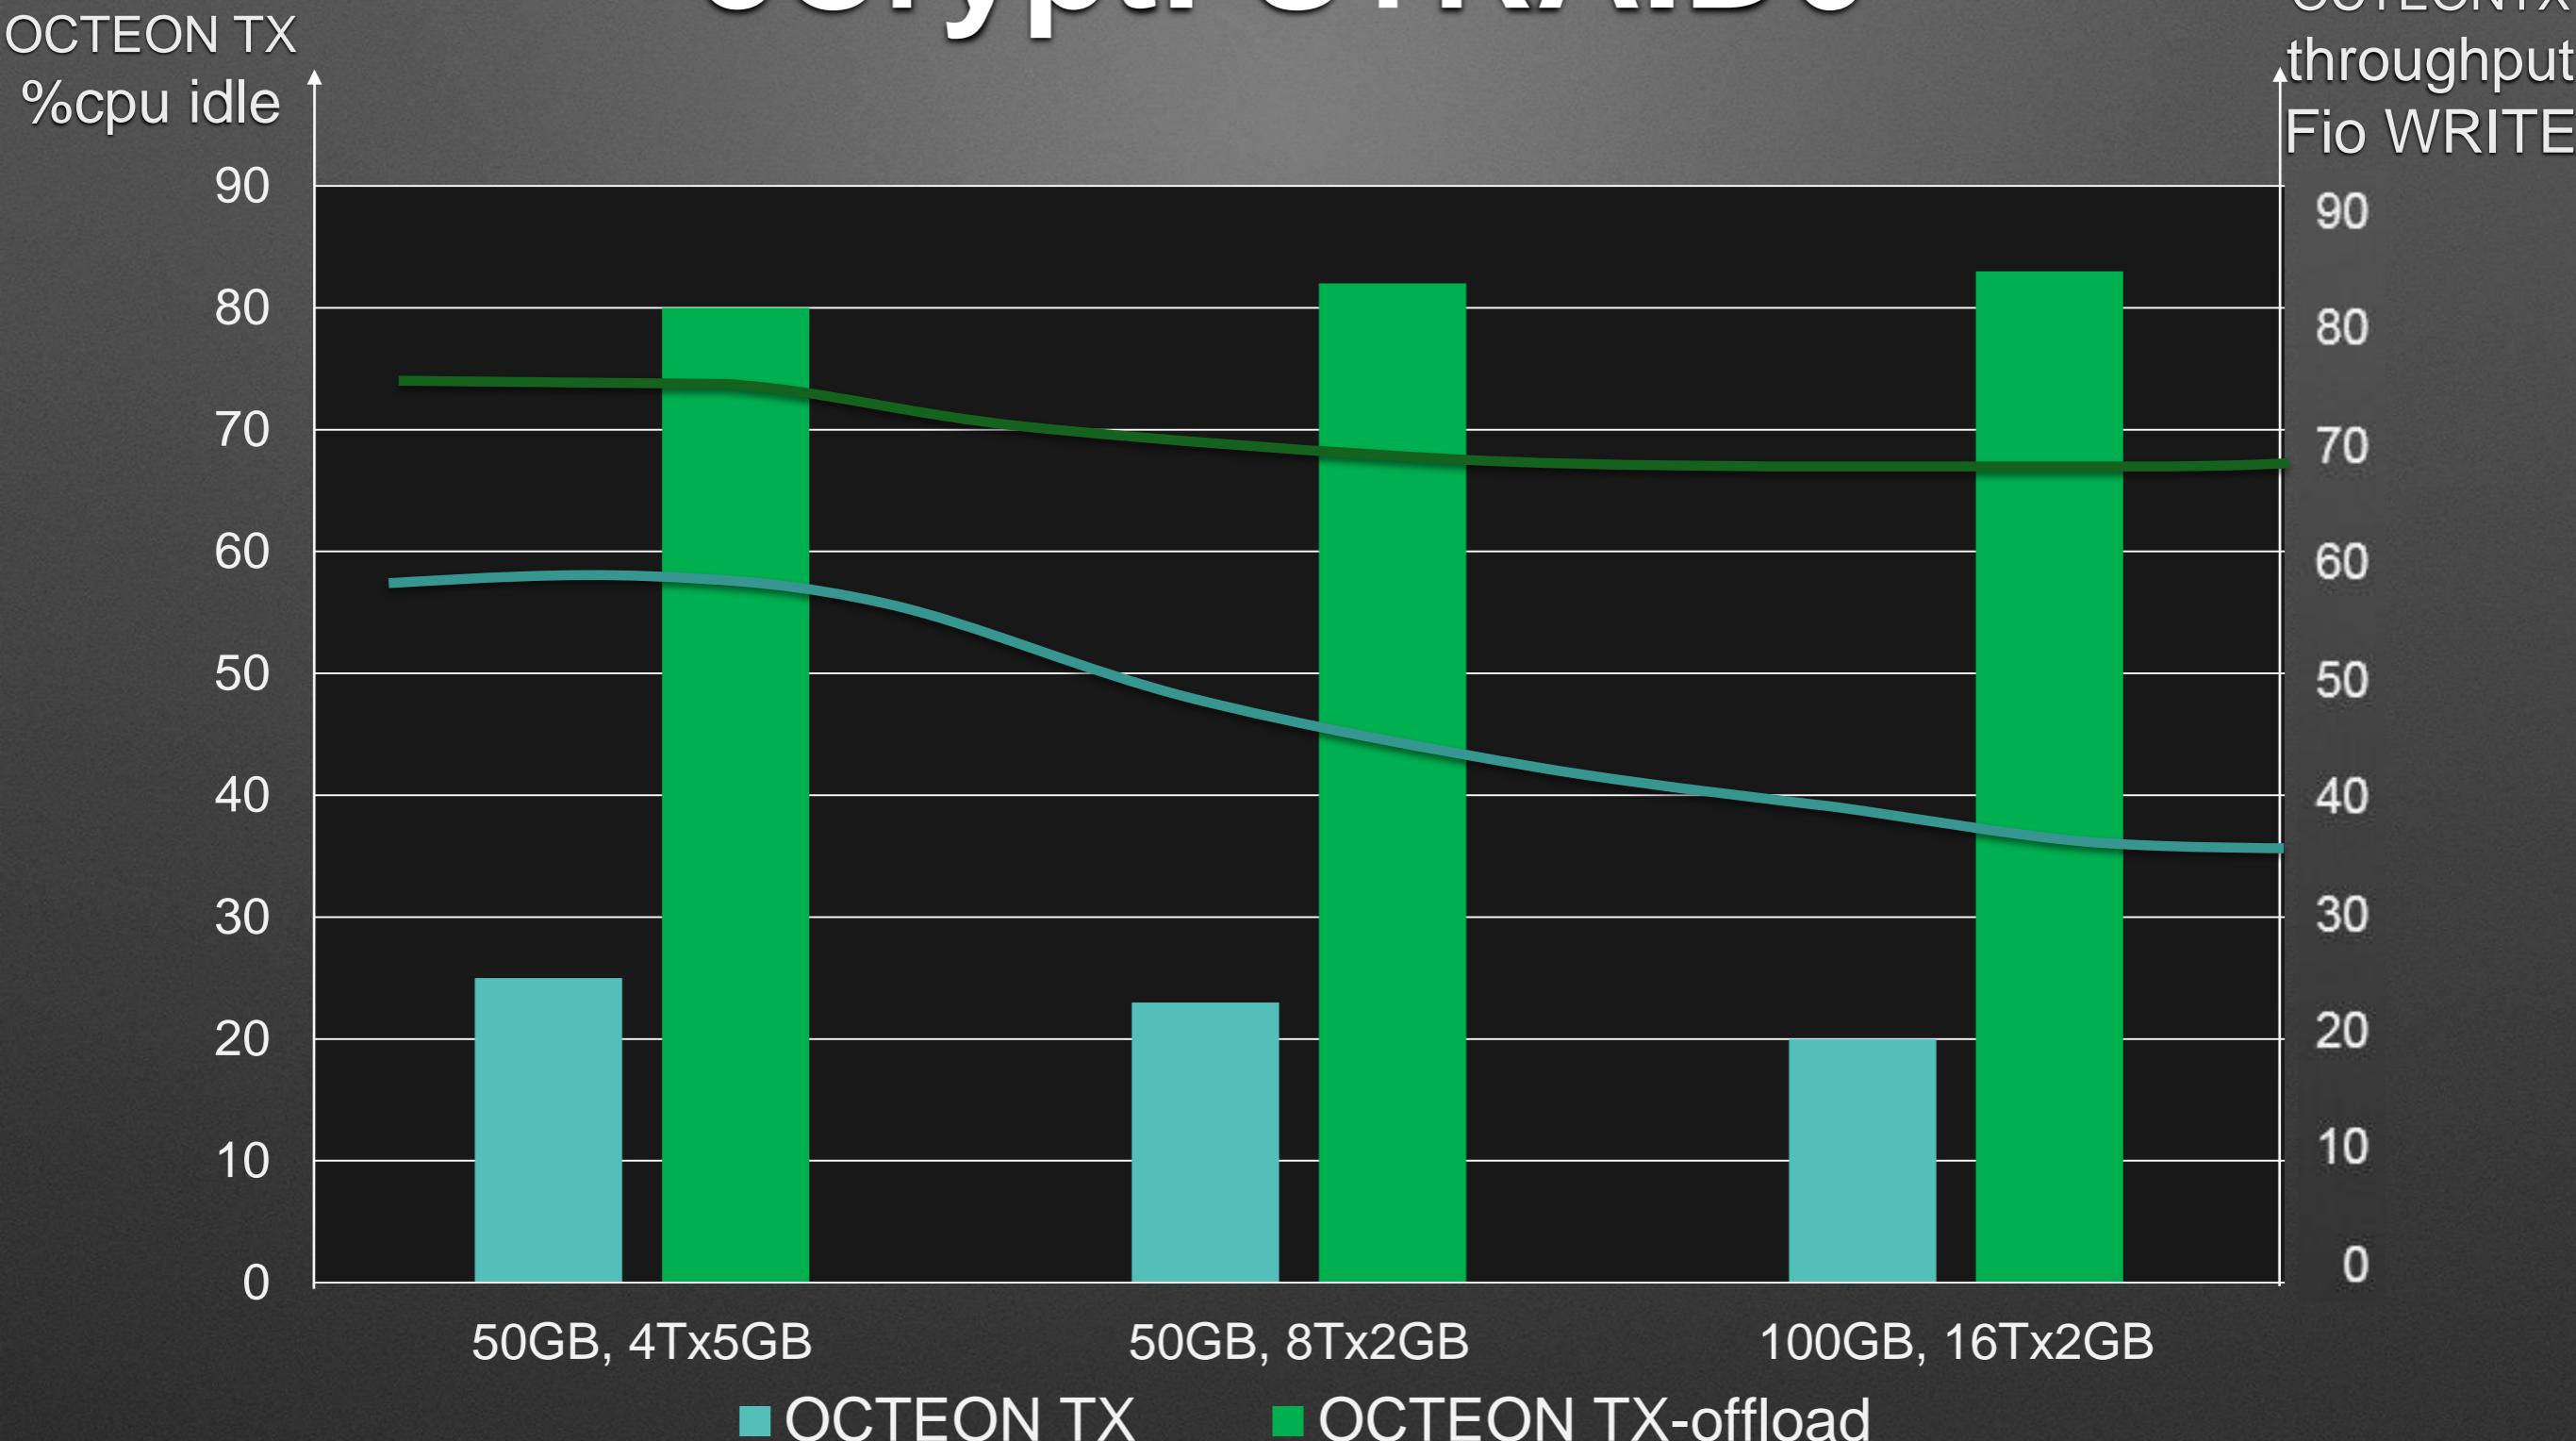
# Dmdedup fio WRITE offload latency



# Ecryptfs



# Dedup + RAID6, eCryptFS+RAID6



# Status

- Status of Work – Preliminary performance results. More in future summits.
- Upstream the drivers.
- Other Platforms:
  - DPDK+SPDK
  - ODP

# Q & A

*Vikas Aggarwal*

*[Vikas.Aggarwal@Cavium.com](mailto:Vikas.Aggarwal@Cavium.com)*

