Nitro: A Capacity-Optimized SSD Cache for Primary Storage

Cheng Li*, Philip Shilane, Fred Douglis, Hyong Shim, Stephen Smaldone, and Grant Wallace

Cheng Li* *Rutgers University*

Philip Shilane, Fred Douglis, Hyong Shim, Stephen Smaldone, and Grant Wallace *EMC Corporation*

Motivation

- Maximize performance and capacity while minimizing cost
- Unified (hybrid) storage is emerging
  - HDD: low performance, low $/GB (.03-.1 for HDD)
  - SSD: high performance, high $/GB (.5-1.2 for SSD)
- SSD caching as performance accelerator
  - Leverage duplicate content
    - Diverse deduplication opportunities (e.g., logs, VM boot-storm)
    - Challenges: duplicate tracking, fingerprint management
  - Leverage compression
    - Large capacity saving opportunities (10-60%)
    - Challenges: fast decompression (LZMA), variable size data

EMC²

RUTGERS

Nitro: A Capacity-Optimized SSD Cache

- Increase effective cache size
  - Leverage deduplication and compression
  - Accelerate two prototype systems
- Support multiple platforms
  - Capacity-Optimized Storage
    - Deduplication and compression
  - Traditional Primary Storage
    - No data reduction technique

- Balance design goals
  - Cost-efficient (SSD as a cache)
  - Performance (~2X over previous work)
  - SSD lifespan (65% fewer erasures)
  - Resources (reduce RAM footprint)

EMC²

RUTGERS

Key Components

- **Extent**:
  - File data caching unit
  - Consecutive file bytes (e.g., 8KB)
- **Write-Evict Unit (WEU)**:
  - Contains compressed extents
  - Inserts into and evicts from SSD (e.g., 2MB)
  - Extends SSD lifespan, may evict useful data

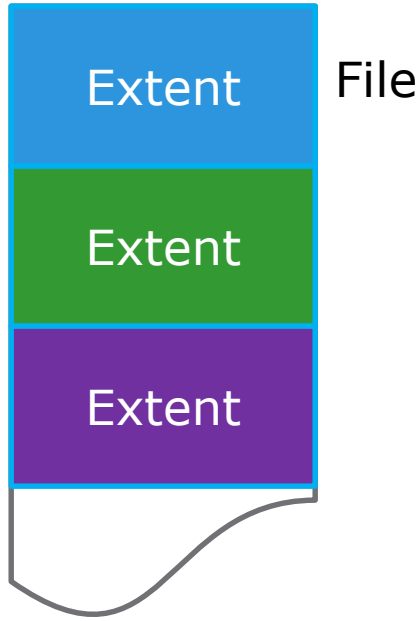

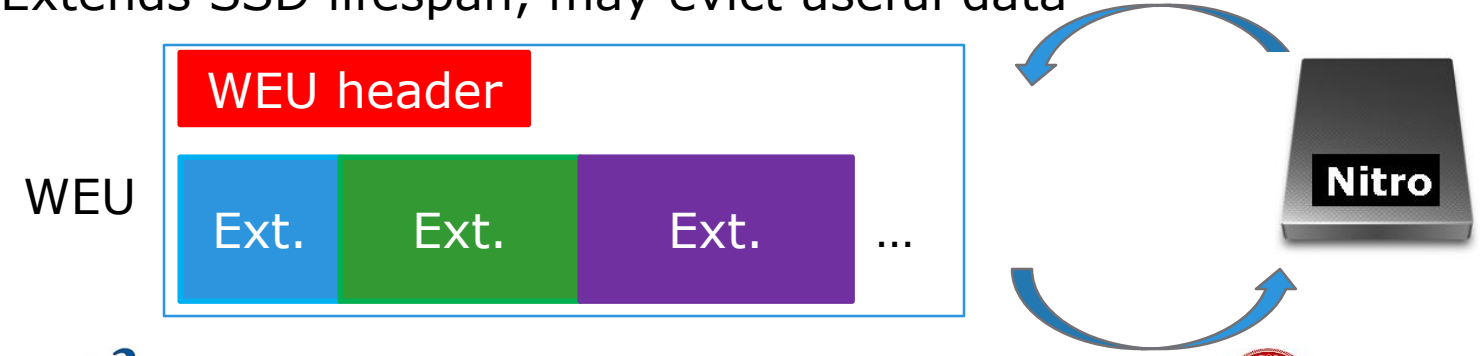

EMC²

RUTGERS

Nitro Insertion Path

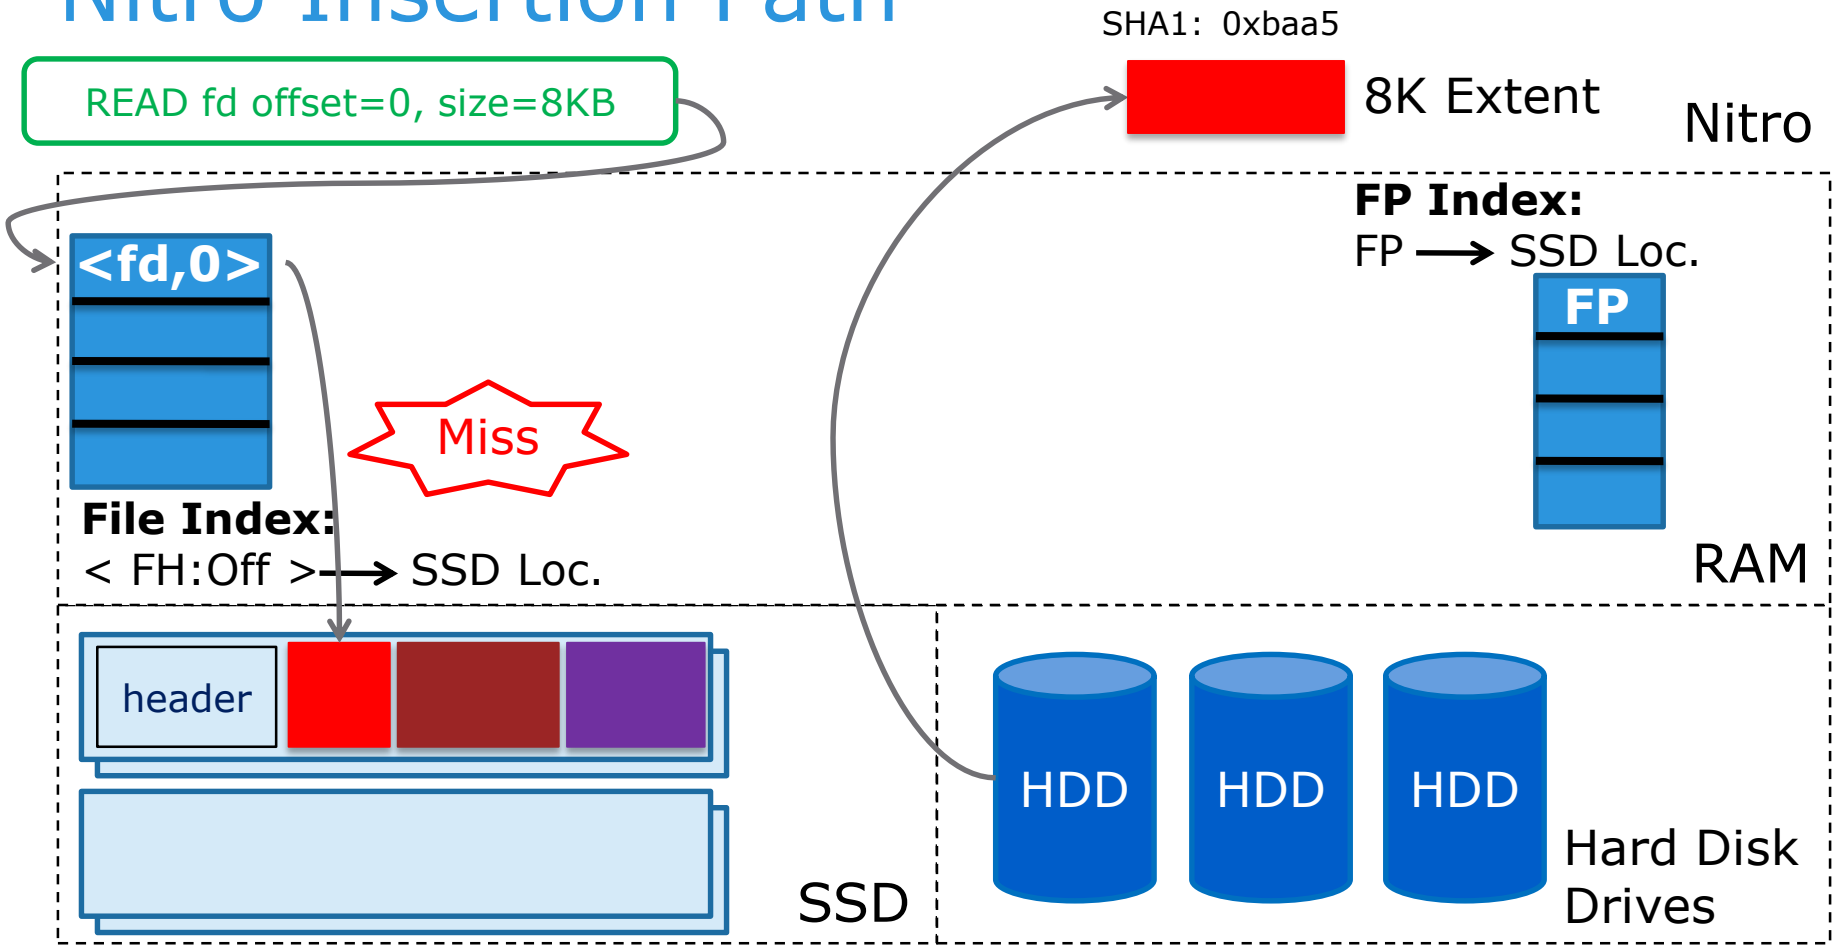

Nitro Insertion Path

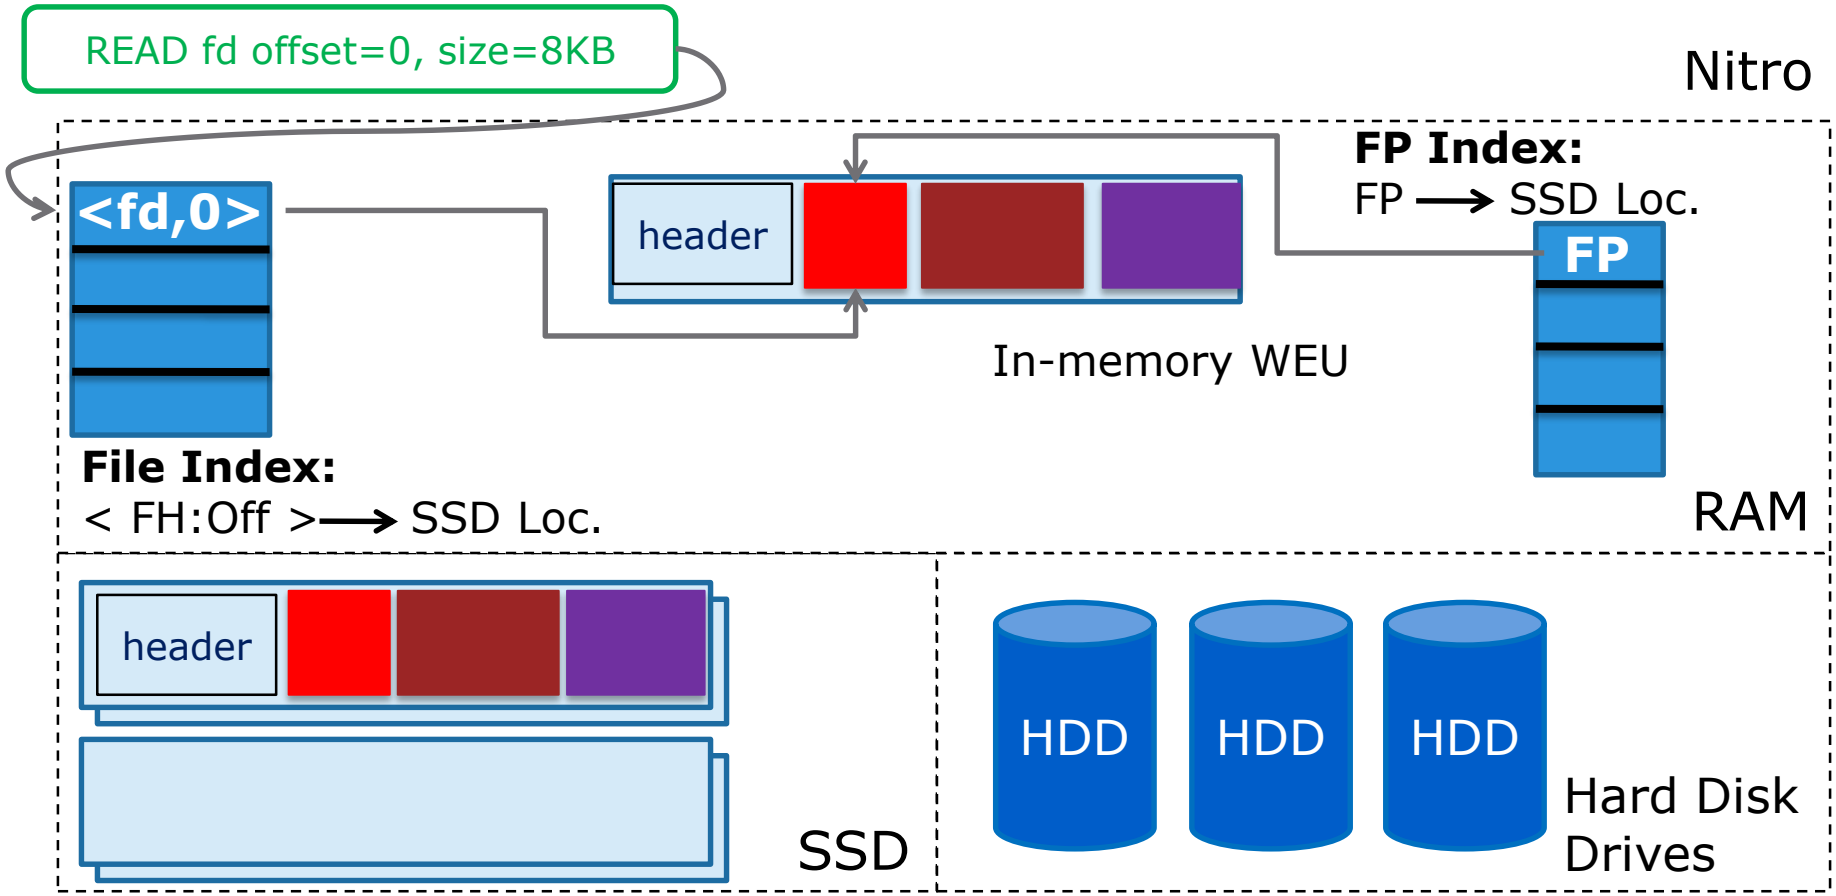

EMC²

RUTGERS

Nitro Read Path

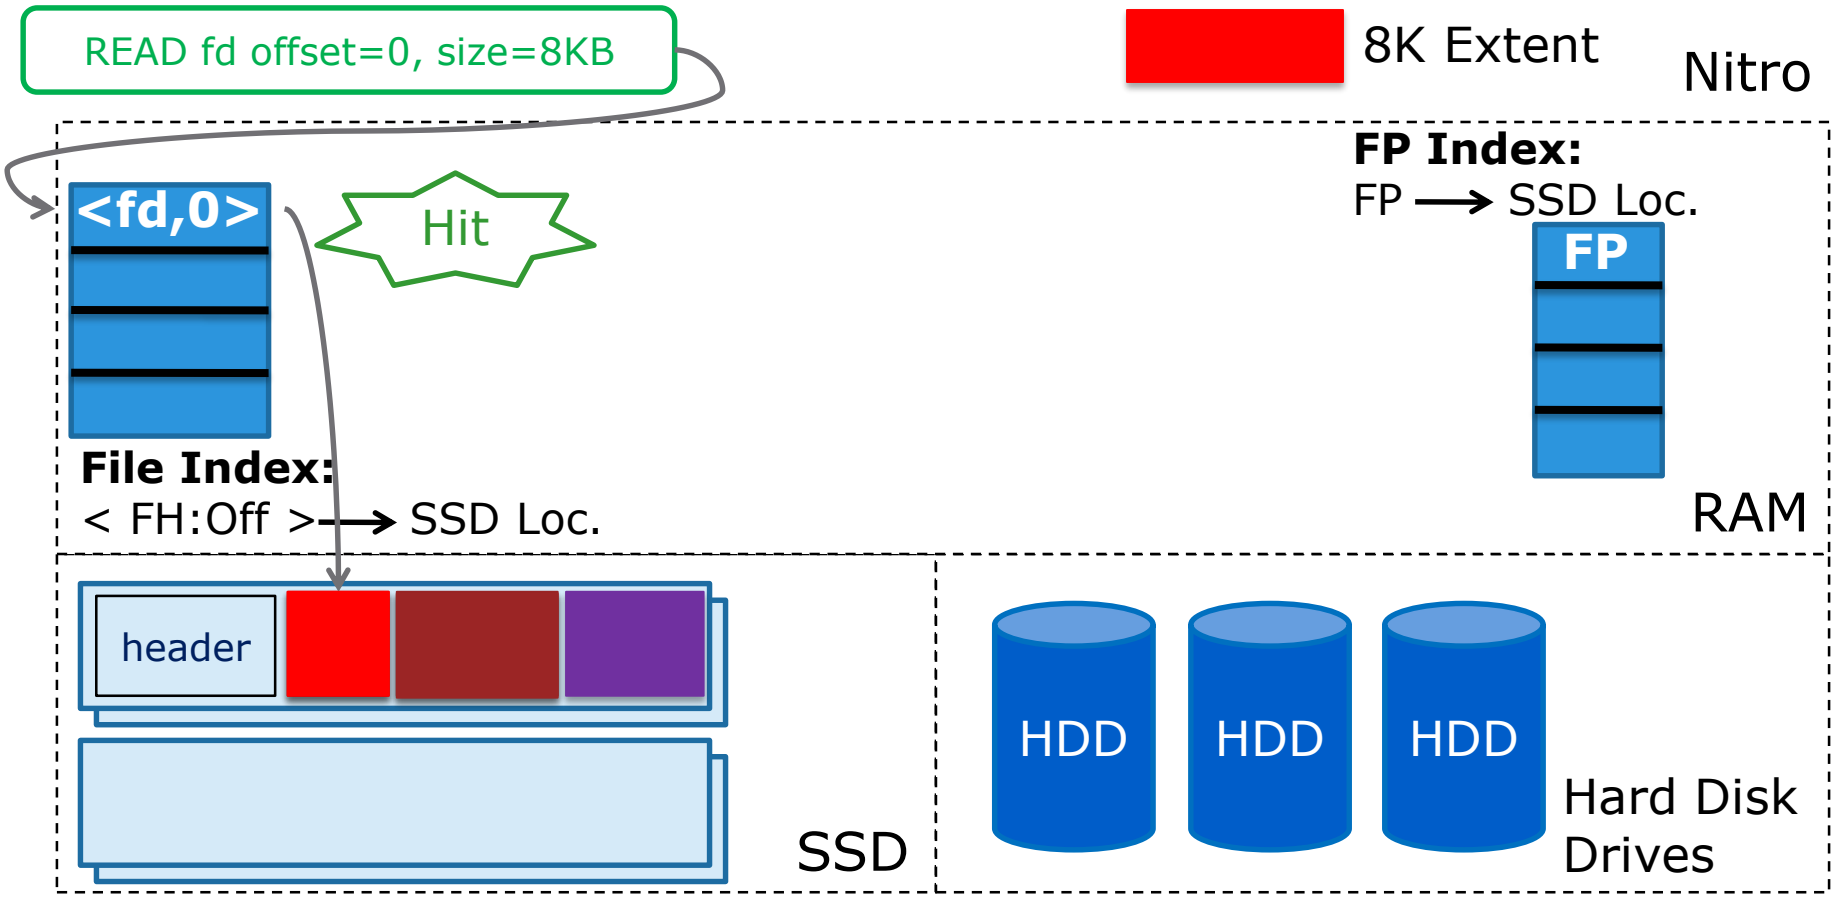

Nitro Read Path

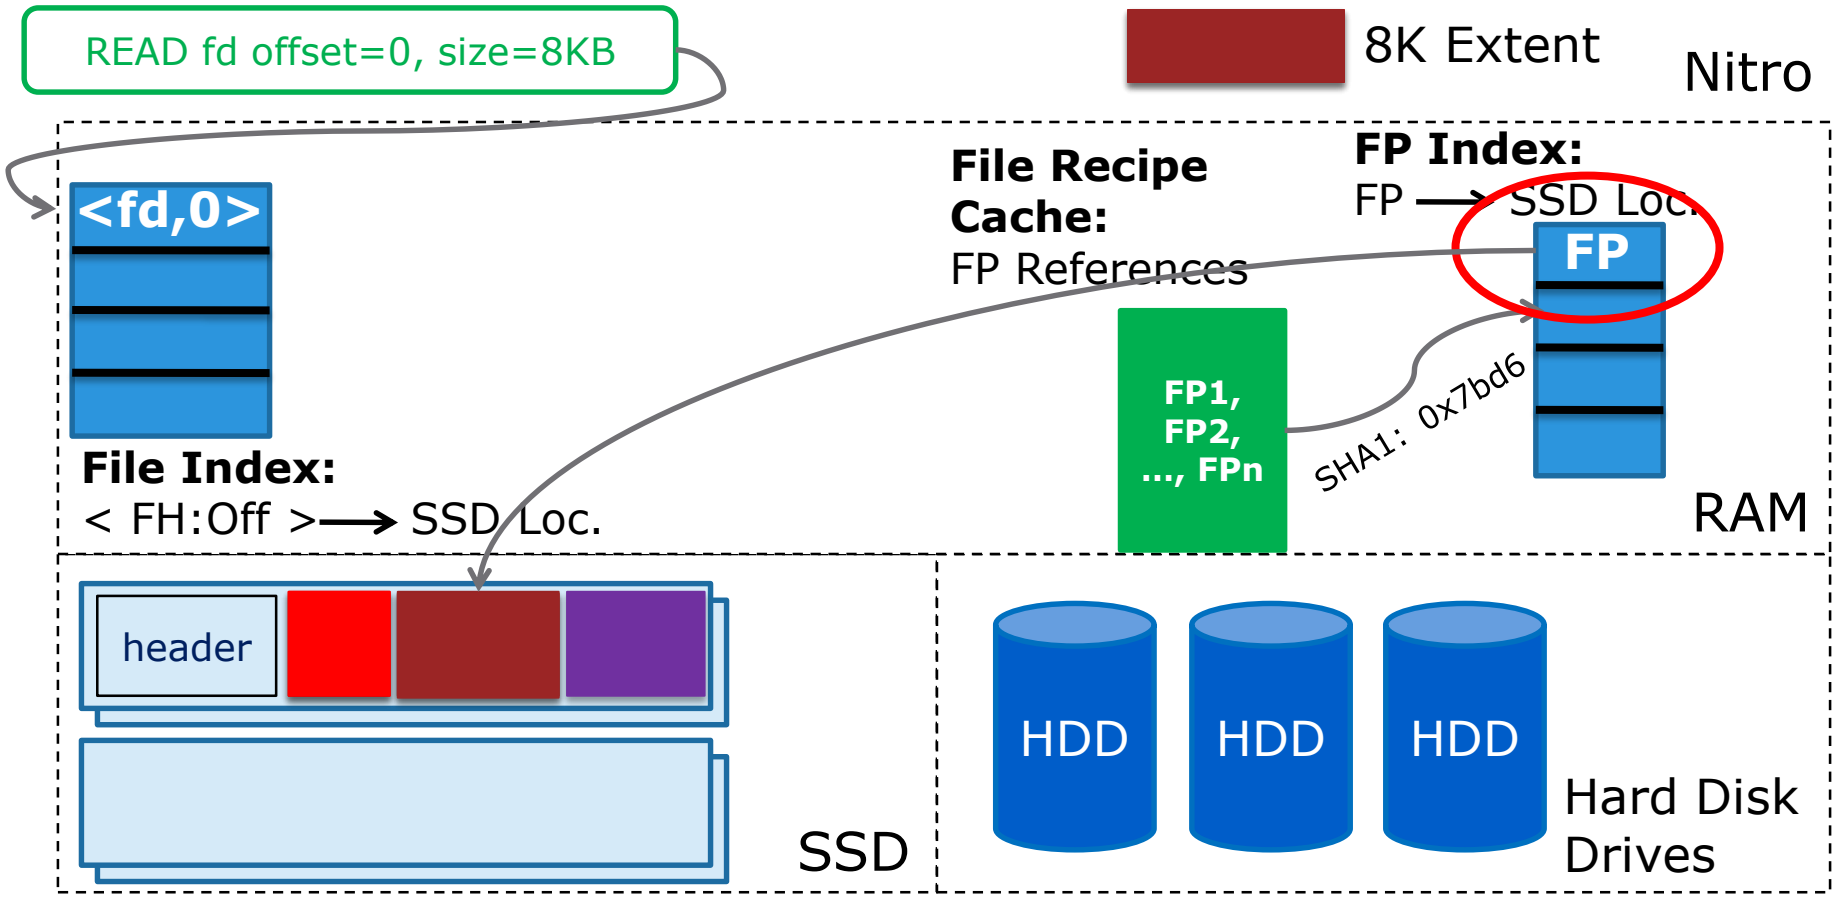

EMC²

RUTGERS

Nitro Deduplication Path

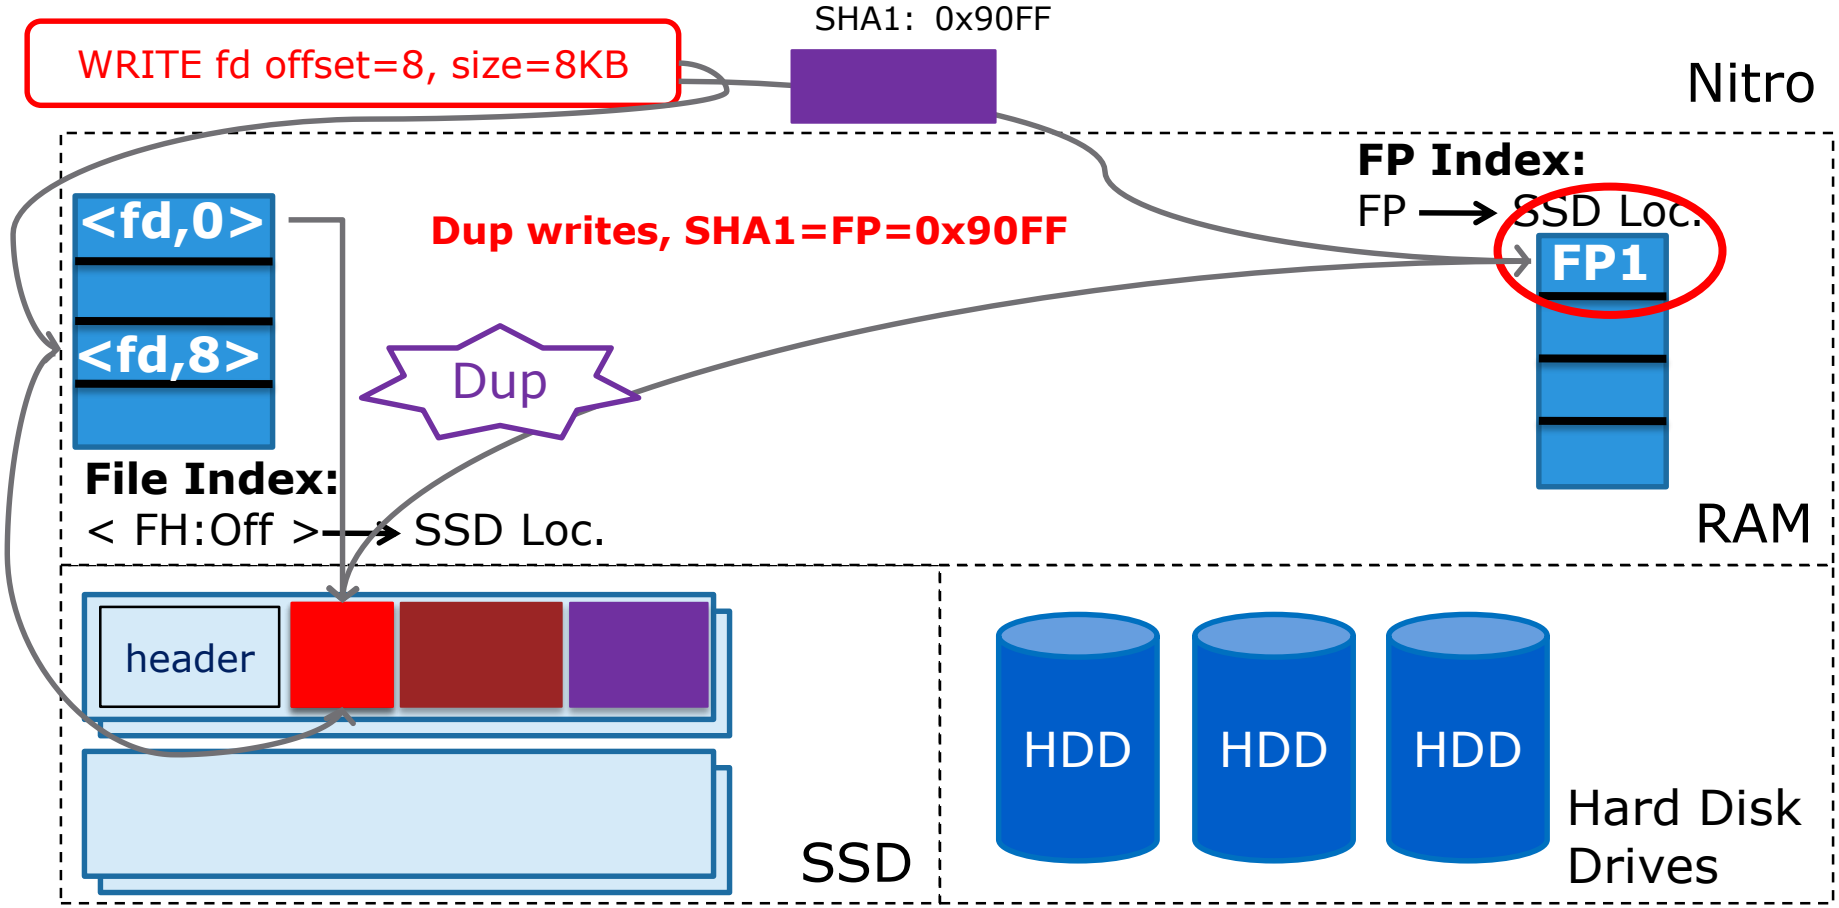

EMC²

RUTGERS

Nitro Deduplication Path

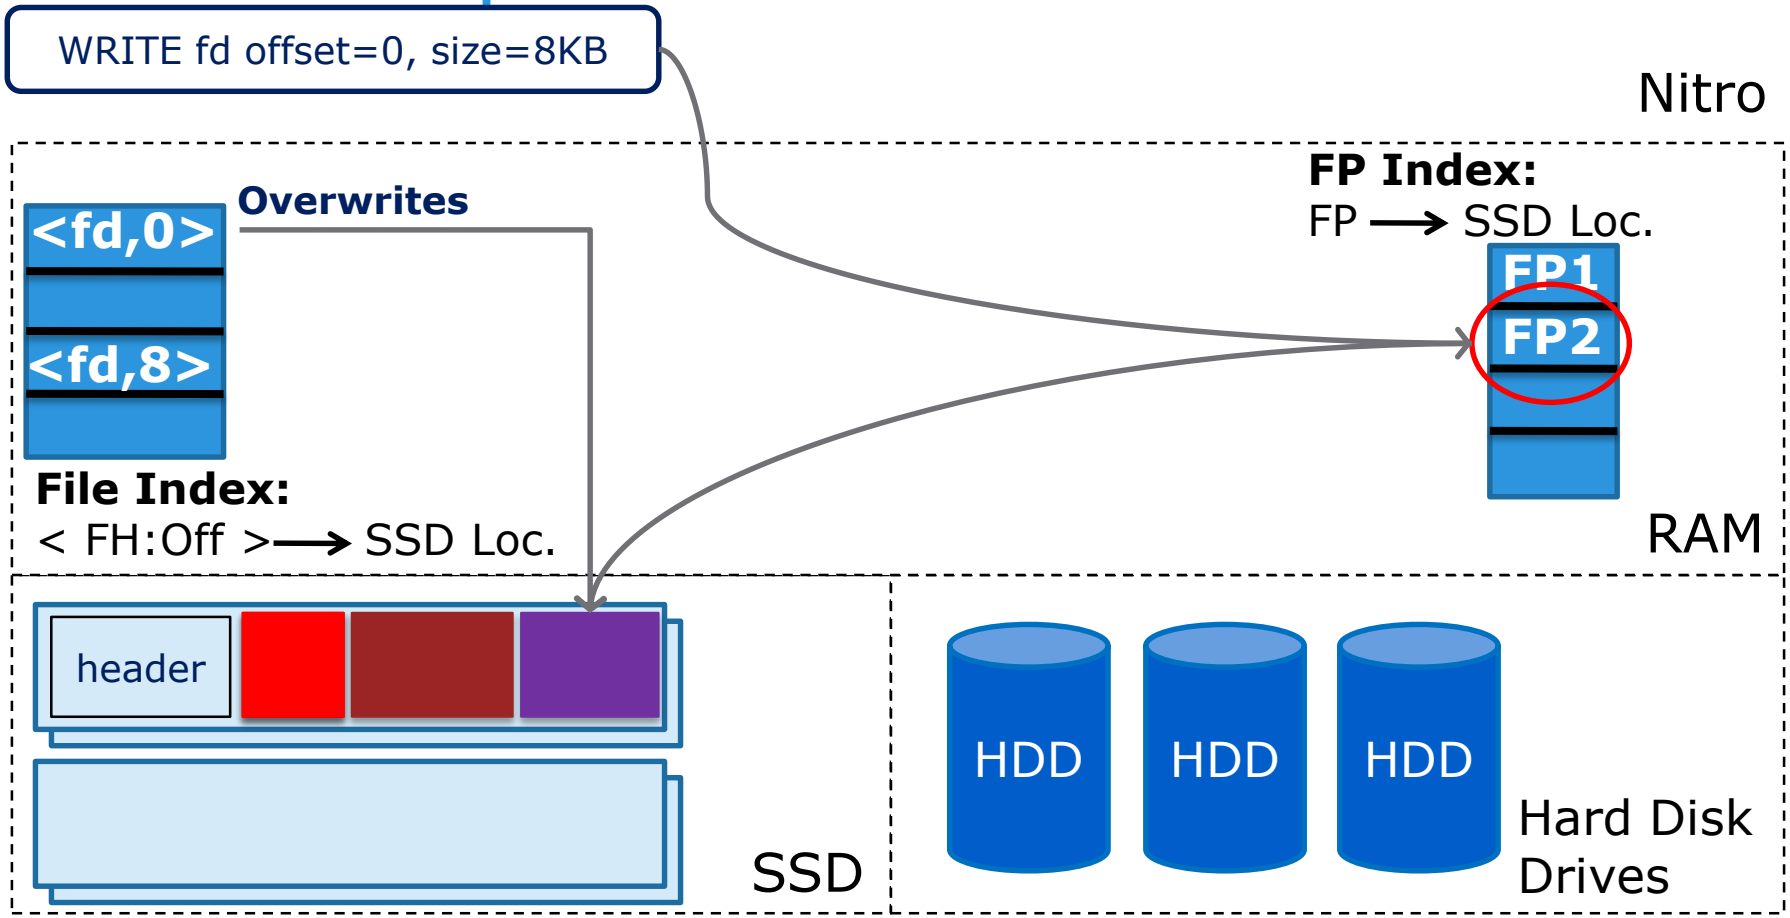

Experimental Methodology

- Storage traces
  - Florida International University (FIU)
    - Homes, WebVM, Mail
  - VM snapshot and restore traces
- Platforms
  - SSD simulator (measure SSD erasures)
  - Prototype with two storage systems
    - Capacity-Optimized Storage (**COS**)
    - Traditional Primary Storage (**TPS**)

EMC²

RUTGERS

Cache Variants

- Explore design variants
  - Extent-based vs. WEU-based

(D=deduplication, C=compression, ND=no deduplication, NC=no compression)

Variants	Write/Evict Granularity	Deduplication	Compression
Extent (ND, NC)	Extent		
Extent (D, NC)	Extent	X	
WEU (ND, NC)	WEU		
WEU (D, NC)	WEU	X	
WEU (D, C) (Nitro)	WEU	X	X

EMC²

RUTGERS

Can Nitro Increase Read-hit Ratio?

Deduplication and compression increase read-hit ratio (2% cache)

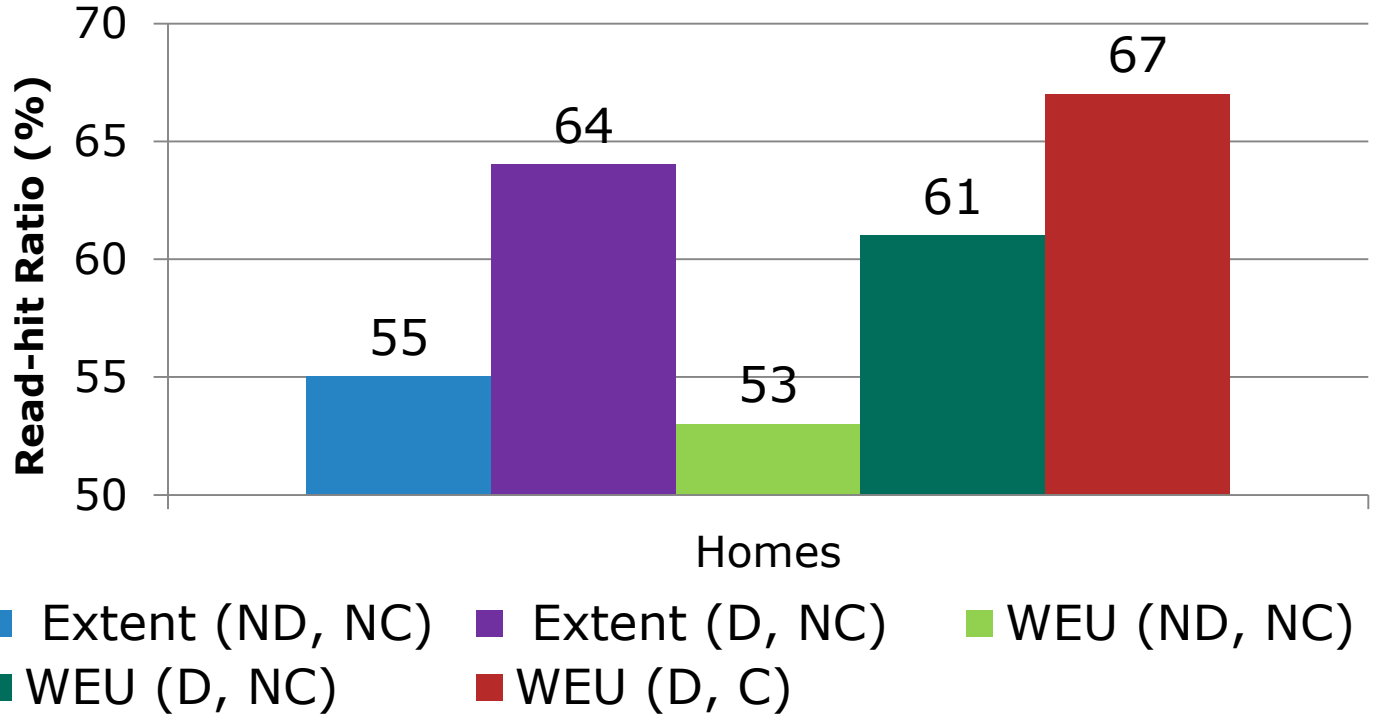

(D=deduplication, C=compression, ND=no deduplication, NC=no compression)

EMC²

RUTGERS

Nitro Extends SSD Lifespan

- WEU eliminates SSD overwrites penalty
- TRIM: SATA command to invalidate addresses
  - Decreases garbage collection copy forward in SSD

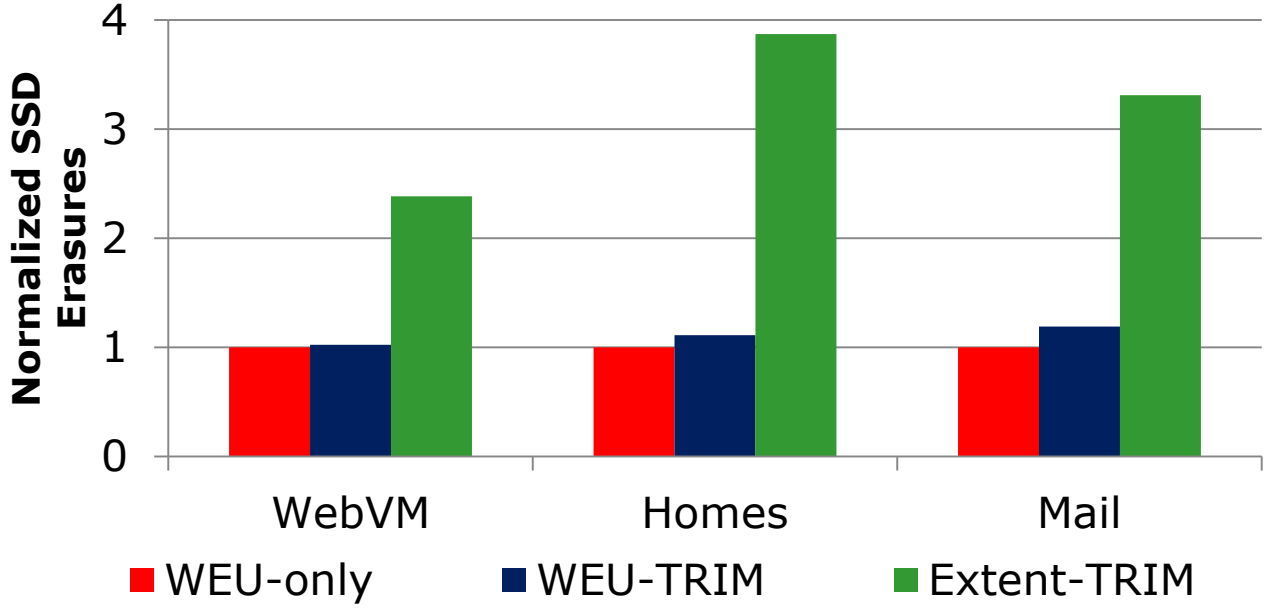

EMC²

RUTGERS

Additional Results

- Small hit-ratio increase leads to large IOPS boost
- Partial fingerprint index
  - Flexibility to trade-off deduplication, performance and RAM
- Nitro decompression has minimal overhead
- Sensitivity analysis
- Leverage content overlap for snapshot restore
- Deduplication reduces writes to SSD

EMC²

RUTGERS

Nitro Benefits

- Leverages deduplication, compression, and SSD performance
- Increases effective cache size
- Improves performance in two systems (**COS** and **TPS**)
- Balances performance, cost, SSD lifespan, and resources
  - Performance: up to 120% improvement
  - Writes to SSD: up to 53% reduction

EMC²

RUTGERS

Q&A

Nitro: A Capacity-Optimized SSD Cache for Primary Storage

Cheng Li*, Philip Shilane, Fred Douglis, Hyong Shim, Stephen Smaldone, and Grant Wallace

**Rutgers University* *EMC corporation*

EMC²

RUTGERS