Advances in SCSI for FLASH

Marty Czekalski

President, SCSI Trade Association

Interface and Emerging Architecture Program Manger, Seagate Technology

Overview

- Extended Copy Feature
- Atomic Writes
- SCSI-SF

Extended Copy Using Tokens

- New feature allowing direct movement of data between storage devices on the same fabric
  - Leverages the ability of SCSI devices to act as both an initiator and a target
- Greatly improves performance
- Greatly reduces overhead
  - Eliminates multiple passes of data over PCIe
  - Eliminates use of system memory as a buffer

Does Offloaded Data Transfer Work?

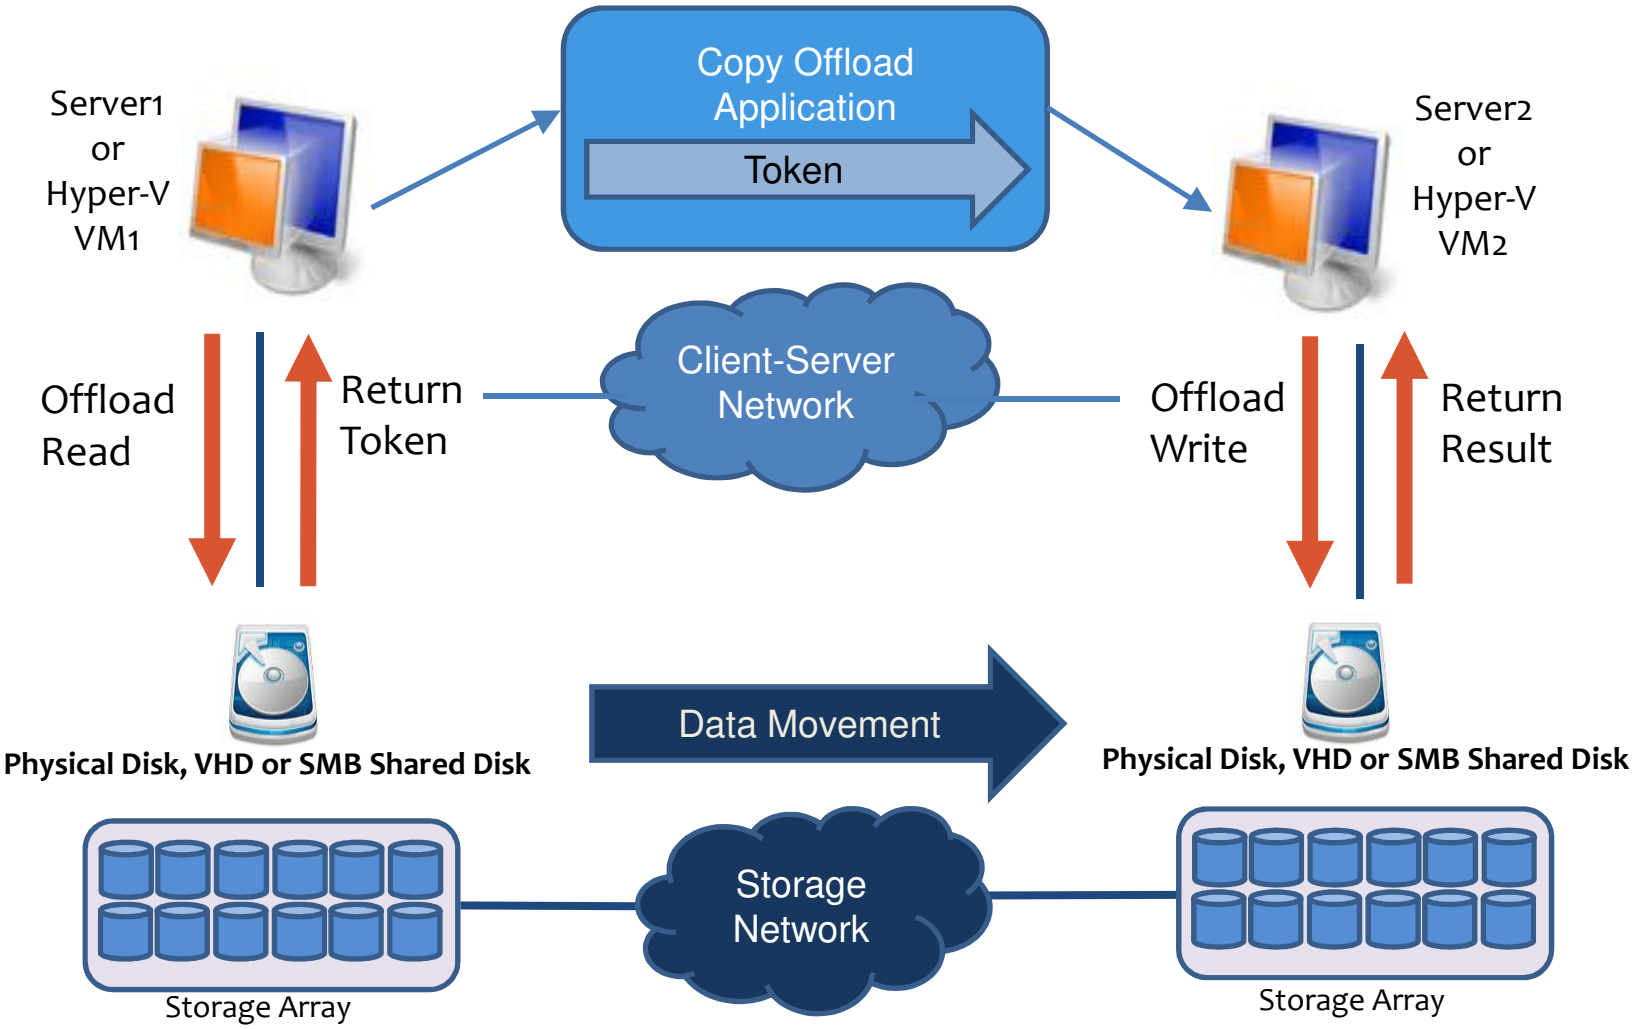

Extended Copy - Connecting the Tiers

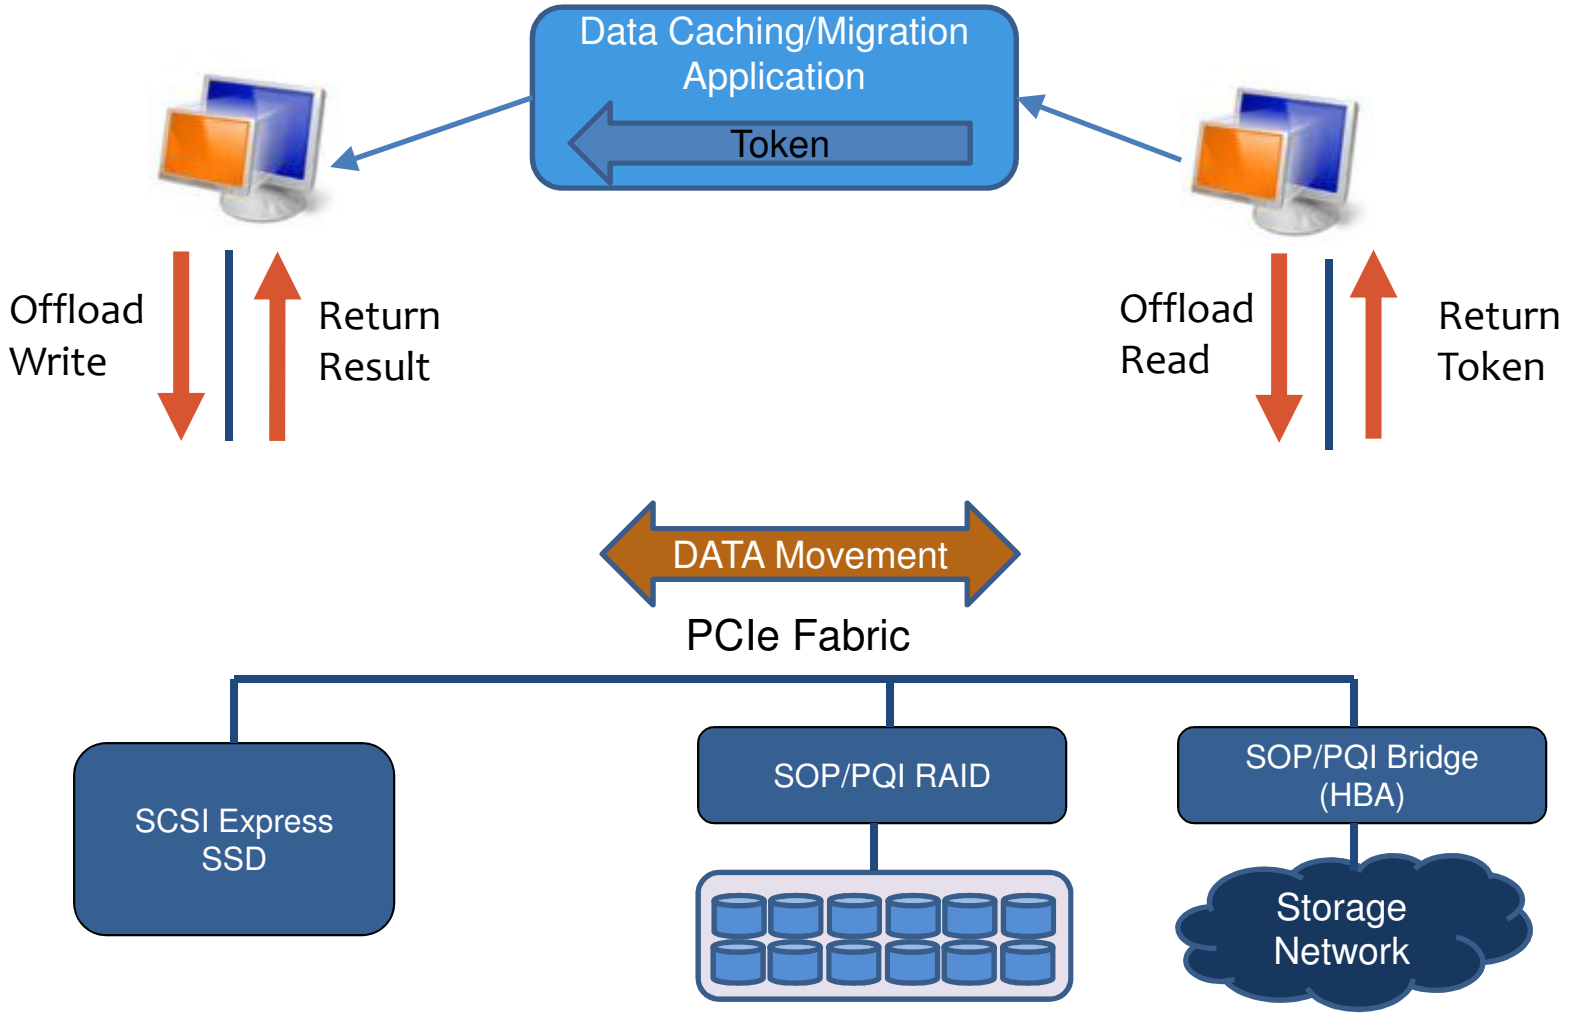

Atomic Operations

- Atomic Write – all or nothing is written
  - For single commands and across non contiguous LBA ranges (Scatter)
- Atomic Read - data read is consistent at a point in time
  - No partial updates in process
  - Multiple extents (Gather)
- Benefits:
  - Simplifies resilient system designs
    - Database, file system, etc.
  - Improves system performance in these applications

Atomic Writes

- Single extent Atomic Writes
  - General agreement on proposal
- Multiple extent (Scatter) Atomic Writes – two versions under discussion
  - Each extent is individually atomic, but no requirement that all extents be completed
    - Reduces overhead
  - All extents must be completed or none are competed
    - Additional programming efficiencies
- Discovery of capabilities for system/application use

SCSI-SF (Simplified Features)

- SCSI contains a rich feature set with multiple methods and options
- SCSI-SF is targeted as a common subset of features for increased efficiency of implementations, qualifications and maintenance
  - Base Feature Set
  - Provisioning Feature Set
  - Maintenance Feature Set

Hinting & Other NVM

- Pass “hints” to devices to make operations more efficient and increase performance
  - Targeted at SSDs and hybrid drives, but also useful for HDDs
    - e.g. UNMAP (TRIM), with variants
- Direct attached devices don’t need to continually OPEN and CLOSE connections
  - Can be implemented within the existing standard
  - Reduces latency on SSDs, SSHDs, and HDDs
- Alignment of commands and features with use cases of NVMP activities in SNIA

Summary

- SAS and SCSI continue to innovate and provide value in the storage ecosystem
- Activity is happening in real time
- Follow/participate in activities at T10 and SFF to make sure your designs take full advantage of these improvements