FlashNode: PetaScale I/O

Computer Architecture and Parallel System Laboratory,
Electrical and Computer Engineering University of Delaware, Newark Delaware

Problem

- HEC compute power advancing faster than I/O performance
  - Widening "I/O performance gap"

Integration into C64

- Connect high capacity Flash memory into every Cyclops-64 node's FPGA
- Flash will serve as:
  - Fast File Cache
  - Extension to shared memory hierarchy
  - Backup to a neighbor's Flash

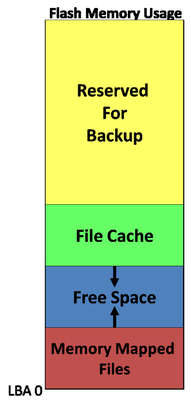

Persistent File Cache

- All files in at least 1 of 3 locations:
  - Primary (Local) Flash
  - Secondary (Backup) Flash
  - File Server (via I/O node)
- Dedicated Service Agent threads handle file ops

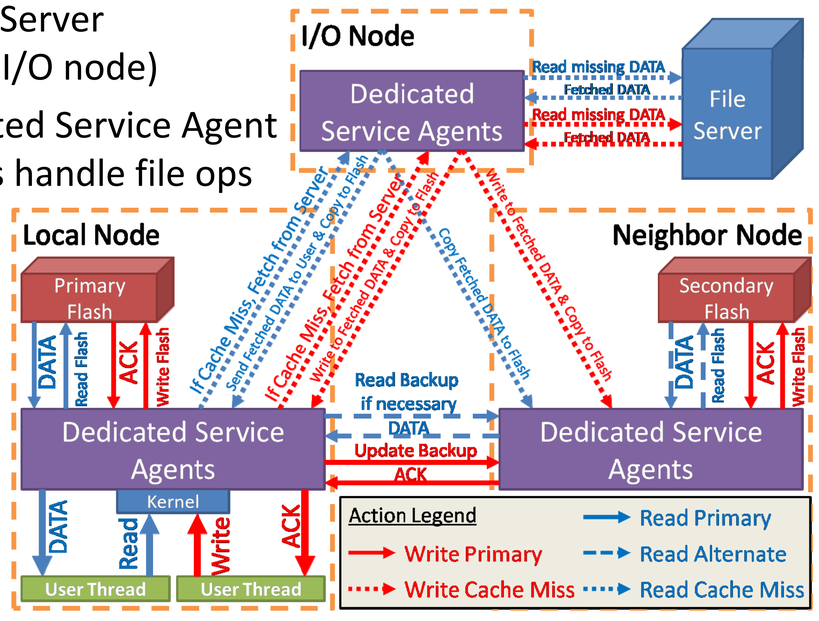

Objectives

- 4 objectives are to develop:
  - I/O architecture model for Peta-scale architecture, which utilizes NAND Flash memory
  - Efficient I/O middleware solution
  - RAS management model
  - Experimental testbed

Memory Extension

- Memory Map files on Flash into shared address space
  - FPGA translates addresses
- Dynamically share space with File Cache

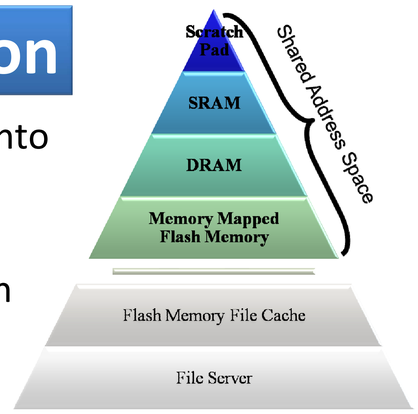

Cyclops-64 Storage Hierarchy

Self-Healing Backup

- Each node has designated backup neighbor
  - Only 2 Flash memories visible to each node
  - 4 Node Cycle backup pattern
- Divert to backup when Flash goes offline
- Self-heal when back online

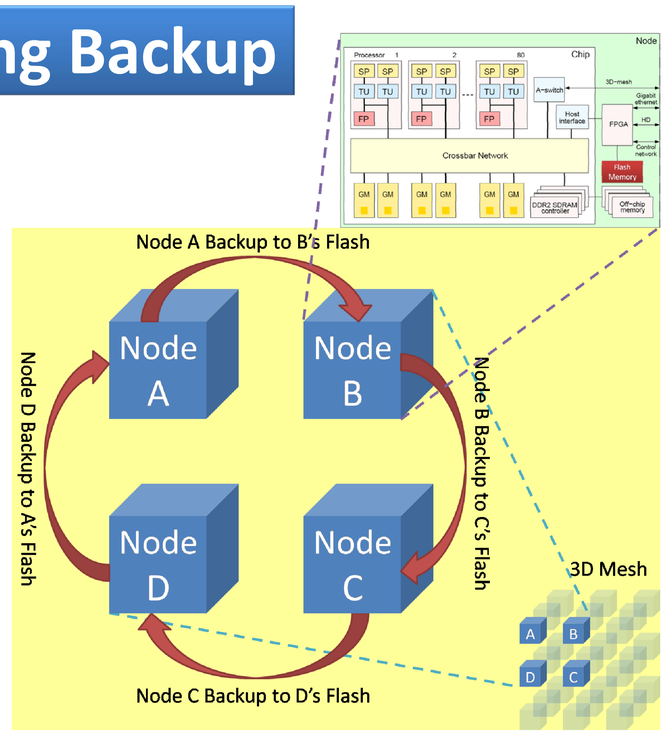

Thanks to ET International and Mihailo Kaplarevic

Poster created by Brian Lucas and Guang R. Gao