



Flash Controller Solutions in Programmable Technology

David McIntyre

Senior Business Unit Manager

Computer and Storage Business Unit

Altera Corp.

dmcintyr@altera.com

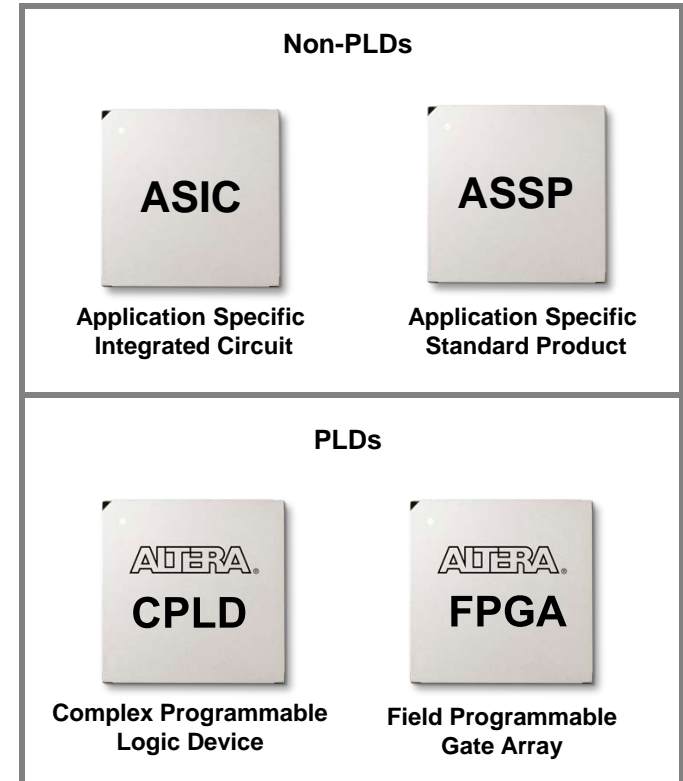


Overview

- What is a PLD?
- Programmable Technology for Flash Controllers
- Example Applications
- Notable Emerging Memory Types


What Is a PLD?

- A programmable logic device (PLD) is a type of semiconductor
- Most semiconductors can be programmed only once to perform a specific function
- PLDs are reprogrammable—functions can be changed or enhanced during development or after manufacturing



Flexibility Makes PLDs Lower Risk and Faster to Design Than Other Types of Semiconductors

PLD Tipping Point vs. ASICs



PLDs Outstripping Traditional ASICs in Technology and Total Cost of Ownership

Source: Altera; data applies to new design starts.

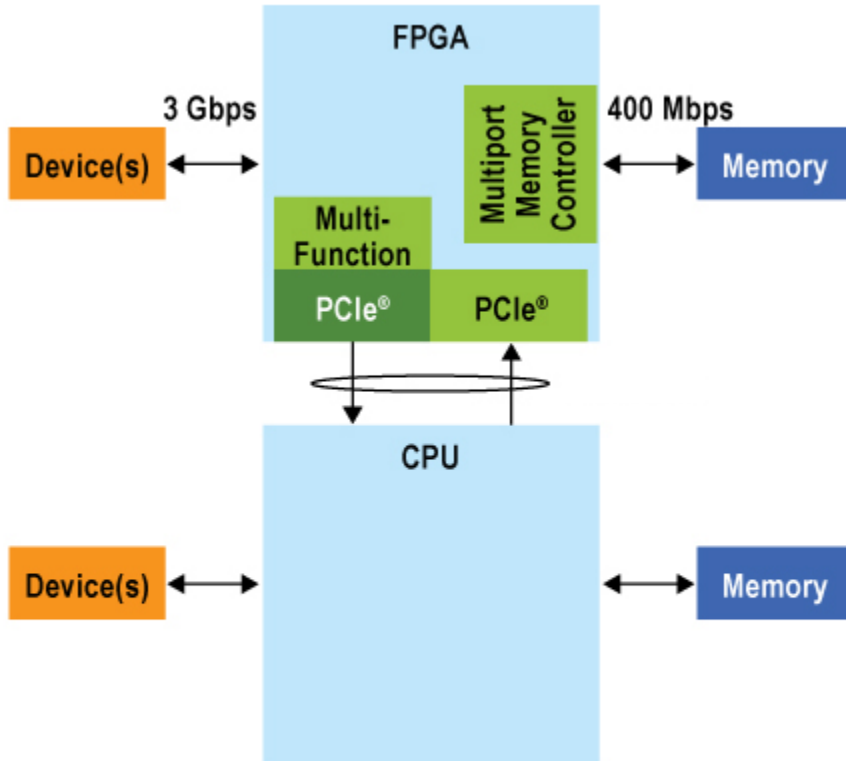


Enabling Technology- Programmable Logic Devices


- Design Logic support
 - Increasing densities to support system on chip (SOC) programmability
- Increased Computational Performance
- Reduced Power
 - Intelligent power management
 - Hardened IP blocks
- High Speed Serial Interface Support
 - Embedded Transceivers

Cost Reduction by Integration

Before

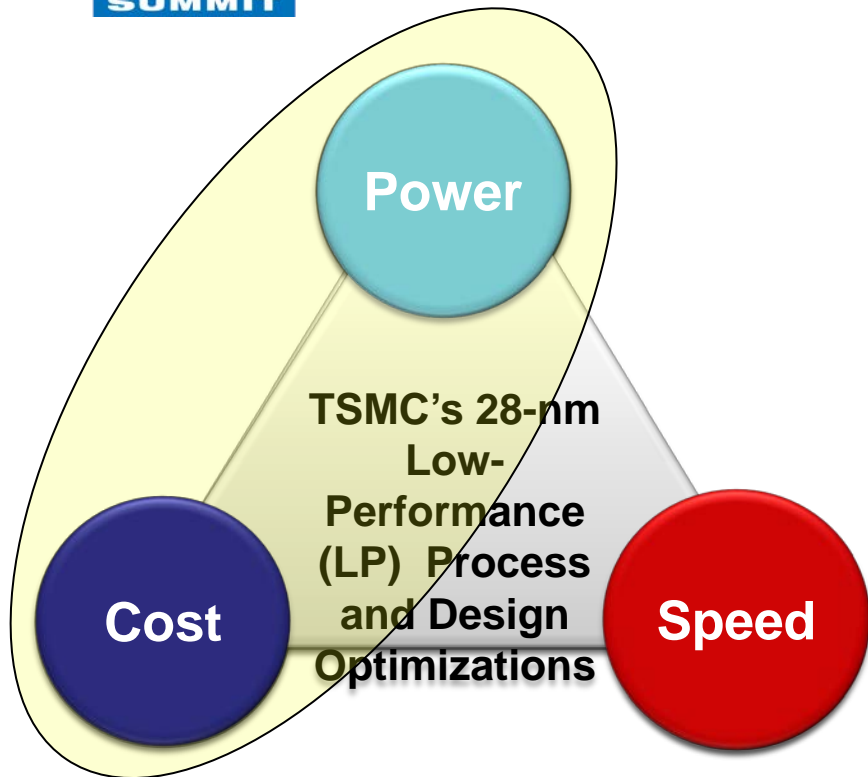


After

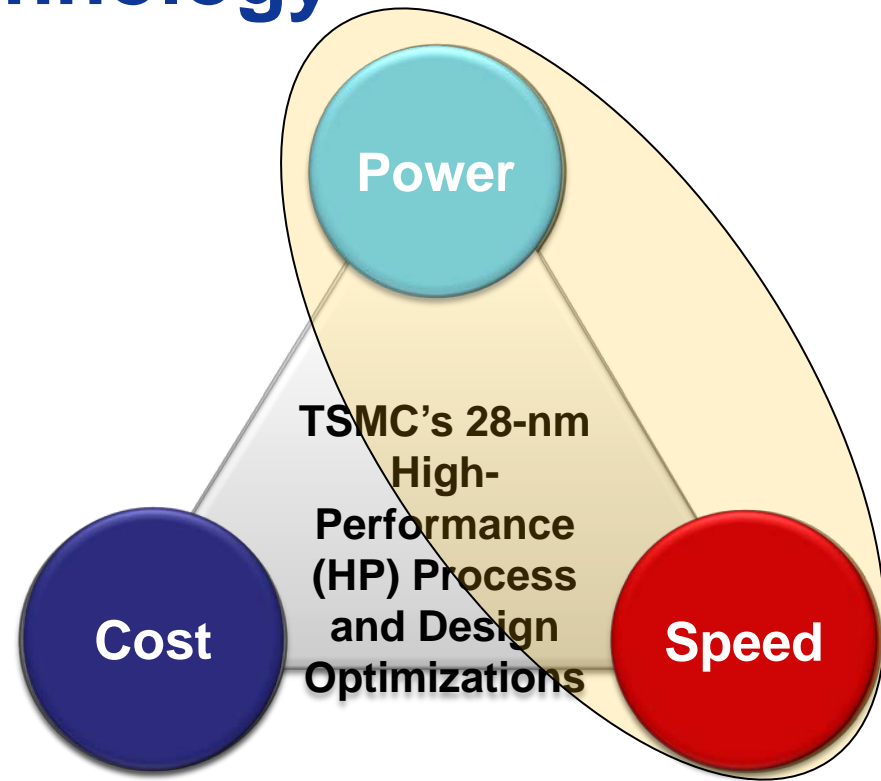


■ Soft IP
■ Hard IP

Balancing a FPGA Family by Process Technology



- The optimal choice for addressing today's power- and cost-constrained applications
- Lowest absolute power



- Highest bandwidth
- 28G transceivers at 200 mW
- Lowest power in high-performance systems



A Complete Solutions Portfolio

ALTERA®



PLD's Have Evolved!

The Lab








Prototyping
1-250 units

The Data Center





Production
10ku-1Mu

Data Center Applications- Servers

Application	Usage Examples
Flash SSD  Flash Controller	PCIe to ONFI bridging, Flash Control
Acceleration  Accelerator Card	Algorithm acceleration for vertical markets
Bridge Plus	Interface bridging with IP function, e.g. compression and encryption, Dedupe
I/O Virtualization (10GbE and PCIe) 	ASIC alternative; low cost with flexibility
Co-ASIC  Mainframe	Features enhancement
Management (BMC, KVM)  Blade Server	IP Flexibility supported with low power

Data Center Applications- Storage

Application	Usage Examples
Flash Cache/SSD  Memory BackUp/Restore	ONFI bridging and RAID adaptor NV DIMM backup, RAID for Flash
RAID Bridging	PCIe Gen 3 x8 best of class signal integrity
Bridge Plus	Interface bridging with IP function
ASIC Replacement  Tape	Lower cost development with flexibility



Flash Controller Applications



Flash Controller Requirements

- Uncertainty Favors PLDs for Flash Control Solutions
- Flash Challenges Continue
 - Data loss, slow writes, wear leveling, write amplification, RAID
- Many Performance Options
 - Write back cache, queuing, interleaving, striping, over provisioning
- Many Flash Cache Opportunities
 - Server, blade and appliance

Flash Controller Design Challenges

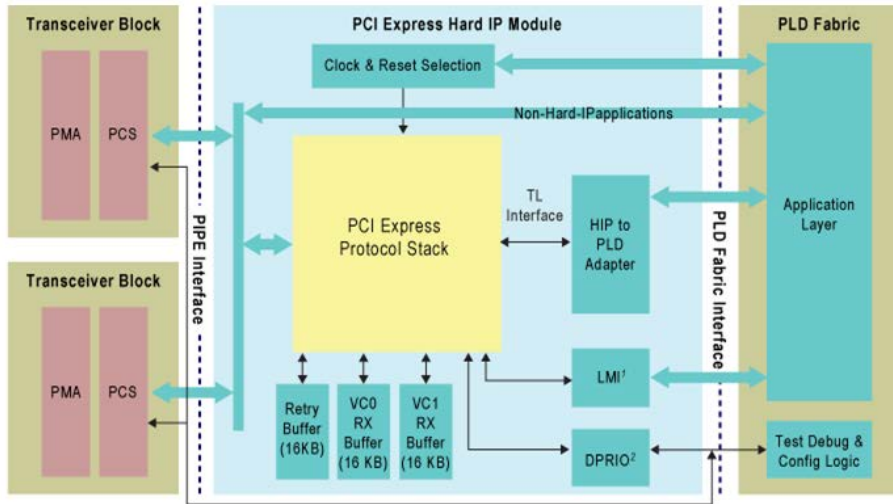
- Emerging memory types
 - ONFI 3.0, Toggle Mode 2.0
 - PCM, MRAM
 - DDR4
- Controller Performance Options
 - Write back cache, queuing, interleaving, striping
- ECC levels
 - BCH encryption
- Data transfer interface support
 - PCI Express, SAS/SATA, FC, IB



Flash Controller Support

IP	Sources	IO	Speed	Logic Density	Comments
ONFI 3.0	SLS, Cadence	40 pins/ch	400 MTps	5KLE/ch	NAND flash control, wear leveling, garbage collection
Toggle Mode 2.0	SLS, Cadence	40 pins/ch	400 MTps	5KLE	Same
DDR3	Altera	72 bit	1066 MHz	10KLE	Flash control modes available for NVDIMM
PCM	Micron			5KLE	PCM- Pending production \$
MRAM	Everspin			5KLE	MRAM- Persistent memory controller (Altera based)
PCIe	Altera	G3x8	64Gbps	HIP	Flash Cache

PCI Express Support



PCIe Mode	Thruput (GT/s per lane)	Production
Gen 2	2.5	Now
Gen 3	5.0	Now
Gen 4	8.0	2016

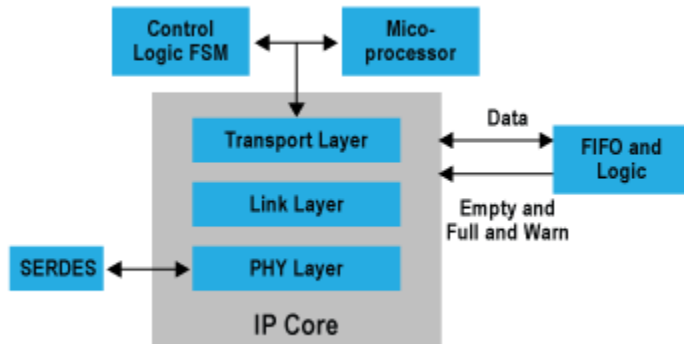
Note:

1. LMI: Local Management Interface
2. DPPIO: Dynamic Partial Reconfigurable Input/Output

Hardened IP (HIP) Advantages

- Resource savings of 8K to 30K logic elements (LEs) per hard IP instance, depending on the initial core configuration mode
- Embedded memory buffers included in the hard IP
- Pre-verified, protocol-compliant complex IP
- Shorter design and compile times with timing closed block
- Substantial power savings relative to a soft IP core with equivalent functionality

SAS/SATA IP



- Flash Controllers can manage SAS/SATA SSD interfaces
- 12Gbps SAS support required for enterprise drives
- FPGA transceivers need to support electrical performance and OOB signaling

Flash Controller Example

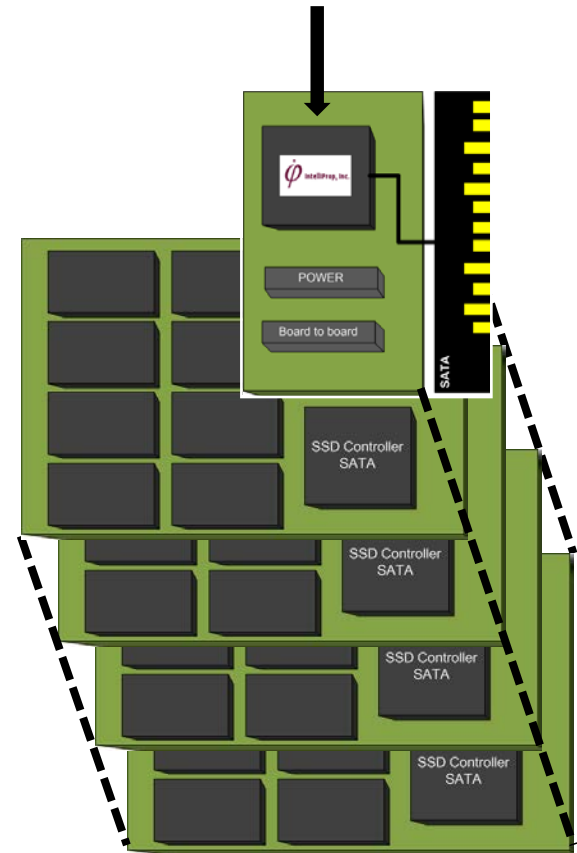
Intelliprop Hydra can provide 4X the capacity and up to 4X the performance of a single drive



Single SATA



HYDRA




**2 or 4 Stacked
drive cards**

Applications

- SSD manufacturers with a need to build larger drives (1TB, 2TB, 4TB)
- Applications requiring stable read/write bandwidth
- Enterprise systems requiring high IOPS and bandwidth
- Host systems writing large amounts of “raw” data
- Systems where economical controllers can be used while still achieving performance of “enterprise” class controllers

MBps
70% Random Read / 30% Rand Write

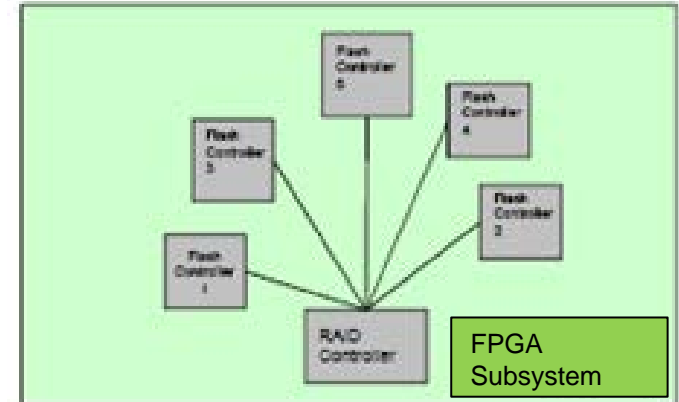


Test Conditions

- Non-compressible data
- Testing done using IOMETER 2010
- Testing with HYDRA in 4 port configuration using 32-bit DRAM
- 4kB page size flash
- Host PC: Intel Core i7-3770K Ivy Bridge 3.5GHz

FPGA RAID Controller for Flash Arrays

- 10+TB plus of Flash Storage
- FPGA Applications
 - Flash Control
 - RAID
 - Data Transfer



PLD Violin Product Portfolio

V3000: Performance Apps

- Any SAN/PCle
- High Reliability
- Low entry cost
- 5TB-20TB per 3U
- 100K – 200K IOPS



Stackable



V6000: Primary Flash Storage

- HA in box
- High bandwidth
- Integrated SAN
- 6-32TB per 3U
- 200K-1M IOPS



Integrated

Technologies



Block Storage with full performance and flexible LUN management



Array management software with Web, CLI, Email and XML & Clustering




Hardware Flash RAID with 4+1 efficiency



Violin Intelligent Memory Modules (VIMMs) with distributed grooming

High Availability

No Single-Point-of-Failure and Hot Swap Everything



Flash Cache

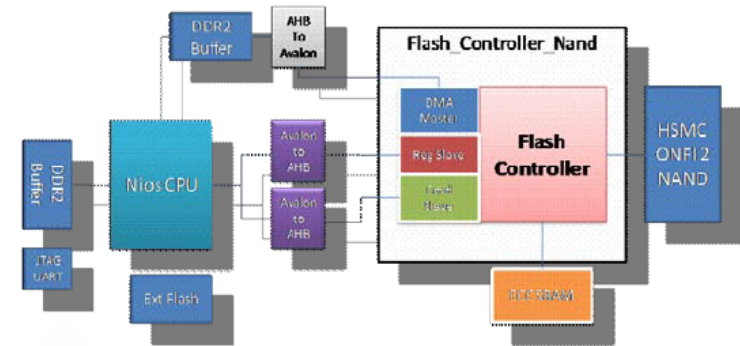
PCIe or Appliance

■ Benefits

- High performance for cache
 - Better bandwidth vs SSD
 - PCIe up to 5X
- Lower power than DRAM
- Less space than disk based
- Hardcopy (structured ASIC) path

■ Solution

- Altera
 - ONFI 2.0 & Toggle Mode cards
 - ONFI 3.0 daughtercard (3Q12)
 - ONFI 3.0 App Note (3Q12)
- Partner
 - Denali/Cadence IP
 - SLS for single channel controller



Denali Flash Controller



Nallatech
PCIe Gen 3x8
Card

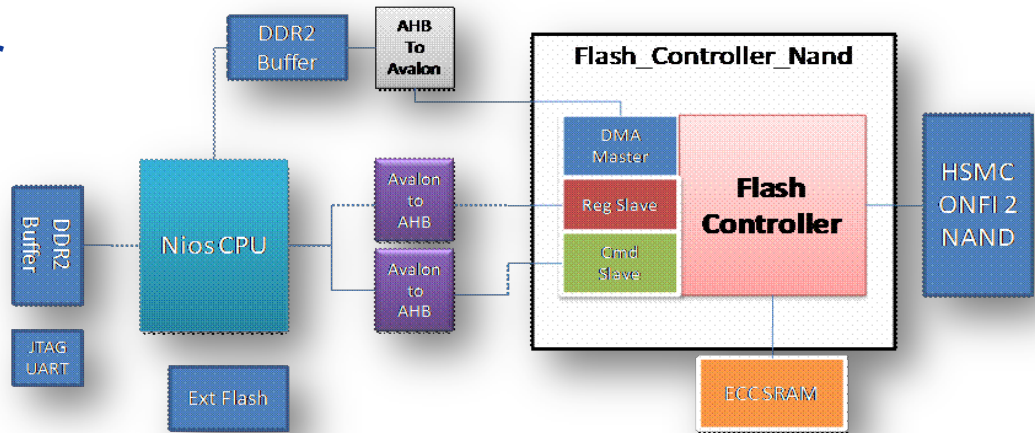


Micron ONFI Daughtercard

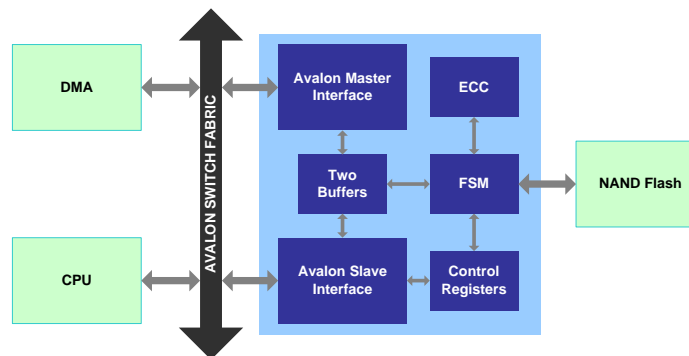
Flash Cache Controllers

Denali Multi Channel Controller

- Single to multi Flash channel capability
- Basic NAND development platform
- Provides High Speed ONFI & Toggle NAND PHY
- ECC of 8 and 15 bits of error correction



Third Party Single Channel Controllers



Data Integrity- The Green Approach

Battery Backed Data Recovery



Add-on modules that protect against data loss in the event of a server or power failure by providing emergency power to the cache memory. When power is restored, the data not yet written to the hard drives can be retrieved from cache memory.



Super Capacitor Based Alternative

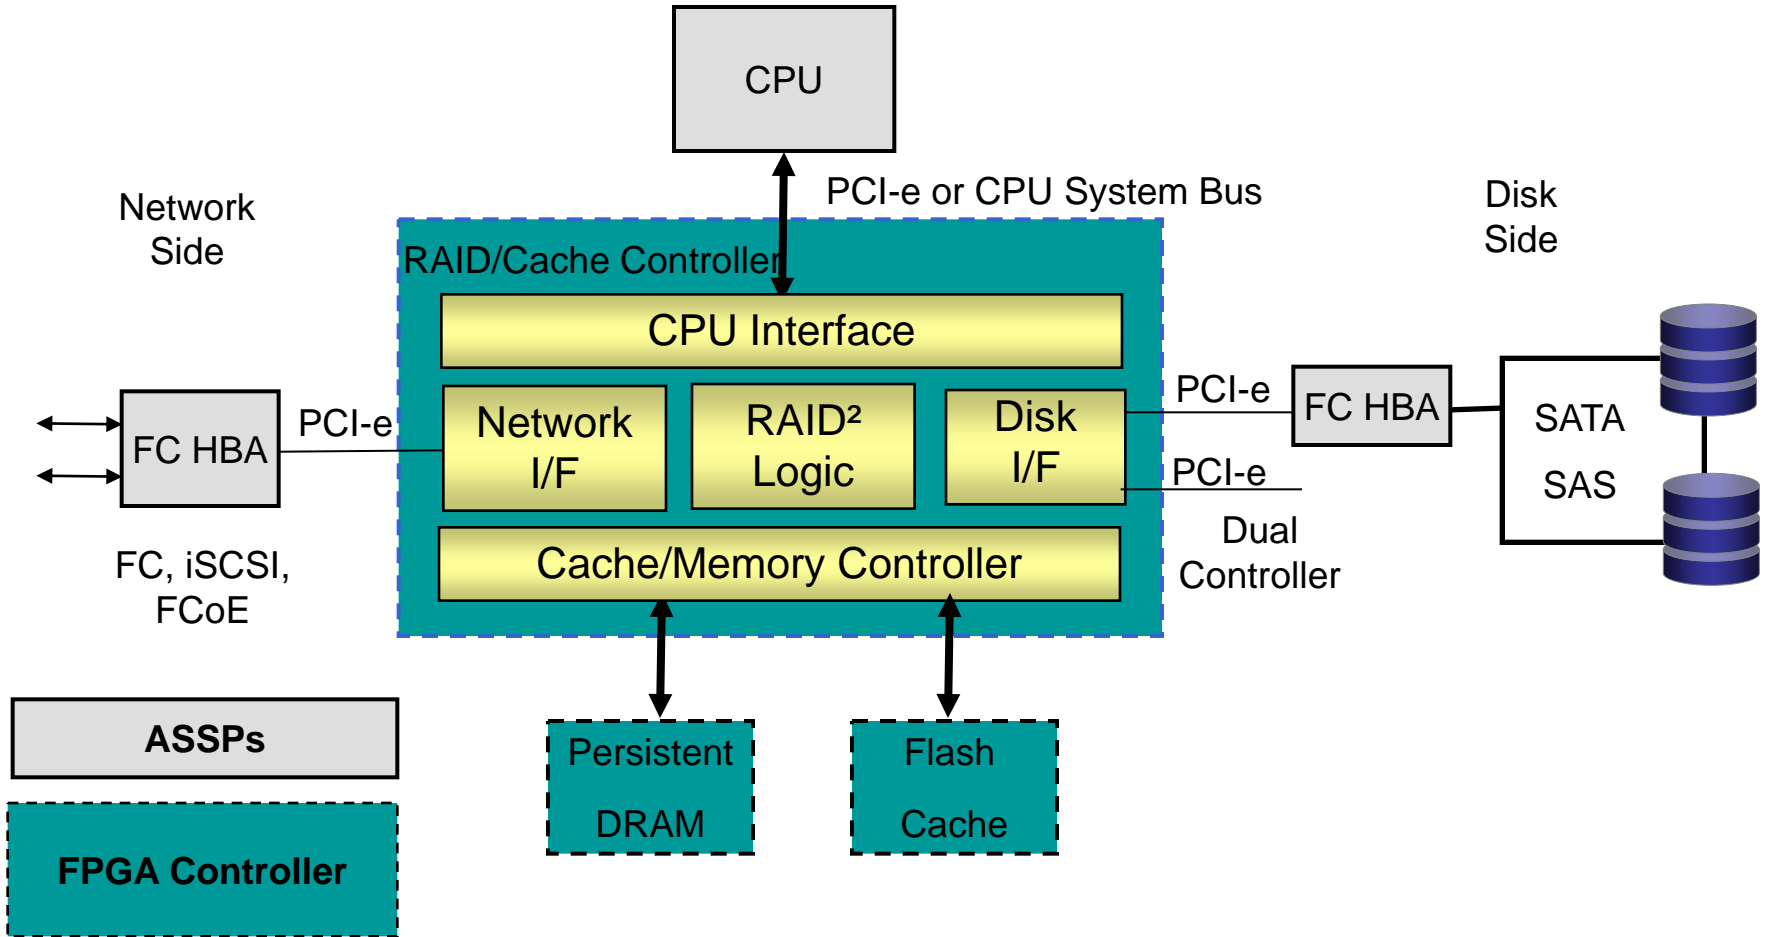


Benefit	Battery Power Source	Super Cap Power Source
Less Cost		<input checked="" type="checkbox"/>
Lower Power		<input checked="" type="checkbox"/>
Smaller Footprint		<input checked="" type="checkbox"/>
Field service required	<input checked="" type="checkbox"/>	
Permanent backup		<input checked="" type="checkbox"/>



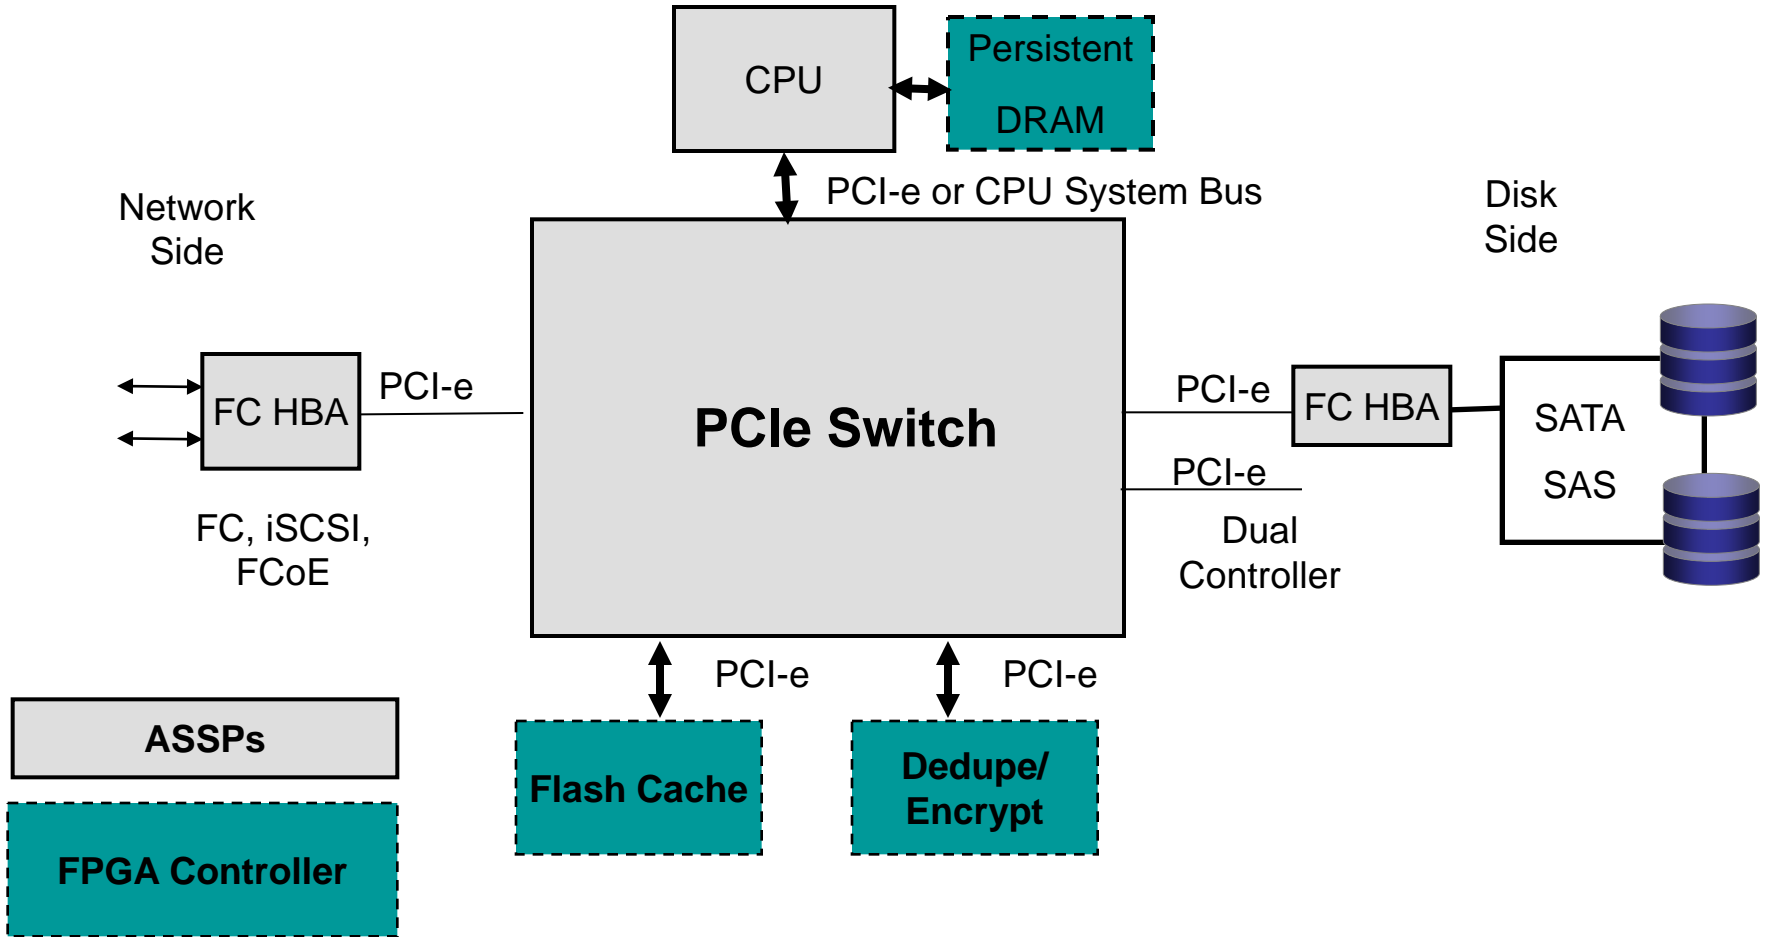
Hybrid RAID System

- Persistent DRAM and Flash Caches





Hybrid RAID System - PCIe Switch Centric





Flashing Forward

- Programmable Logic Devices (PLDs) are a well suited technology for emerging flash memory controller requirements.
 - Multiple memory types
 - BCH enhancements
 - Emerging technologies (MRAM)
 - High Speed Interface Requirements (PCIe, SAS)

Visit the Altera Exhibit Booth to Learn More!