# Low Latency Hardware Accelerators for File Systems

Zacodi Labs
OpenZFS Summit 2014

# Motivation

- Compute intensive tasks
  - Cryptographic hash SHA256
  - Data compression
  - Encryption
- Used in many applications: storage, networking, etc.
- Existing hardware acceleration solutions
  - High overheads and latency
  - A lot of parallelism/concurrency required to achieve throughput
  - Not usable for latency sensitive applications

# Compute Intensive Tasks

ZFS example: SHA256

- Cryptographically strong hash function
- Provides checksum for data integrity
- Aids redundancy detection and deduplication
- Uses almost all CPU time:
  - 6-core Sandy Bridge
  - 0.6-0.7 GB/s

# Hardware Acceleration Approach

Offload compute intensive tasks to FPGA

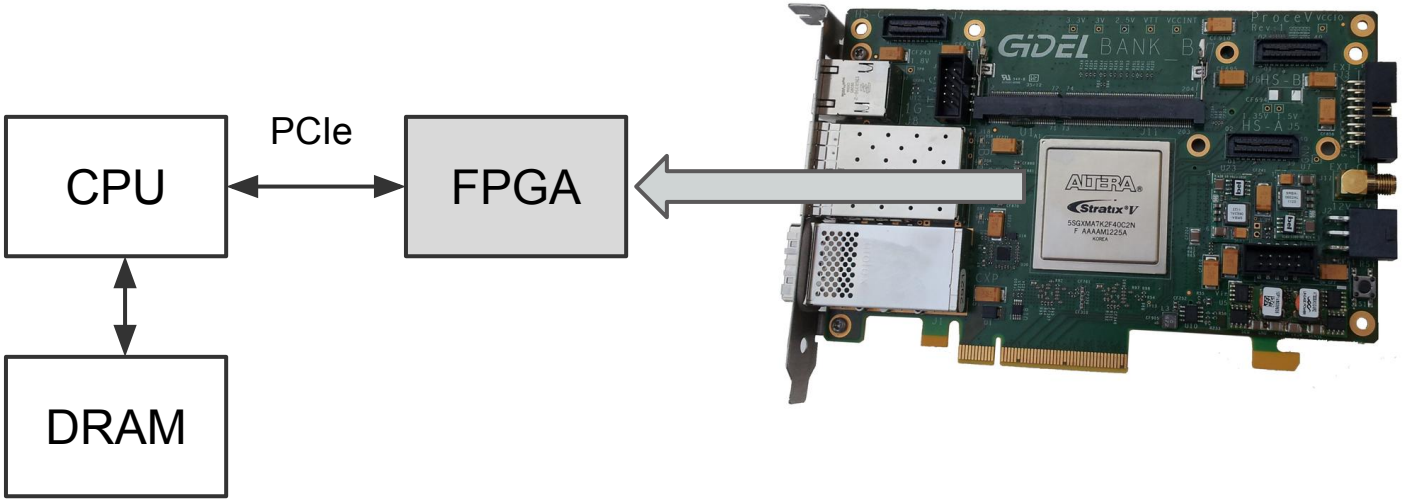

# Hardware/Software Interface

- **Low overhead/Low latency**
  - minimum number of PCIe round trips
  - overlap DMA and SHA256 computation
- Zero copy
  - no additional buffers/copies
  - no buffer management, allocation, de-allocation, etc.
- No DMA descriptors in memory
  - no descriptor management, synchronization, etc.
- Interface optimized for multi-core CPU architecture

# Integration with ZFS

- Open source ZFS (Linux, Illumos)
- PCIe driver for FPGA - both Illumos and Linux are supported
- Driver modifies SHA256 function pointer in ZFS
- SHA256 function is a callback into driver
- **No ZFS source code changes required**

# ZFS performance comparison

CPU only

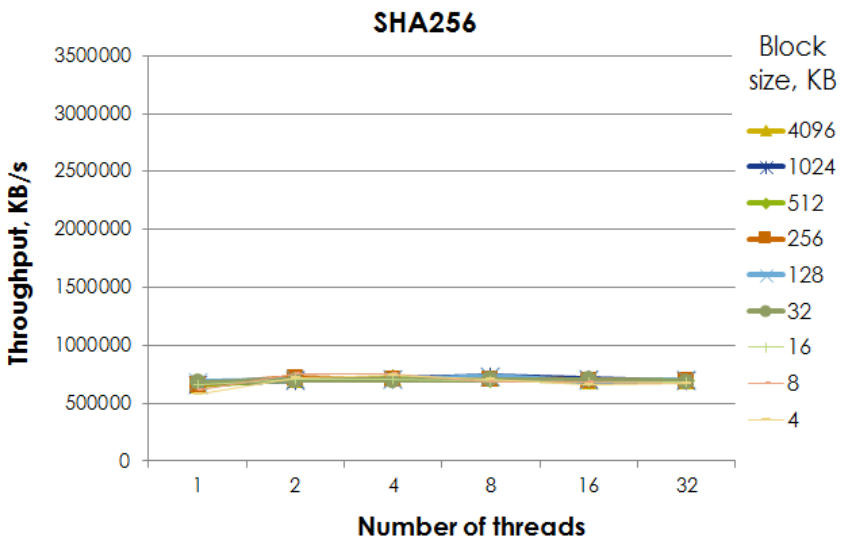


CPU + Arria V FPGA

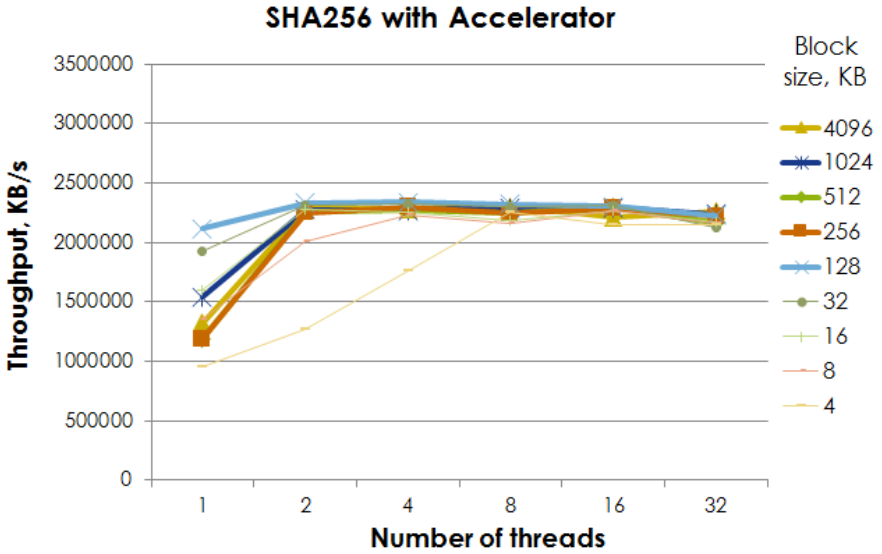

# Latency comparison

- SHA256 computation latency of 4KB and 64KB blocks of data
- Comparing CPU and accelerator

| | **4KB** | **64KB** |
|---|---|---|
| **Sandy Bridge Xeon CPU** | 40.5 usec | 650 usec |
| **Arria V FPGA** | 38.5 usec | 500 usec |
| **Stratix V FPGA** | 21.5 usec | 250 usec |

# Latency vs. Throughput

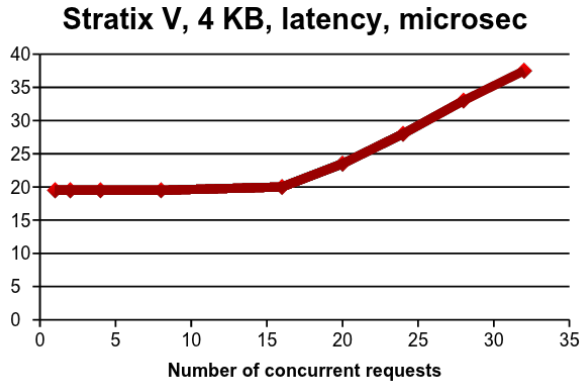


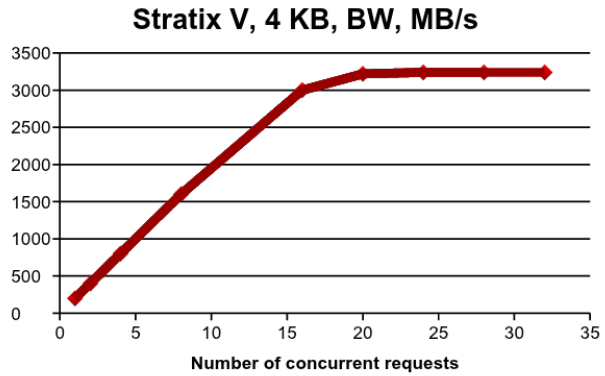


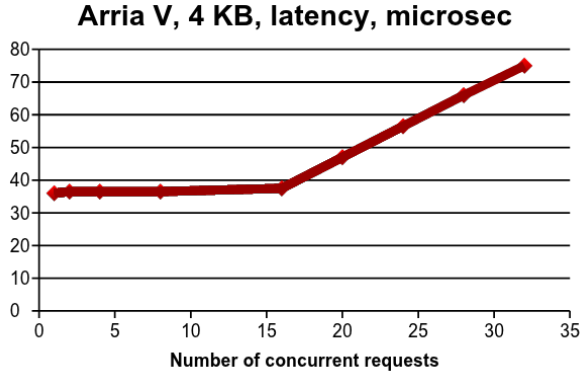


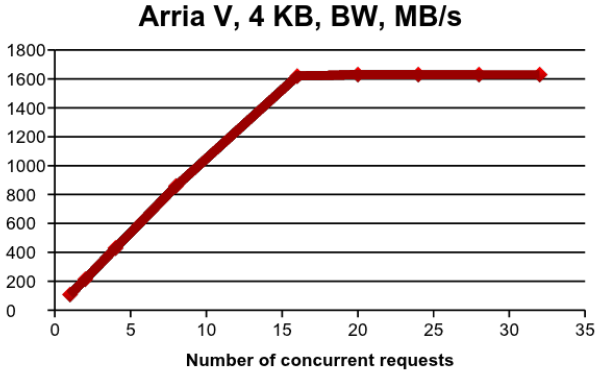

# Further extensions

- Compression
- Sliding window hashes
  - Rabin fingerprints, used for advanced deduplication
- Encryption
- Combinations of the above
  - e.g. parallel checksumming and compression

# Conclusions

- High throughput hardware accelerator
- Low overhead, low latency
- Simple integration with ZFS
  - no changes in file system necessary
- Demonstrated 3x throughput improvement
- Interested? Write us at *info@zacodi.com*