

Platforms for SSD and IoT Applications

Amit Saxena, VP Engineering

4/9/2019

The Mobiveil Team



Vision

Provide Technology, Platform and value added services to develop storage solutions

- RTL IP
- FPGA Platform
- Engineering services for Custom Designs

Leadership

Management with 30+ years experience in Semiconductor/ Silicon IP/ Product Engineering Services

Team working together developing Silicon IP & Engineering Services for 15+ years


- Silicon valley “Fast 50” in 2018
- 2018 Inc. 5000 “Fastest Growing Private Companies in America”
- 10 Most Promising Solution Providers in Storage by CIO Magazine
- 4 Patents in Storage and Flash Reliability

Location

Headquarters in Milpitas, Engineering Centers in Chennai, Bangalore and Hyderabad. Total headcount ~225


Mobiveil Platform Focus Domain

Flash Storage




- Enterprise
- Data Centers
- Laptop/Consumer
- Mobile

AI/ML/IoT



- Data Centers
- Smart Cities
- Industrial IOT
- Consumer/Automotive

Communications



- Base Stations
- HPC
- Industrial
- Aerospace

Mobiveil Vision for Application Platforms

- What are platforms
 - Bundle of highly configurable Silicon IP blocks, Pre- integrated
 - Pre-Validated FPGA platform
 - Operating Firmware
 - Pre-verified hooks to 3rd party IPs like Verification IP, PHYs
 - Pre-verified environment for the design flow like the Emulation
- Benefits of Platform
 - True acceleration of product development as much of the integration and verification is already completed
 - Reduction in cost and schedule, Product development risk minimized
- SSD Platform Example for customers developing SSD SOC
 - IP Blocks: PCIe Gen4, NVMe, DDR4, ONFI, LDPC, Custom Blocks
 - Xilinx/Intel FPGA based Hardware platform
 - Operating Firmware for ARM and RISC-V*
 - Hooks to Avery VIP
 - Standard Interfaces to PCIe,DDR4,ONFI PHYs

Understanding Our Customer's Challenge

To enable early availability of IPs and Platforms

Energy

- Energy cost due to data movement will be comparable, if not more than the computing cost.


"Exascale Computing Study: Technology Challenges in Achieving Exascale Systems", Peter Kogge, editor and study lead, 2008.

- Energy efficiency will become the primary metric for system design since we need to increase the computing power by 1000x with only 10x increase in the power envelope.

U.S. Department of Energy. DOE Exascale Initiative Technical Roadmap, December 2009.

I/O Bottleneck

- Multi flash chip architecture accelerating SSD bandwidth while host I/F lags and CPU performance hits Moore's Law
- Current direction is PCIe Gen 3→4→5
- More and faster pipes to bring data to the host



Challenges and Constraints



Compute and Storage Server Bottlenecks

- **IO Bottleneck:** Memory speeds \gg host speeds
 - Increasing flash die further increasing memory speeds
- **Cycles and power** spent moving data from storage to server
 - IO power sending from storage to server for processing
 - Network and protocol overhead (NFS, iSCSI, rDMA, etc)
- **Power** for moving data within server
 - NIC \rightarrow DRAM \rightarrow L2/L1 cache \rightarrow uP ... and back out
- **Inefficiency** \rightarrow CPU and power to move “sparse” data

Network Bottlenecks


- **Power** increasing with more data and data movement
- **Latency and congestion** increasing
 - High data movement
 - Storage \rightarrow compute \rightarrow storage
 - Latency of on-chip buffering in switches and NICs
 - Over subscribed core \rightarrow dropped packets and retransmission
- **Increasing data** driving bandwidth
- **Inefficiency** \rightarrow High bandwidth to move “sparse” data
- Data overhead: NFS, iSCSI, security, etc

Analogy

Current solution: build more roads, bigger roads, buy more gas

Our solution: Enable People to work at home

Example - HPC Scientific Modeling Application



"Jaguar"

**224,256 cores
(1.75 petaflops)**

**240 GB/s total
= 1MB/s per
core**

10 PB storage

**Data moved three
times**

- 1 Generate data**
- 2 Analyze data**
- 3 Update data**

**Power: 5-10
megawatts**

Other Applications

- Climate
- Genetic
- Finance
- Security
- Medical
- Fraud
- Retail
- Automotive

Boboila, Simona, et al. "Active flash: Out-of-core data analytics on flash storage." Mass Storage Systems and Technologies (MSST), 2012 IEEE 28th Symposium on. IEEE, 2012.

Nearly 70% Power Savings Running on SSD vs Server

Data analysis using spare cycles of existing SSD controller

Additional offload possible with optimized Storage IQ architecture

Additional power savings if data is processed in DRAM before moving to flash

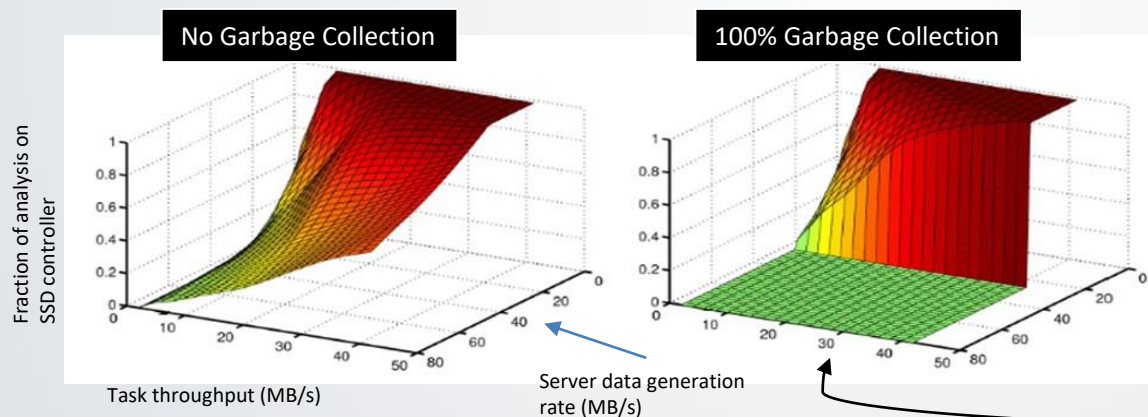
Observations:

Nearly 70% power savings vs sending back to server: reduced IO, reduced re-read and re-write of data

No performance impact: analysis runs in parallel with compute application running on server

Many tasks possible to offload on SSD controller

App performance will go down over time (i.e. as drive fills) due to garbage collection



Task

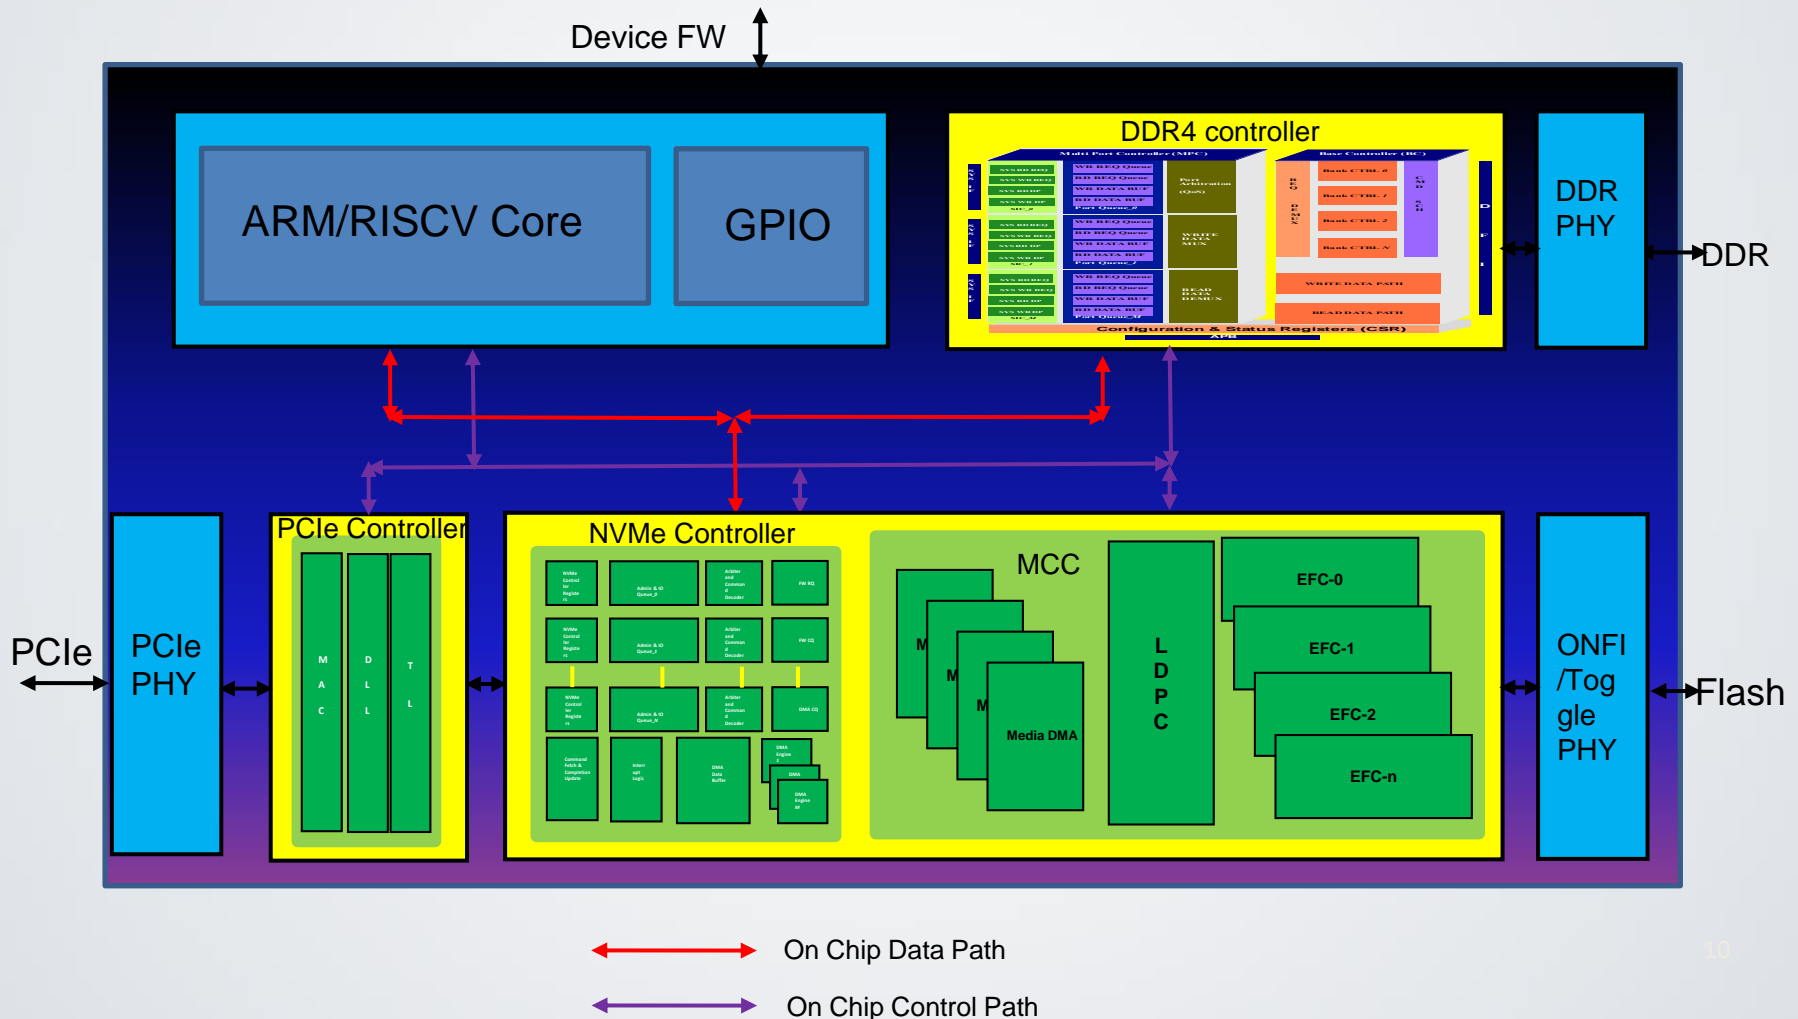
| Task throughput on SSD ARM Cortex A9 (bigger numbers are easier tasks) | |
|--|----------|
| Statistics (mean) | 416 MB/s |
| Pattern Matching (grep) | 123 MB/s |
| Data formatting (transpose) | 76 MB/s |
| Compression (LZO) | 41MB/s |
| Dim. reduction (PCA) | 10 MB/s |
| Edge detection (vision) | 7 MB/s |
| Compression (gzip) | 4 MB/s |
| Dedup. (Rabin fingerprint) | 2 MB/s |
| Clustering (k-means) | 1 MB/s |

Sparse Data Well Suited to StorageIQ SSDs

- Sparse data: Important but rare
 - Benign, 0s, empty cells, much can be ignored
- Processing sparse data
 - Divide and conquer large databases
 - Find interesting data and process further
 - Apache MapReduce, MPI, Dryad, Pregel, GraphLab, Piccolo
- Find interesting data using basic commands
 - Scan / filter / sort / group / hash
 - Eliminate trivial items
 - Compress (e.g., empty cells, 0s, etc.)
 - Clean / pre-process (e.g., normalize data, etc.)
- Numerous applications in machine learning and data management

| | |
|-------------------|--|
| Security | Network intrusions, attacks, misuse traffic = small fraction of total network traffic |
| Fraud | Credit card, telecom, subscription, e-commerce, insurance fraud = small fraction of legit transactions |
| Medical | Tumorous pixels in overall image ("Video CoSeg") |
| | Medication errors (0.3% - 5%) |
| | Adverse drug events (0.05% – 6.7%) |
| | Infection analysis (5.9 of 1000 days) |
| Retail | Recommending products out of millions available |
| Movies | Individual movie recommendations out of millions of movies available on Netflix, Amazon, Vudu |
| Scientific | Analyzing genetic signatures across petabytes |
| IOT | ARM projecting 1 trillion IOT nodes; most will be reporting "Situation normal" |
| ADAS | Accidents in 270+B / month vehicle miles in US |
| Photos | Hundreds of millions of photos uploaded each week |

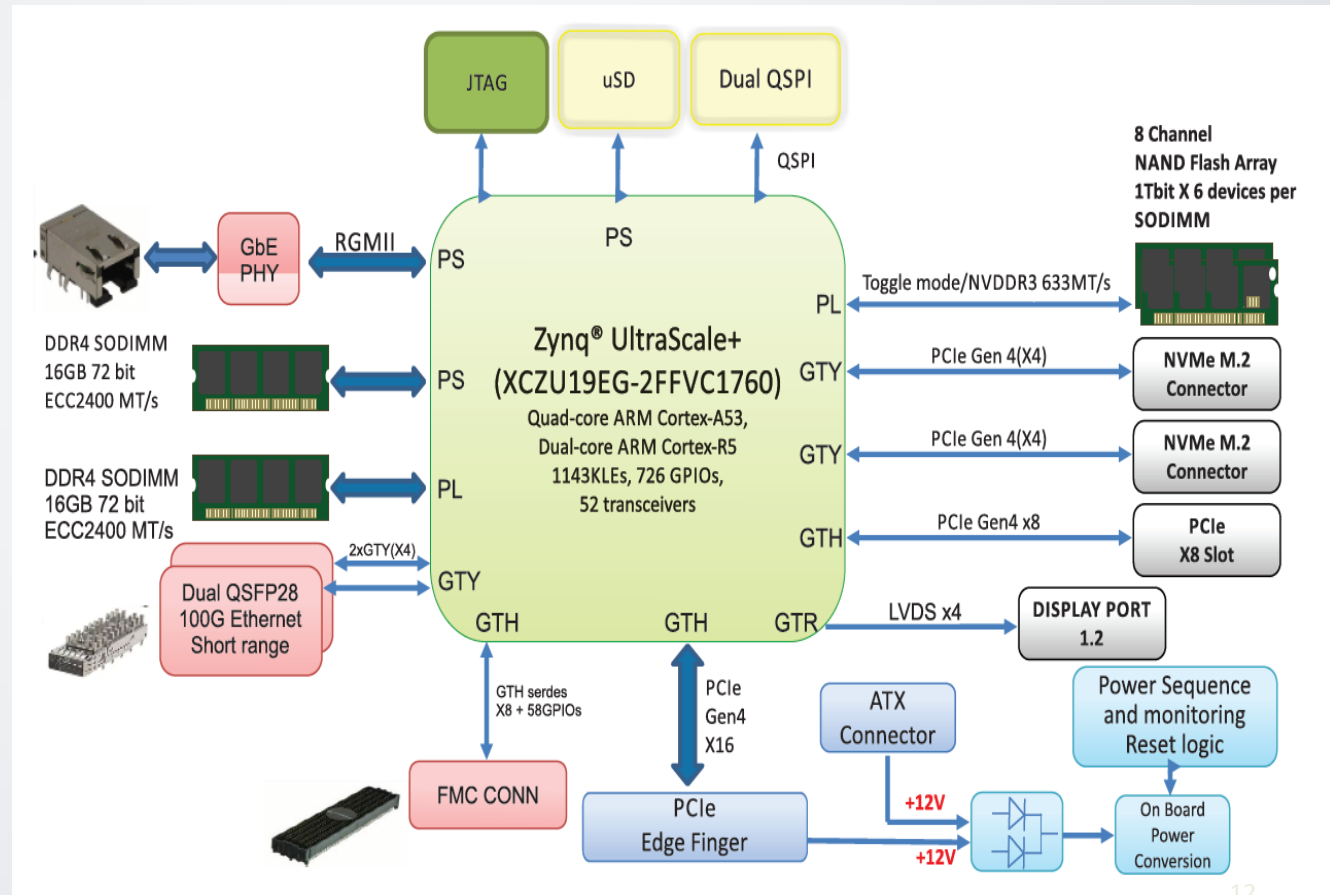
NVMStor-Ultra Configurable NVMe SSDC Platform



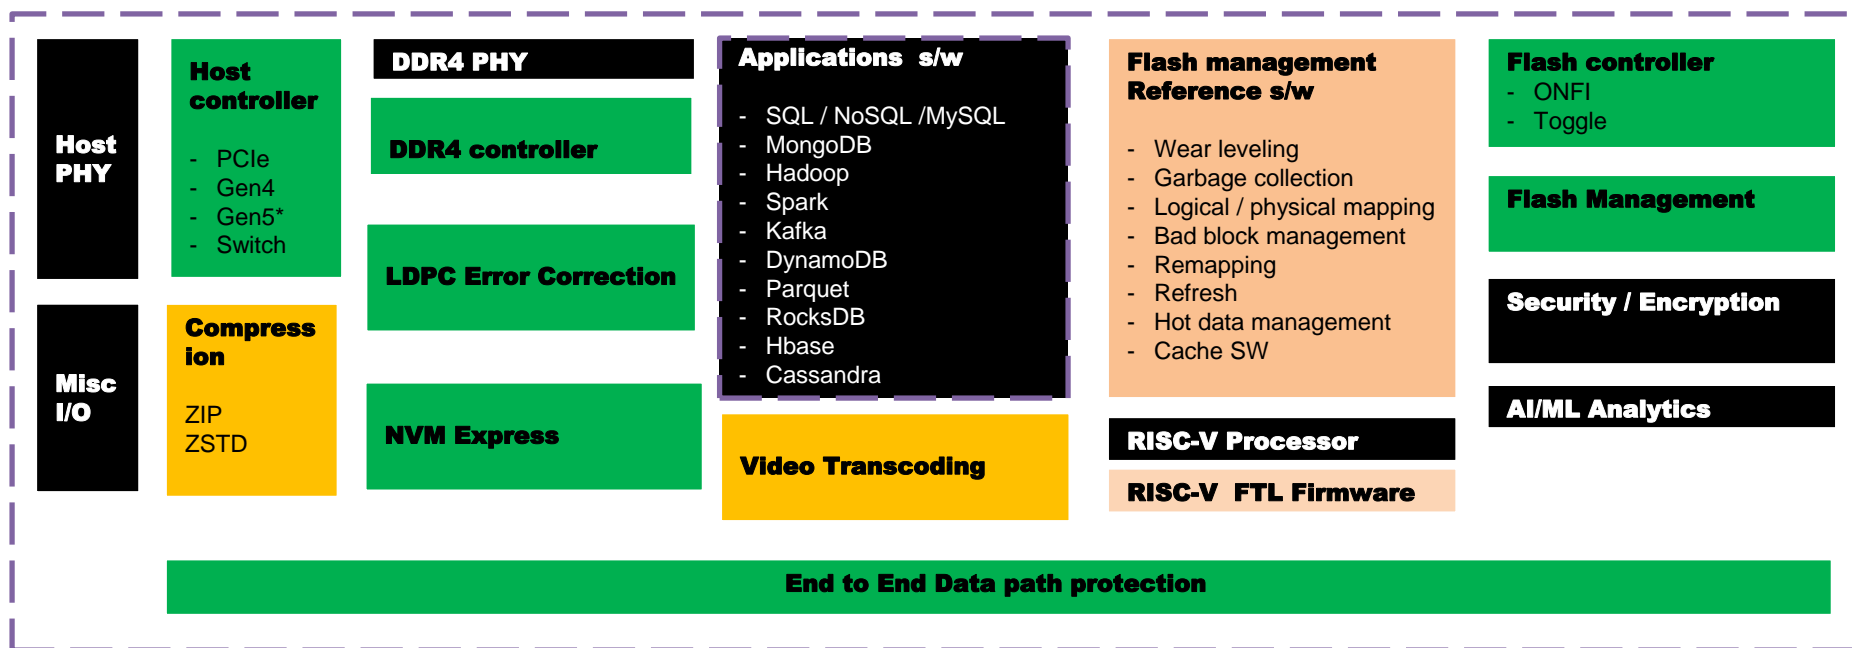
Unique Subsystem Development Solution

- **Provides Full NVMe Based Reference Design Using Mobiveil's Controllers**
 - **PCIe Gen4.0 PCIe Controller (GPEX)**
 - **Multiport NVMe (UNEX)**
 - **Flash Reliability – (LDPC)**
 - **Enterprise Flash Controller (EFC)**
 - **UMMC**
 - **Media Control Cluster**
- **Reference Firmware is also provided**
- **Allows various Flash parts to be used**
- **Customer can add their custom value add in SW or HW**

Mobiveil Configurable NVMe SSDC Platform



Computational Storage Platform



In-house: Completed (FPGA)

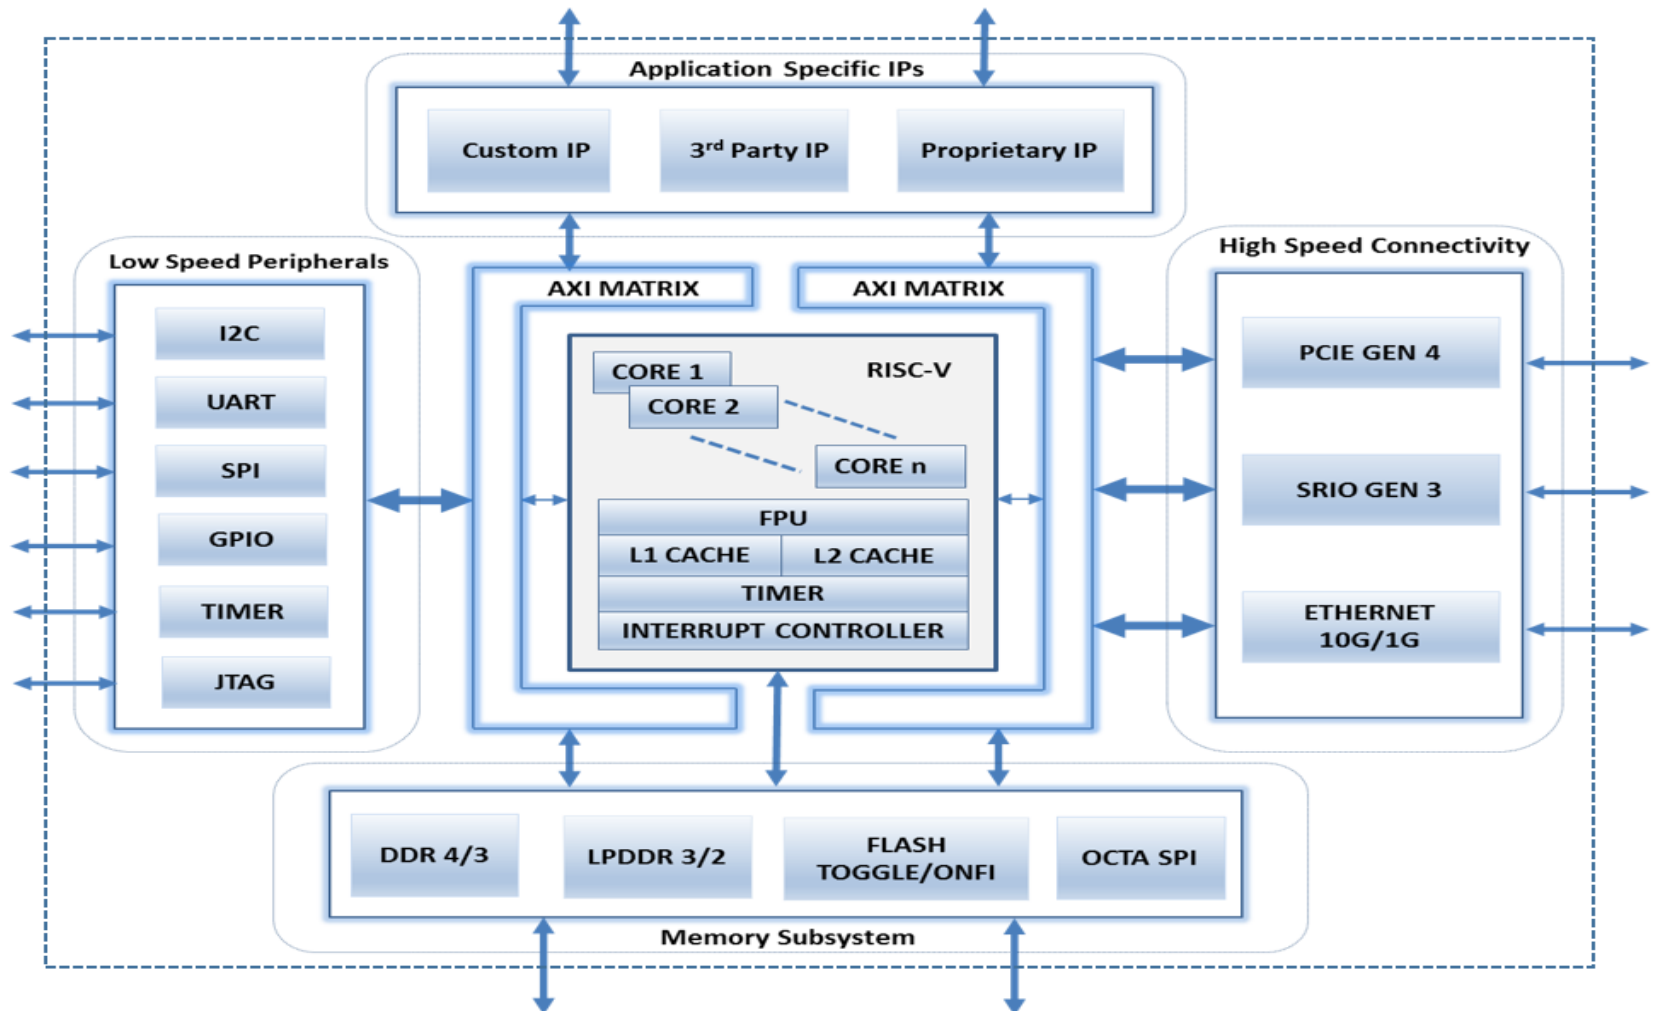
In-house: New development

In-house: Ongoing

3rd Party

RISC V based IoT SOC Platform

RISC-V SOC Block Diagram:



MV-IoT Platform



Further Engagement with Mobiveil



- Visit Mobiveil Booth for
 - SSD related IP blocks
 - NVMStor-Ultra Platform Details
 - How to accelerate your product development
- Contact us at ip@Mobiveil.com