The Changing Landscape of Compute

with RISC-V Open Architecture and P4

Richard New, VP Research
Western Digital

Western Digital®

Forward-Looking Statements

Safe Harbor | Disclaimers

This presentation contains forward-looking statements that involve risks and uncertainties, including, but not limited to, statements regarding our products and technologies, business strategies, product development efforts and growth opportunities, emerging storage and memory technologies, our investments in and contributions to the RISC-V ecosystem, industry and market trends, and data growth and its drivers. Forward-looking statements should not be read as a guarantee of future performance or results, and will not necessarily be accurate indications of the times at, or by, which such performance or results will be achieved, if at all. Forward-looking statements are subject to risks and uncertainties that could cause actual performance or results to differ materially from those expressed in or suggested by the forward-looking statements.

Key risks and uncertainties include volatility in global economic conditions; business conditions and growth in the storage ecosystem; impact of competitive products and pricing; market acceptance and cost of commodity materials and specialized product components; actions by competitors; unexpected advances in competing technologies; our development and introduction of products based on new technologies and expansion into new data storage markets; risks associated with acquisitions, mergers and joint ventures; difficulties or delays in manufacturing; and other risks and uncertainties listed in the company's filings with the Securities and Exchange Commission (the "SEC") and available on the SEC's website at www.sec.gov, including our most recently filed periodic report, to which your attention is directed. We do not undertake any obligation to publicly update or revise any forward-looking statement, whether as a result of new information, future developments or otherwise, except as required by law.

A Leading Data Infrastructure Company

PORTFOLIO BREADTH

TECHNOLOGY ENGINE

~14K active patents

GLOBAL SCALE

~62K employees worldwide

CUSTOMER VALUE

Western Digital. WD G-Technology SanDisk

Cache Coherence

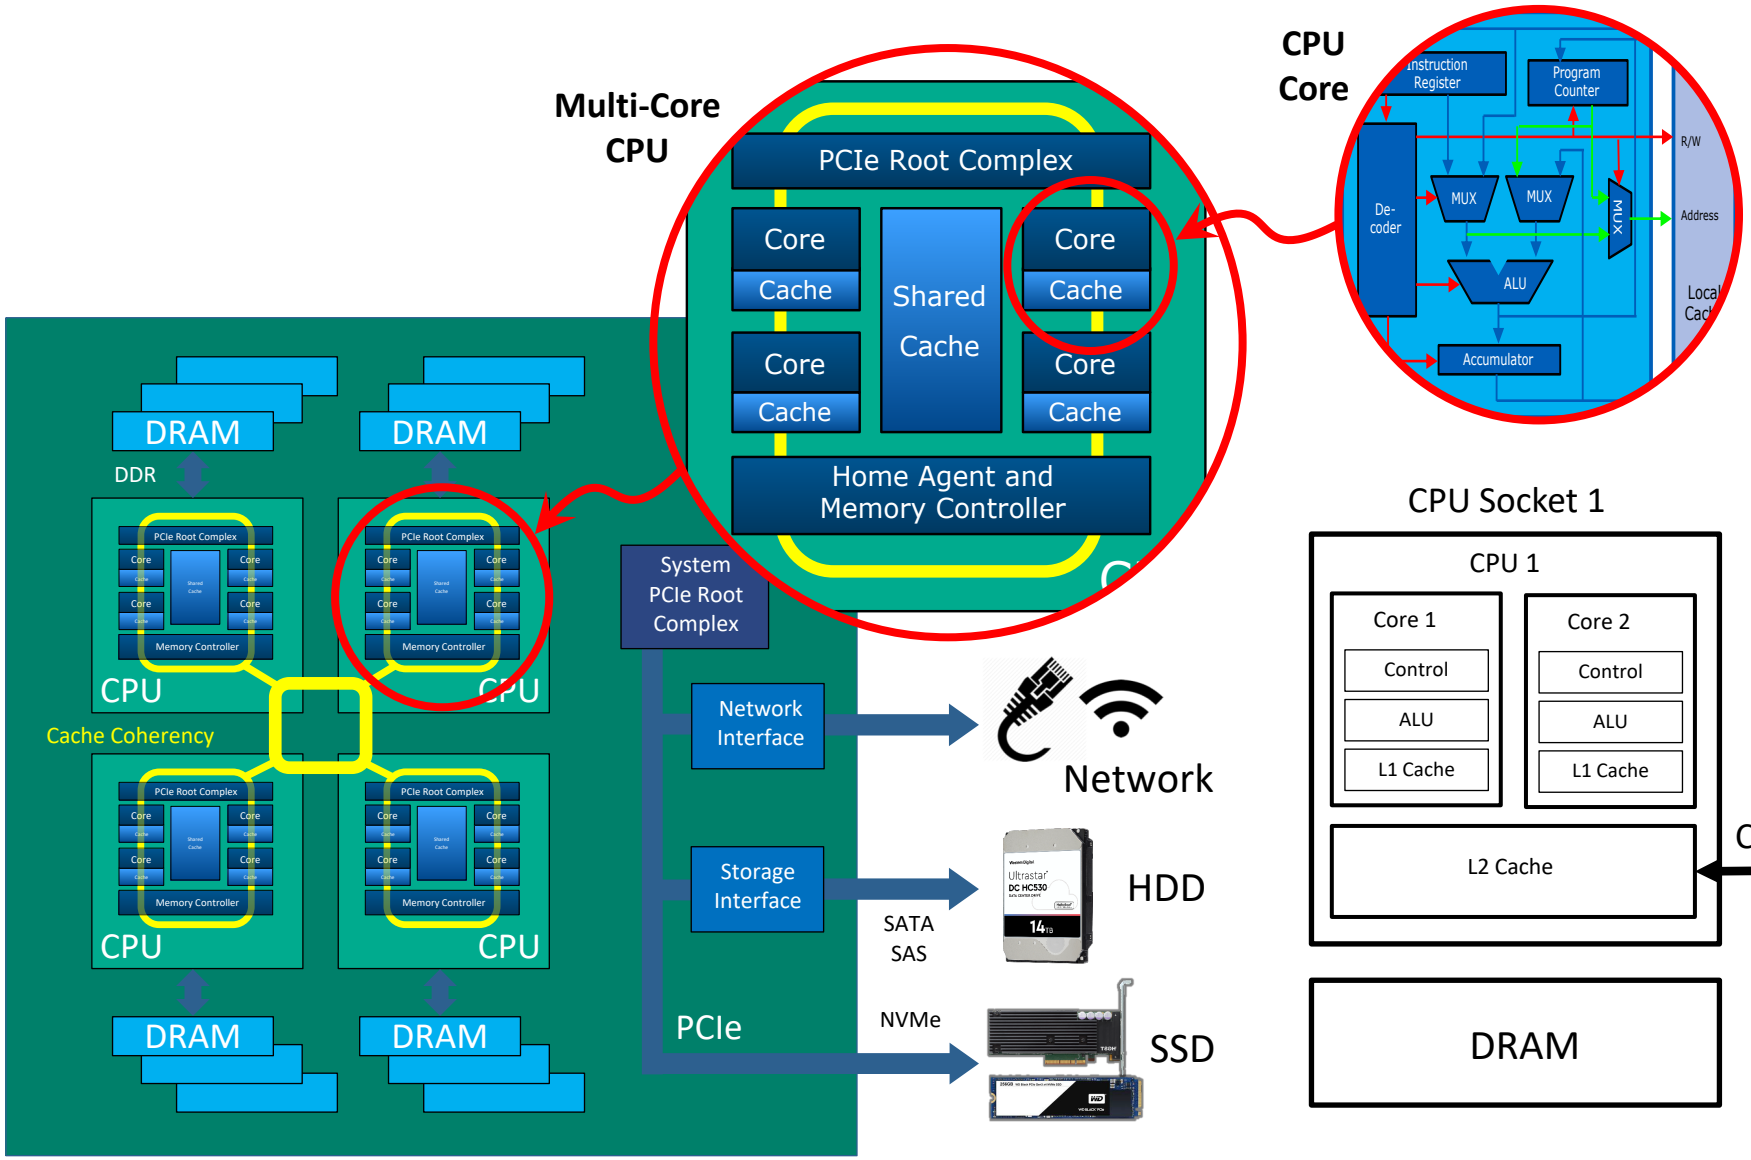

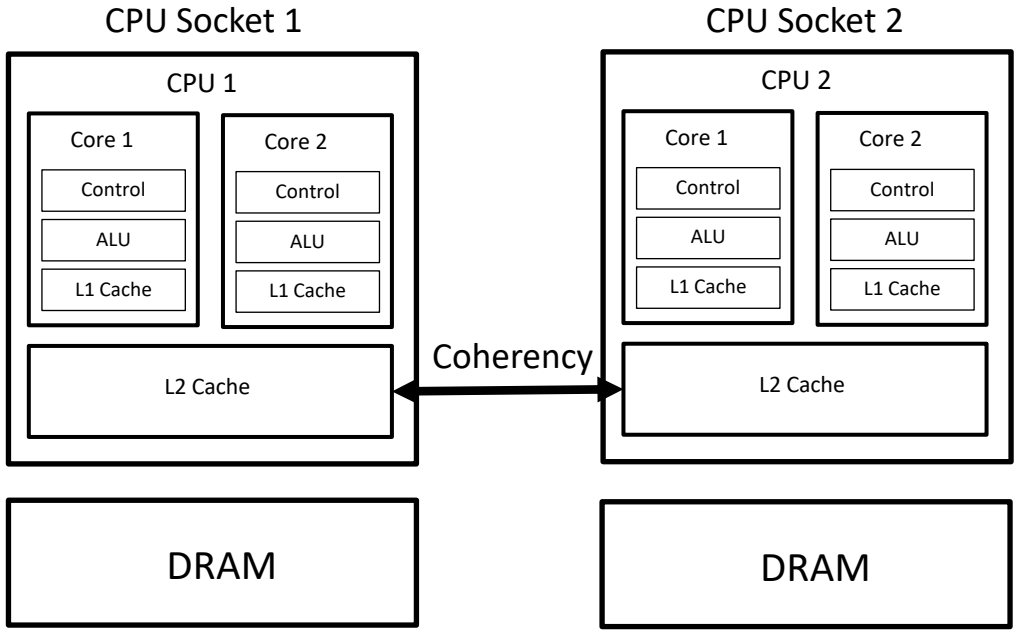

Cache Coherence Protocols

- Enables multiple processors to meaningfully share the same memory
- Coherence protocol ensures that all the processors see synchronized copies of data
- Key building block for creating a useful memory consistency model
- Enables single operating system to run across multiple cores

MESI Protocol

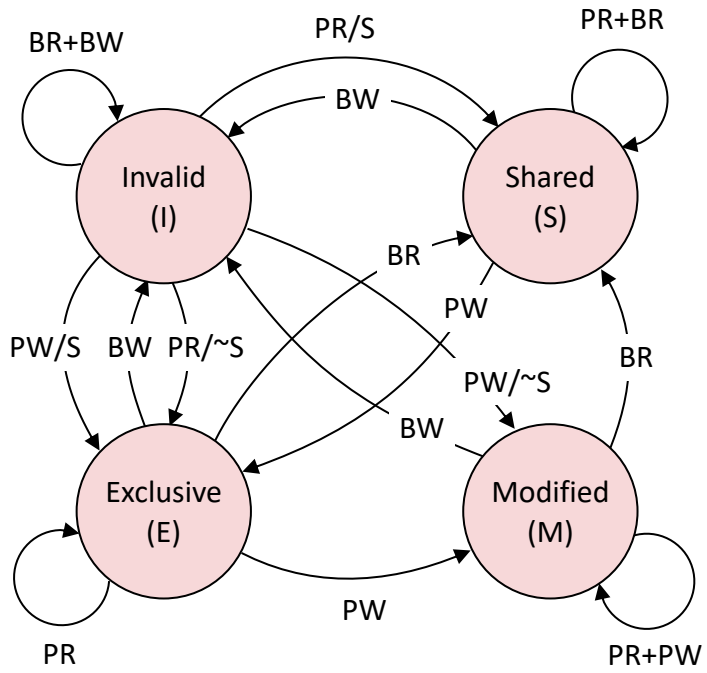

PR = processor read
PW = processor write
BR = observed bus read
BW = observed bus write

S/~S = shared or not shared

Modified (M) Cache line present only in the current cache, and is dirty.
Exclusive (E) Cache line present only in the current cache, and is clean.
Shared (S) Cache line may be stored in other caches and is clean.
Invalid (I) Cache line is invalid (unused).

Other Alternatives

Memory Extension

Page fault trap leading to RDMA request (incurs context switch and SW overhead) (e.g.mmap across network)

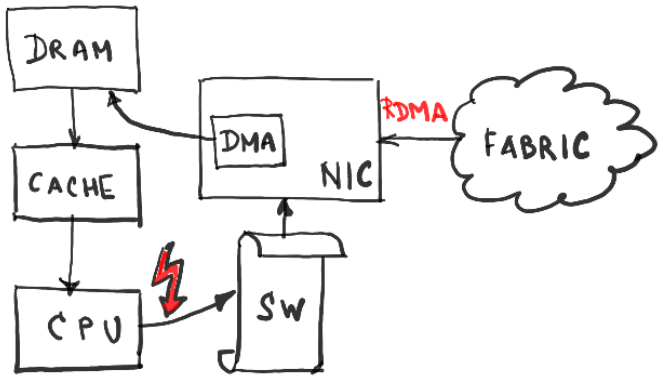

Context switch cost comparable to remote memory access latency; wasted wall power

Software Address Translation

Global address translation management in SW, leading to LD/ST across global memory fabric (e.g. LD/ST to MMIO PCIe device that translates address to network destination).

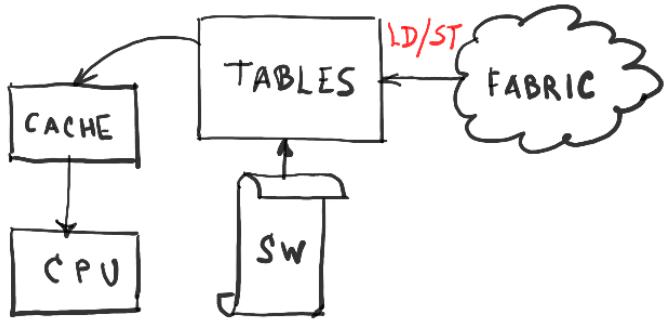

Require software/kernel support and/or rewriting of applications

Distributed Software

Abandon the concept of shared memory, and write distributed software which explicitly passes messages back and forth or copies data with RDMA between nodes.

RISC-V Open Instruction Set Architecture

Open

- Completely open source ISA that can be applied for any modern computing devices

Free

- Develop independent IP
- RISC-V Foundation does not charge authorization fee
- BSD license allows developers to decide if they want to share their own work

Safe

- Clear separation between user and privileged ISA

Regulation

- Prevent fragmentation of RISC-V in China
- Small standard base ISA with only 40+ base instructions

Extensibility/Specialization

- Variable-length instruction encoding
- Vast opcode space available for instruction-set extensions

Stable

- Base and standard extensions are frozen
- Additions via optional extensions, not new versions

Western Digital SweRV Core™

- First production-grade open source RISC-V core
- 2-way, superscalar, in-order core with 9 stage pipeline
- 1 GHz operation @ 28nm
- Support for RV32IMC
- Ideal for high-performance embedded applications
- Download it at: https://github.com/westerndigitalcorporation/swerv

Western Digital ships in excess of **1 Billion** cores per year

…and we expect to **double that.**

SoC with OmniXtend™ Support

Western Digital and SiFive collaborating on extension of TileLink cache coherency protocol over ethernet.

- Leverages ubiquitous Ethernet
- Routable and switchable fabric
- Scalable as Ethernet speeds increase

OmniXtend initial specification and FPGA bitstream are now live at https://github.com/westerndigitalcorporation/omnixtend

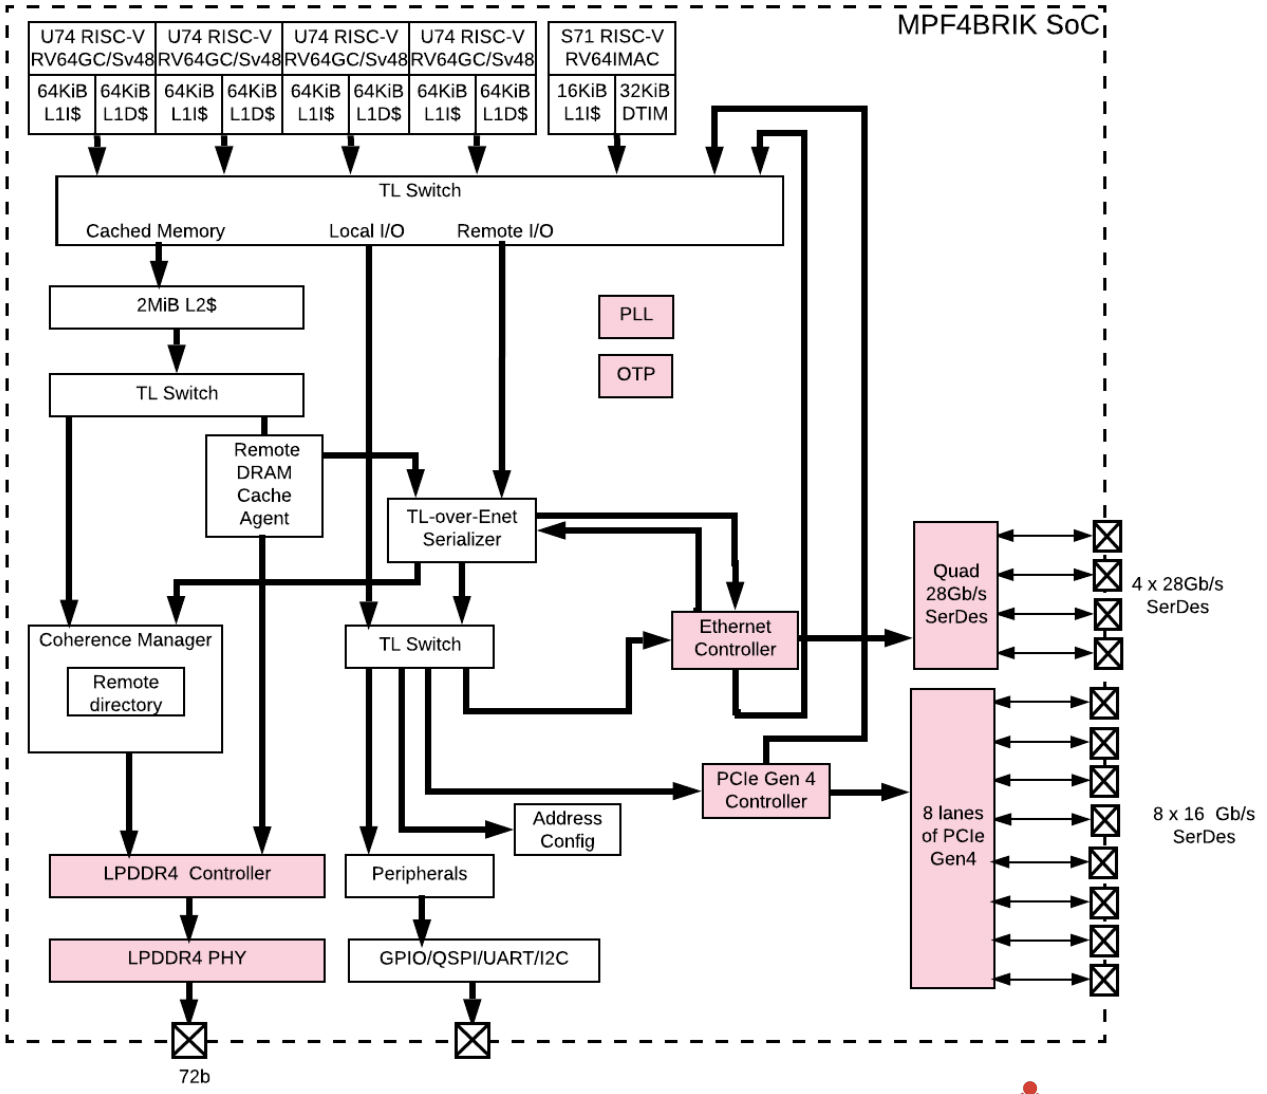

OmniXtend Ethernet Implementation

- Replace L2 Ethernet with Tilelink on top of 802.3 L1
  - Data plane frame processing programmed in P4.
  - Supports innovation required for RAS
- Barefoot Tofino™ ASIC
  - 256-lane 25 GT/s Ethernet switch, 6.4 Tbit/s throughput
  - Supports P4 HDL, successor to OpenFlow
  - Match-Action Pipeline enables line-rate performance

802.3 Ethernet packet and frame structure

Layer	Preamble	Start of frame delimiter	MAC destination	MAC source	802.1Q tag (optional)	Ethertype (Ethernet II) or length (IEEE 802.3)	Payload	Frame check sequence (32-bit CRC)	Interpacket gap
	7 octets	1 octet	6 octets	6 octets	(4 octets)	2 octets	46–1500 octets	4 octets	12 octets
Layer 2 Ethernet frame			OmniXtend uses these bits						
Layer 1 Ethernet packet & IPG	← 72–1530 octets →								← 12 oct. →

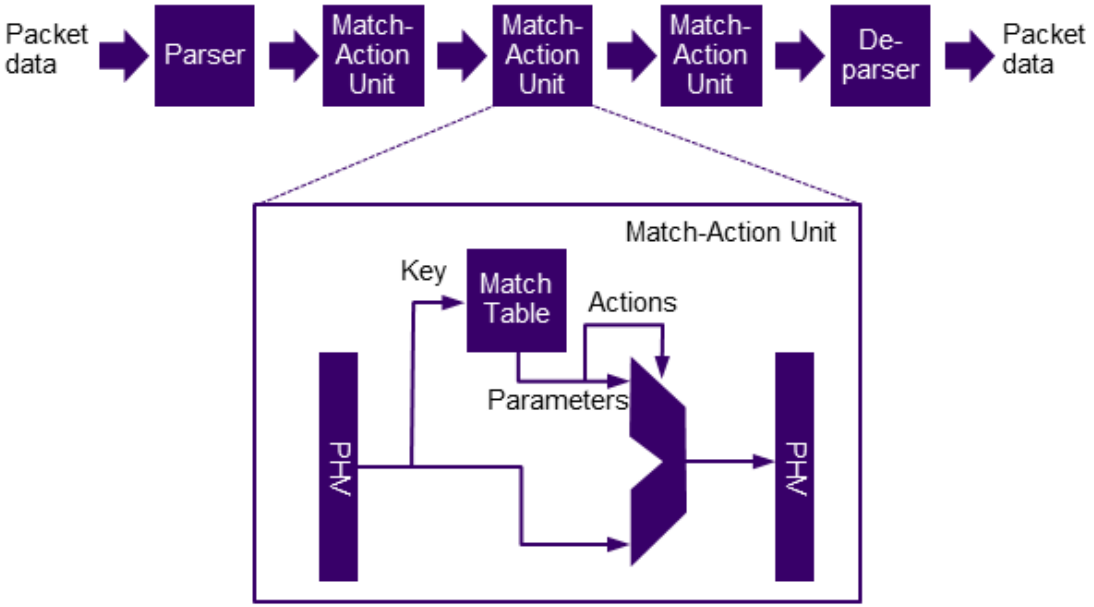


```
/* -*- P4_16 -*- */

#include <core.p4>
#include <v1model.p4>

/*
 * Define the headers the program will recognize
 */

/*
 * This is a custom protocol header for the Gen-Z packets.
 * We'll use ethertype 0x1234 for is (see parser)
 */

const bit<16> P4GENZ_P2P_CORE_ETYPE = 0x1234;
const bit<1>  P4GENZ_P2P_END2END_ETYPE = 0x0;
const bit<1>  P4GENZ_P2P_LINKLOCAL_ETYPE = 0x1;

/*
 * Defined the headers of Gen-Z packets
 *
 */

header p4genz_p2p_corereq_t {
    bit<5>   p4genz_pcrc;
    bit<1>   p4genz_oc_0;
    bit<1>   p4genz_o;
```

```
parser MyParser(
    packet_in               packet,
    out   my_headers_t    hdr,
    inout my_metadata_t   meta,
    inout standard_metadata_t standard_metadata)
{
    state start {
        packet.extract(hdr.ethernet);
        transition select(hdr.ethernet.etherType) {
            P4GENZ_P2P_CORE_ETYPE : check_p4genz;
            default               : accept;
        }
    }


    //check L of Gen-Z packet
    state check_p4genz {
        transition
select(packet.lookahead<p4genz_p2p_corereq_t>().p4genz_l) {
            P4GENZ_P2P_END2END_ETYPE  : parse_p4genz;
            default                   : accept;
        }
    }

    state parse_p4genz {
        packet.extract(hdr.p4genz_p2p_corereq);
        transition accept;
    }
}
```

OmniXtend Architectures

Two Servers Point to Point

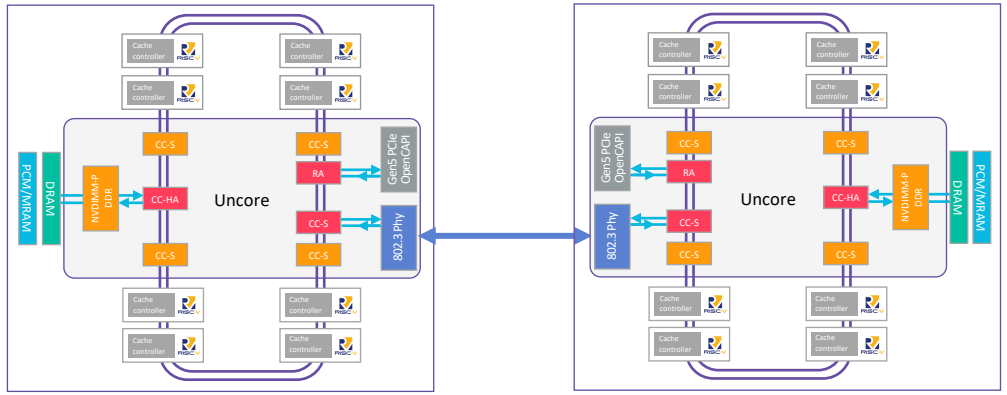

Looks like single machine with more CPUs

Multiple Devices Through an Ethernet Switch

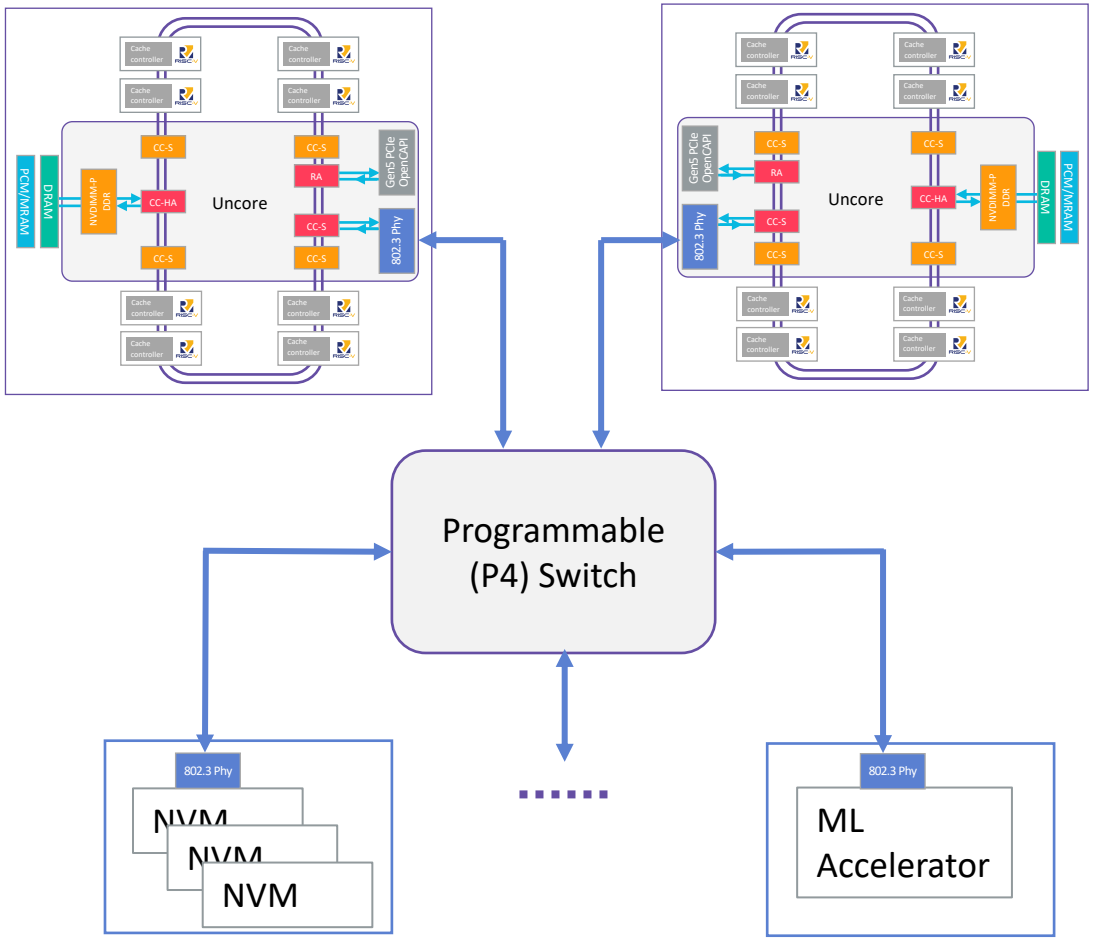

First OmniXtend 0.1.1 measurements

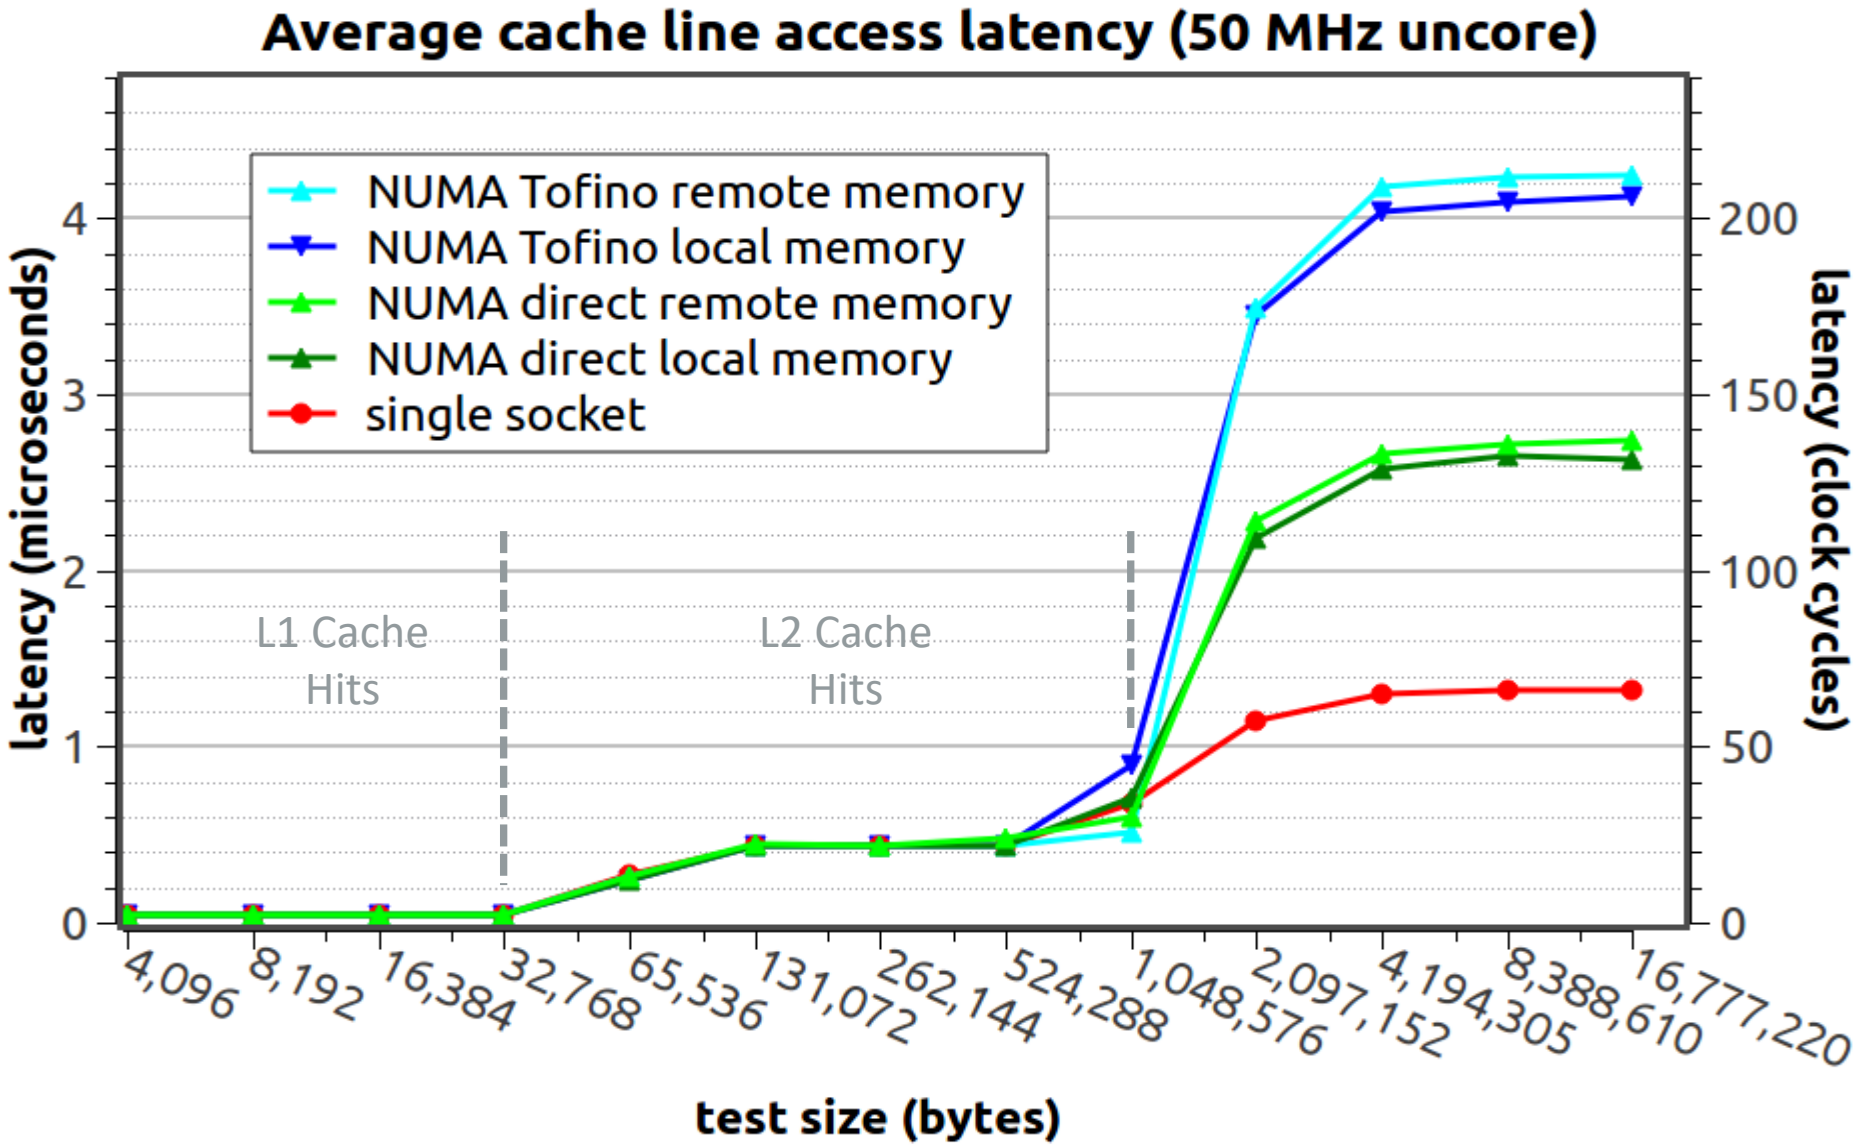

A Truly Open Memory Fabric -- OmniXtend

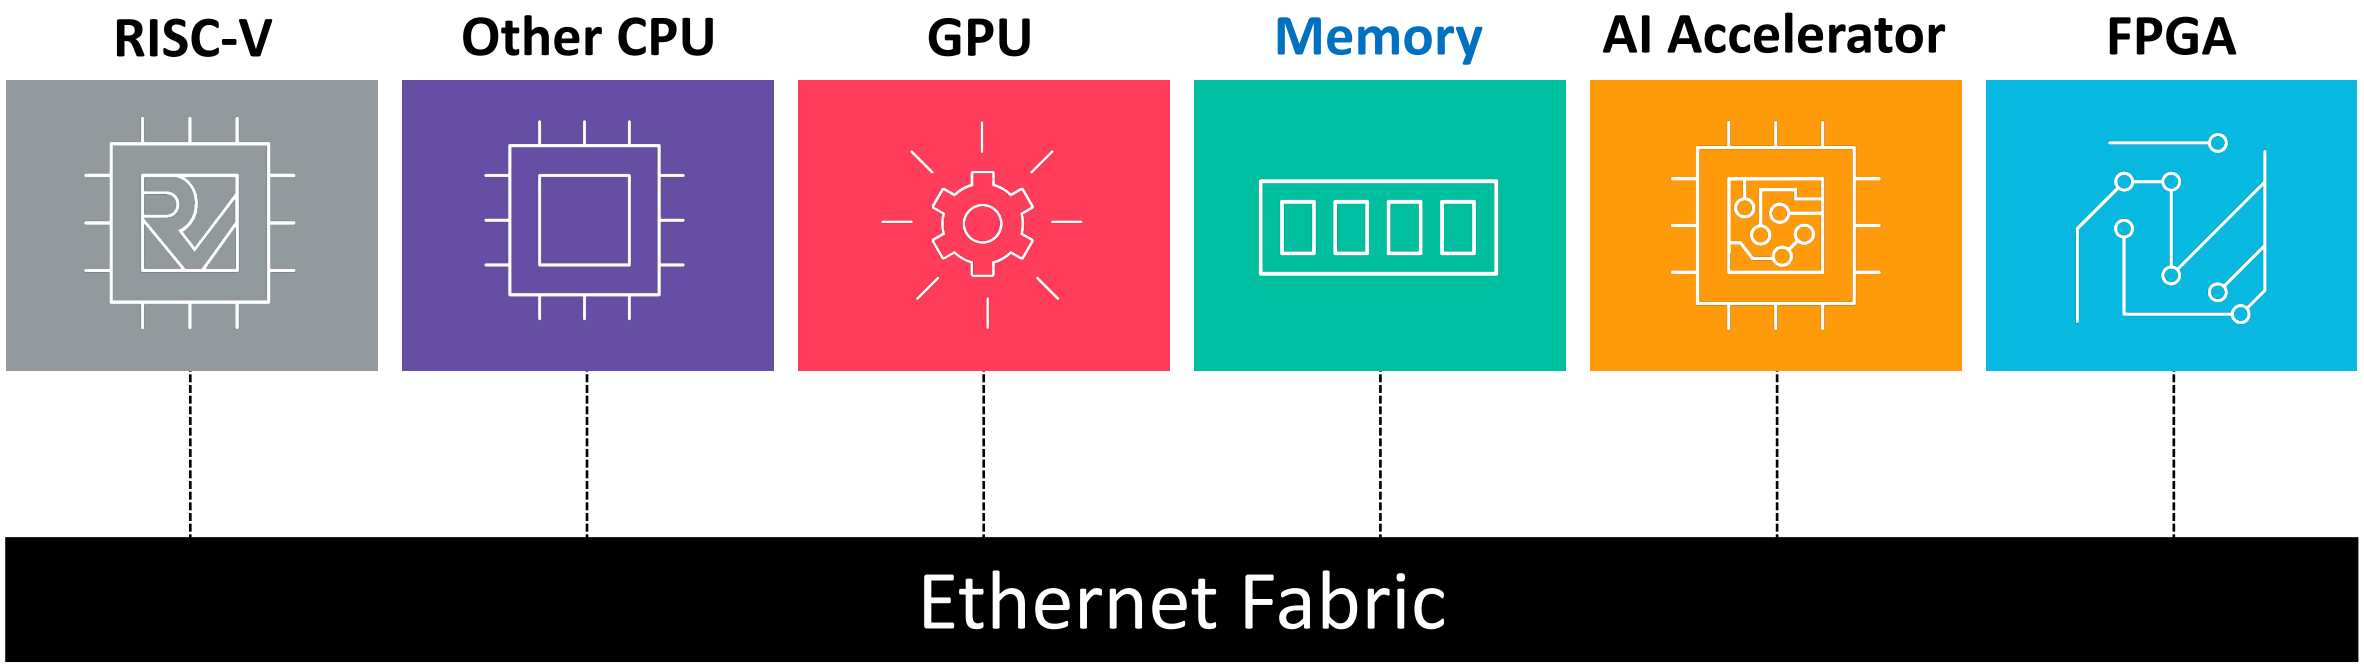

Data is the center of the architecture

No established hierarchy – CPU doesn't 'own' the GPU or the Memory

Cache Coherency preserved system-wide over the Fabric

Contacts

- Check out RISC-V as a candidate for embedded networking cores.
  Join RISC-V Foundation and Chips Alliance.
  https://riscv.org/
  https://chipsalliance.org/

- Western Digital SweRV core on github:
  https://github.com/westerndigitalcorporation/swerv

- If interested in memory-fabric applications of P4 and programmable switch, lease contact Western Digital,
  or see github site.
  https://github.com/westerndigitalcorporation/omnixtend

Thank You

Western Digital®