Boost your high bandwidth data acquisition by combining OpenCAPI + memory coherency + FPGAs

Filip LEONARSKI – *Paul Scherrer Institut – beamline data scientist – Project Leader*
Bruno MESNET – *IBM Systems – Hardware Acceleration*
Toshaan BHARVANI – *Van Tosch – OpenPower Foundation Technical Chair*

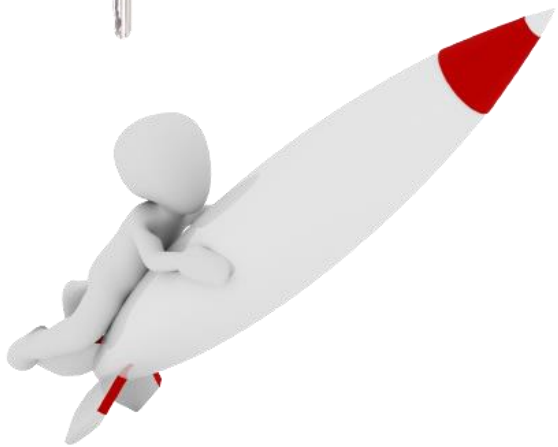

POWER9

OpenCAPI™

Future ?

OpenCAPI™

System Composability: PowerAXON & Open Memory Interfaces

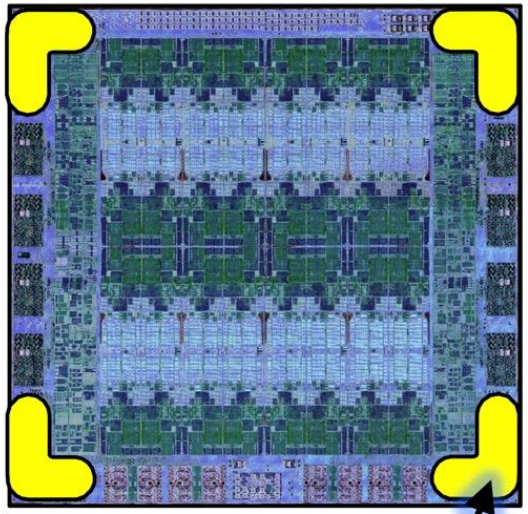

PowerAXON corner
4x8 @ 32 GT/s

Multi-protocol
"Swiss-army-knife"
Flexible / Modular Interfaces

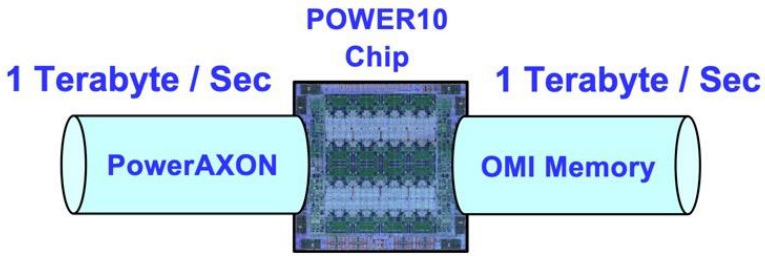

Built on best-of-breed
Low Power, Low Latency,
High Bandwidth
Signaling Technology

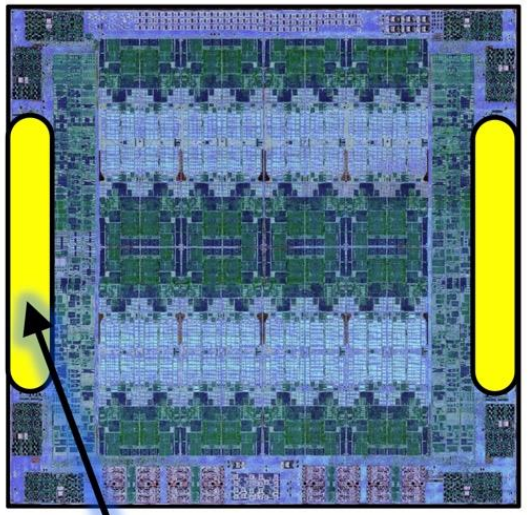

OMI edge
8x8 @ 32 GT/s

6x bandwidth / mm^2
compared to DDR4
signaling

IBM POWER10

Data Plane Bandwidth and Capacity: Open Memory Interfaces

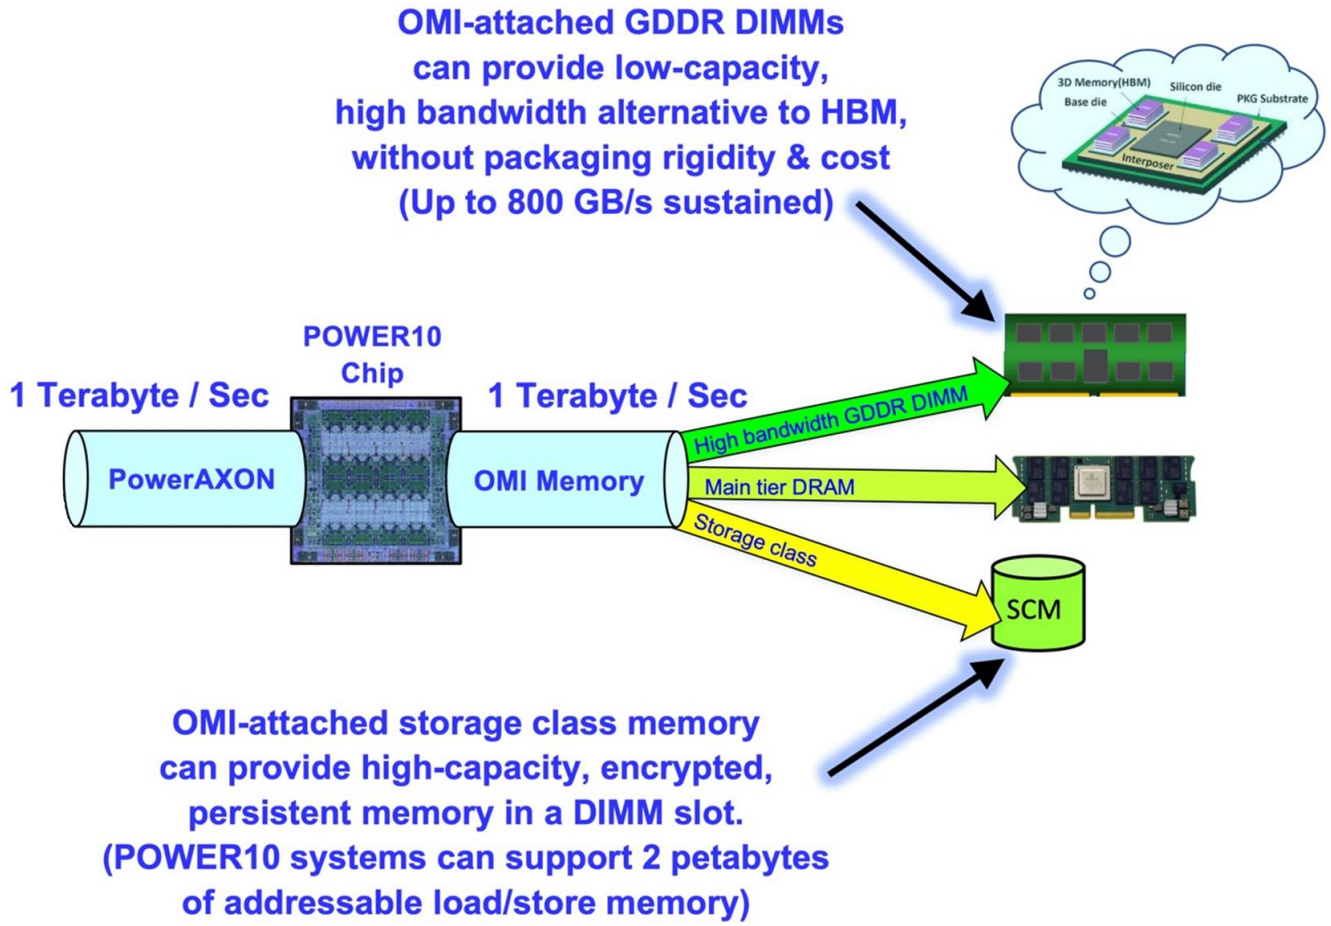

(PowerAXON and OMI Memory configurations show processor capability only, and do not imply system product offerings)

IBM POWER10

System Enterprise Scale and Bandwidth: SMP & Main Memory

Multi-protocol
"Swiss-army-knife"
Flexible / Modular Interfaces

Build up to 16 SCM socket
Robustly Scalable
High Bisection Bandwidth
"Glueless" SMP

Allocate the bandwidth
however you need to use it

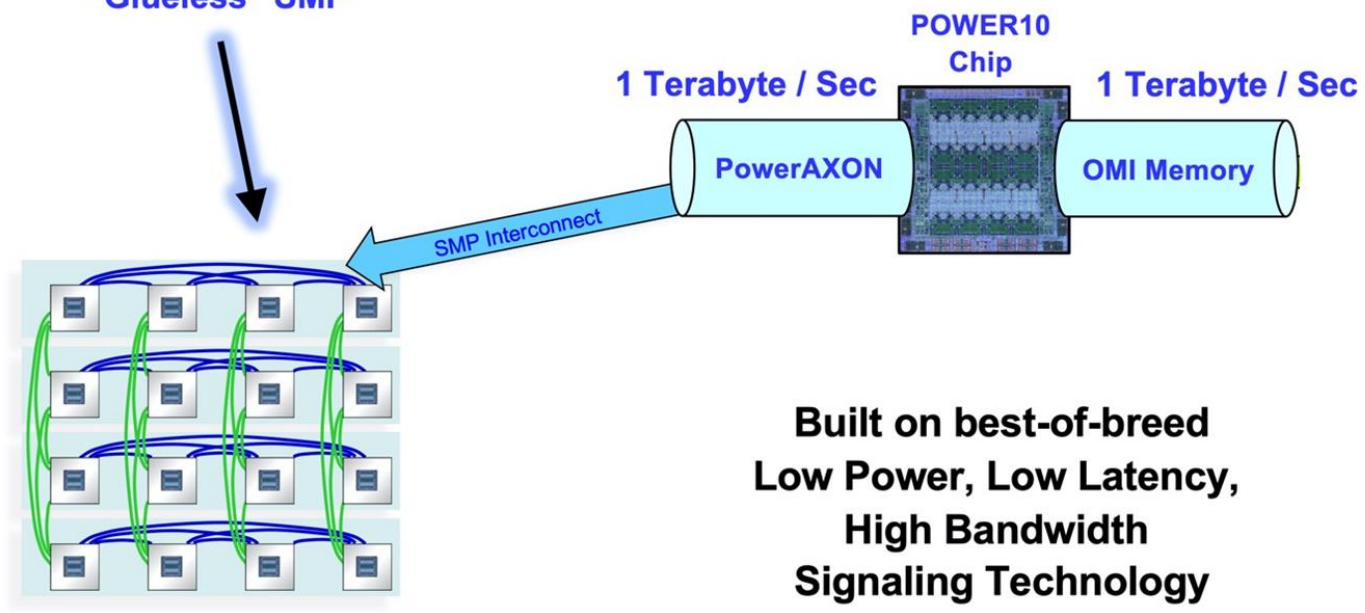

Built on best-of-breed
Low Power, Low Latency,
High Bandwidth
Signaling Technology

(PowerAXON and OMI Memory configurations show processor capability only, and do not imply system product offerings)

IBM POWER10

System Heterogeneity and Data Plane Capacity: OpenCAPI

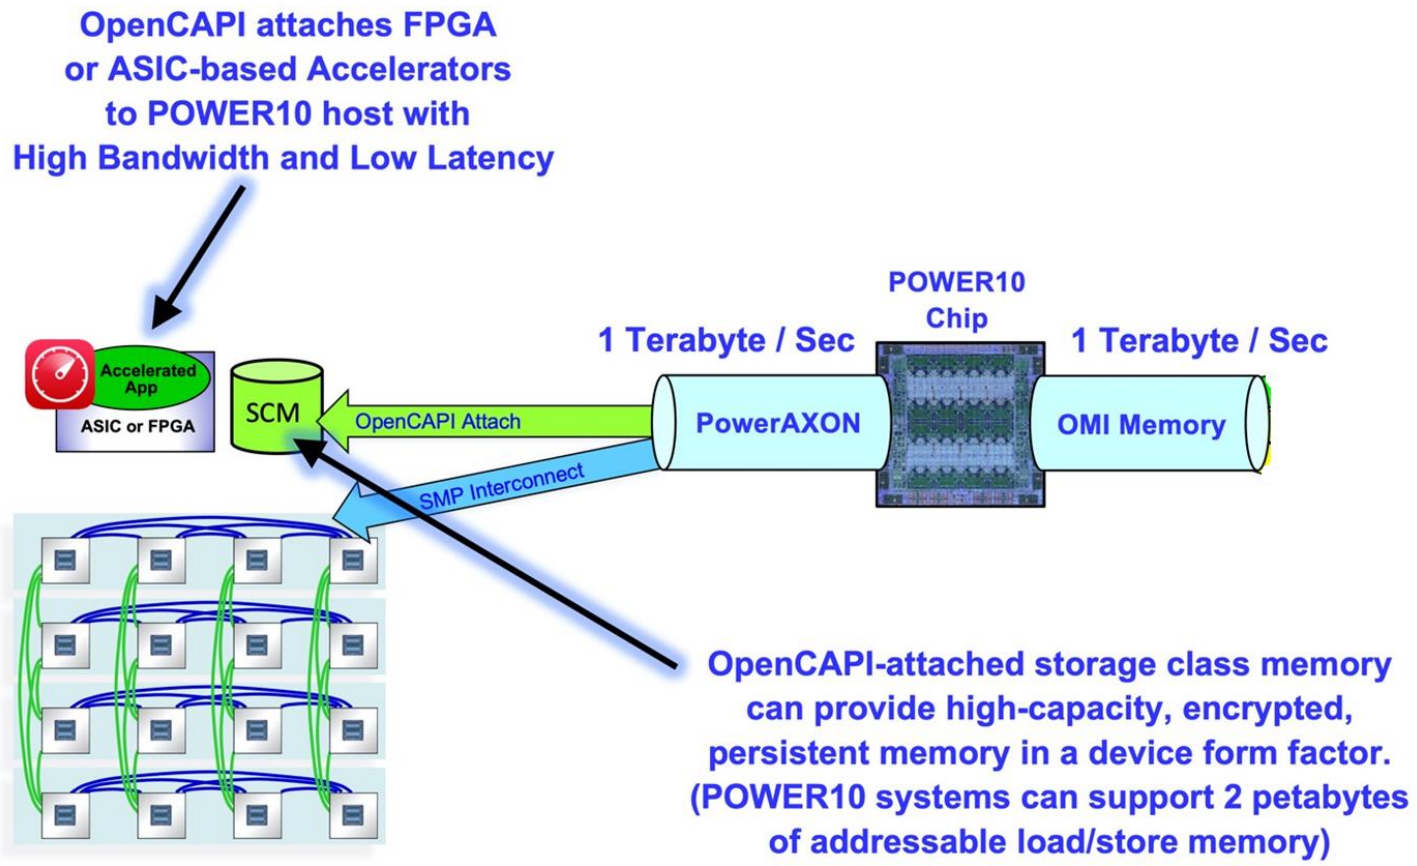

(PowerAXON and OMI Memory configurations show processor capability only, and do not imply system product offerings)

IBM POWER10

Pod Composability: PowerAXON Memory Clustering

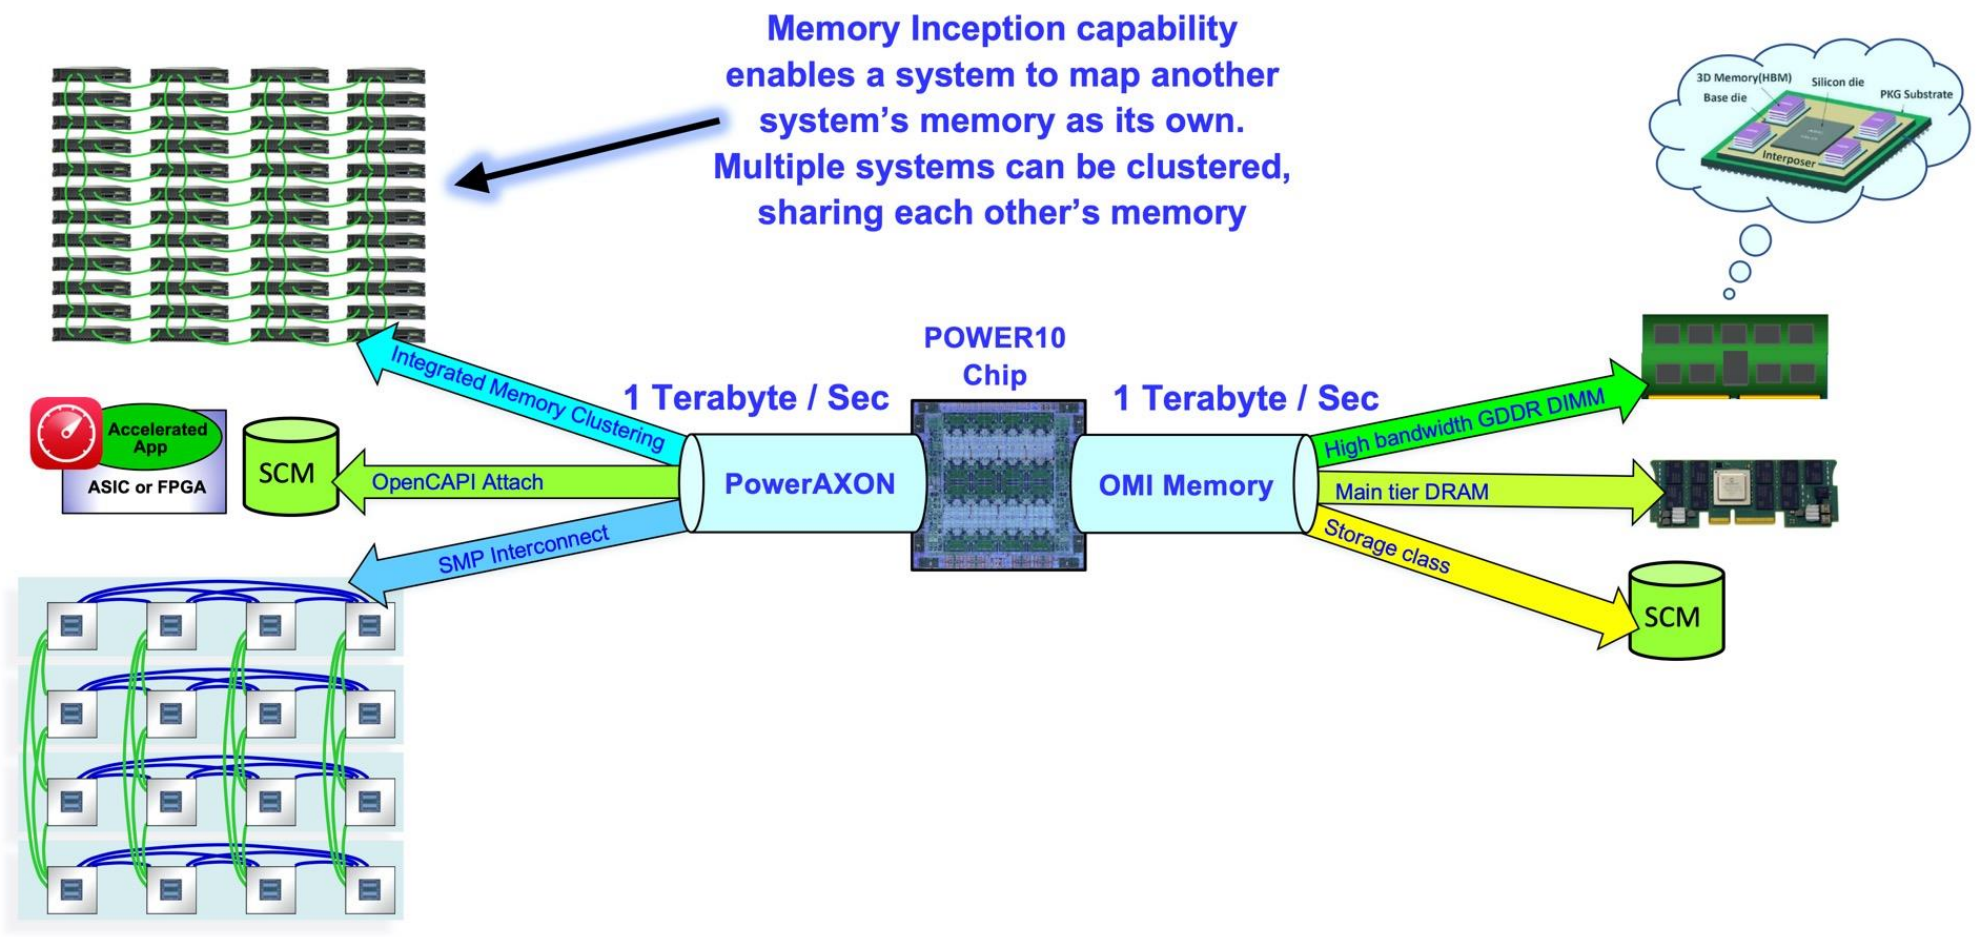

(PowerAXON and OMI Memory configurations show processor capability only, and do not imply system product offerings)

IBM POWER10

- ***Know more about accelerators ?***
- ***See a live demonstration?***
- ***Access to a server?***
- ***Do a benchmark ?***
- ***Get answers to your questions?***

CERN Openlab CTO for an access to an IC922 + AC922: maria.girone@cern.ch

IBM local partner: lclavien@inno-boost.com

IBM OpenCAPI team :
alexandre.castellane@fr.ibm.com - bruno.mesnet@fr.ibm.com - fabrice_moyen@fr.ibm.com

OpenCAPI Repository: https://github.com/OpenCAPI/oc-accel

More about decoupling compute with OMI: https://youtu.be/c0DuGSwDpqY

OpenCAPI™

IBM

Backup

Servers

2 servers available in Openlab

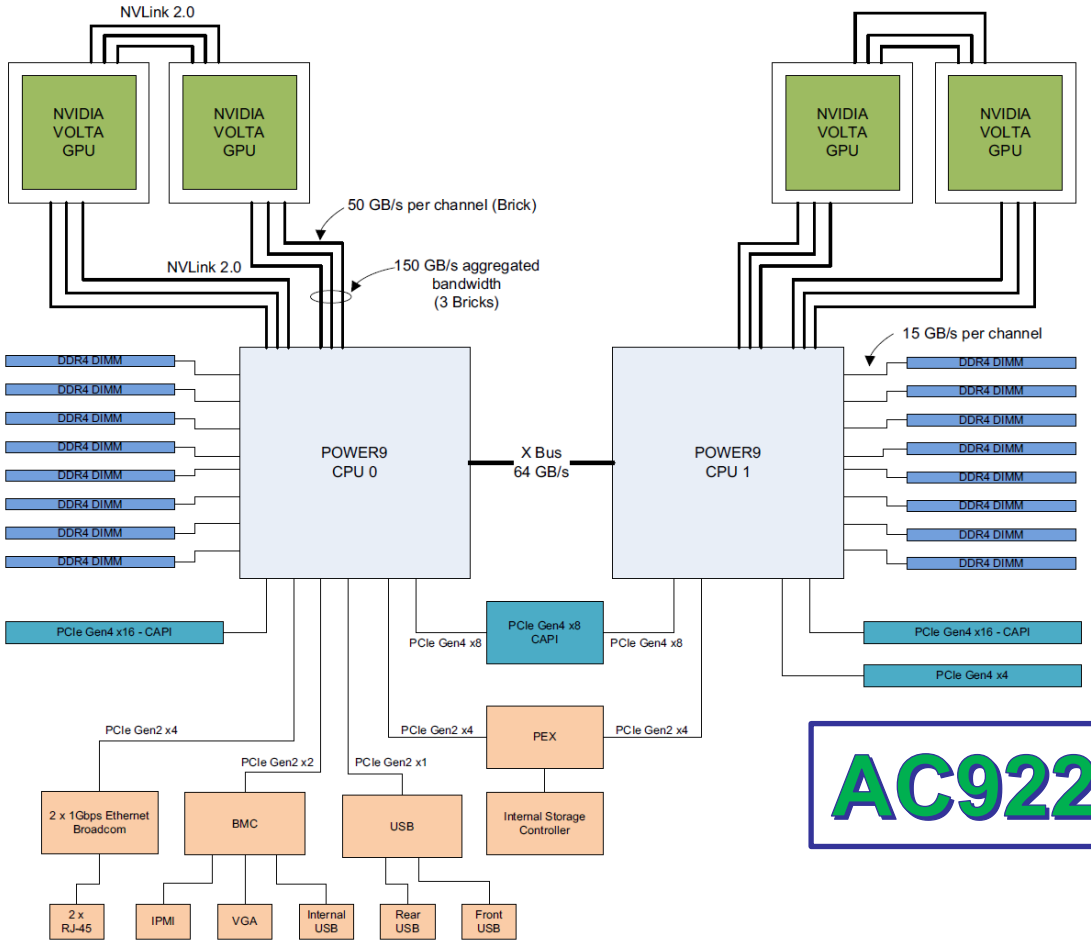

AC922

Mechanical		CAPI2.0	Processor	Shared Slots	Characteristics
PCIe Slot1	P1-C5	NO	CPU0		PCIeG4 x4, LP
PCIe Slot2	P1-C4	YES	CPU1 (Shared)	YES	PCIeG4 x8, LP
PCIe Slot3	P1-C3	YES	CPU1		PCIeG4 x16, LP
PCIe Slot4	P1-C2	YES	CPU0		PCIeG4 x16, LP

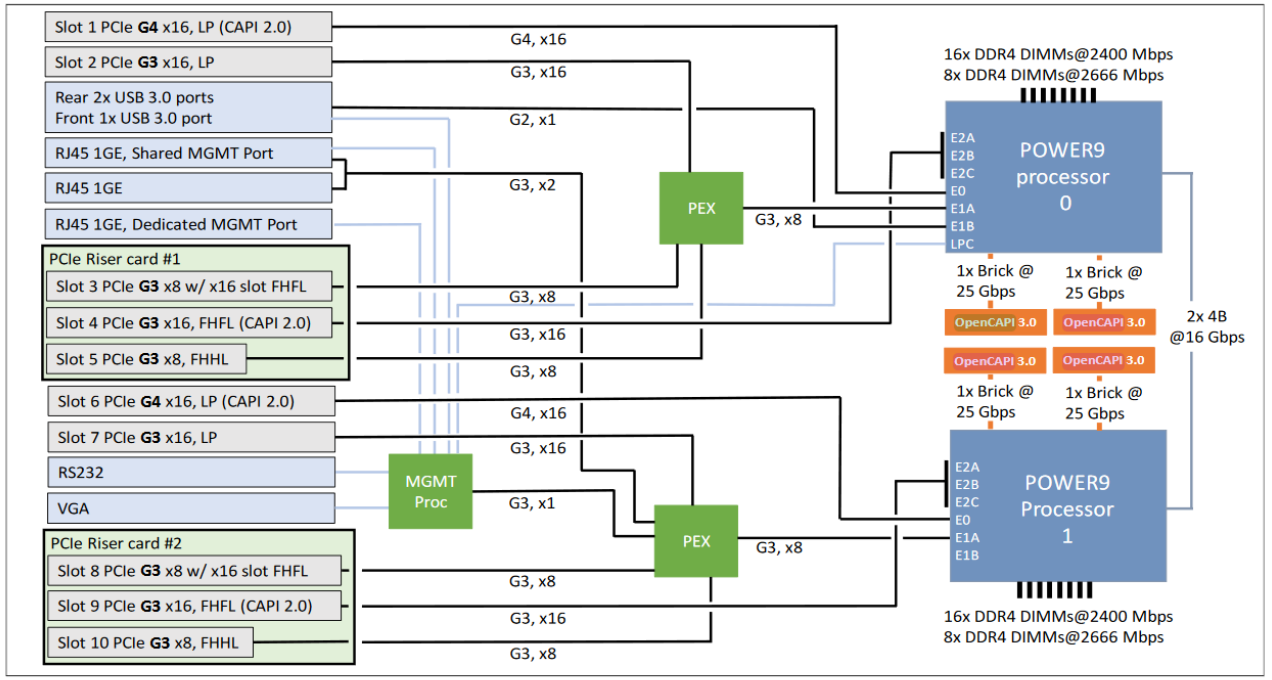

Figure 1-3 Power IC922 server logical system diagram with standard riser cards

IC922

Mechanical	CAPI2.0	Processor	Shared Slots	Characteristics
PCIe Slot1	YES	CPU0		PCIe G4 x16, LP
PCIe Slot2	NO	CPU0	2-3-5	PCIe G3x16, LP
PCIe Slot3	NO	CPU0	2-3-5	PCIe G3x8 w/ x16 slot FHFL
PCIe Slot4	YES	CPU0		PCIe G3 x16, FHFL
PCIe Slot5	NO	CPU0	2-3-5	PCIe G3x8, FHHL
PCIe Slot6	YES	CPU1		PCIe G4 x16, LP
PCIe Slot7	NO	CPU1	7-8-10	PCIe G3x16, LP
PCIe Slot8	NO	CPU1	7-8-10	PCIe G3x8 w/ x16 slot FHFL
PCIe Slot9	YES	CPU1		PCIe G3x16, FHFL
PCIe Slot10	YES	CPU1	7-8-10	PCIe G3x8, FHHL

FPGA boards

FPGA boards supported today on OC-Accel

Alpha-Data boards

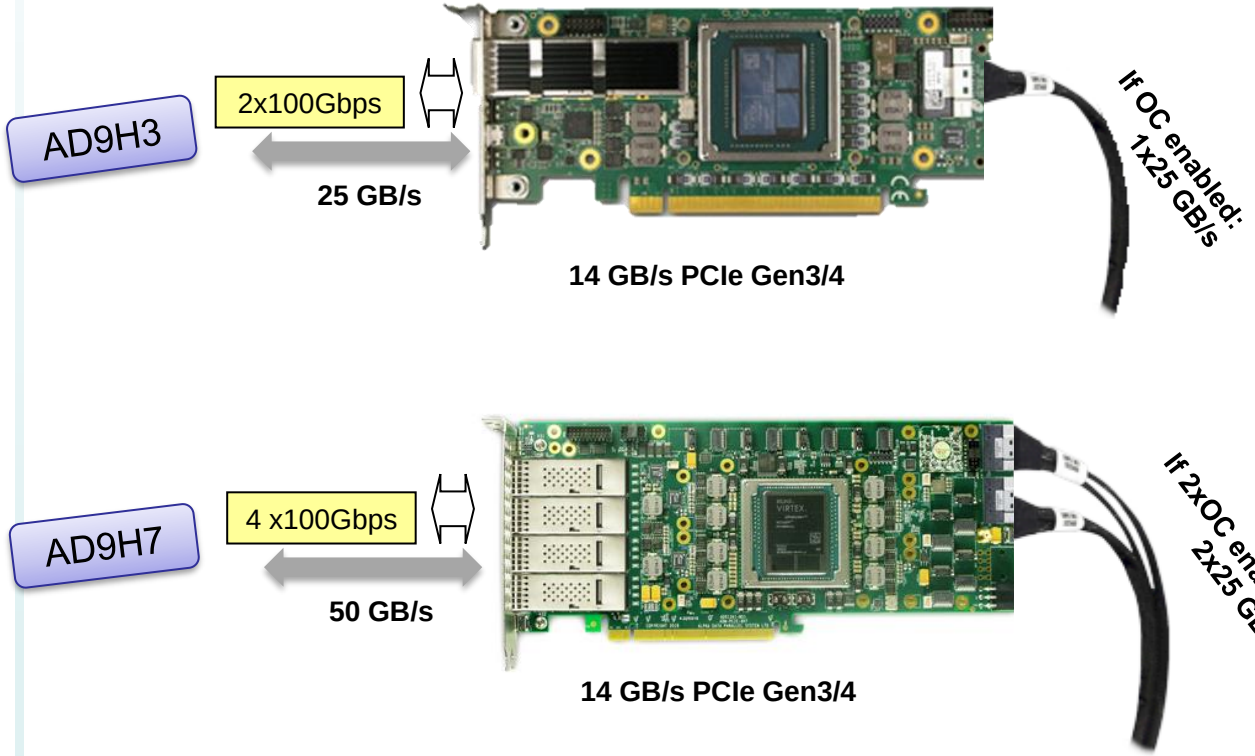

Bittware boards

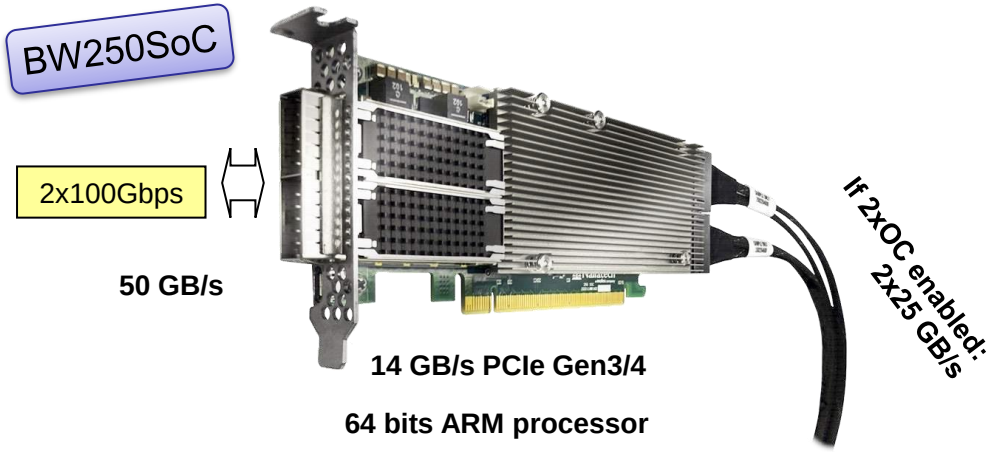

But also, AD9V3, AD9V5,... and more to come

Accelerator Product guide for CAPI/OpenCAPI

OpenPOWER™

Card Name	FPGA	Memory	Ethernet	IO	Format	Mode	Power	Server	Comments	Framework Status
Alpha Data 9V3	XCVU3P 394 k LUTs	2x 8GB DDR4	2x QSFP28 = 2x 100GbE	1x 25GBps OpenCAPI	HH-HL-S Slot x16 PCIe	CAPI2 / OpenCA PI	75W	IC922 AC922	Old Card Functionally ok but OC connector badly located constraints specific PCIe slots use	SNAP: ENABLED OC-Accel: ENABLED
Alpha Data 9H3	XCVU33P 440 k LUTs	2x 4GB HBM	1x QSFP- DD = 2x100GbE	1x 25GBps OpenCAPI	HH-HL-S Slot X16 PCIe	CAPI2 / OpenCA PI	75W	IC922 AC922		SNAP: ENABLED OC-Accel: ENABLED
Alpha Data 9H7	XCVU37P 1340 k LUTs	2x 4GB HBM	4x QSFP28 = 4x 100GbE	2x 25GBps OpenCAPI	FH-FL-D Slot x16 PCIe	CAPI2 / OpenCA PI	225W	IC922 ~~AC922~~	*Requires a specific riser*	SNAP: ENABLED OC-Accel: ENABLED
Bittware 250S+	KU15P 523 k LUTs	1x 8GB DDR4 4x 1TB M2 NVMe sticks	None	4x OCuLink to NVME drives	HH-HL-S Slot x8 PCIe	CAPI2	75W	IC922 AC922	Old card –PCIe IP for Gen4x8 IP not supported any more by Vivado > 2018.2 and PCIe IP for NVME drives to be adapted	SNAP: not fully enabled anymore OC-Accel: N/A
Bittware 250-SoC	ZU19EG 523 k LUTs Quad-core ARM	1x 4GB DDR4 (FPGA) 1x 4GB DDR4 (ARM)	2x QSFP28 = 2x 100GbE	2x 25GBps OpenCAPI	HH-HL-S Slot x16 PCIe	CAPI2 / OpenCA PI	75W	IC922 AC922		SNAP: not in plan OC-Accel: ENABLED
Xilinx Alveo U50	U50 872 k LUTs	2x 4GB HBM	1x QSFP28 = 1x 100GbE		HH-HL-S Slot x16 PCIe	CAPI2	75W	IC922 AC922		SNAP: ETA 4Q21 OC-Accel: N/A
Xilinx Alveo U200	U200 892 k LUT	4x 16GB DDR4	2x QSFP28 = 2x 100GbE		FH-FL-D Slot x16 PCIe	CAPI2	225W	IC922 ~~AC922~~		SNAP: ETA 4Q21 OC-Accel: N/A