



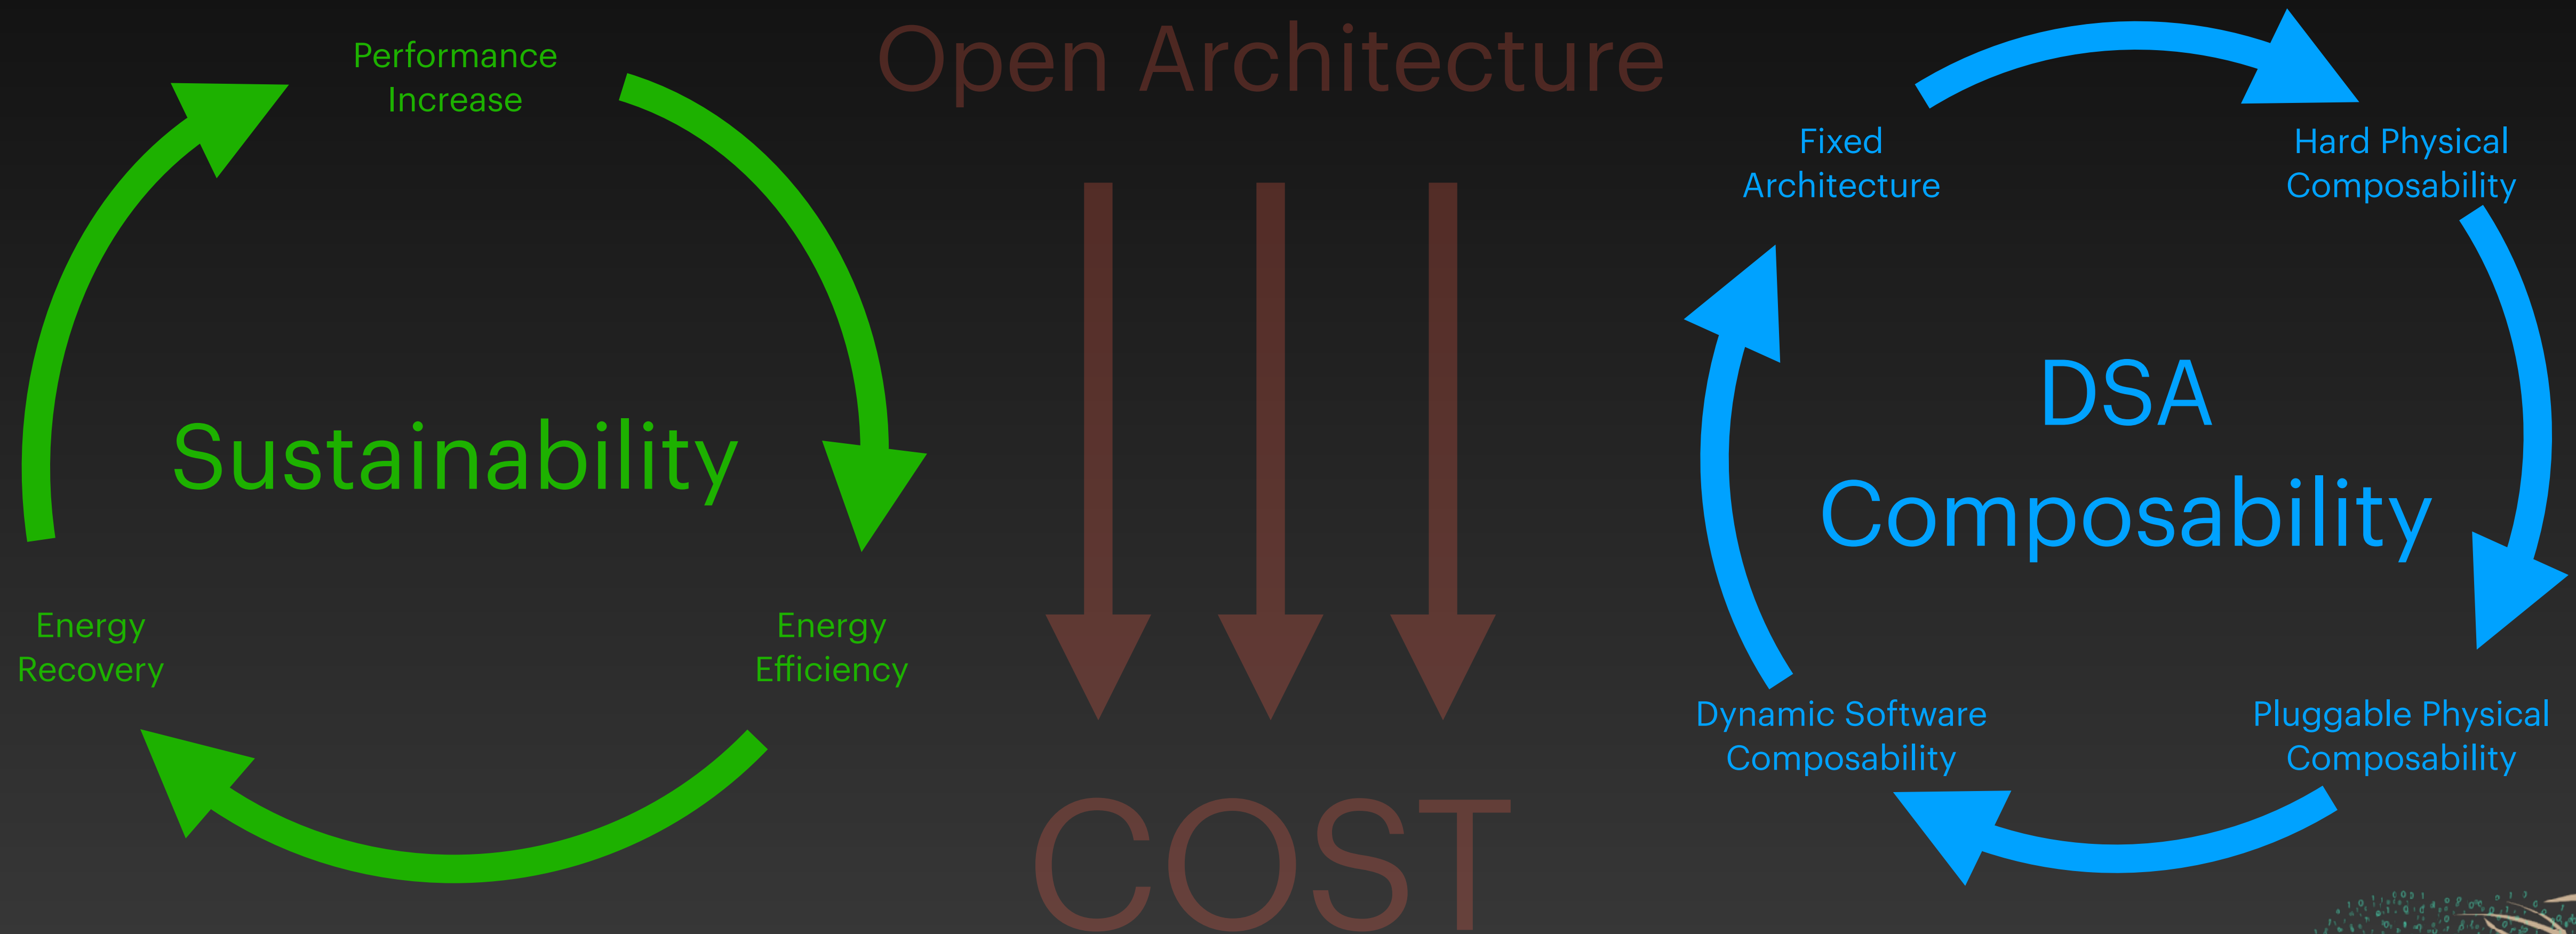
OMI & CXL Composable Memory Centric Computing Vision

OMI Seminar

Allan Cante - 6/16/2022

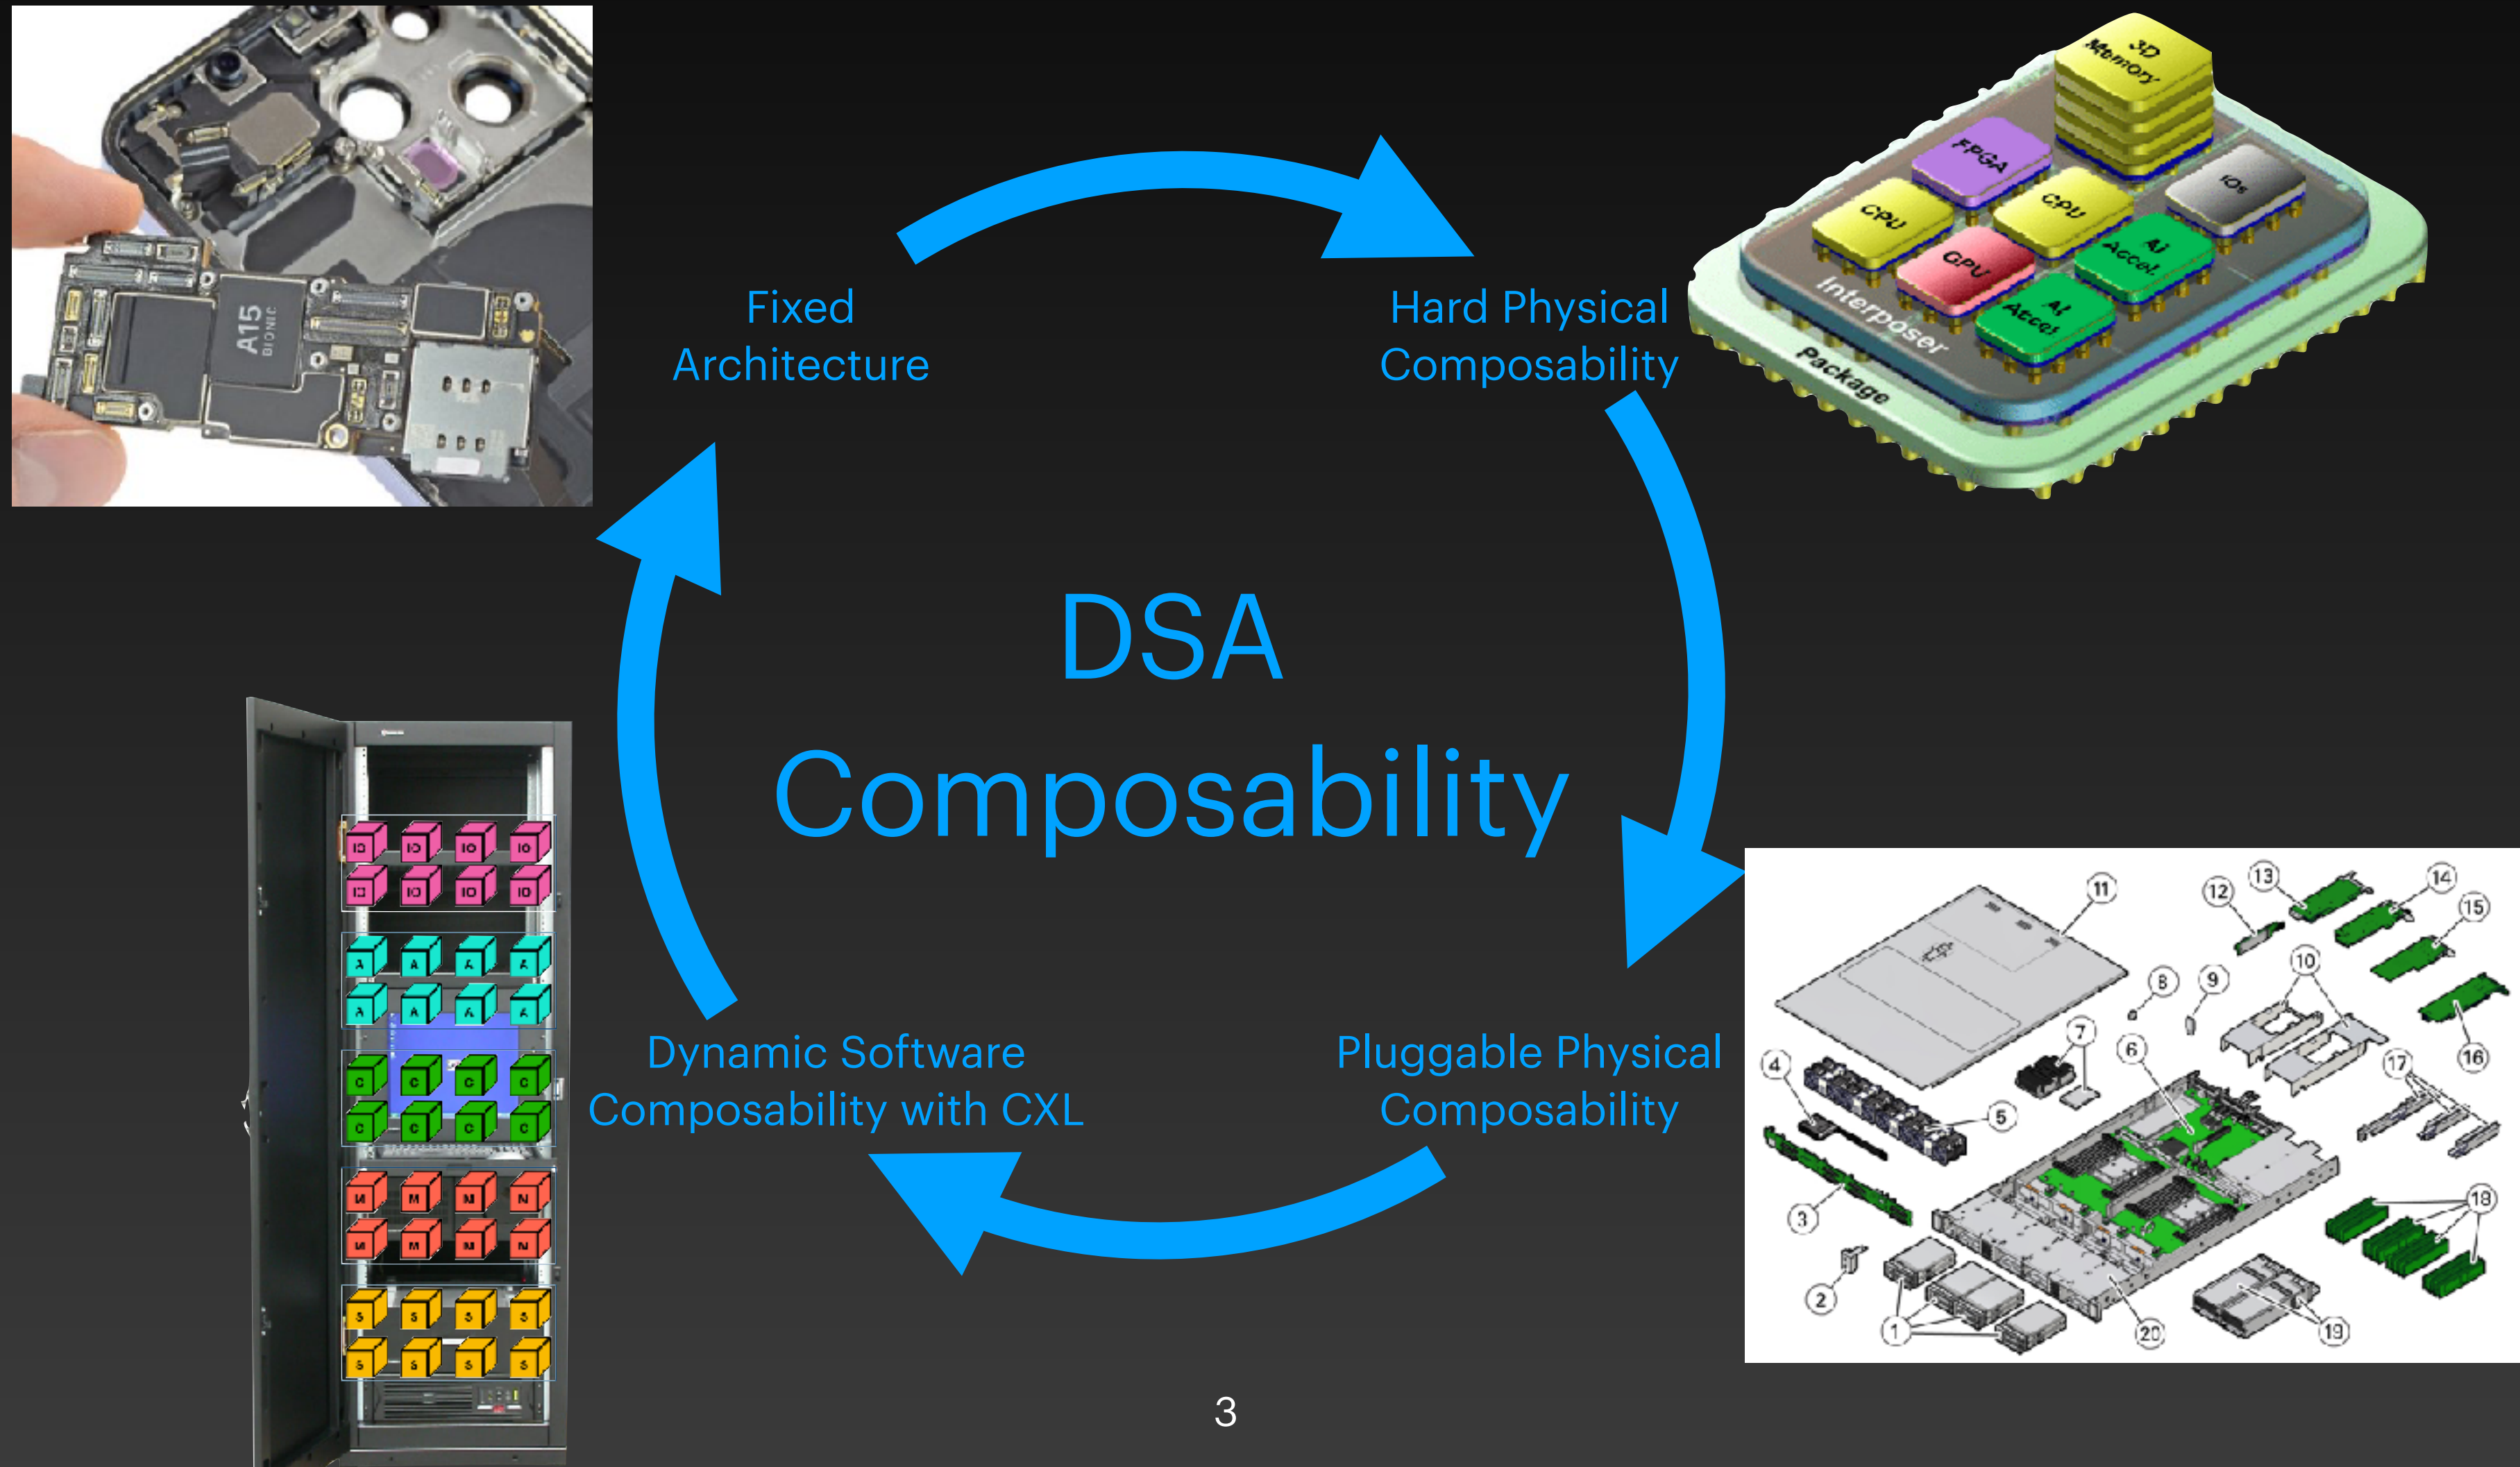
What's driving Computing Architecture

Power & Cost Efficient Domain Specific Architectures, DSA



What's driving Computing Architecture



Power & Cost Efficient Domain Specific Architectures, DSA



What is the Motivation for Chiplets?

Efficiently Combining Compute & Data in Pursuit of Performance


- Computer Architecture is JUST Compute & Data!
 - & Circumventing the Von-Neumann Bottleneck
- Improving Data to Compute Movement Efficiency
 - Moores Law.....shrinking everything
 - Bringing Memory on chip - Caching
 - Heterogeneous Processors
 - Moving everything closer together
 - Chiplets in a package.....2D, 2.5D & 3D



But what about Off Package IO?

System size rarely shrinks with higher density packaging!

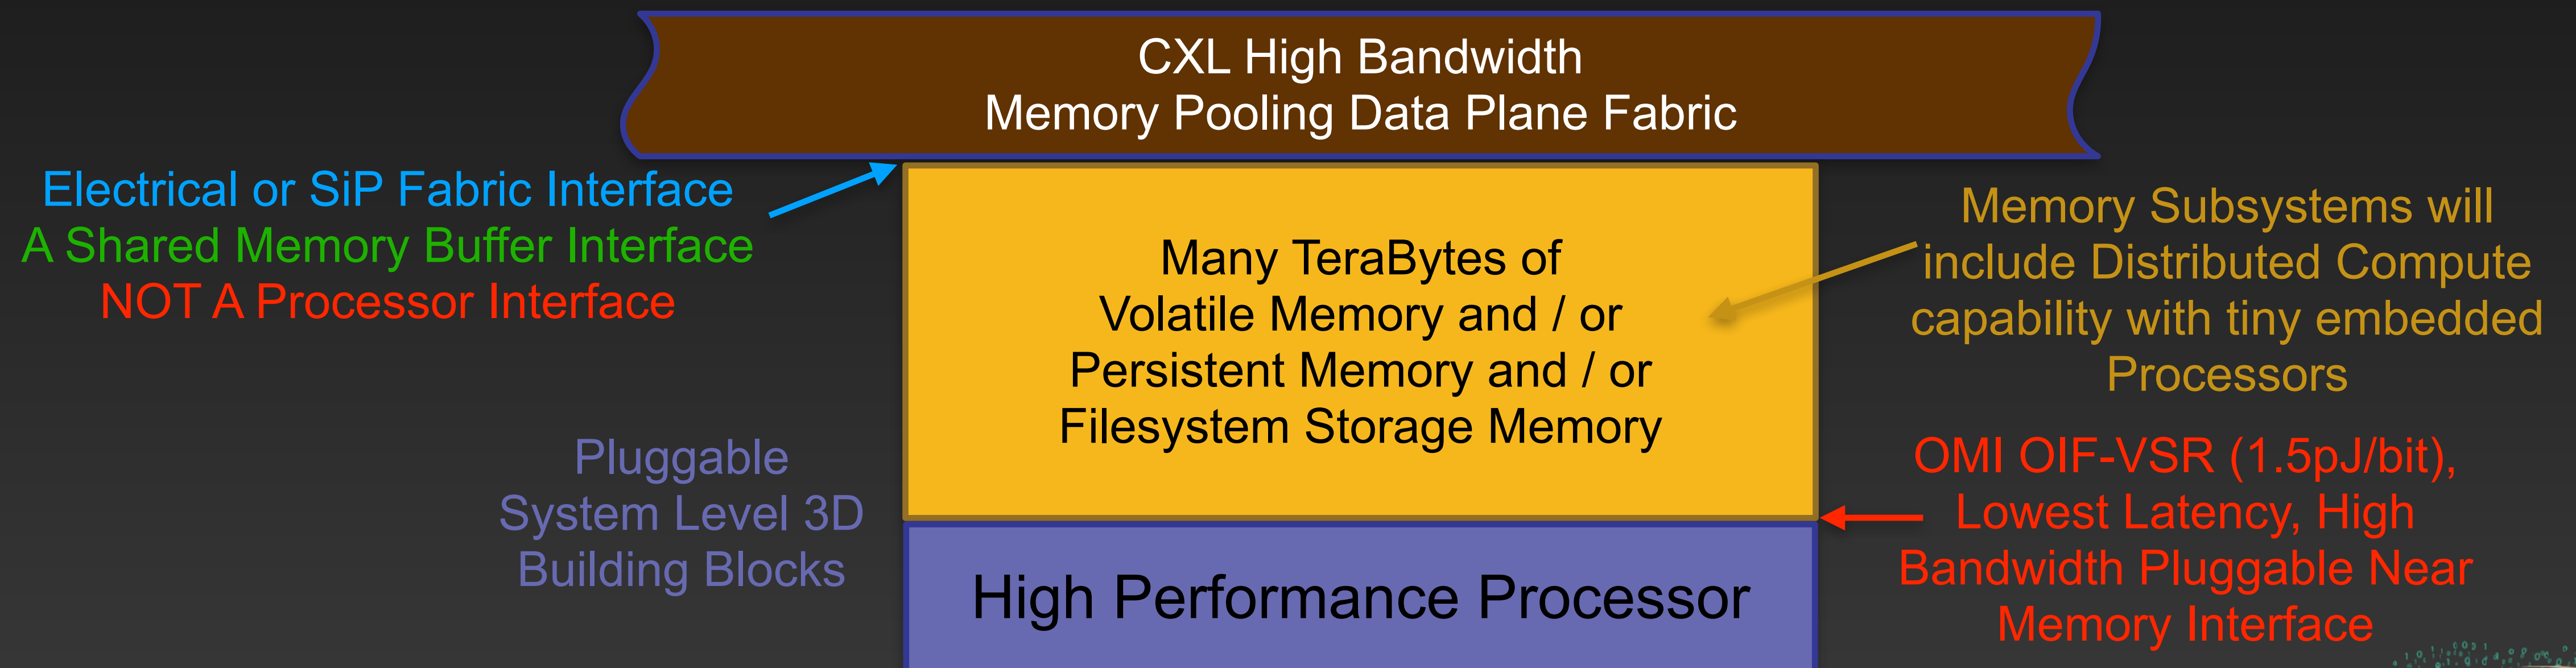
- From ~0.5pJ/bit chiplet IO Energy
 - To ~10pJ/bit System IO Energy
 - A 20x off package increase!
- But there is the Promise of Co-Packaged Optics
- Ultimately we want Highest System Performance at the Lowest Energy & Cost
- Is there are more Graceful Approach?



Pluggable Chiplet 3D Architectural Concept

Energy Efficient Memory Centric Architecture

- Tightly Couple Compute with ALL/ANY Memory Types
 - using OIF-VSR PHYs with the Low Latency Open Memory Interface, OMI, Protocols
- Efficiently share Processors Near Memory with Other Processors over CXL



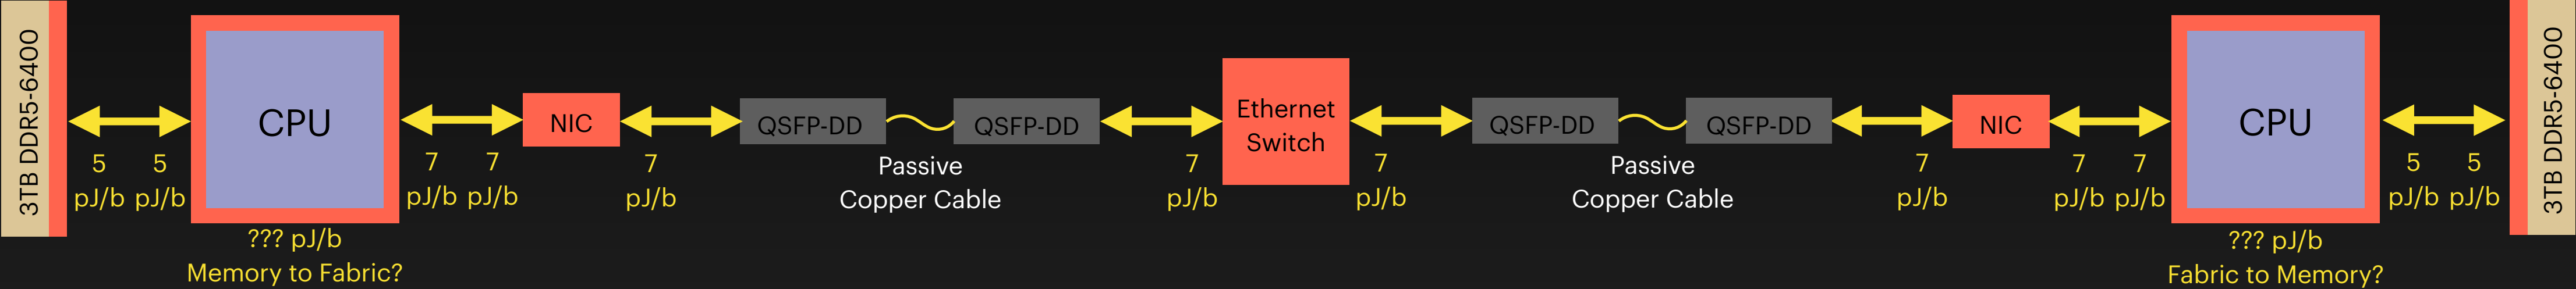
Popular PHY IO Standards

PHY Name	Substrate Type	Ch. Speed Gb/s	Tbits/s/mm Beachfront Tbits/s	Reach mm	Energy pJ/bit	Possible Chiplet IO	Possible Off Package IO
❖ Traditional QSFPxx Optical IO	Organic	112	Same as PHY	100,000's	15	✗	✓
* PCIe-G5 Reference	Organic	32	0.22	973	7	✗	✓
▲ DDR Memory Ch. Reference	Organic	6.4	0.06	300	5	✗	✓
Silicon Photonics	Any	Tb/s	Same as PHY	100,000's	1	✗	✓
* OIF - VSR	Organic	40	1.1	160	1.5	✓	✓
* BoW Fast	Organic	16	1.0	25	<0.5	✓	✗
† UCle - Standard	Organic	32	1.8	<25	0.5	✓	✗
* OpenHBI	Silicon Interposer	8	3.4	8	0.5	✓	✗
† UCle - Advanced	Silicon Interposer	32	10.5	<2	0.25	✓	✗

Potential
Off Package
Chiplet PHY

On and Off Package Memory Configurations

Traditional Local Memory + RDMA Example

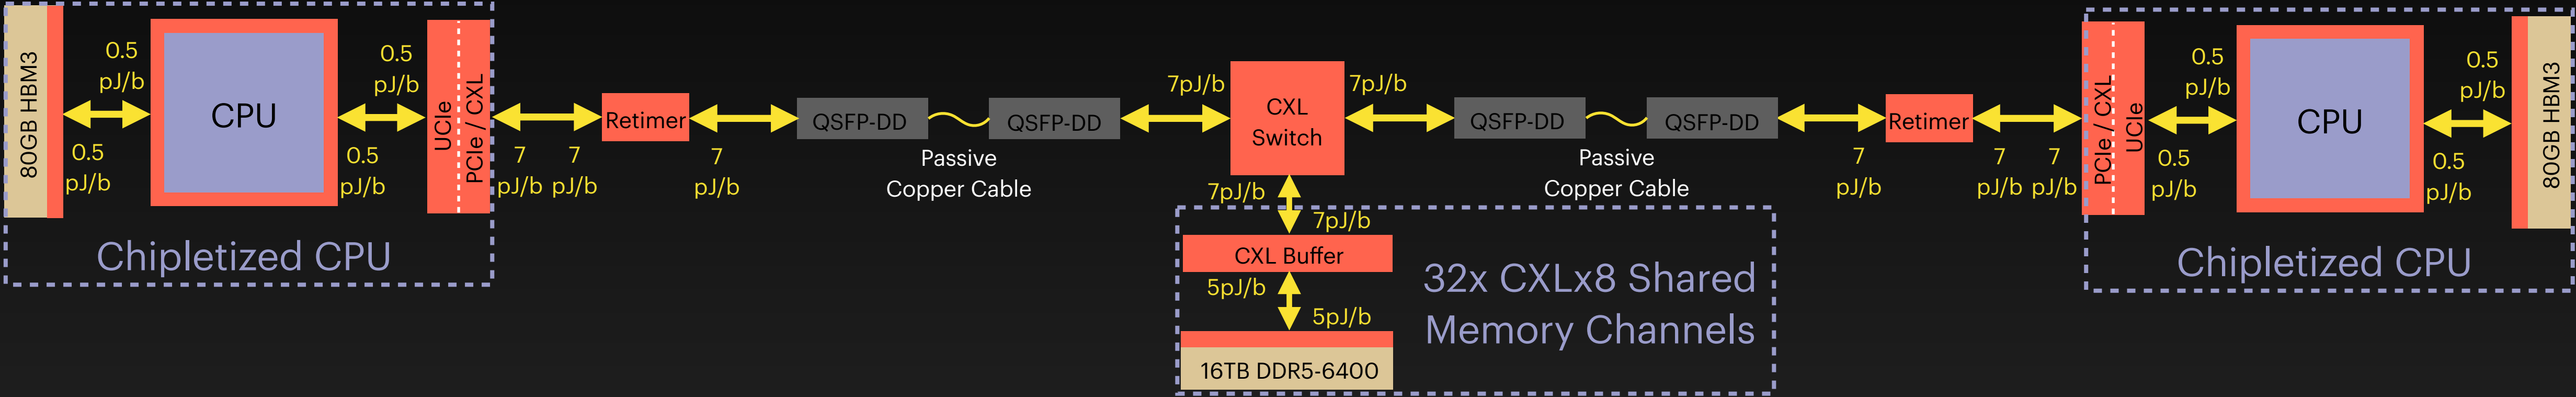


Example	80GB		3TB		8TB		16TB+	
	Bandwidth TBytes/s	Power Watts	Bandwidth TBytes/s	Power Watts	Bandwidth TBytes/s	Power Watts	Bandwidth TBytes/s	Power Watts
Traditional Local Memory + RDMA Example	0.6	49	0.6	49	0.6	373+	0.6	373+

Data Movement Energy shown in Red - Assume NIC & Ethernet Switch Power insignificant compared to their IO power

On and Off Package Memory Configurations

CXL Shared Memory Example

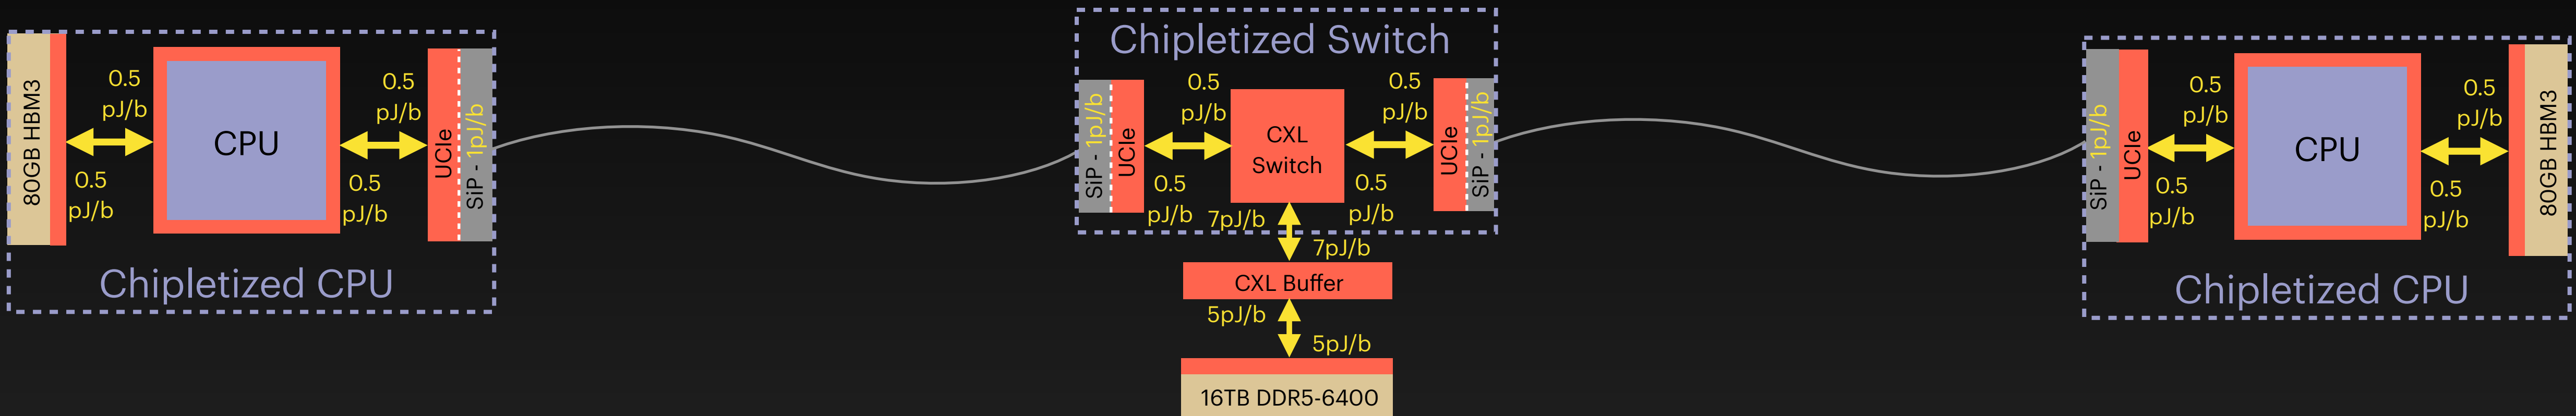


Example	80GB		3TB		8TB		16TB+	
	Bandwidth TBytes/s	Power Watts	Bandwidth TBytes/s	Power Watts	Bandwidth TBytes/s	Power Watts	Bandwidth TBytes/s	Power Watts
Traditional Local Memory + RDMA Example	0.6	49	0.6	49	0.6	373+	0.6	373+
CXL Shared Memory Example	3	24	2	848	2	848	2	848

Data Movement Energy shown in Red - Assume Retimer & CXL Switch Power insignificant compared to their IO power

On and Off Package Memory Configurations

CXL Shared Memory Example with Co-Packaged Optics

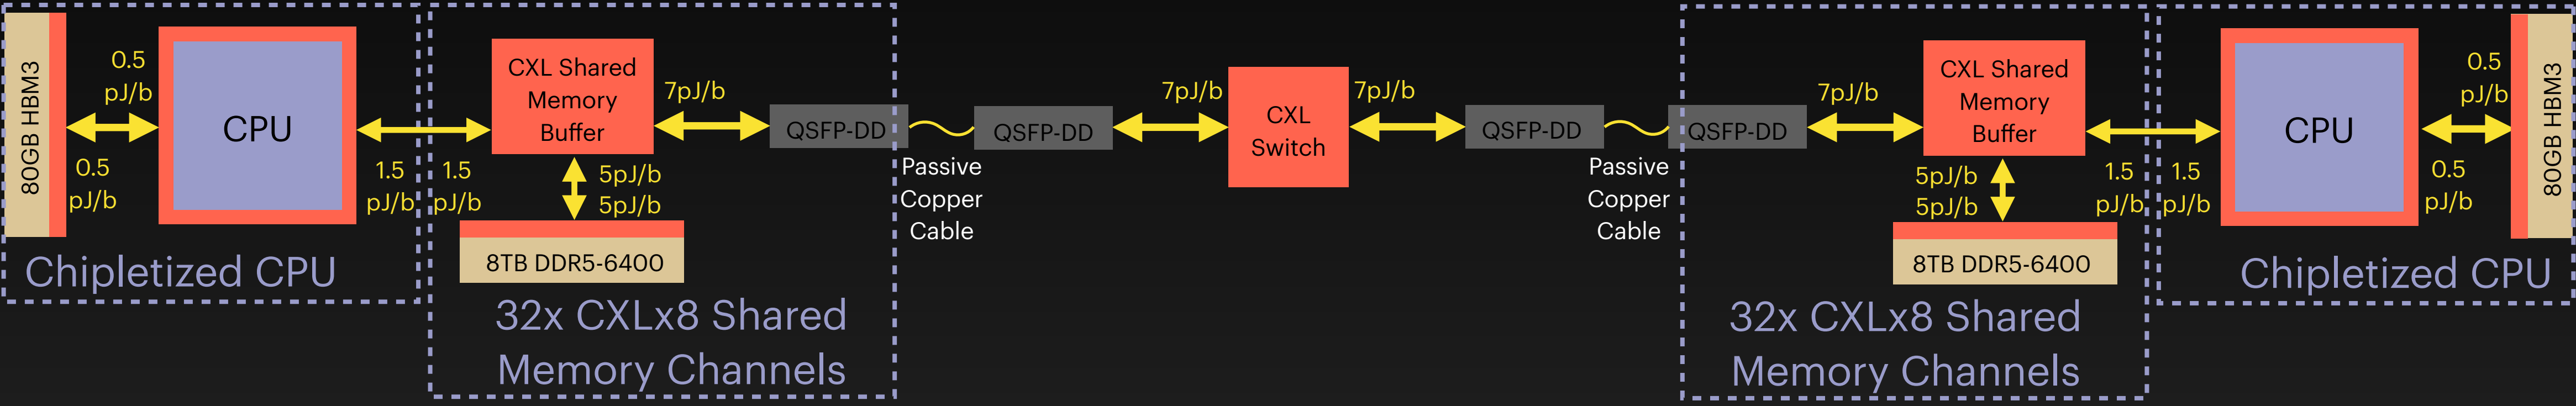


Example	80GB		3TB		8TB		16TB+	
	Bandwidth TBytes/s	Power Watts	Bandwidth TBytes/s	Power Watts	Bandwidth TBytes/s	Power Watts	Bandwidth TBytes/s	Power Watts
Traditional Local Memory + RDMA Example	0.6	49	0.6	49	0.6	373+	0.6	373+
CXL Shared Memory Example	3	24	2	848	2	848	2	848
SiP CXL Shared Memory Example	3	24	2	448	2	448	2	448

Data Movement Energy shown in Red - Assume CXL Switch Power insignificant compared to its IO power

On and Off Package Memory Configurations

CXL Shareable, OMI Local Memory using OIF-VSR Example



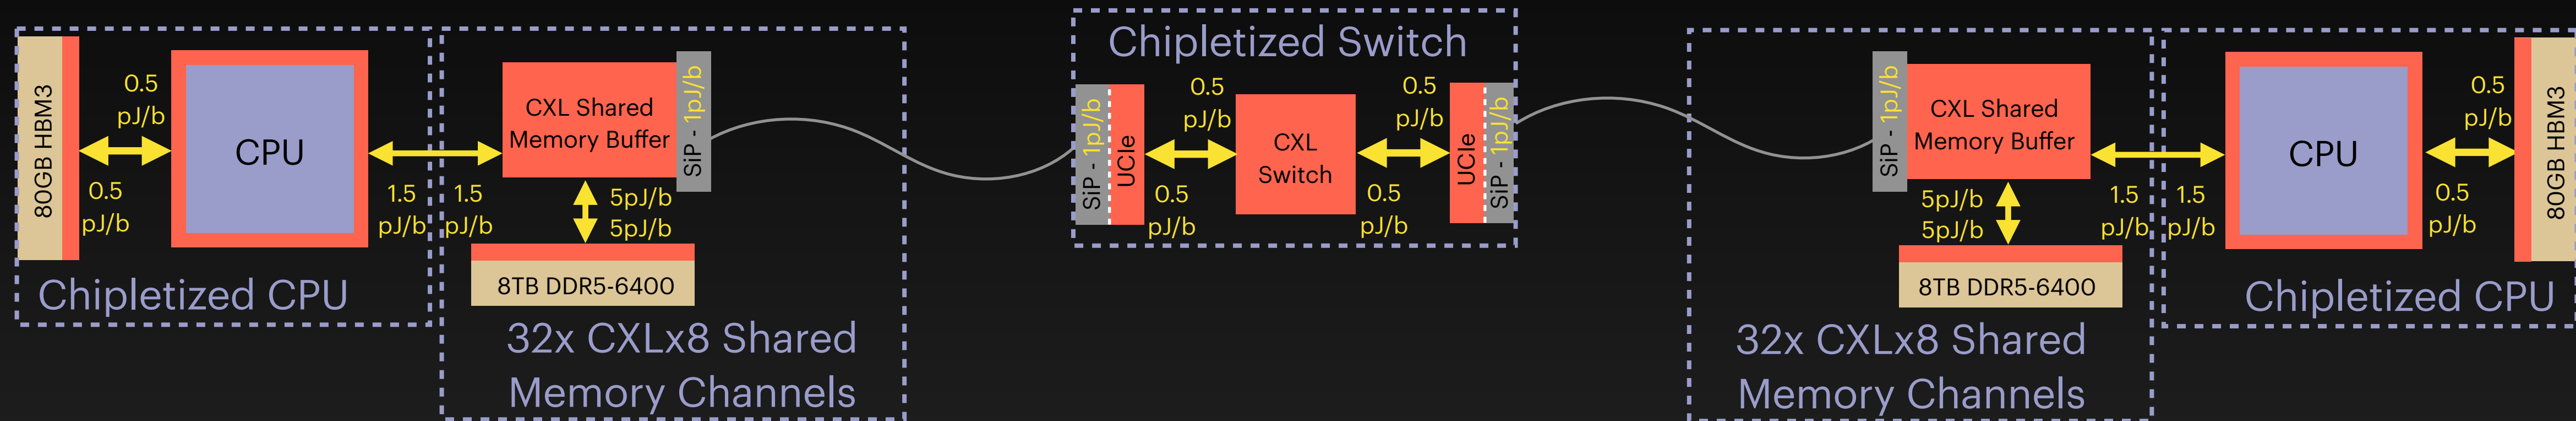
Example	80GB		3TB		8TB		16TB+	
	Bandwidth TBytes/s	Power Watts	Bandwidth TBytes/s	Power Watts	Bandwidth TBytes/s	Power Watts	Bandwidth TBytes/s	Power Watts
Traditional Local Memory + RDMA Example	0.6	49	0.6	49	0.6	373+	0.6	373+
CXL Shared Memory Example	3	24	2	848	2	848	2	848
SiP CXL Shared Memory Example	3	24	2	448	2	448	2	448
CXL Shareable, Local Memory Example	3	24	2	208	2	208	2	656

Data Movement Energy shown in Red - Assume CXL Shared Memory Buffer & CXL Switch Power insignificant compared to their IO power



On and Off Package Memory Configurations

Optical, CXL Shareable, OMI Local Memory using OIF-VSR Example



Example	80GB		3TB		8TB		16TB+	
	Bandwidth TBytes/s	Power Watts	Bandwidth TBytes/s	Power Watts	Bandwidth TBytes/s	Power Watts	Bandwidth TBytes/s	Power Watts
Traditional Local Memory + RDMA Example	0.6	49	0.6	49	0.6	373+	0.6	373+
CXL Shared Memory Example	3	24	2	848	2	848	2	848
SiP CXL Shared Memory Example	3	24	2	448	2	448	2	448
CXL Shareable, Local Memory Example	3	24	2	208	2	208	2	656
SiP CXL Shareable, Local Memory Example	3	24	2	208	2	208	2	304


Data Movement Energy shown in Red - Assume CXL Shared Memory Buffer & CXL Switch Power insignificant compared to their IO power

Watts/TByte/s Memory Sharing Comparison


Energy Advantage of Off-Package Chiplet Shared Memory Buffers

Passive Copper Interconnect

Lower is Better




Silicon Photonic Interconnect



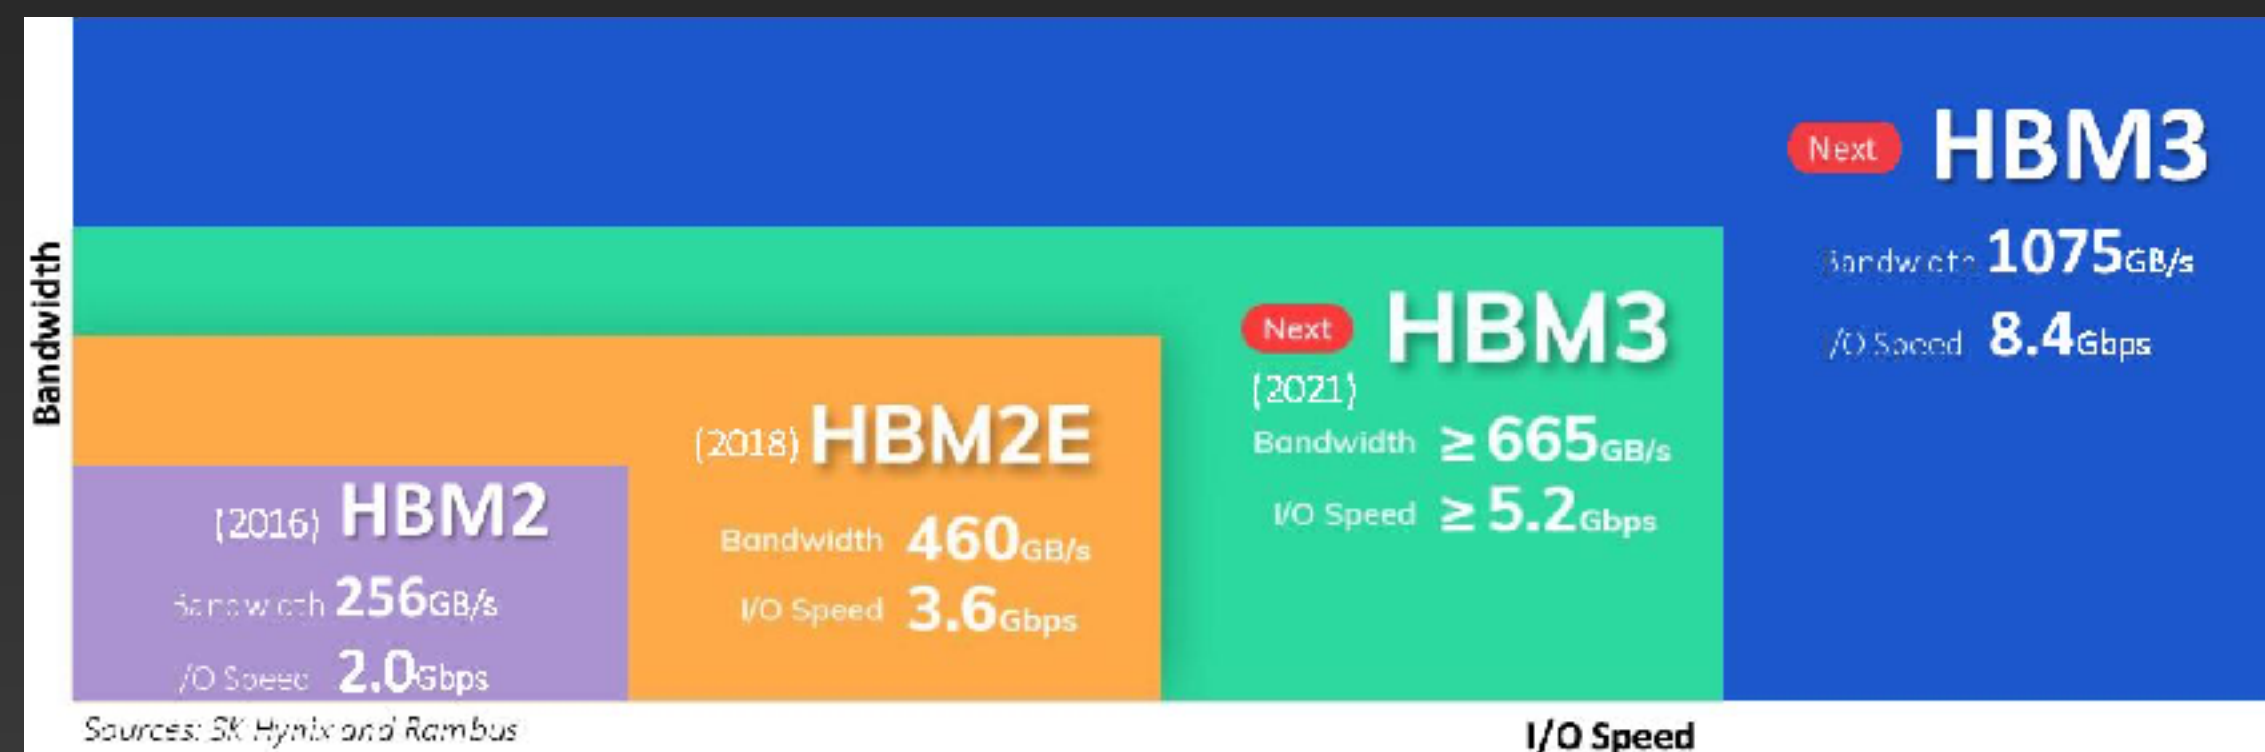

Can UCle Standard support HBM?

UCle-Standard for HBM with Organic packages

- HBM = 11mm beachfront + roadmap to 1.1TB/s
- Requires $1.1\text{TB/s} * 8 / 11\text{mm} = 0.8\text{Tb/s/mm}$
- UCle-Standard & BoW can support HBM using organic substrates





UCle-S / BoW Beachfront



Proposed Chiplet Industry Standardization

Accessible to Tier 2 and 3 Industry Players


- Utilize BoW & OIF-VSR PHYs for On & Off Package Chiplet IO's respectively
- Leverage Cost effective Organic Packaging
- Unified Off Package IO for ALL Processors that will Democratize Compute



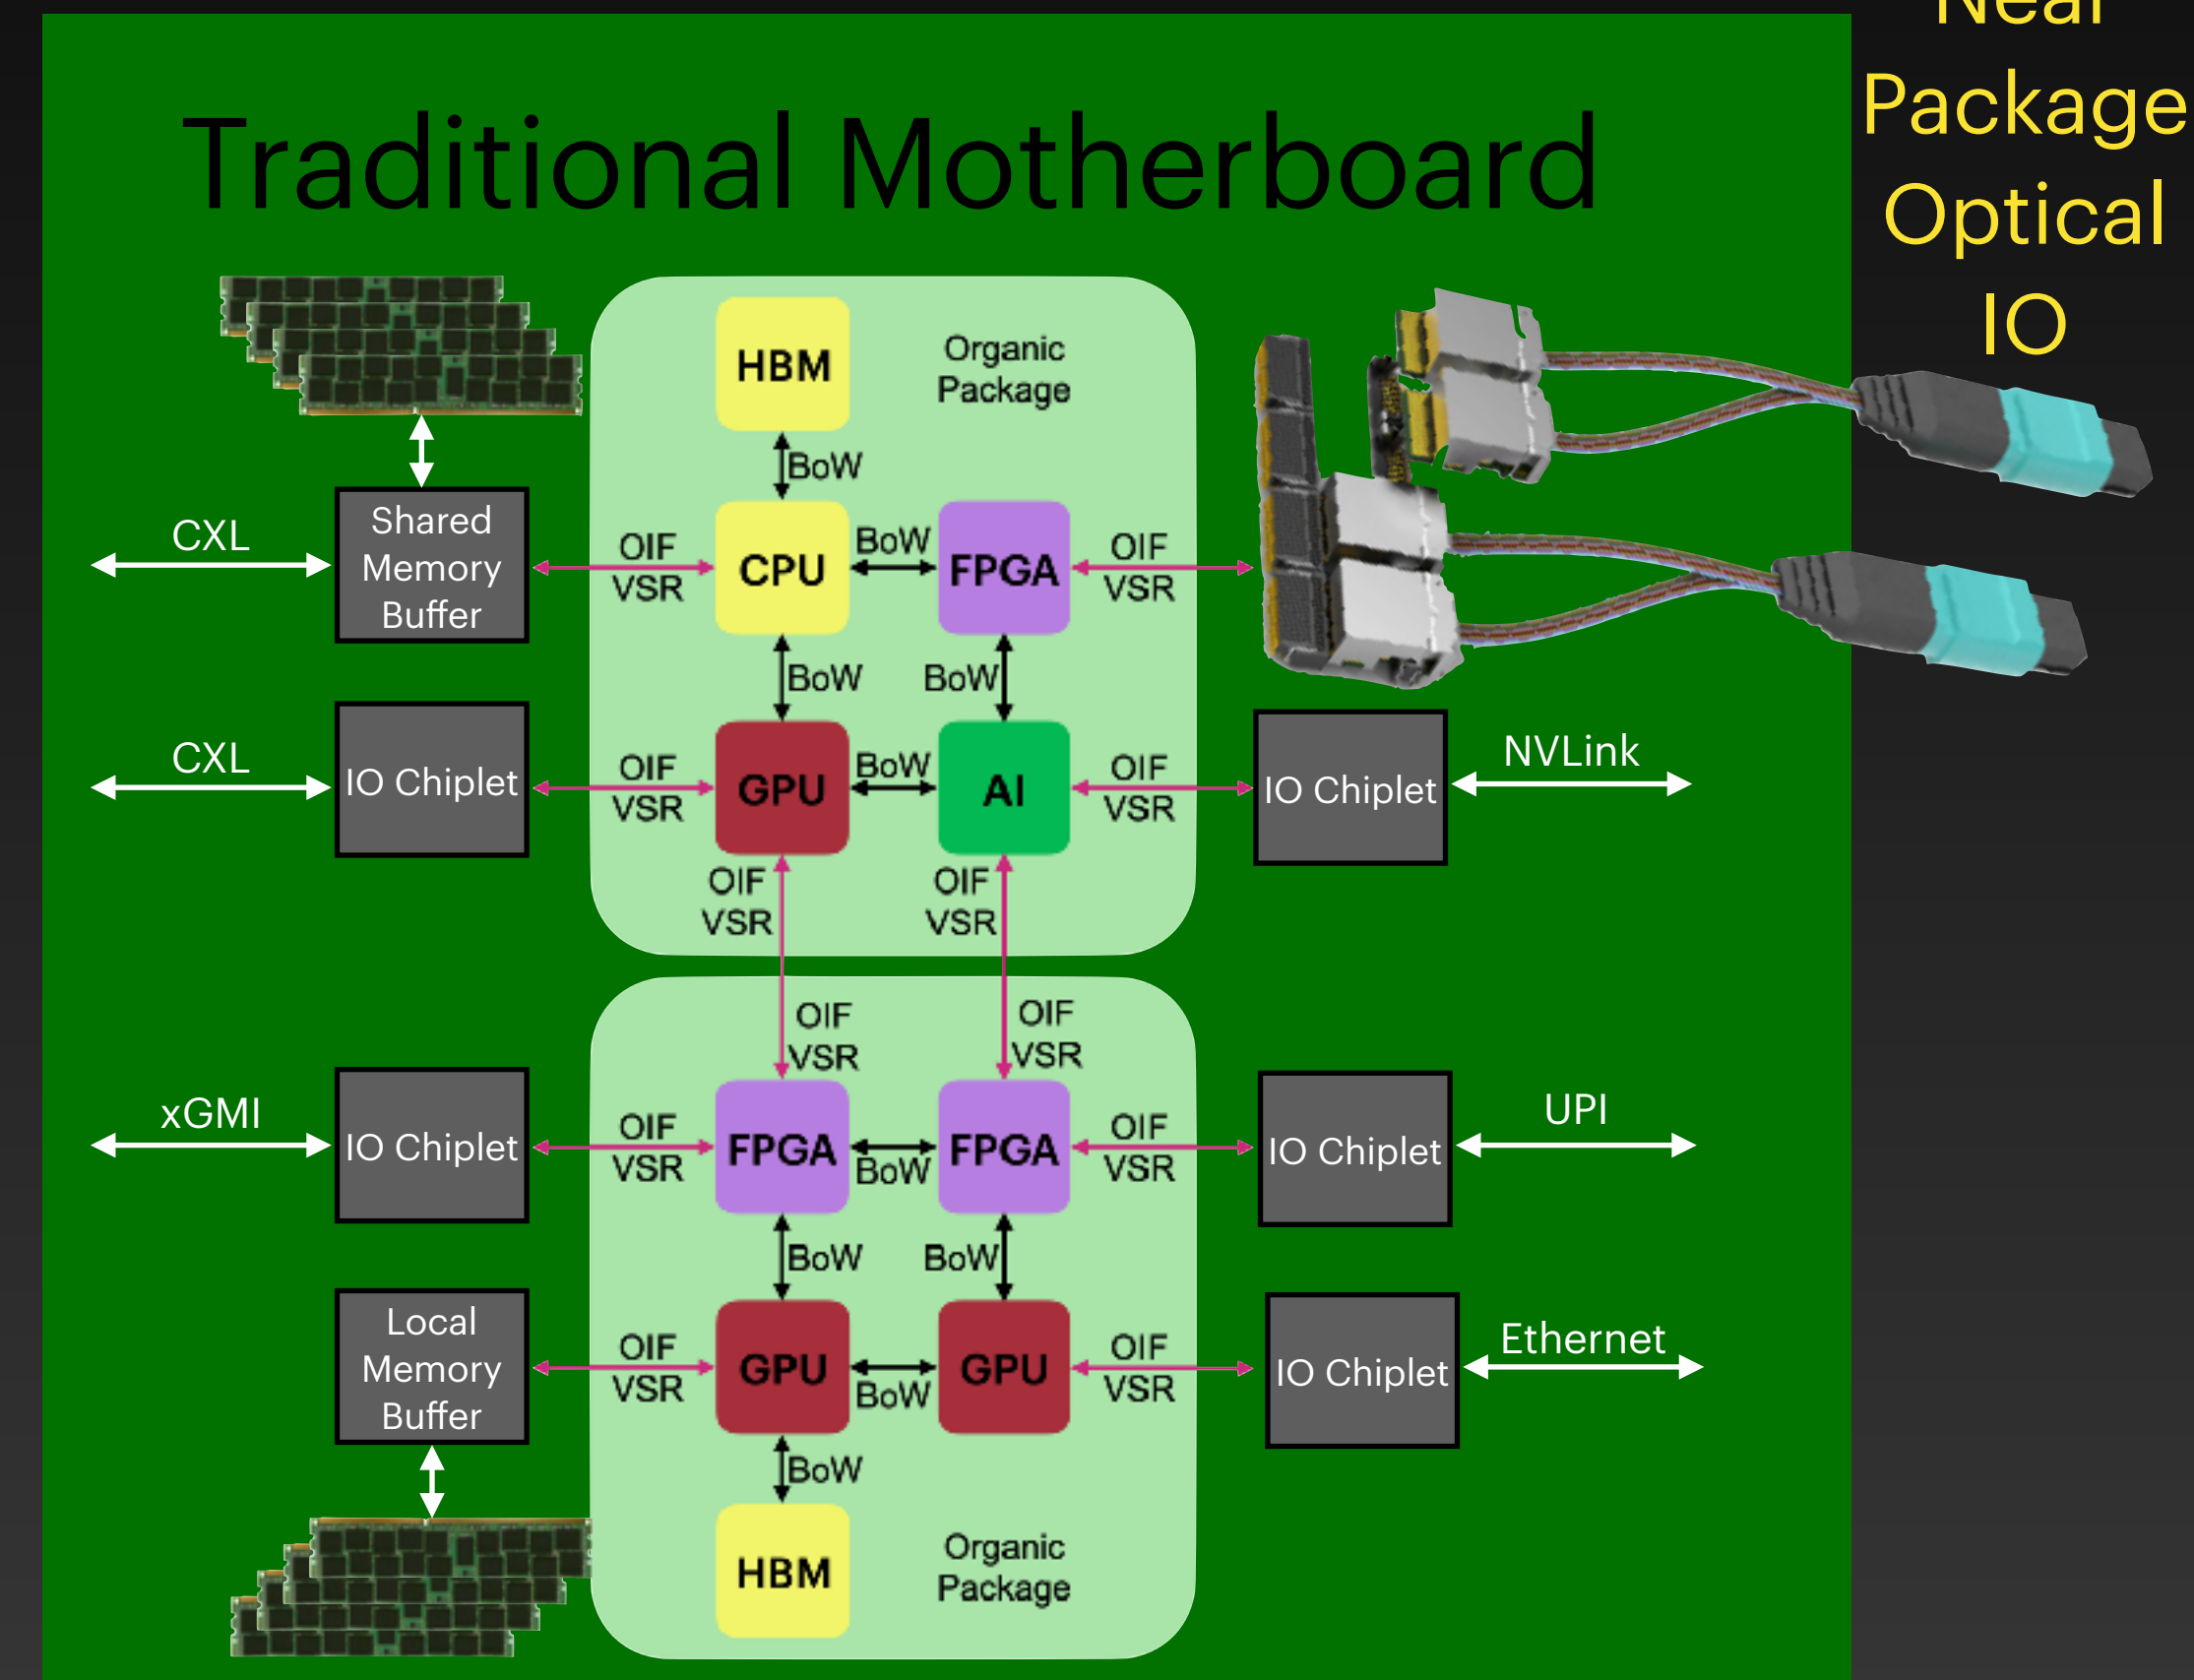
Proposed Chiplet Industry Standardization

Accessible to Tier 2 and 3 Industry Players

- Off-Package Chiplets are simply IO Translators!
 - For differing IO Use Cases
- OCP HPCM - 3D System Level Packaging



Novel Pluggable Chiplet 3D System Level Packaging₆



Summary

- Today's Disaggregation Models are driving : -
 - Energy Efficiency in Wrong Direction
- Supporting an off-package Chiplet IO Standard allows :-
 - A Graceful increase in Power & Latency when moving data
 - A more Flexible & Composable Computing Architecture
 - The potential for IO Unification across All Heterogeneous Processors
 - A Lower cost, Sustainable, and Competitive Ecosystem