ODSA: Technical Introduction

Bapi Vinnakota, Netronome

ODSA Project Workshop

June 10th, 2019

Consume. Collaborate. Contribute.

ODSA: A New Server Subgroup (Incubation)

- Extending Moore's Law
  - Domain-Specific Architectures: Programmable ASICs to accelerate high-intensity workloads (e.g. Tensorflow, Network Flow Processor, Antminer…)
  - Chiplets: Build complex ASICs from multiple die, instead of as monolithic devices, to reduce development time/costs and manufacturing costs.
- Open Domain-Specific Architecture: An architecture to build domain-specific products
  - Today: All multi-chiplet products are based on proprietary interfaces
  - Tomorrow: Select best-of-breed chiplets from multiple vendors
  - Incubating a new group, to define a new open interface, build a PoC

Open Interface for Chiplet-Based Design

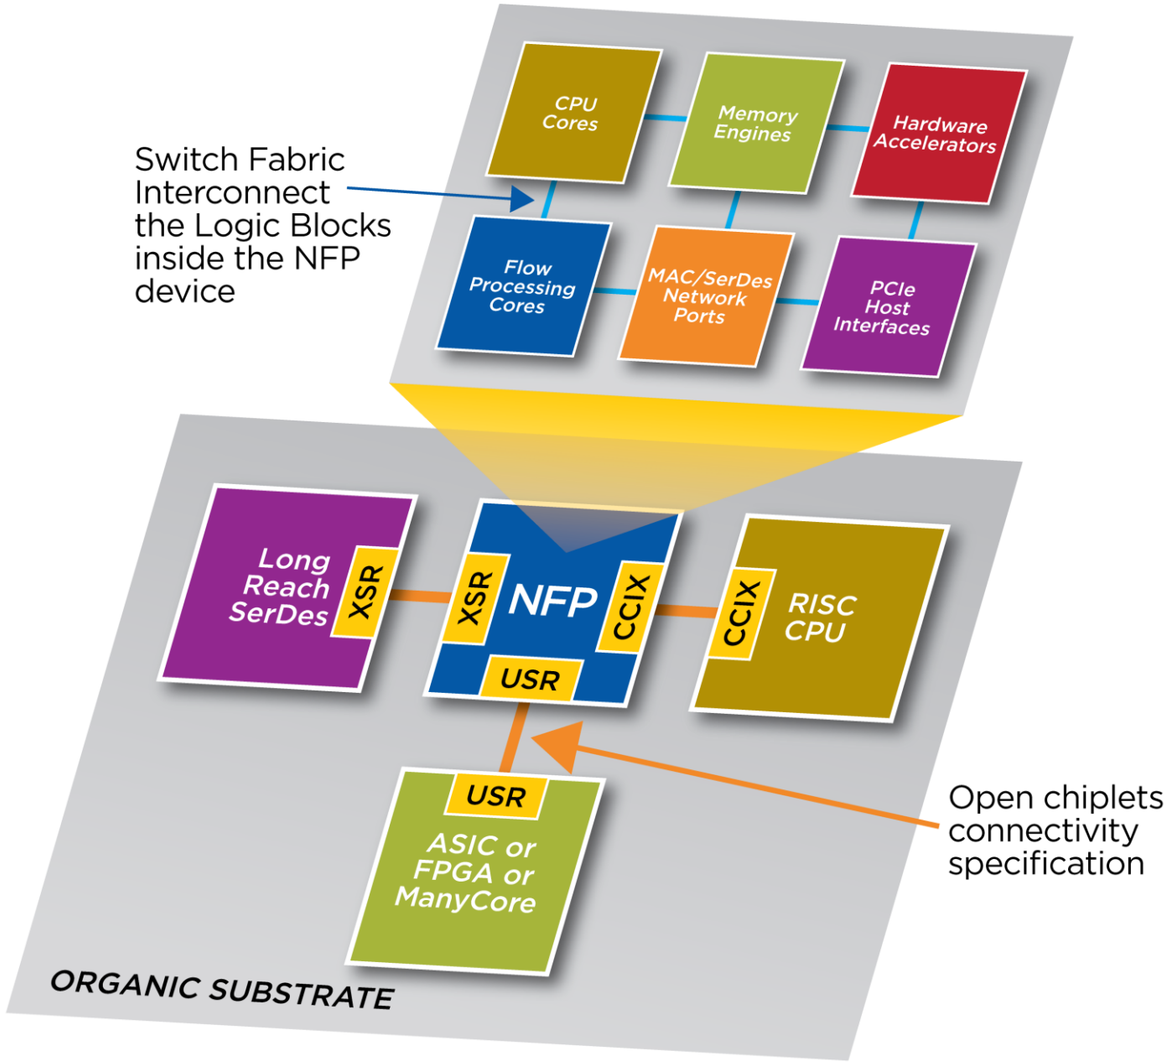

Multiple chiplets need to function as though they are on one die

Application	Firmware Infrastructure for: Accelerator Use, Host Integration[1]			
Transport	Coherent Protocol (CCIX[2] or TileLink[2] or...)	Instruction-Driven Transfer (ISF Transport Layer)[1]	DMA	
Network	Routing (To Be Developed From ISF Routing Layer)			
Link	Inter-Chiplet (e.g. PCIe[2], TileLink[2] or new Link Layer[1])	Intra-Chiplet (From ISF Link Layer[1])		
PCIe PIPE	PIPE Adapter	PIPE Adapter	PIPE Adapter	PIPE Adapter
PHY	PCIe[2] PHY	XSR[2] SerDes	USR[2] SerDes	Parallel BoW[1]/AIB[2]
Substrate	Organic Substrate Packaging[1] + Interconnect	Interposer		
Chiplet	NFP	RISC[2]	FPGA	ML ASIC	SerDes	•••

Architecture Interface

[1] New Open IP/Specification
[2] Existing Open Standard

Source: ODSA

~~How to Participate~~
Please Help! : Join a Workstream

Join the PoC, Build fast:
(Quinn Jacobson/Jawad Nasrullah)

Join Interface/Standards:
(Mark Kuemerle/Aaron Sullivan)

Join Business, IP and workflow:
(Sam Fuller)

Develop software

Provide ODSA chiplets

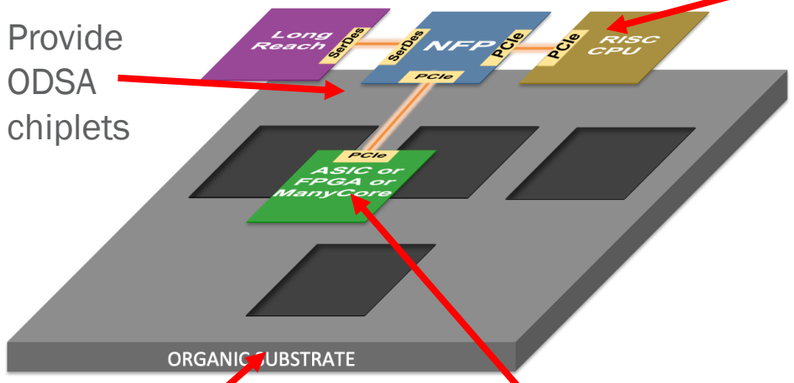

Develop Packaging + Socket, Dev Board

Provide FPGA IP

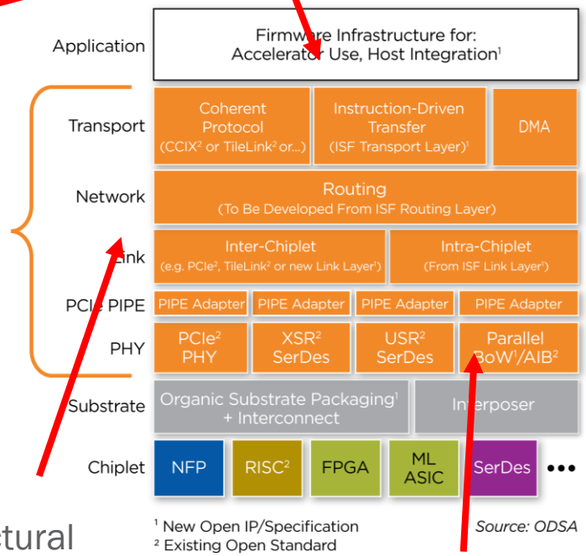

Define Architectural Interface

Provide PHY technology

Define test and assembly workflow

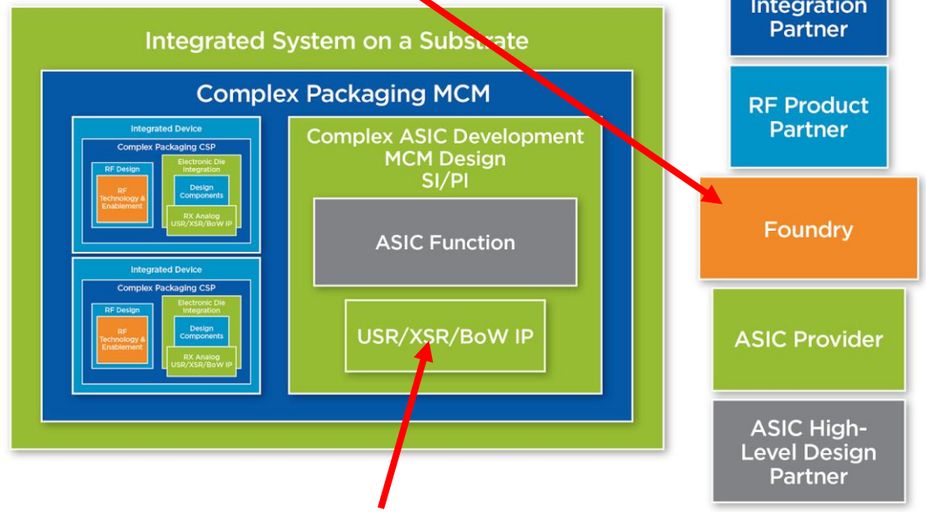

Provide Chiplet IP

Workstream contact information at the ODSA wiki

Consume. Collaborate. Contribute.

Domain-Specific Architectures

Tailor architecture to a domain*

- Server-attached devices — programmable, not hardwired
- Integrated application and deployment-aware development of devices, firmware, systems, software
- 5-10X power performance improvement

- Big - more of a processor to I/O mismatch => more memory
- Each serves a smaller market

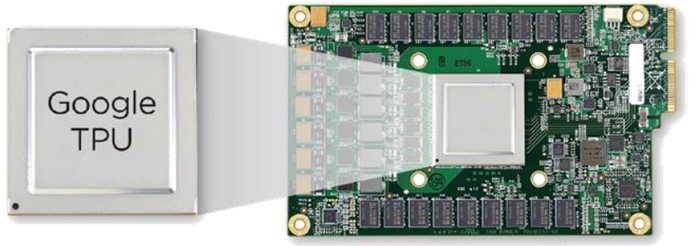

Domain-Specific for Machine Learning and AI

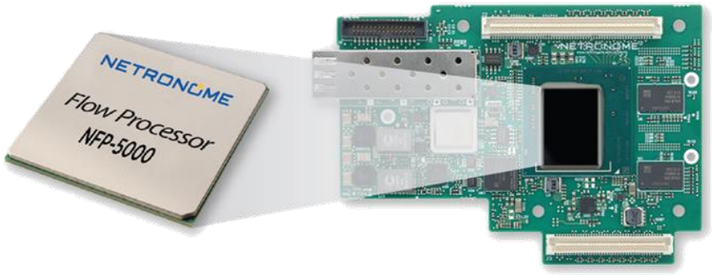

Domain-Specific for Networking and Security

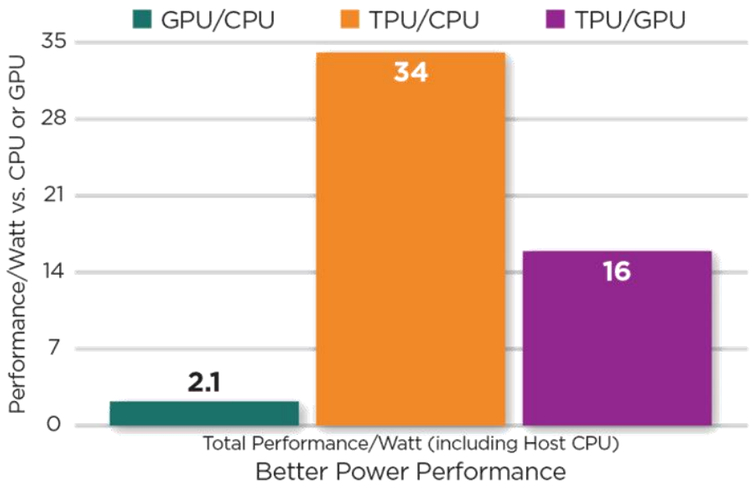

Google TPU vs. CPU and GPU

Source: "An in-depth look at Google's first Tensor Processing Unit (TPU)," Google Cloud, May 2017

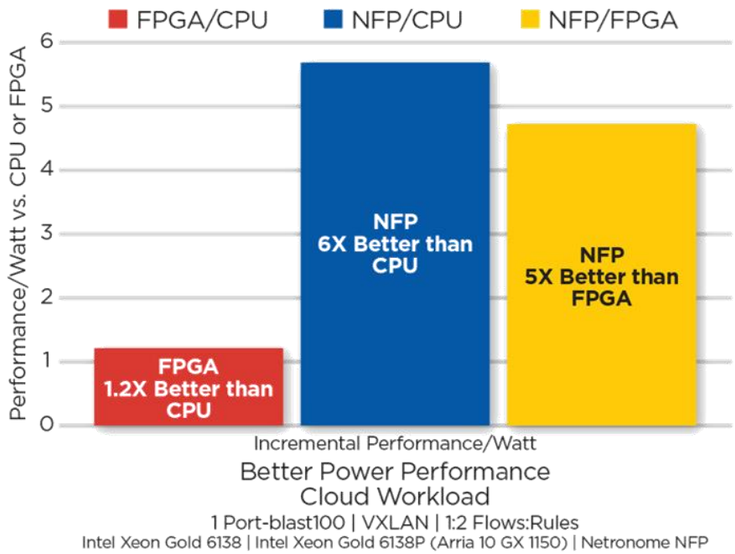

Netronome NFP vs. CPU and FPGA

Source: Netronome, based on internal benchmarks and industry reports related to Xeon CPUs and Arria FPGAs

*A New Golden Age for Computer Architecture
John L. Hennessy, David A. Patterson
Communications of the ACM, February 2019, Vol. 62 No. 2, Pages 48-60

Monolithic vs Chiplets

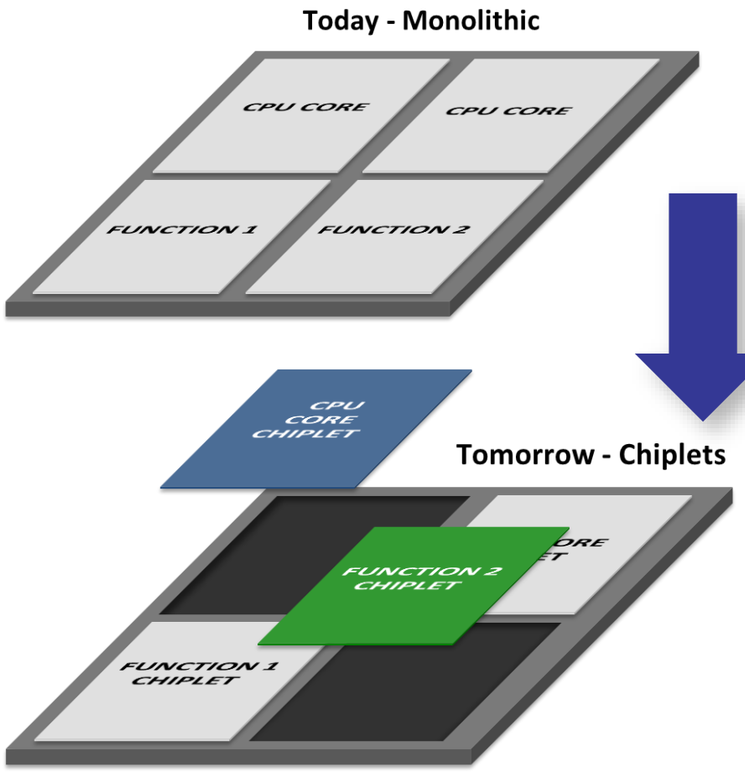

AMD Data

4 Die are ~30% cheaper than a single large die

Shrink: Monolithic process shrink

Integration: Multi-chip on same process

Integration provides nearly all the benefits of a shrink at a fraction of the cost, because of efficient inter-chiplet interconnect

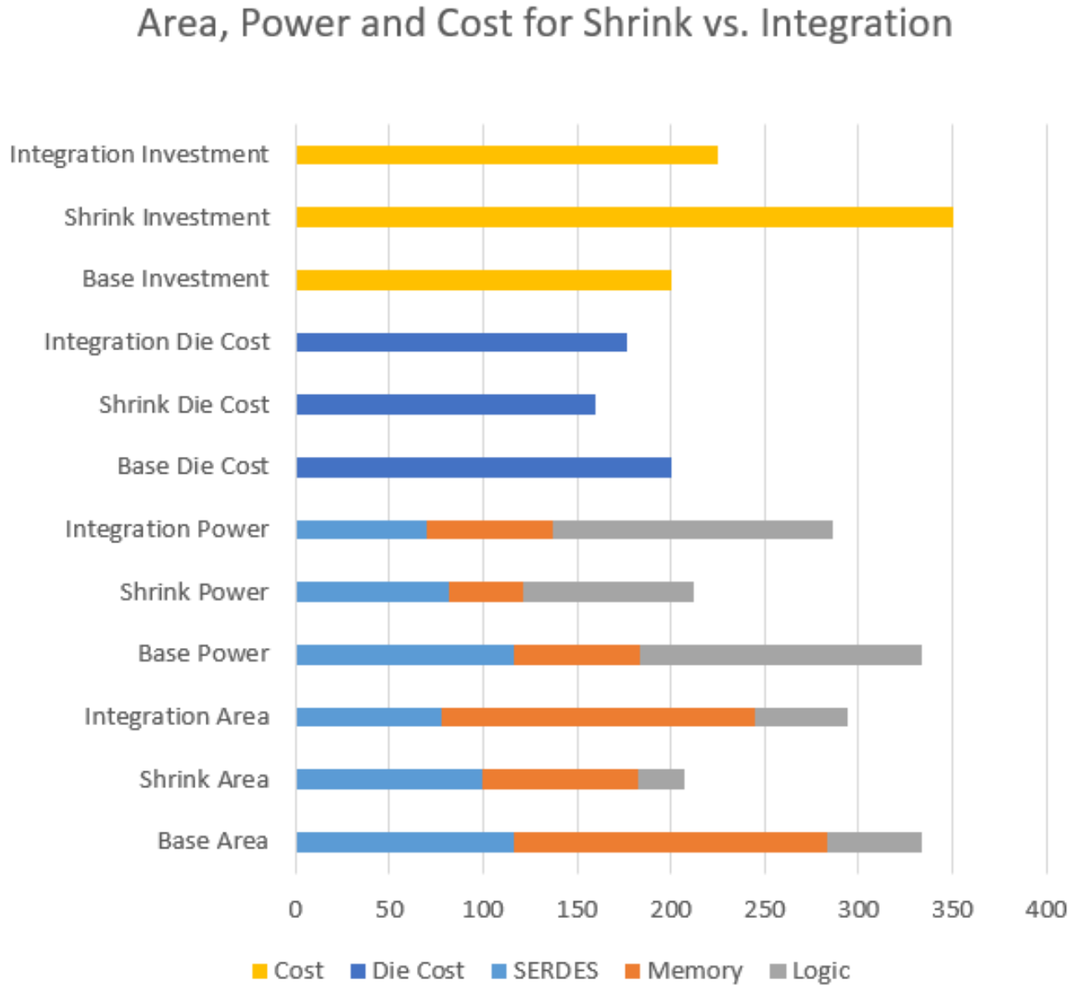

https://www.netronome.com/media/documents/WP_ODSA_Open_Accelerator_Architecture.pdf

Chiplets for DSAs

Design Function	Value
IP Qualification	Verified IP for inter-chiplet communication
Architecture	Leverage reference architecture. Focus investment on domain-specific logic. Reuse chiplets instead of IP for 40% of the functions in a monolithic design
Verification	
Physical	
Software	Open source firmware and software for host-attached operation
Prototype	Aim for reference package design with area, power budgets and pinouts for components
Test and Validation	Develop workflow for chiplets

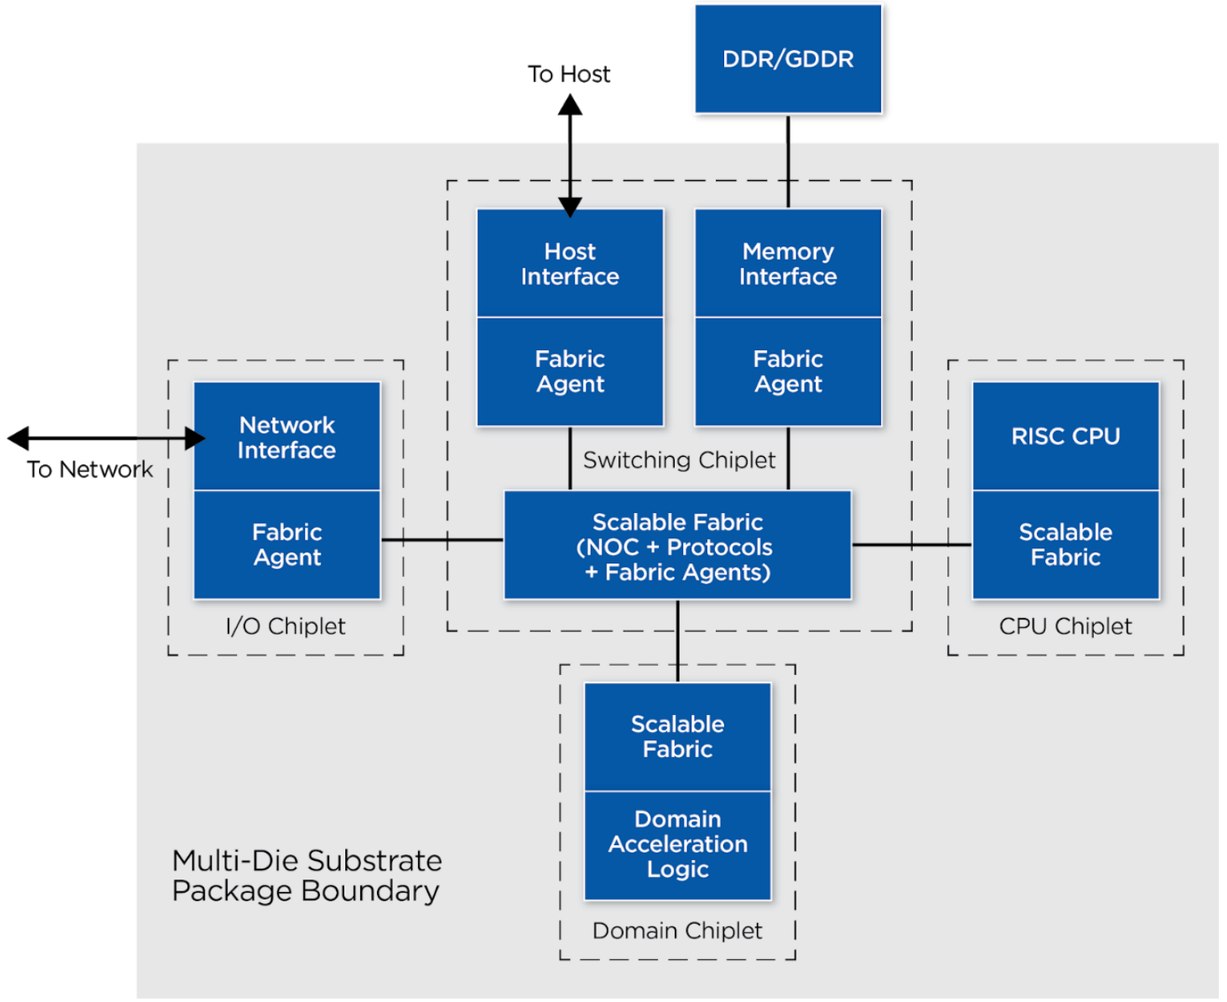

Chiplet reuse reduces development costs
Partition large devices into smaller devices with better yield.

Consume. Collaborate. Contribute.

Cross-Chiplet ODSA fabric

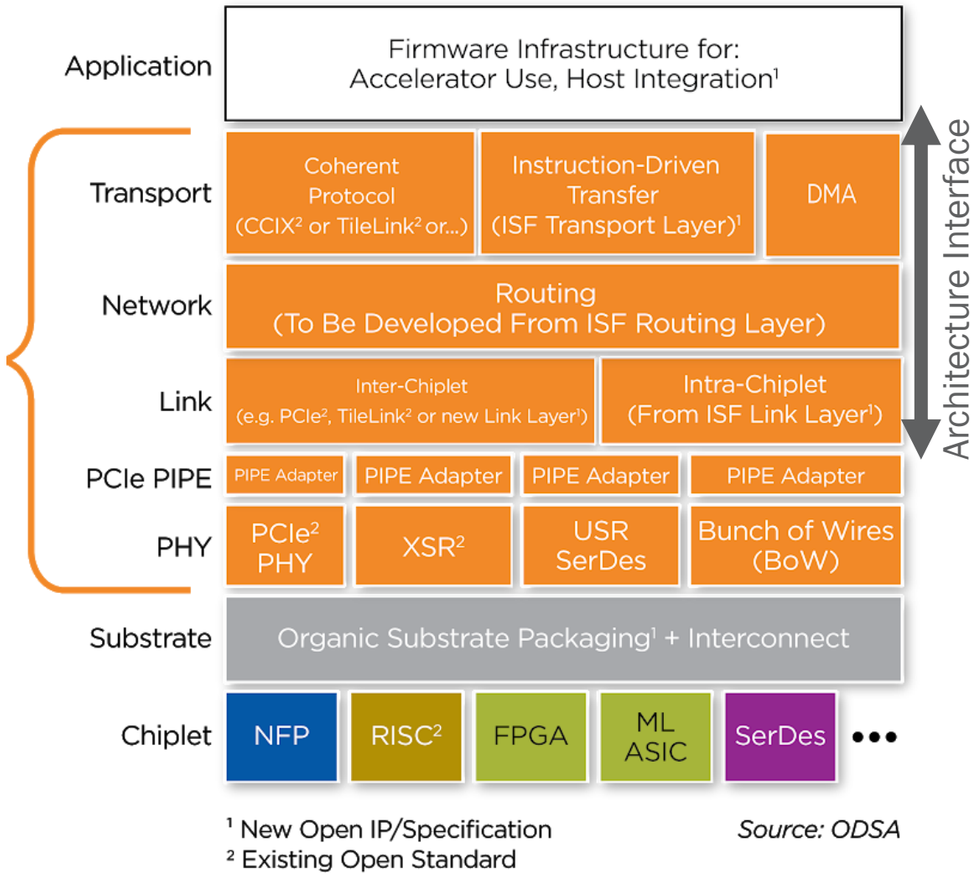

[1] New Open IP/Specification
[2] Existing Open Standard

Source: ODSA

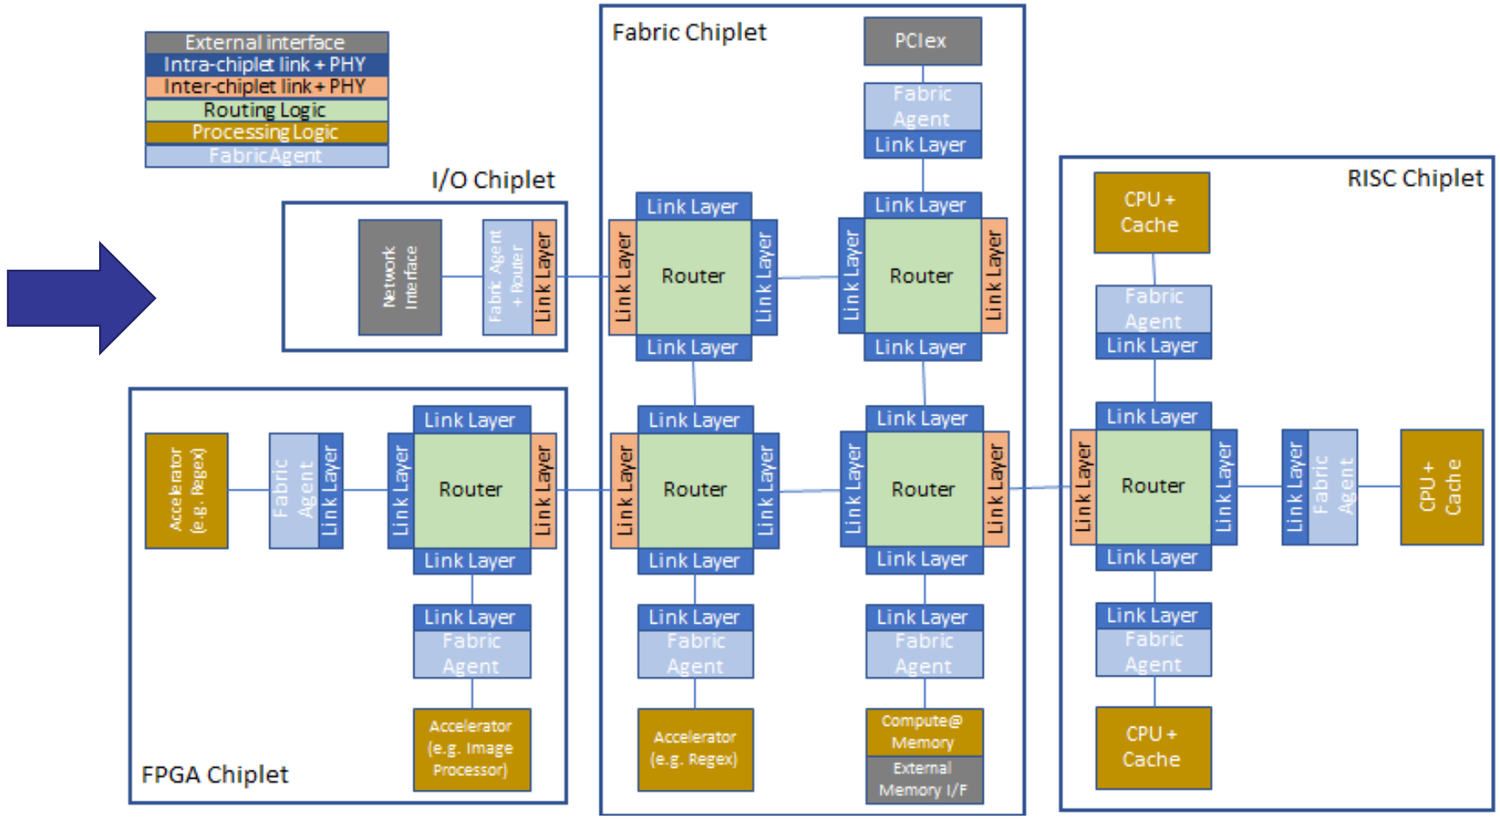

ODSA Scope

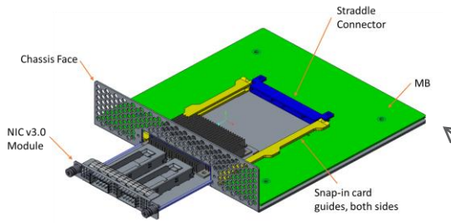

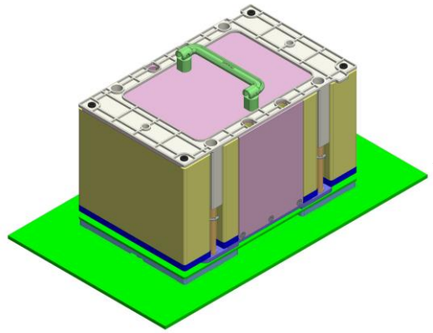

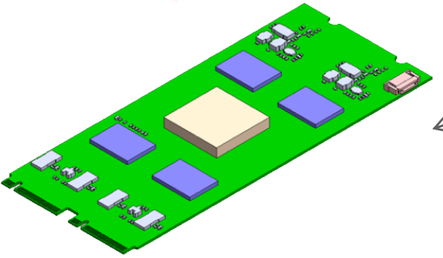

Accelerators drive the requirements
Chiplets are a means to meet requirements

Assemble DSAs from a library of ODSA-compliant chiplets.

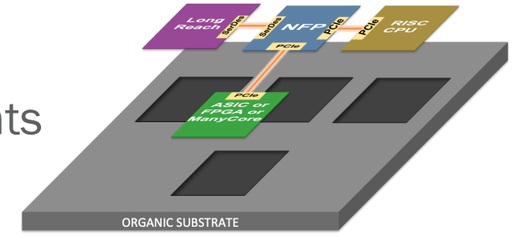

Reference architectures, PoCs for Networking, Storage, Inferencing, Training, Video and Image processing

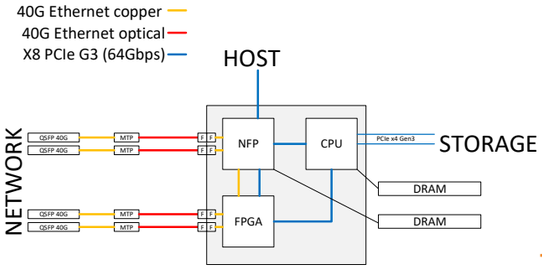

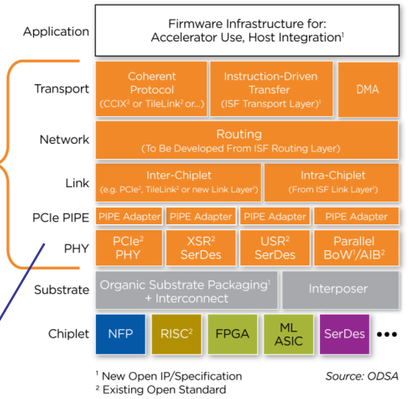

Multi-technology ODSA stack

OCP Form Factors drive Power, I/O Footprint, Performance

Business, Tools Workflow to assemble Product

Stack Compliant Interoperable Chiplet Marketplace

Consume. Collaborate. Contribute.

Timeline

• ODSA Announced	10/1/18	7 companies
• White Paper	12/5/18	10 companies
• First Workshop @Global Foundries	01/28/19	35 companies
• Joined OCP	03/15/19	
• Second Workshop @Samsung	03/28/19	53 companies
• Today @Intel	06/10/19	65 companies

Meet weekly on Fridays. Status updates, new project proposals, guest speakers. All the content on the ODSA wiki - https://www.opencompute.org/wiki/Server/ODSA

We may not have the right solution, we likely have the right problem.

Our Progress - How You Can Participate

Project	Objective	Participants	Recent Results	Upcoming Milestones	Needs
PHY Analysis	PHY requirements PHY analysis Cross-PHY abstraction	Alphawave,Aquantia, Avera Semi, Facebook, Intel, Kandou, Netronome, zGlue,	PHY Analysis paper (to be published at Hot Interconnect in August)	PIPE abstraction Operations, test and management	
BoW Interface	No technology license fee, easy to port inter-chiplet interface spec	Aquantia, Avera Semi, Netronome	New BoW Interface (to be published at Hot Interconnect in August)	Data i/f spec, Aug, 2019 0.9 spec, Sep, 2019	Foundry support for test chips. Chiplet library with interface Open source implementation
Prototype	Device that integrates existing die from multiple companies into one package	Achronix, Cisco, Netronome, NXP, Samtec, Sarcina, zGlue	Decomposable design flow.	Committed schedule	End users End user participation ~30% funding is open
Chiplet Design eXchange	Open chiplet physical description format starting with zGlue format. Information normally confidential.	Ayar, NXP, zGlue	Open chiplet survey.	ZEF Exchange format draft specification	EDA participation
Inter-chiplet Link Layer	Interface and implementations – requirements and definition	Achronix, Avera Semi, Intel, Netronome, more needed			

Listing participation does not imply official endorsement by employer

Projects Requested

Project	Objective	Participants	Recent Results	Upcoming Milestones	Needs
Cross-chiplet network layer, fabric agents	Scalable network layer. Netronome offers a starting point	Netronome			
Pchiplet design flow	Chip/chiplet open design flow to integrate across companies				
Reference Architectures	I/O, Compute, Memory, functional partition for SmartNIC, Inferencing, Storage, Learning, Image/Video				Leads to assess requirements, create reference designs
Chiplet proposals	Proposals for chiplets for common functions – I/O, CPU, Memory				Chiplet vendors supporting a PHY interface
Business workflow	Leverage learnings from prototype effort