UCIe™ (Universal Chiplet Interconnect Express™)

Building an open ecosystem of chiplets for on-package innovations

Tutorial for Hot Chips, 2023

Presented by: Debendra Das Sharma[1], Nathan Kalyanasundaram[2], Swadesh Choudhary[1], Marvin Denman[3], Gerald Pasdast[1], Anwar Kashem[2], Jerome Giselle[4], and Sridhar Muthrasanallur[1]

[1] : Intel Corporation, [2]: AMD Corporation, [3]: Nvidia Corporation, [4]: Google Inc

Tutorial Sessions

- Overview
- Protocol
- Electrical

Coffee Break

2

- Form-Factor and Compliance
- Software, Management and Security

UCIe Overview

120+ Member Companies and growing!

Board Members

Leaders in semiconductors, packaging, IP suppliers, foundries, and cloud service providers are joining together to drive
The open chiplet ecosystem.

JOIN US!

UCIe Consortium is open for membership

- UCIe Consortium welcomes interested companies and institutions to join the organization at the **Contributor and Adopter level.**

- **UCIe** was founded in March 2022, incorporated in June 2022. Two levels of memberships: Contributor and Adopter

- **Contributor Membership**
  - Access the Final Specifications (ex: 1.0, 1.1, 2.0, etc.)
  - Implement with the IP protections as outlined in the Agreements
  - Right to attend Corporation trade shows or other industry events as determined by the Board
  - Participate in the technical working groups
  - Influence the direction of the technology
  - Access the intermediate (dot level) specifications
  - Election to get to the Promoter Class/ Board every year when the term of half the board completes

- **Adopter Membership**
  - Access the Final Specifications (ex: 1.0, 1.1, 2.0, etc.), but not intermediate level specifications
  - Implement with the IP protections as outlined in the Agreements
  - Right to attend Corporation trade shows or other industry events as determined by the Board

On-Package Interconnects:

Opportunities and Challenges

Moore Predicted "Day of Reckoning"

*"It may prove to be more economical to build large systems out of smaller functions, which are separately packaged and interconnected."**

- Gordon E. Moore

*"Cramming more components onto integrated circuits," Electronics, Volume 38, Number 8, April 19, 1965

Drivers for On-Package Chiplets

- Reticle Limit, yield optimization, scalable performance
  - →Same dies on package (Scale-up)
- Increasing design costs at leading edge process nodes
  - →Disaggregate dies across different nodes
  - →Deploy latest process node for advanced functionality
- Time to Market (Late binding)
- Easily enables Custom silicon for different customers leveraging a common base product
  - →E.g., Different acceleration functions with common compute
- Different process nodes optimized for different functions
  - →E.g., Memory, logic, analog, co-packaged optics
- Enables high, power-efficient bandwidth with low-latency access (e.g., HBM memory)

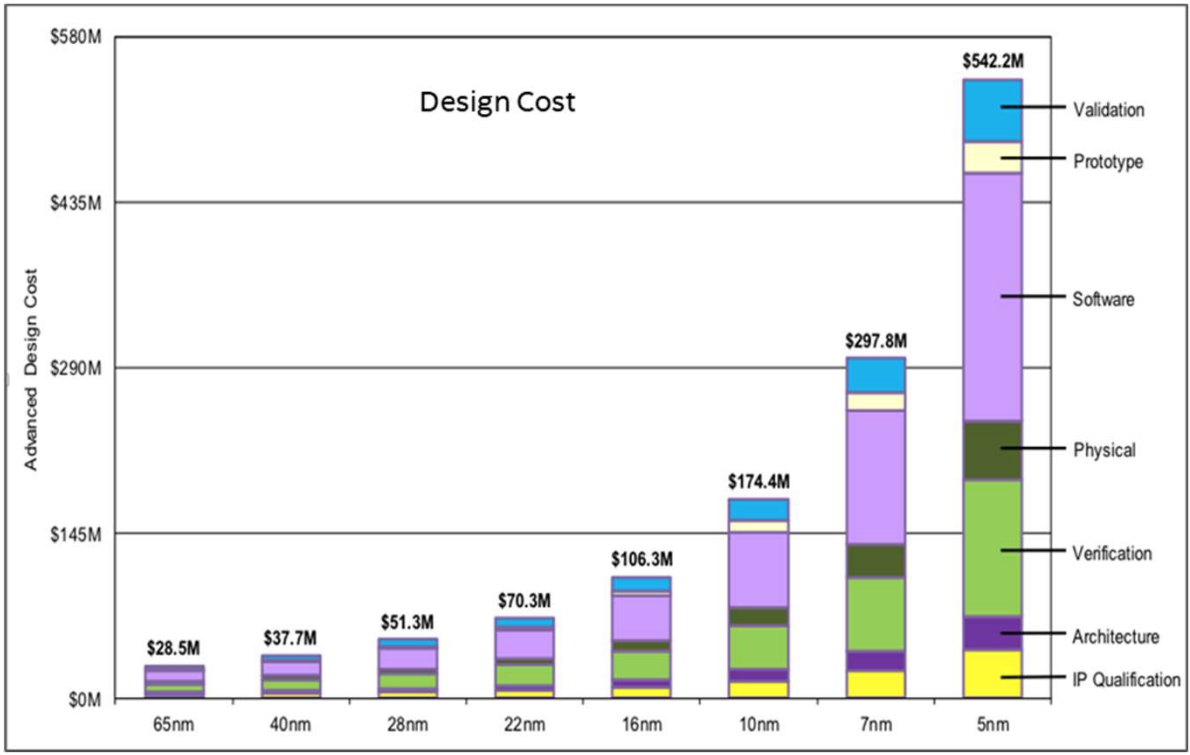

Source: IBS (as cited in IEEE Heterogeneous Integration Roadmap)

Components of Chiplet Interoperability

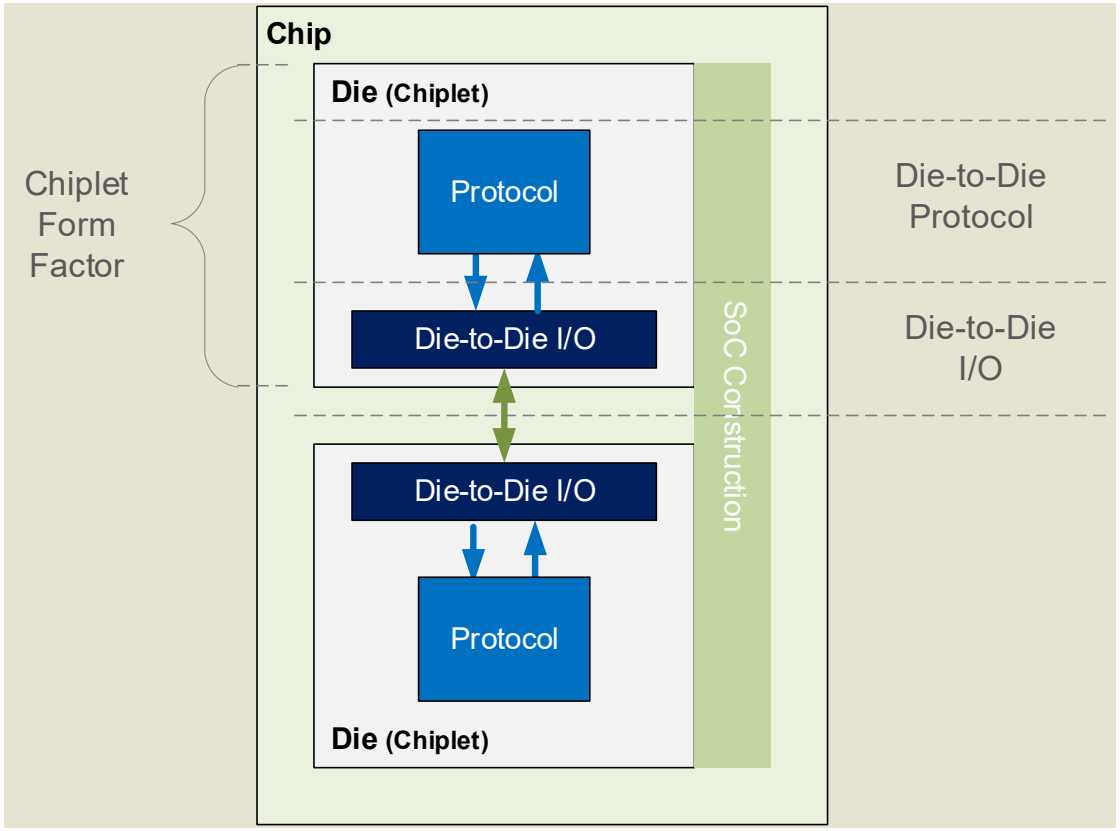

(Example SoC showing two chiplets only)

- **Chiplet Form Factor**
  - Die Size / bump location
  - Power delivery
- **SoC Construction** (Application Layer)
  - Reset and Initialization
  - Register access
  - Security
- **Die-to-Die Protocols** (Data Link to Transaction Layer)
  - PCIe/ CXL/ Streaming
  - Plug and play IPs
- **Die-to-Die I/O** (Physical Layer)
  - Electrical, bump arrangement, channel, reset, initialization, power, latency, test repair, technology transition

Design Choice: Seamless Integration from Node → Package → On-die Enables Reuse, Better User Experience

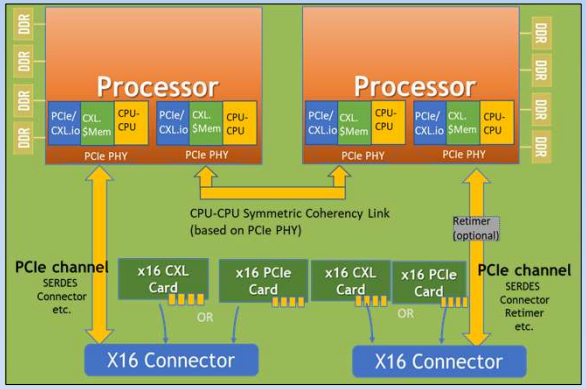

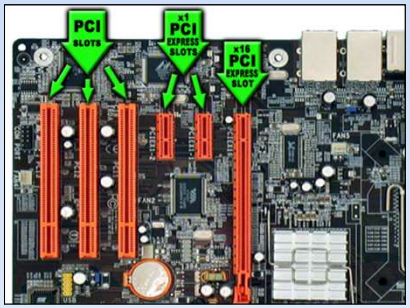

Node / Board Level Integration

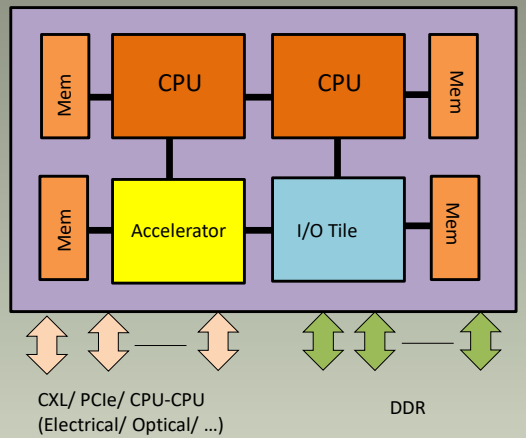

Package Level Integration (with on-package interconnects)

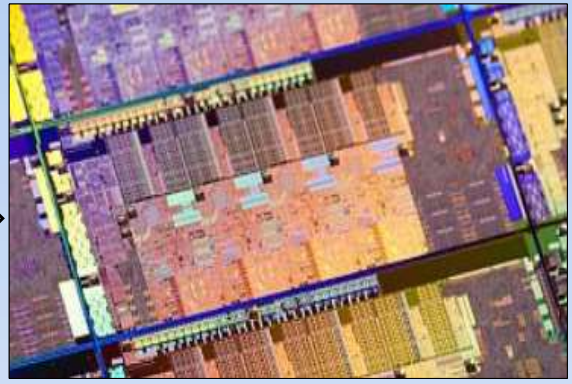

On-die Integration

Same Software, IP, and Subsystem to build scalable solutions offers economies of scale , time to market advantage, and seamless user experience. Innovations at the open slot in board level needs to migrate to package level for multiple usages!

Universal Chiplet Interconnect Express (UCIe):

An Open Standard for Chiplets

<u>Guiding principles of UCIe</u>

1. Open Ecosystem with Plug-and-play
2. Backward compatible evolution when appropriate to ensure investment protection
3. Best power, performance, and cost metrics across the industry applicable across the entire compute continuum
4. Continuously innovate to meet the needs of evolving compute landscape

(Leveraging decades of experience driving successful industry standards at the board level: PCIe, CXL, USB, etc.)

Motivation

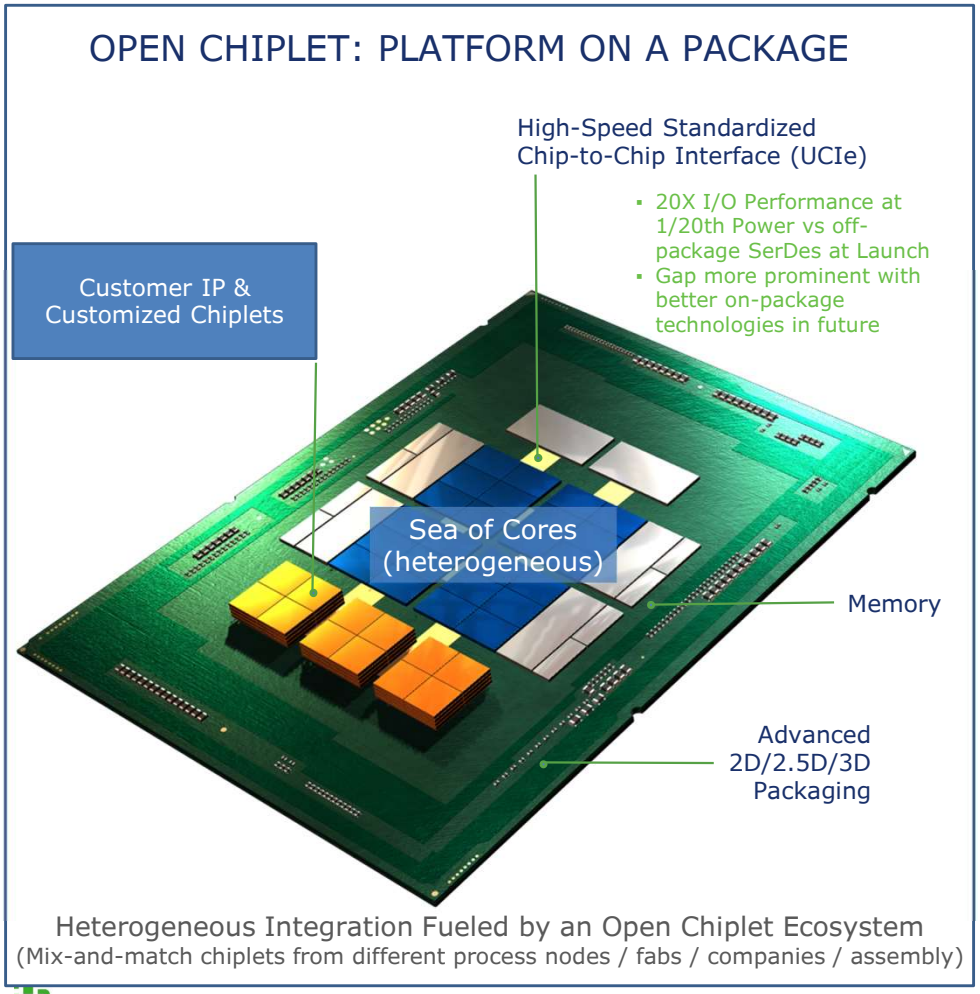

Align Industry around an open platform to enable chiplet based solutions

- Enables construction of SoCs that exceed maximum reticle size
  - Package becomes new System-on-a-Chip (SoC) with same dies (Scale Up)
- Reduces time-to-solution (e.g., enables die reuse)
- Lowers portfolio cost (product & project)
  - Enables optimal process technologies
  - Smaller (better yield)
  - Reduces IP porting costs
  - Lowers product SKU cost
- Enables a customizable, standard-based product for specific use cases (bespoke solutions)
- Scales innovation (manufacturing/ process locked IPs)

Key Metrics and Adoption Criteria

Key Technology Metrics

- Bandwidth density (linear & area)
  - Data Rate & Bump Pitch
- Energy Efficiency (pJ/b)
  - Scalable energy consumption
  - Low idle power (entry/exit time)
- Latency (end-to-end: Tx+Rx)
- Channel Reach
- Technology, frequency, & BER
- Reliability & Availability
- Cost (Standard vs advanced packaging)

Factors Affecting Wide Adoption

- Interoperability
- Full-stack, plug-and-play with existing s/w is+
- Different usages/segments
- Technology
  - Across process nodes & packaging options
  - Power delivery & cooling
  - Repair strategy (failure/yield improvement)
  - Debug – controllability & observability
- Broad industry support / Open ecosystem
  - Learnings from other standards efforts

UCIe - Architected and specified from the ground-up to deliver the best KPIs while meeting wide adoption criteria to drive innovations at package level

UCIe 1.0 Specification

- **Layered Approach with industry-leading KPIs**
- **Physical Layer:** Die-to-Die I/O
- **Die to Die Adapter:** Reliable delivery
  - Support for multiple protocols: bypassed in raw mode
- **Protocol:** CXL/PCIe and Streaming
  - **CXL™/PCIe® for volume attach and plug-and-play**
    - SoC construction issues are addressed w/ CXL/PCIe
    - CXL/PCIe addresses common use cases
      - I/O attach, Memory, Accelerator
  - **Streaming for other protocols**
    - Scale-up (e.g., CPU/ GP-GPU/Switch from smaller dies)
    - Protocol can be anything (e.g., AXI/CHI/SFI/CPI/ etc)
- **Well defined specification**: interoperability and future evolution
  - Configuration register for discovery and run-time
    - control and status reporting in each layer
    - transparent to existing drivers
  - Form-factor and Management
  - Compliance for interoperability
  - Plug-and-play IPs with RDI/ FDI interface

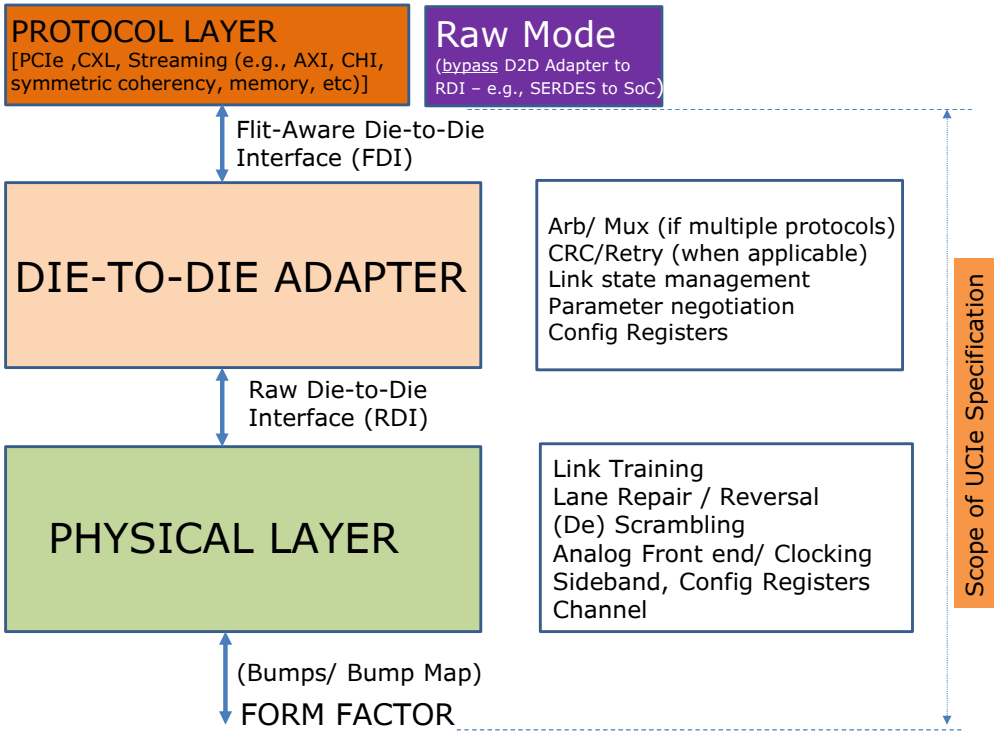

UCIe 1.0: Supports Standard and Advanced Packages

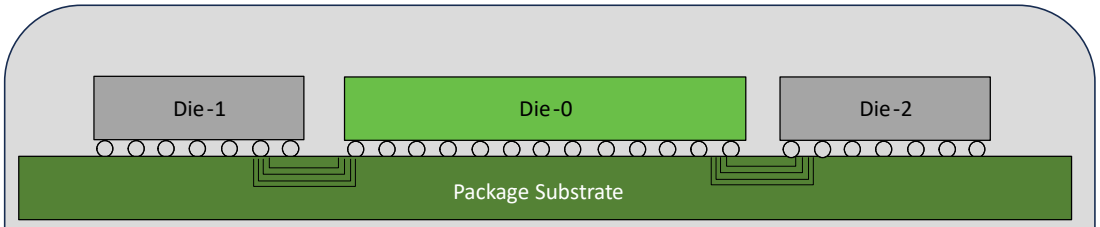

(Standard Package)

Standard Package: 2D – cost effective, longer distance

Advanced Package: 2.5D – power-efficient, high bandwidth density

Dies can be manufactured anywhere and assembled anywhere – can mix 2D and 2.5D in same package – Flexibility for SoC designer

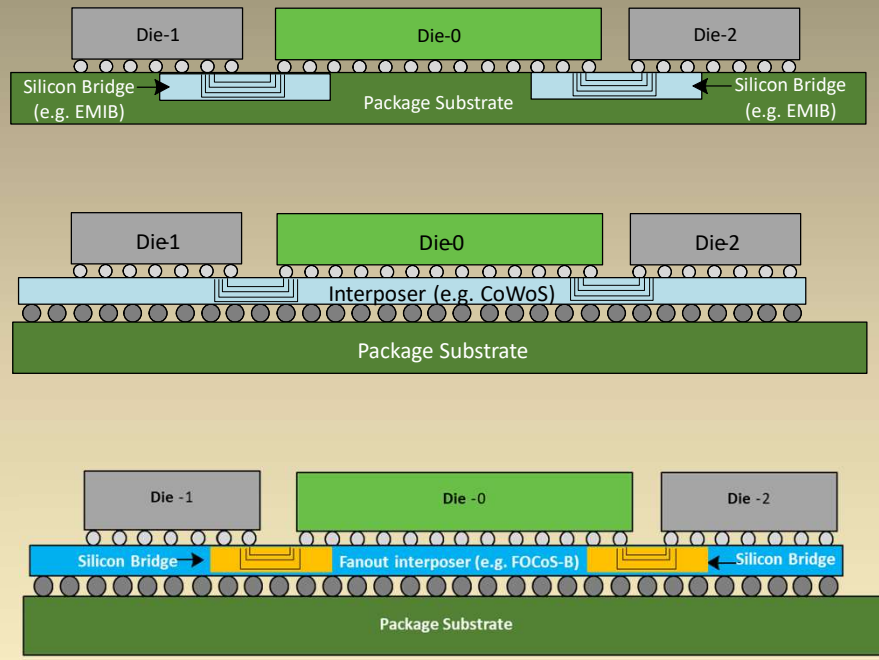

(Multiple Advanced Package Choices)

One UCIe 1.0 spec supports different flavors of packaging options to build an open ecosystem

UCIe PHY: Bump-out for Interoperability

- UCIe architected with process portability in mind
  - Circuit components can be built with common digital/ analog structures
- Bump-out specified in the specification for interoperability even with future bump-pitch reductions
  - Die rotation and mirroring supported

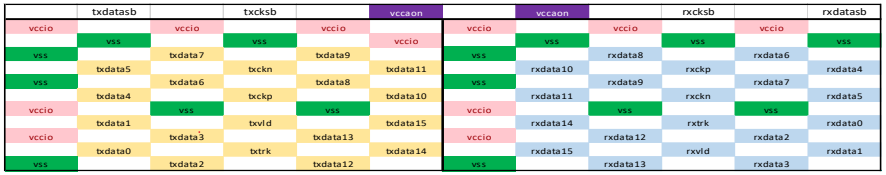

(UCIe-S Unstacked Bump-out)

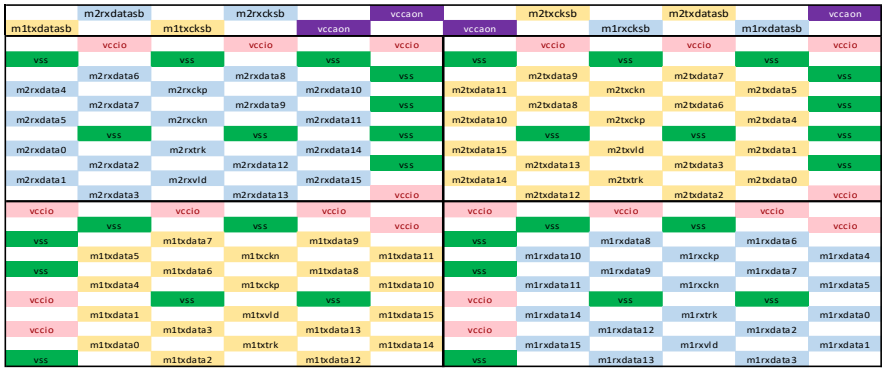

(UCIe-S Stacked Bump-out)

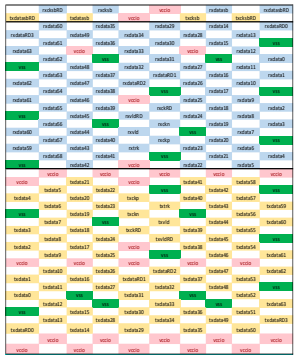

(UCIe-A Bump-out)

Fixed beachfront allows for
Multi-generational compatibility
As bump pitches decrease

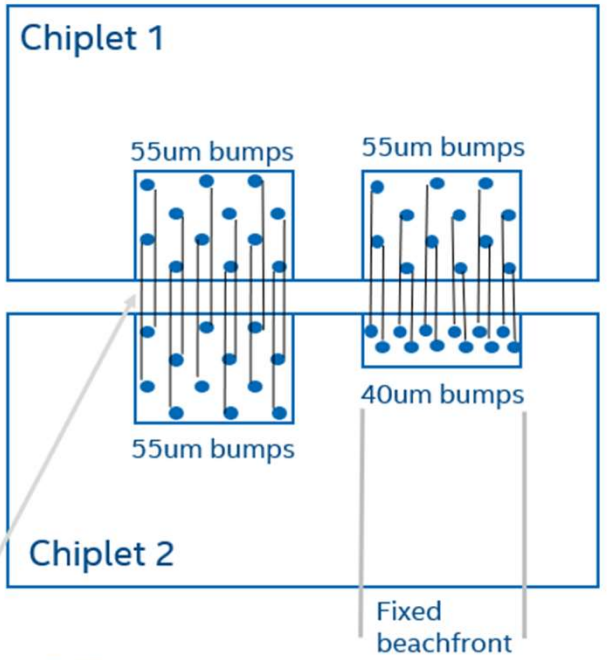

CoWos or EMIB or FoCoS or similar tight-pitch tech

Physical Layer

- Unit is One Module: uni-directional: 1, 2, or 4 modules form a Link
  - 16 (64) SE Lanes for Std (Adv)
  - 1 SE Lane of valid
  - 1 differential pair of forwarded clock
  - 1 lane (SE) calibration - Track
  - Lane reversal on Transmit side
  - Reliability: Spare Lanes in Adv; degradation in Std
  - Supported frequencies: 4, 8, 12, 16, 24, 32 GHz
  - A component must support all data rates up to its advertised maximum data rate for interoperability
  - B/W per module/ dir: 64 GB/s Std, 256 GB/s Adv: Two module gets 2X, 4-module gets 4X
- Sideband: always on; 2 Lanes/ direction @ 800 MHz – data and clock
  - Used for training, debug, management, etc; Leverages depopulated bumps to ensure no extra shore-line
- Valid used for effective dynamic power management

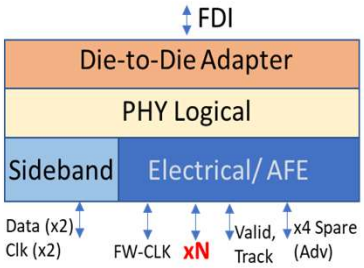

(Single Module Configuration)

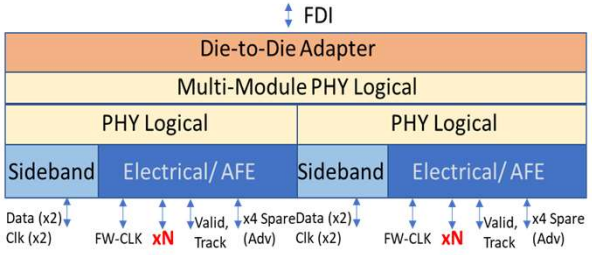

(Two Module Configuration)

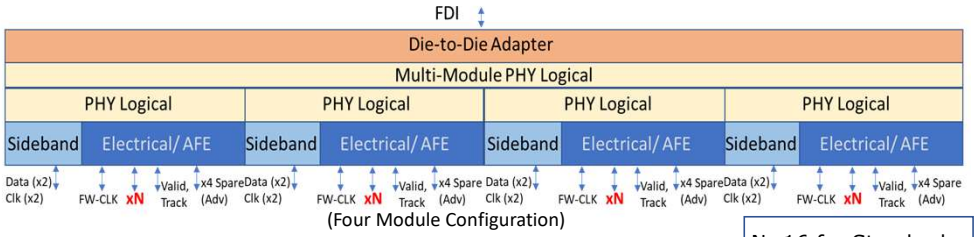

(Four Module Configuration)

(4-Module Configuration with UCIe)

N=16 for Standard
N=64 for Advanced

D2D Adapter and Flit Mapping through FDI

- Responsible for packetization
  - Adds Flit Header (2B) and CRC (2B)
- Supported Flit Sizes: 68B and two flavors of 256B
  - Decided at negotiation
- Flit Hdr (2B): Protocol ID (3b), Credit (1b), Flit Ack/Nak management (2b command + 8b sequence number), Rsvd (2b)
- CRC: Covers 128B payload (smaller payloads are 0-extended)
  - Triple bit flip detection guarantee with 16 bits
  - Replay if CRC fails
  - Sample RTL code for CRC provided in the spec

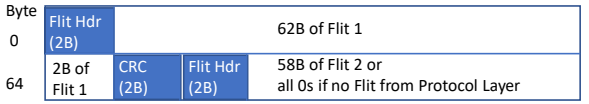

(a. 68-Byte Flit – usage CXL 2.0/ PCIe Non-Flit Mode/ Streaming)

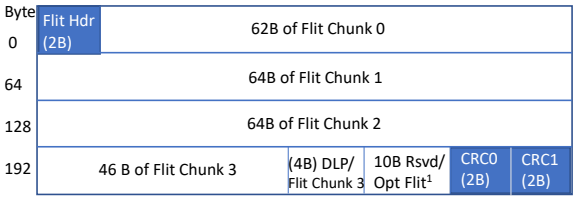

(b. 256-Byte Flit – usage CXL 3.0/ PCIe 6.0)

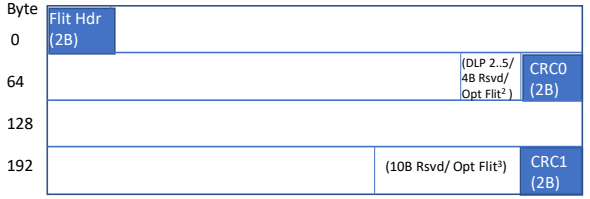

(c. 256-Byte Latency-Optimized Flit – usage CXL 3.0/ Streaming)

(Opt Flit is for better link efficiency to use the unused CRC/ FEC bytes in PCIe/ CXL)

Usage Models for UCIe: SoC at Package level

- SoC as a Package level construct
  - Standard and/ or Advanced package
  - Homogeneous and/or heterogeneous chiplets
  - Mix and match chiplets from multiple suppliers
- Across segments: Hand-held, Client, Server, Workstation, Comms, HPC, Automotive, IoT, etc
- UCIe PHY and D2D adapter common
  - PCIe/CXL protocol for plug-and-play
  - Streaming for others (similar to board level connectivity today where scale-up systems are on PCIe PHY)
  - Similar to PCIe/ CXL at board level

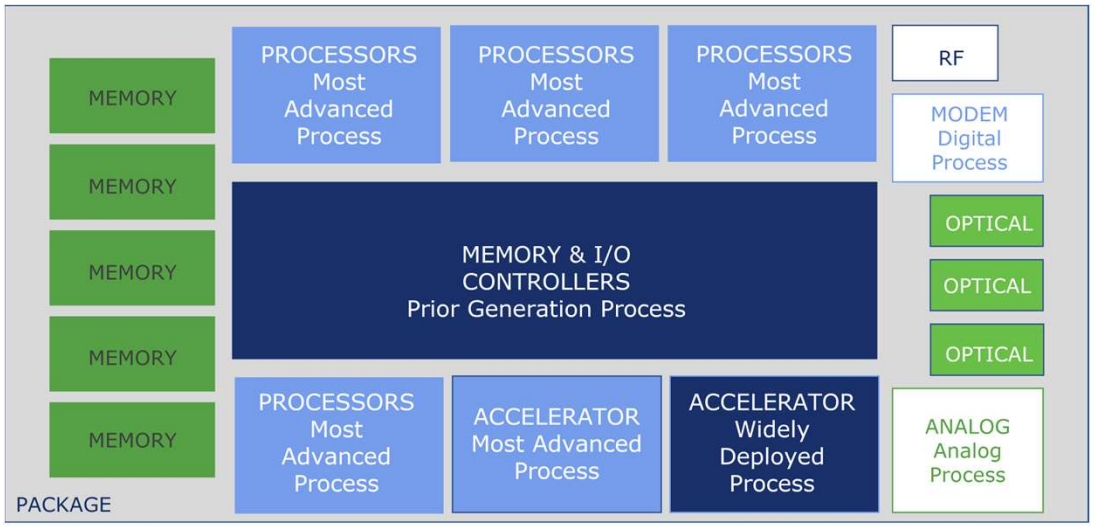

Processors: symmetric coherency protocol mapped on UCIe through FDI
Memory: CXL.Mem mapped on UCIe through FDI
Accelerators: PCIe/ CXL mapped on UCIe through FDI
Modem/ RF/ Optical: Raw mode on UCIe

Example Scale-up SoC from homogeneous dies: Large Switch with on-die protocol as streaming over UCIe

- Need large radix CXL switches – challenges: reticle limit, cost, etc.
- UCIe based Chiplets should help with scalable products
- 64G Gen6 x16b CXL links
- UCIe as d2d interconnect – while this is a scale-up CXL switch , a switch vendor may prefer to have their on-die interconnect protocol be transported over UCIe rather than create a hierarchy of switches which will not work for CXL 2.0 tree-based topology

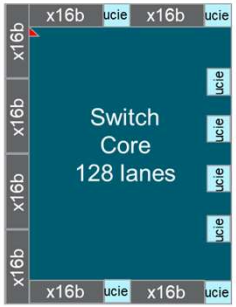

Small CXL Switch (128 lanes)

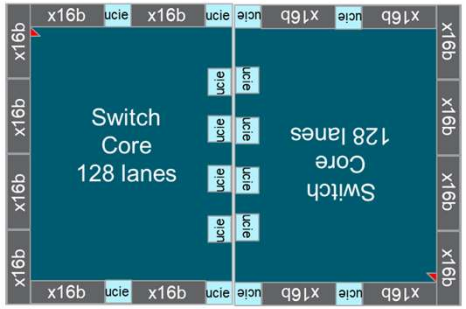

Medium-sized CXL Switch (256 lanes)

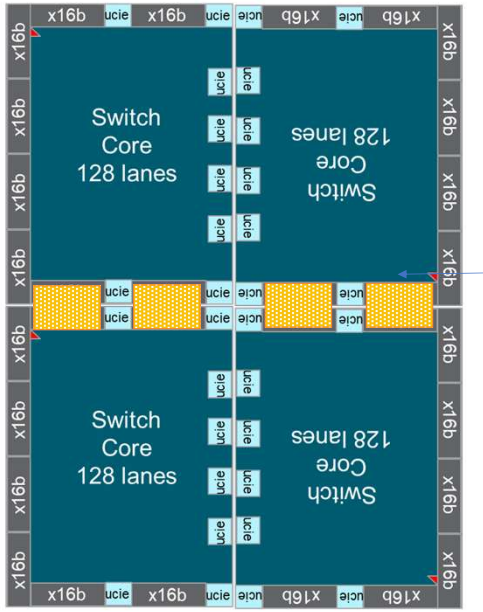

Large CXL switch (512 lanes)

One can construct CPUs (low, medium, large core-count CPUs) from smaller dies connected through UCIe using the same principle
Here the UCIe PHY and D2D adapter will carry the packetized version of internal CPU interconnect fabric

Example Scale-up Package using Streaming and open-plug-in using PCIe/ CXL

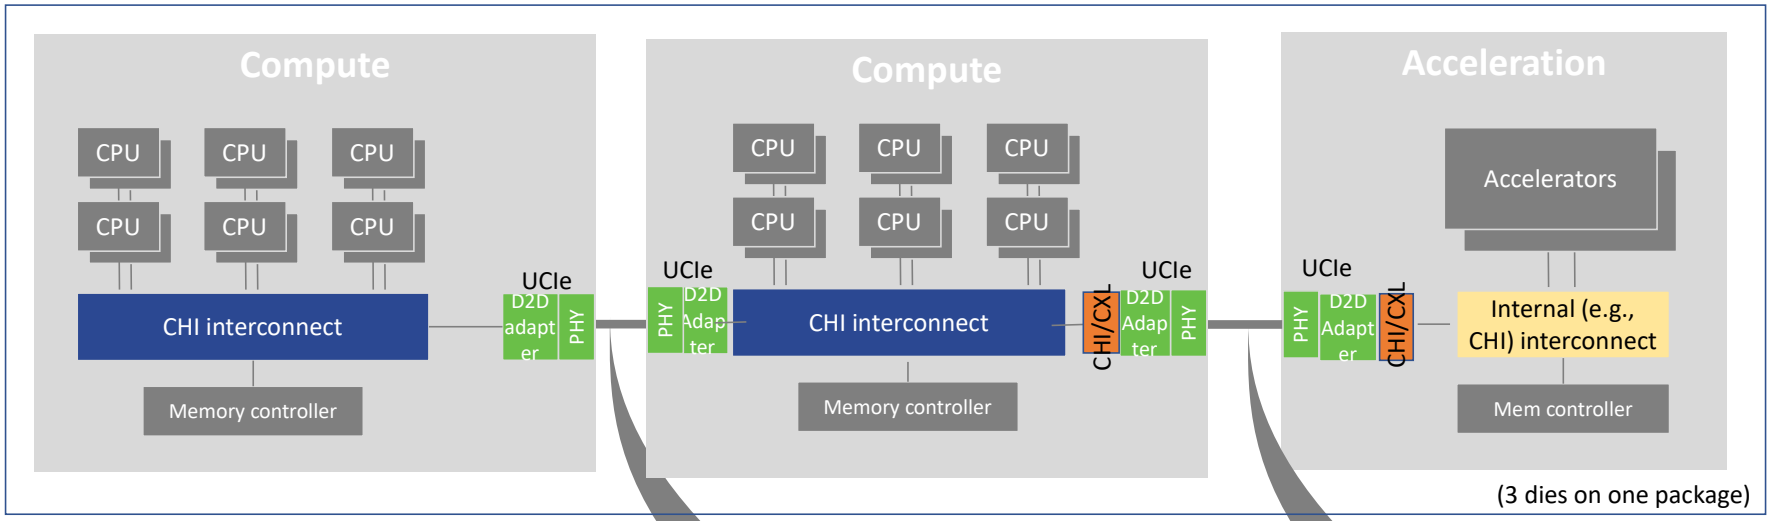

- Transporting the same on-chip protocol allows seamless use of architecture specific features without protocol conversion
- Streaming interface with additional flit formats provide link robustness using UCIe defined data-link CRC and retry

- Any device type in this open plug-in slot with CXL (or CHI if both support it)

UCIe Usage: Off-package connectivity w/ Retimers

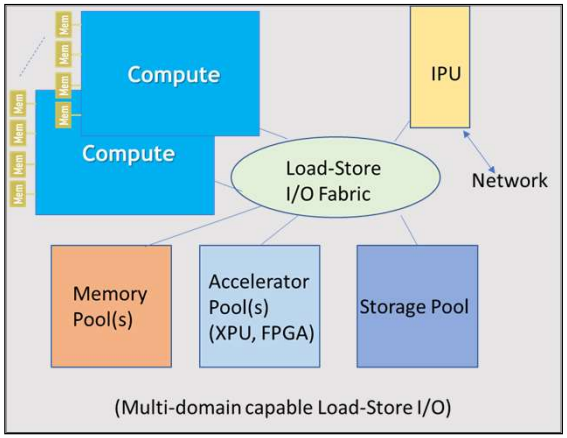

(Use Case: Load-Store I/O (CXL) as the fabric across the Pod providing low-latency and high bandwidth resource pooling/ sharing as well as message passing)

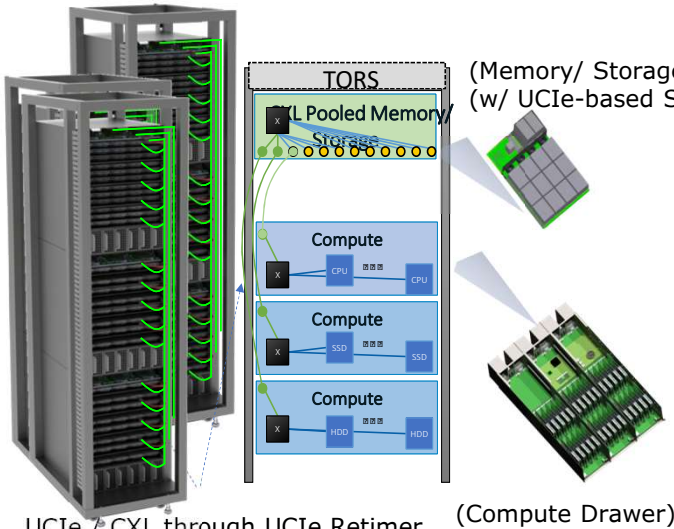

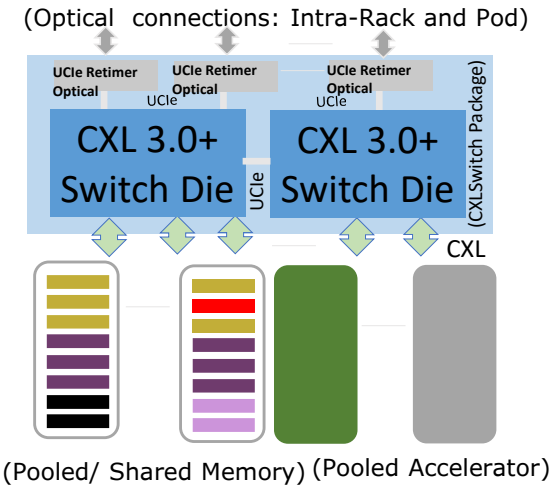

(Switch dies connected through UCIe PHY + Adapter Running a proprietary switch internal protocol)

Provision to extend off-package with UCIe Retimers connecting to other media (e.g., optics)

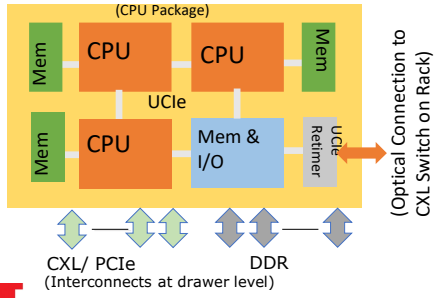

(Another example can be multi-terabit networking switches Constructed from UCIe-based co-packaged optics and partitionable networking switch dies connected through UCIe on package

UCIe 1.0: Characteristics and Key Metrics

CHARACTERISTICS	STANDARD PACKAGE	ADVANCED PACKAGE	COMMENTS
Data Rate (GT/s)	4, 8, 12, 16, 24, 32		Lower speeds must be supported -interop (e.g., 4, 8, 12 for 12G device)
Width (each cluster)	16	64	Width degradation in Standard, spare lanes in Advanced
Bump Pitch (um)	100 – 130	25 - 55	Interoperate across bump pitches in each package type across nodes
Channel Reach (mm)	<= 25	<=2	

KPIs / TARGET FOR KEY METRICS	STANDARD PACKAGE	ADVANCED PACKAGE	COMMENTS
B/W Shoreline (GB/s/mm)	28 – 224	165 – 1317	Conservatively estimated: AP: 45u; Standard: 110u; Proportionate to data rate (4G – 32G)
B/W Density (GB/s/mm^2)	22-125	188-1350	
Power Efficiency target (pJ/b)	0.5	0.25	
Low-power entry/exit latency	0.5ns <=16G, 0.5-1ns >=24G		Power savings estimated at >= 85%
Latency (Tx + Rx)	< 2ns		Includes D2D Adapter and PHY (FDI to bump and back)
Reliability (FIT)	0 < FIT (Failure In Time) << 1		FIT: #failures in a billion hours (expecting ~1E-10) w/ UCIe Flit Mode

UCIe 1.0 delivers the best KPIs while meeting the projected needs for the next 5-6 years across the compute continuum.

Ingredients for a broad inter-operable chiplet ecosystem

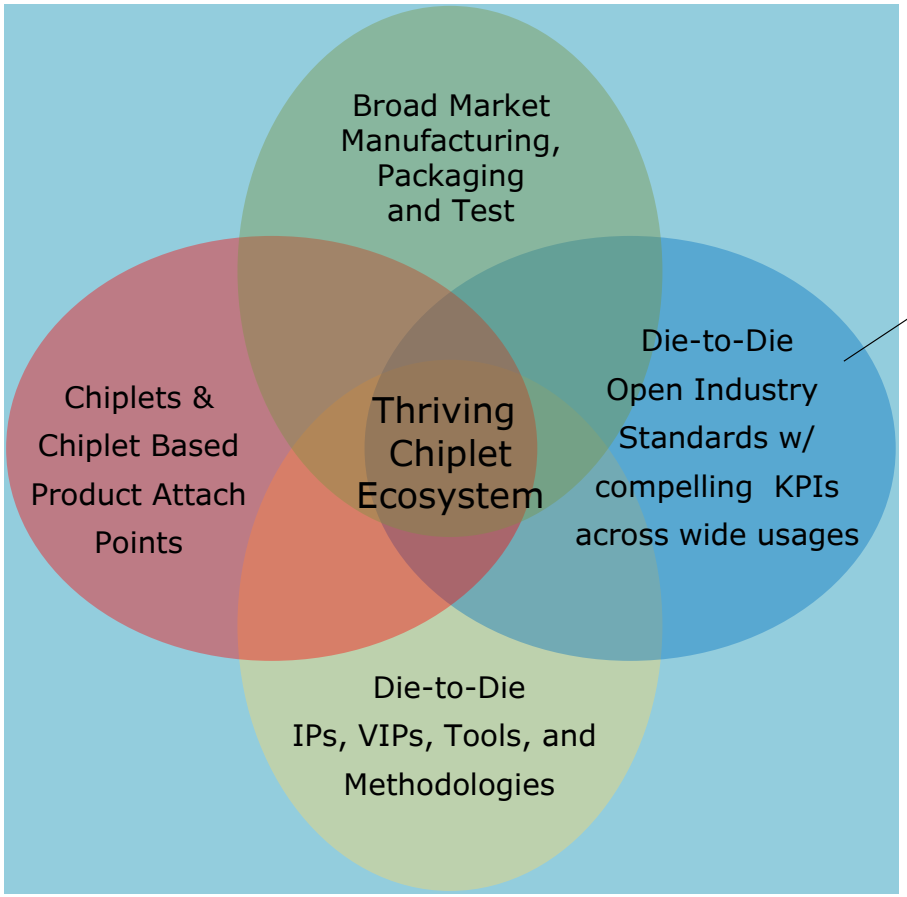

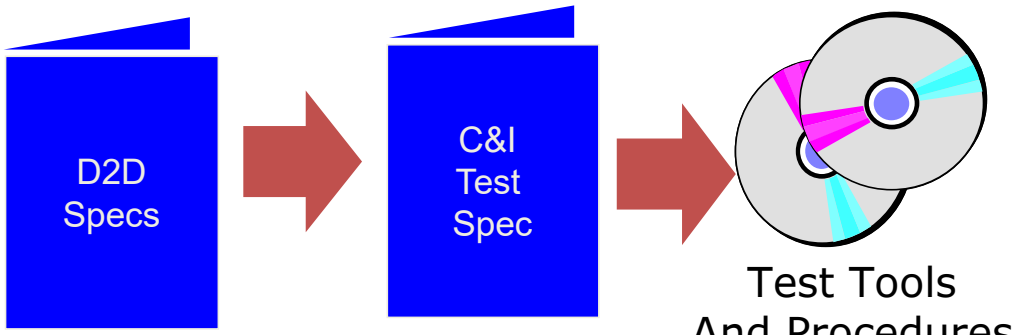

Well-defined Specs

(Electrical, Logical, Protocol (e.g., PCIe/ CXL) Software, Form-Factor, Management)

Test criteria based on Specs

(Test Definitions, Pass/Fail Criteria: Electrical, Logical, Protocol, Software)

Test Tools And Procedures

Test H/W & S/W *Validates*

Test criteria

- Compliance
- Interoperability

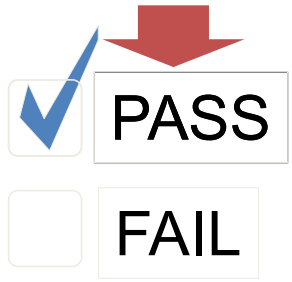

Predictable path to design compliance with UCIe