

# UCIe 3.0 Specification: Driving Innovation for Efficient, Scalable, and Reliable Chiplet Integration

Presented by:

Brian Rea, MWG Chair, UCIe Consortium

Swadesh Choudhary, Protocol WG Co-Chair, UCIe Consortium

# Agenda

- UCIE Consortium Update
- UCIE Specification Recap: UCIE 1.0 – UCIE 2.0 Highlights
- Introducing the UCIE 3.0 Specification
- Conclusions

# Universal Chiplet Interconnect Express™ (UCIe™)

## An Open Standard for Chiplet Development

### ■ UCIe Guiding Principles

- Open chiplet ecosystem
- Backward-compatible evolution to ensure investment protection
- Optimized power, performance, and cost metrics applicable across the entire compute continuum
- Continuously innovate to meet the needs of the evolving ecosystem



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

High-bandwidth, Low-latency, Power-efficient, Cost-effective Interconnects for  
AI, HPC, Cloud, Edge, Enterprise, 5G, Automotive, Handhelds



# Board Members

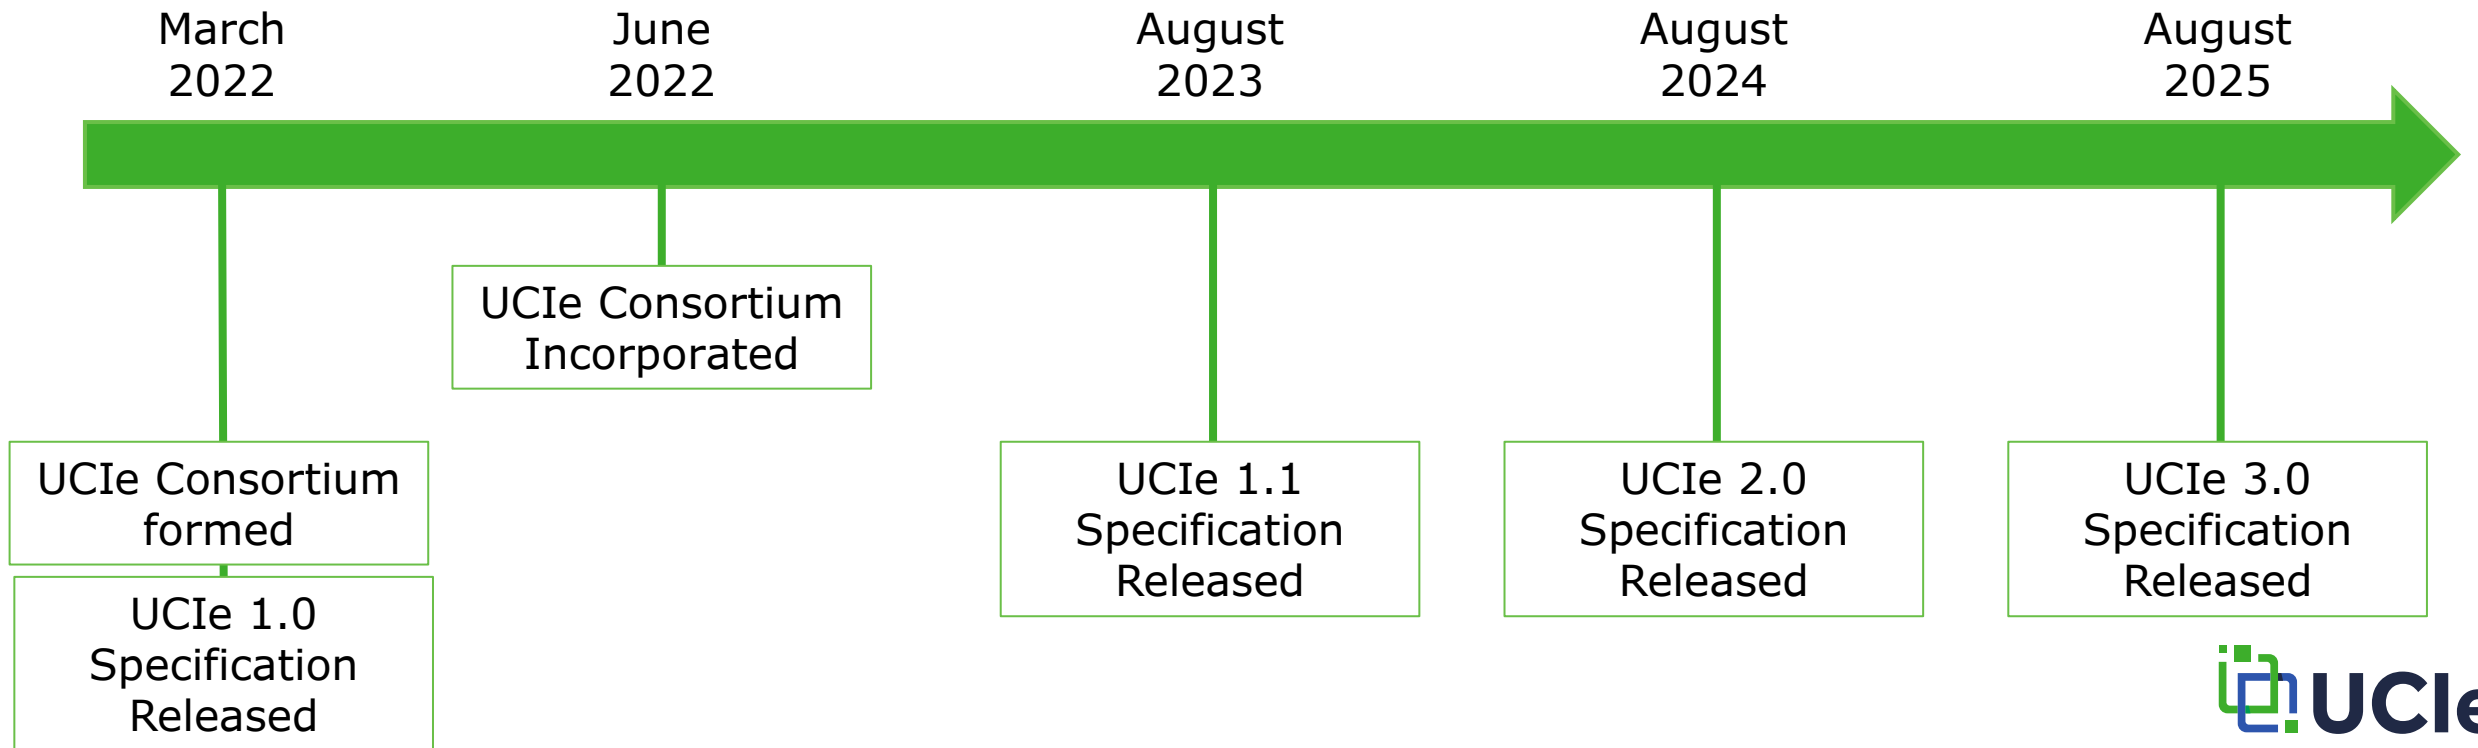
Leaders in semiconductors, packaging, IP suppliers, foundries, and cloud service providers are joining together to drive the open chiplet ecosystem.



140+ Member Companies...and growing!



# Member-Driven Evolution

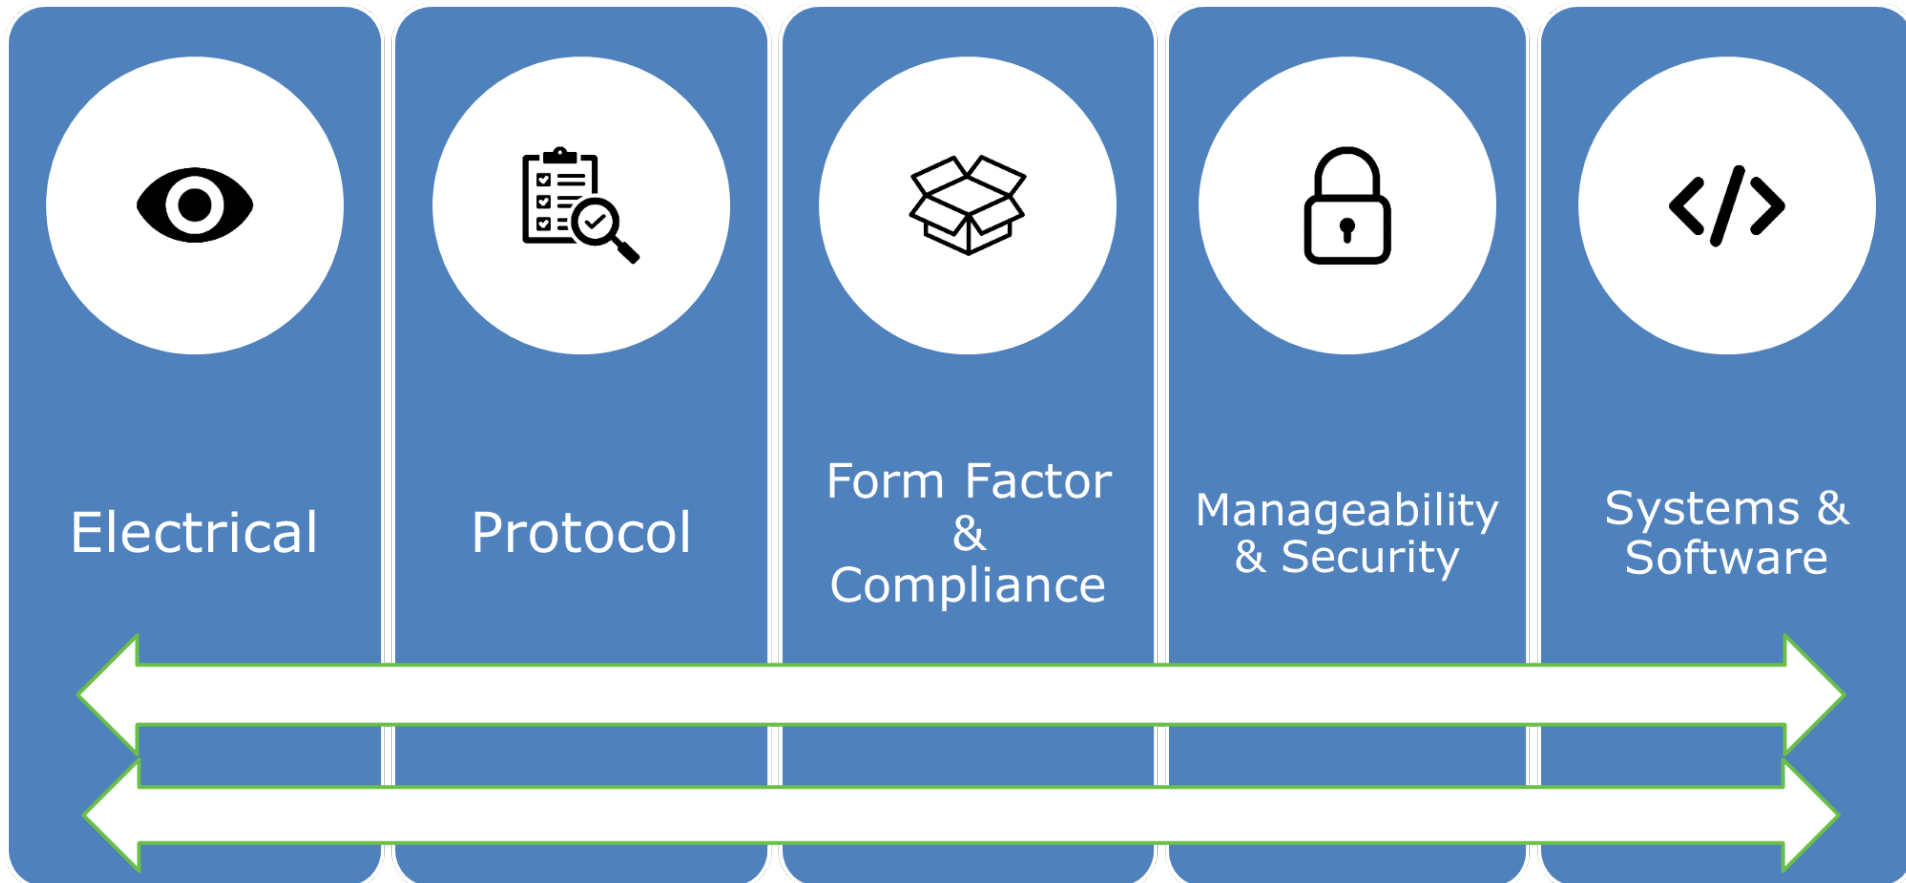


# UCIe Consortium is Open for Membership

- 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

# UCIe Consortium Working Groups

Working Groups are identifying and addressing the demands of a complete, full-stack solution for strengthening the open standards-based ecosystem.




# UCIe Specification Recap

## UCIe 1.0 – UCIe 2.0 Highlights



# Motivation

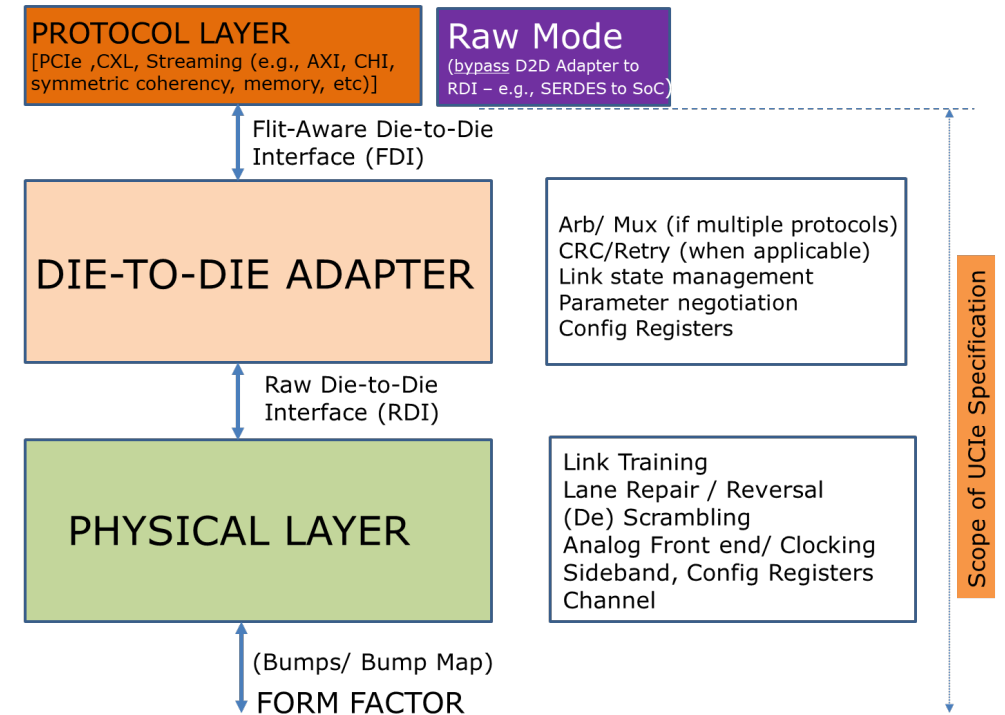


## Aligning the industry around an open platform to enable chiplet based solutions

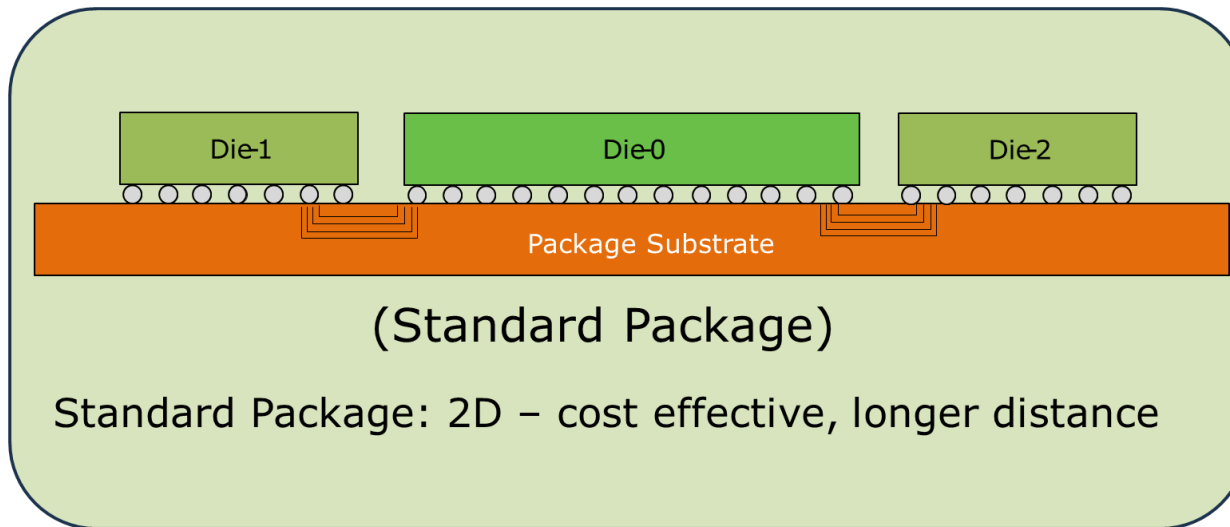
- Enables construction of SoCs that exceed maximum reticle size
  - Package becomes new System-on-a-Chip (SoC) with same dies (Scale Up)
- Reduces time-to-solution (e.g., enables die reuse)
- Lowers portfolio cost (product & project)
  - Enables optimal process technologies
  - Smaller (better yield)
  - 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)

# UCIe 1.0: Building an Open Chiplet Ecosystem

- 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)
    - Raw Mode only
- 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



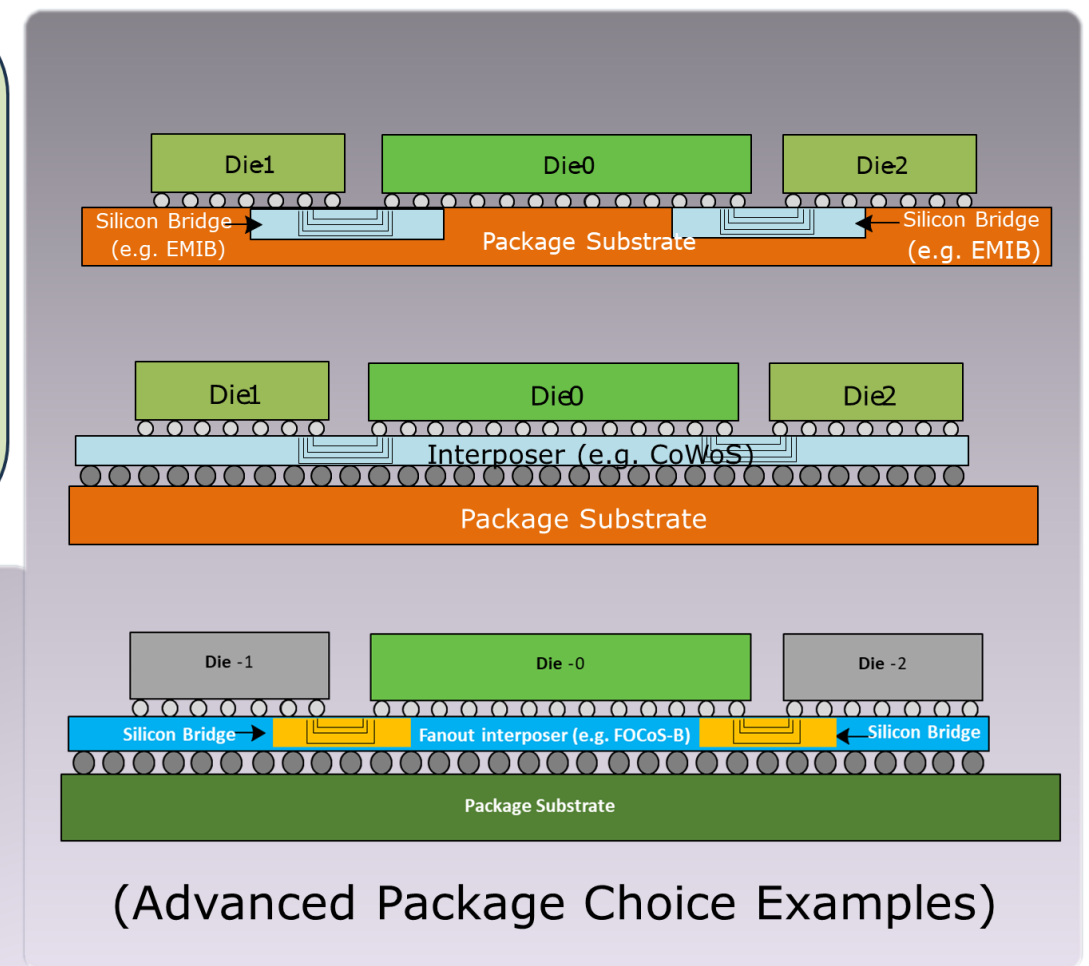
# UCIe 1.0: Support for Standard and Advanced Packages



Standard Package: 2D – cost effective, longer distance

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

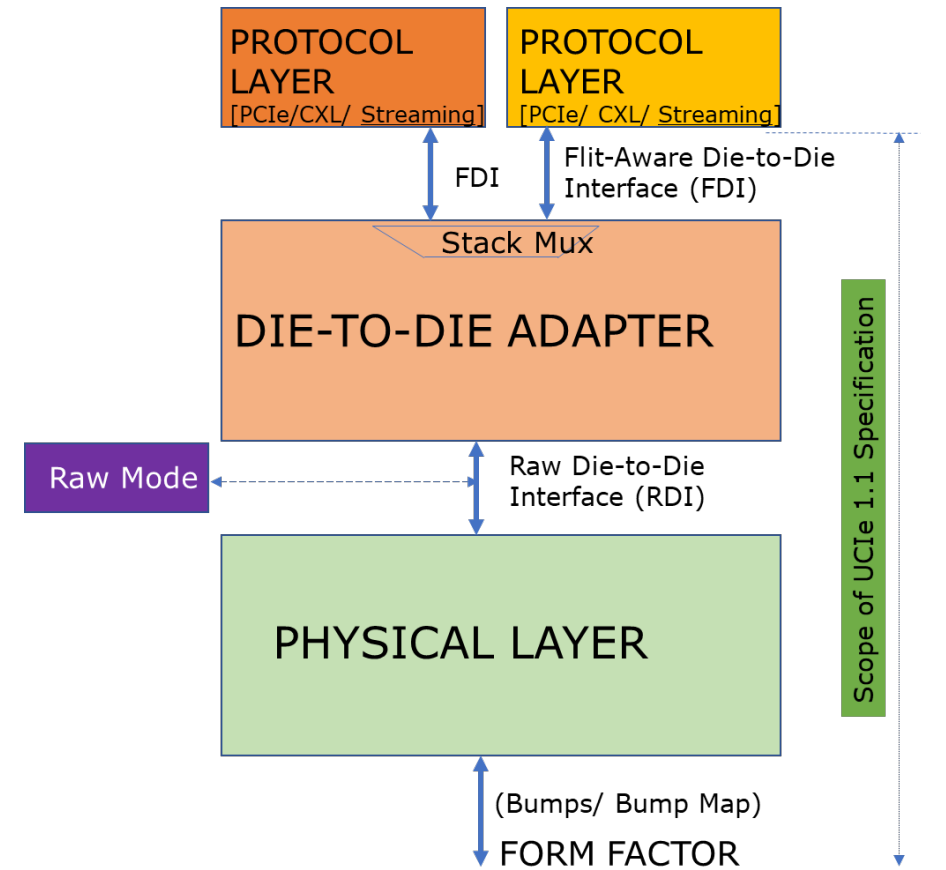
Dies can be manufactured anywhere and assembled anywhere – can mix 2D and 2.5D in same package – Flexibility for SoC designer



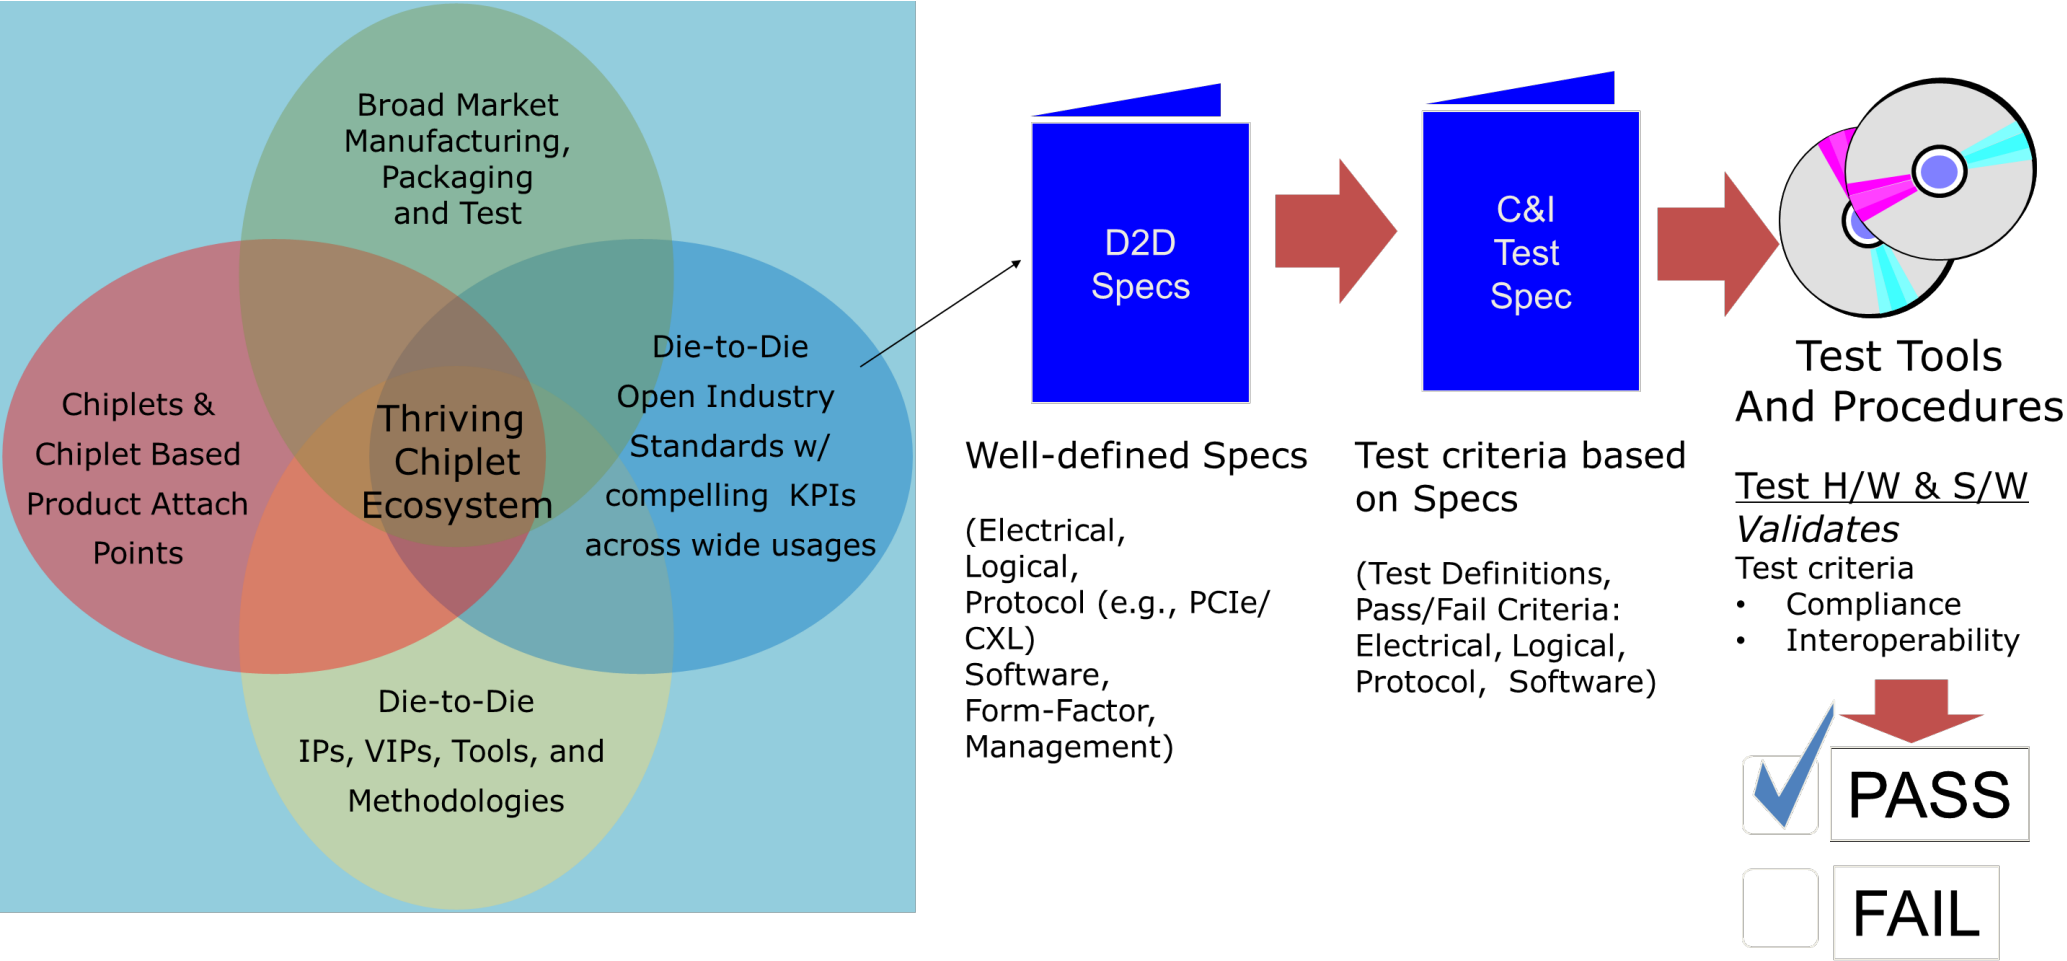
One UCIe 1.0 spec supports **different flavors** of packaging options to build an open ecosystem

# UCIe 1.1: Enhancements for Automotive and Compliance Testing

- Enhancements for Automotive Segment Usage
  - Preventive monitoring
  - Run-time testability of link health
  - Field repairability
- New Usages: Streaming Protocols with Full Stack
  - Enables D2D adapter for streaming protocols
  - Streaming protocols can multiplex with other protocols with on-demand interleaving
- Cost Optimization for Advanced Packaging
- Enhancements for Compliance Testing




# Ingredients for Broad Interoperable Chiplet Ecosystem



Predictable path to design compliance with UCle

# UCIe 2.0: Vertical Chiplets with UCIe-3D

- 3D deployed in commercial offerings (Memory, CPU)
  - Hybrid bonding (HB) looks promising
  - Standardize for constrained interop (e.g., bump pitch match)
- High bandwidth density
  - 3D → areal connectivity (vs shore-line in 2D/ 2.x D)
  - Bump pitches aggressively shrinking
    - Number of wires increases inversely as the square of bump pitch
  - Must ensure we continue to be bump-limited
- Low power
  - Reduced interconnect distance ( $\sim 0$ ) between dies, electrical parsitics
  - Simple circuits and lower frequency are essential
- Better power, bandwidth, and latency than UCIe 2.5D



3D can deliver power-efficient performance comparable/better than large monolithic die

# Introducing the UCIE 3.0 Specification

# UCIe 3.0 Specification Feature Overview

- **Higher bandwidth density:** 48 GT/s and 64 GT/s for UCIe-S and UCIe-A
  - Doubling the data rate to power next-gen multi-chip systems such as AI and HPC while maintaining low power.
- **New Usages:** Added support for continuous transmission protocols
  - Enables uninterrupted data flow in Raw Mode for new applications such as connectivity between SoC and DSP chiplets.
- **Power Savings:** Runtime recalibration and L2 Optimization
  - Enable power-efficient link tuning during operation by reusing initialization states.
  - Reduces Idle Power on the sideband.
- **Manageability Infrastructure Enhancements:**
  - Enhancements for Early Firmware Download, Sideband Priority Packets, Extending Sideband Reach, Open Drain Pin, and Fast Throttle/Shutdown.





# Doubling Data Rates for UCIe-A and UCIe-S

- **Motivation**: Continued demand for higher linear bandwidth density for SoCs used in applications such as AI, HPC, etc., with shore-line constraints
- **Solution**: Increase the data rate from maximum 32 GT/s to 48 and 64 GT/s
- **UCIe's Approach**:
  - Full backwards compatibility – same sideband, valid, track, data, training, etc.
  - Signaling: NRZ Uni-directional
  - Clocking: Quarter rate for 48/64 GT/s; free running
  - BER:  $10^{-15}$  for 48 GT/s and  $10^{-12}$  for 64 GT/s
  - Termination: RX Termination required for both UCIe-S and UCIe-A
  - Enhanced Equalization: 3-tap TX FFE (1-pre + 1-post); 1<sup>st</sup> order (passive) RX CTLE : can possibly be combined with T-coil network; Optional 1-tap RX DFE
  - B/W Density target: 1.7-2x linear, 1.3-1.6x areal.
  - Power Target: 0.5-0.75pJ/b
    - Breakdown: ~ 40% TX, 40% RX, 20% common circuits
- **Result**: Linear B/W Density increases 1.65x/2x for UCIe-S/UCIe-A with similar power efficiency

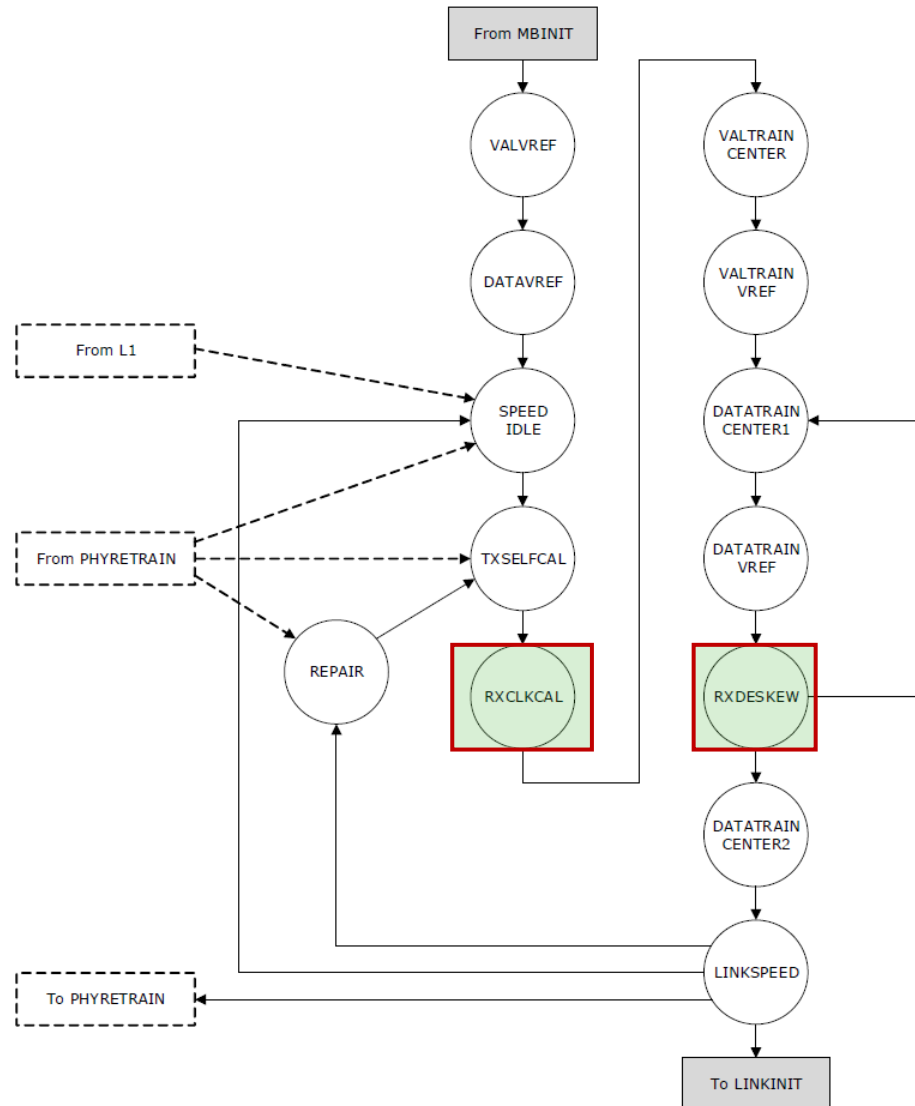
# Clocking

- **Quarter rate** and **free running clock** only for 48 and 64 GT/s
- Valid Framing and Fast Idle Entry/Exit through Valid Gating remain the same

## Forward Clock Frequency and Phase:

| Data Rate<br>(GT/s) | Clock freq. (fCK)<br>(GHz) | Phase-1 | Phase-2 | Deskew<br>(Req/Opt) |
|---------------------|----------------------------|---------|---------|---------------------|
| 64                  | 16                         | 45      | 135     | Required            |
| 48                  | 12                         | 45      | 135     | Required            |
| 32                  | 16                         | 90      | 270     | Required            |
|                     | 8                          | 45      | 135     | Required            |
| 24                  | 12                         | 90      | 270     | Required            |
|                     | 6                          | 45      | 135     | Required            |
| 16                  | 8                          | 90      | 270     | Required            |
| 12                  | 6                          | 90      | 270     | Required            |
| 8                   | 4                          | 90      | 270     | Optional            |
| 4                   | 2                          | 90      | 270     | Optional            |

# Training and Equalization at 48/64 GT/s



- I/Q training done in RXCLKCAL phase
- EQ adjustments done in RXDESKEW
  - TX Preset selection (of 6) is accomplished in RXDESKEW phase
  - Pick best Preset based on RX Eye Margin
  - Can go back to DATATRAIN CENTER1 if more training time is needed

Preset Table:

|           | C(-1) | C(0) | C(+1) | Accuracy  |
|-----------|-------|------|-------|-----------|
| <b>P0</b> | 0     | 1    | 0     |           |
| <b>P1</b> | -0.05 | 0.95 | 0     | +/- 0.025 |
| <b>P2</b> | 0     | 0.9  | -0.1  | +/- 0.025 |
| <b>P3</b> | -0.05 | 0.85 | -0.1  | +/- 0.025 |
| <b>P4</b> | 0     | 0.8  | -0.2  | +/- 0.025 |
| <b>P5</b> | -0.05 | 0.75 | -0.2  | +/- 0.025 |

# Overall KPI after adding 48GT/s and 64GT/s support

Table 1-4. UCle 2D and 2.5D Key Performance Targets

| Metric  | Link Speed/<br>Voltage    | Advanced Package<br>(x64)           | Standard Package  |
|---|---------------------------|-------------------------------------|-------------------|
| Die Edge Bandwidth<br>Density <sup>a</sup><br>(GB/s per mm) | 4 GT/s                    | 165                                 | 28                |
|   | 8 GT/s                    | 329                                 | 56                |
|   | 12 GT/s                   | 494                                 | 84                |
|   | 16 GT/s                   | 658                                 | 112               |
|   | 24 GT/s                   | 988                                 | 168               |
|   | 32 GT/s                   | 1317                                | 224               |
|   | 48 GT/s                   | 1975                                | 278 <sup>b</sup>  |
|   | 64 GT/s                   | 2634                                | 370 <sup>b</sup>  |
| Energy Efficiency <sup>c</sup><br>(pJ/bit)                  | 0.7 V<br>(Supply Voltage) | 0.5 (<= 12 GT/s)                    | 0.5 (4 GT/s)      |
|   |                           | 0.6 (>= 16 GT/s)                    | 1.0 (<= 16 GT/s)  |
|   |                           | -                                   | 1.25 (>= 24 GT/s) |
|   | 0.5 V<br>(Supply Voltage) | 0.25 (<= 12 GT/s)                   | 0.5 (<= 16 GT/s)  |
|   |                           | 0.3 (>= 16 GT/s<br>and <= 32 GT/s)) | 0.75 (>= 32 GT/s) |
|   |                           | 0.5 (>= 48 GT/s)                    |                   |
| Latency Target <sup>d</sup>                                 |                           | <=2ns                               |                   |

- a. Die edge bandwidth density is defined as total I/O bandwidth in GB per sec per mm silicon die edge, with 45-um (Advanced Package) and 110-um (Standard Package) bump pitch. For a x32 Advanced Package module, the Die Edge Bandwidth Density is 50% of the corresponding value for x64.
- b. Die edge bandwidth density for Standard Package at 48 GT/s and 64 GT/s is less than 2x of that at 24 GT/s and 32 GT/s, respectively. This is because of increased die edge to improve signal integrity at the higher data rates. Future revisions of the specification will look at addressing this.
- c. Energy Efficiency (energy consumed per bit to traverse from FDI to bump and back to FDI) includes all the Adapter and Physical Layer-related circuitry including, but not limited to, Tx, Rx, PLL, Clock Distribution, etc. Channel reach and termination are discussed in [Chapter 5.0](#).
- d. Latency includes the latency of the Adapter and the Physical Layer (FDI to bump delay) on Tx and Rx. See [Chapter 5.0](#) for details of Physical Layer latency. Latency target is based on 16 GT/s. Latency at other data rates may differ due to data rate-dependent aspects such as data accumulation and transfer time. Note that the latency target does not include the accumulation of bits required for processing; either within or across Flits.

# Key Metrics with UCIe 3.0

| Characteristics / KPIs             | UCIe-S (2D)                    | UCIe-A (2.5D)                                 | UCIe 3D                                    | Comments  |
|------------------------------------|--------------------------------|---|--|---|
| <b>Characteristics</b>             |                                |   |  |   |
| Data Rate (GT/s)                   | 4, 8, 12, 16, 24, 32, 48, 64   |   | Up to 4                                    | UCIe 3D SoC Logic frequency – power efficiency is critical<br>Added 48G and 64G with UCIe 3.0 |
| Width (each cluster)               | 16                             | 64  | 80   | UCIe 3D: Options or reduced width to 70, 60...  |
| Bump Pitch (μm)                    | 100 – 130                      | 25 – 55                                       | ≤ 10 (optimized)<br>> 10 – 25 (functional) | Must scale so that UCIe fits within the bump area, UCIe-3D must support hybrid bonding        |
| Channel Reach (mm)                 | ≤ 25                           | ≤ 2   | 3D vertical                                | UCIe-3D: FtF, FtB, BtB, multi-stack possible  |
| <b>Target for Key Metrics</b>      |                                |   |  |   |
| BW Shoreline (GB/s/mm)             | 28 – 224<br>278, 370           | 165 – 1317<br>1975, 2634                      | N/A (vertical)                             | For UCIe-S and UCIe-A: First row is for 4-32G. Second Row is for 48G and 64G respectively     |
| BW Density (GB/s/mm <sup>2</sup> ) | 22 – 125<br>144,192            | 188 – 1350<br>1235, 1646                      | 4,000 (9μm) –<br>300,000 (1μm)             | For UCIe-S and UCIe-A: First row is for 4-32G. Second Row is for 48G and 64G respectively     |
| Power Efficiency Target (pJ/b)     | 0.5 (≤16 G)<br>0.75 (≥ 32 G)   | 0.25 (≤12G)<br>0.3 (16G – 32G)<br>0.5 (≥ 48G) | <0.05 at 9μm -><br>0.01 at 1 μm            |   |
| Low-Power Entry/Exit               | 0.5nS ≤ 16G, 0.5-1nS ≥ 24G     |   | 0nS  | No preamble or post-amble   |
| Reliability (FIT)                  | 0 < FIT (Failure in Time) << 1 |   | 0 < FIT << 1                               |   |
| ESD                                | 30V CDM                        |   | 5V CDM → ≤3V                               | UCIe-3D: 5V CDM at introduction, no ESD for W2W hybrid bonding possible                       |

UCIe continues to deliver compelling power-efficient and cost-effective performance

# Continuous Transmission Protocols

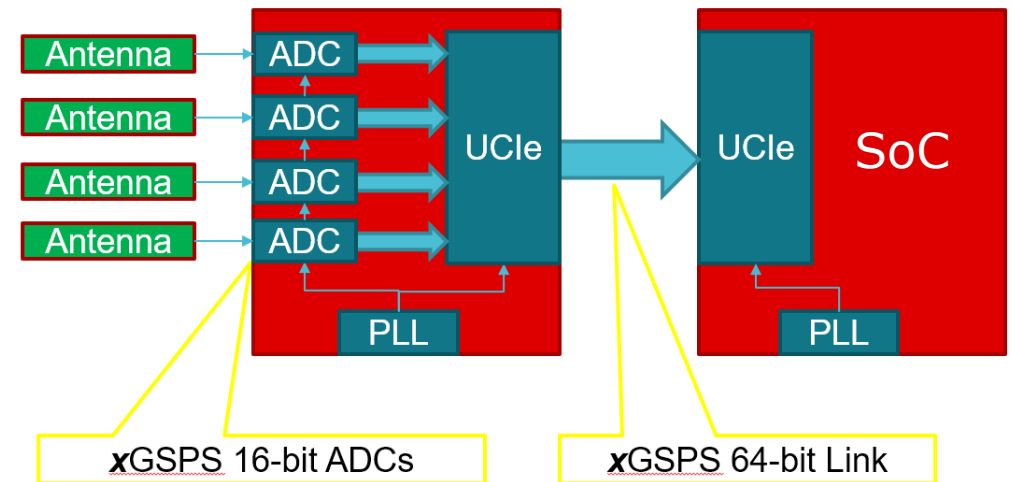
- **Function:** High-speed data transmission protocols between data converters (ADCs and DACs) can be mapped to UCIe Raw Mode of operation

- Leading DSP companies want to use UCIe standard
- Run link at same data rate as data generation/ consumption
- No need for separate PLLs
- Avoids introducing additional frequency noise in sensitive analog circuits
- Need periodic synchronization markers, parity

- **UCIe's Approach:**

- Use existing raw mode with enhancements to the internal RDI / FDI interface
- Reuse the UCIe Retimer encodings in Valid to send periodic synchronization markers and parity (full use of all data lanes which is desired)
- Support range of frequencies
  - System designer controls the data rate by varying the REFCLK
  - Compliance only performed for the UCIe data rates

- **Benefit:** Addresses new market segments such as the wireless infrastructure, software-defined radio, rada systems, and more



# Continuous mode transmission

- System designer controls the data rate by varying the REFCLK provided to the PLLs in the UCIE IP (see table for range)
  - IP will work since the change is with Refclk and it is within the interoperability range
- Compliance only performed for the UCIE data rates supported by the UCIE IP (i.e. Link Speed Setting in the table below)

| Link Speed Setting | Min Adjusted Operating Speed | Max Adjusted Operating Speed |
|--------------------|------------------------------|------------------------------|
| 4 GT/s             | 2 GT/s                       | 4 GT/s                       |
| 8 GT/s             | 4 GT/s                       | 8 GT/s                       |
| 12 GT/s            | 8 GT/s                       | 12 GT/s                      |
| 16 GT/s            | 12 GT/s                      | 16 GT/s                      |
| 24 GT/s            | 16 GT/s                      | 24 GT/s                      |
| 32 GT/s            | 24 GT/s                      | 32 GT/s                      |
| 48 GT/s            | 32 GT/s                      | 48 GT/s                      |
| 64 GT/s            | 48 GT/s                      | 64 GT/s                      |

## Enhanced Runtime Recalibration

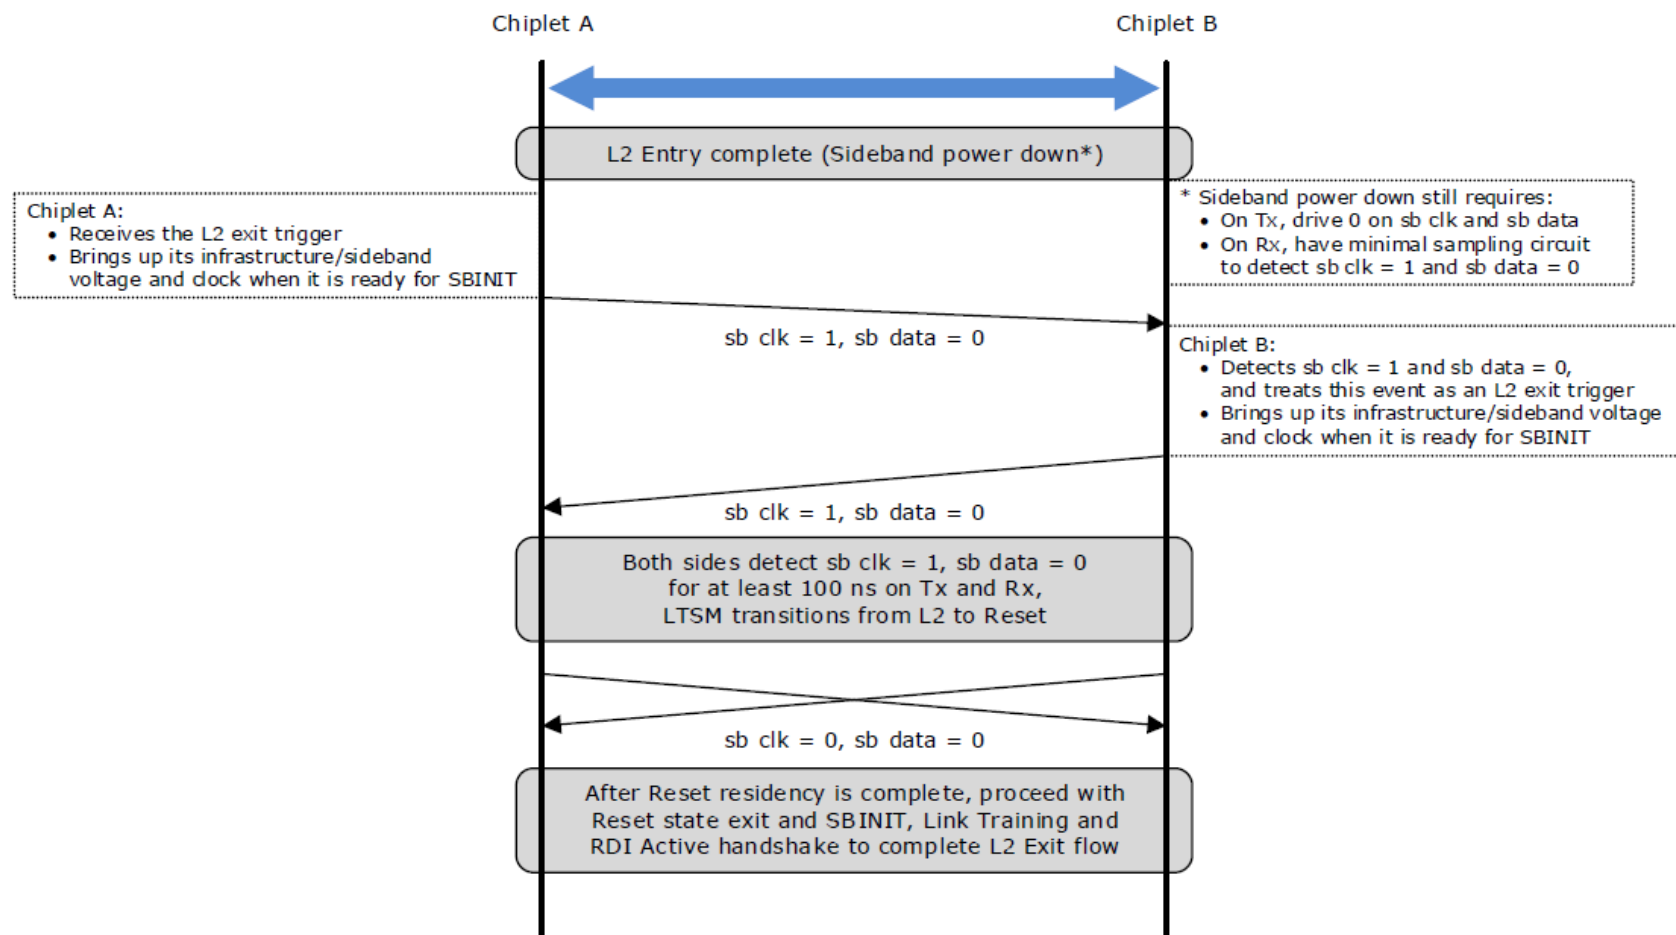
- **Function:** Allows Transmitter (Tx) adjustment of clock during runtime recalibration of the Link
- **Benefit:** Decreases the impact on Receivers (Rx) and enables power saving features for the UCIe Physical Layer by giving Tx a wider adjustment range during Link Initialization flows that can repurposed during runtime recalibration.



## L2 Exit Handshake

- **Motivation:** Deeper power savings by turning off power and clock to sideband infrastructure in L2, the deep power saving state
  - The Main band is already off
- **Solution:** A mechanism to wake up the sideband infrastructure on L2 exit and initialize it before sideband packets can be exchanged
- **Our Approach:**
  - Use the existing sideband clock and data pins (no new wires) to indicate L2 exit using DC signal levels (see flow on next slide)
  - A small amount of logic is active while the rest of sideband is powered off/ clock gated to detect exit from L2 and wake up the rest of sideband
  - Rules are provided such that the exit can be symmetric or one sided

# L2 Optimization Flow



# Optimized Manageability Framework

## ■ Early Firmware Download

- **Function:** Standardize data structures and capabilities for firmware download
- **Benefit:** Enable chiplet use of firmware without each chiplet in SiP needing its own flash or firmware loading mechanisms

## ■ Priority Packets Over Sideband

- **Function:** Permit low-latency (bounded) transmission of sideband messages for notification events
- **Benefit:** High-priority events are not blocked by low-priority traffic

## ■ Extended Reach Sideband (UCIe-S only)

- **Function:** Permit 100mm sideband channel to minimize hops/daisy chaining SiP
- **Benefit:** Enables star topology with sideband – director chiplet connected to each chiplet

## ■ Open Drain Pins

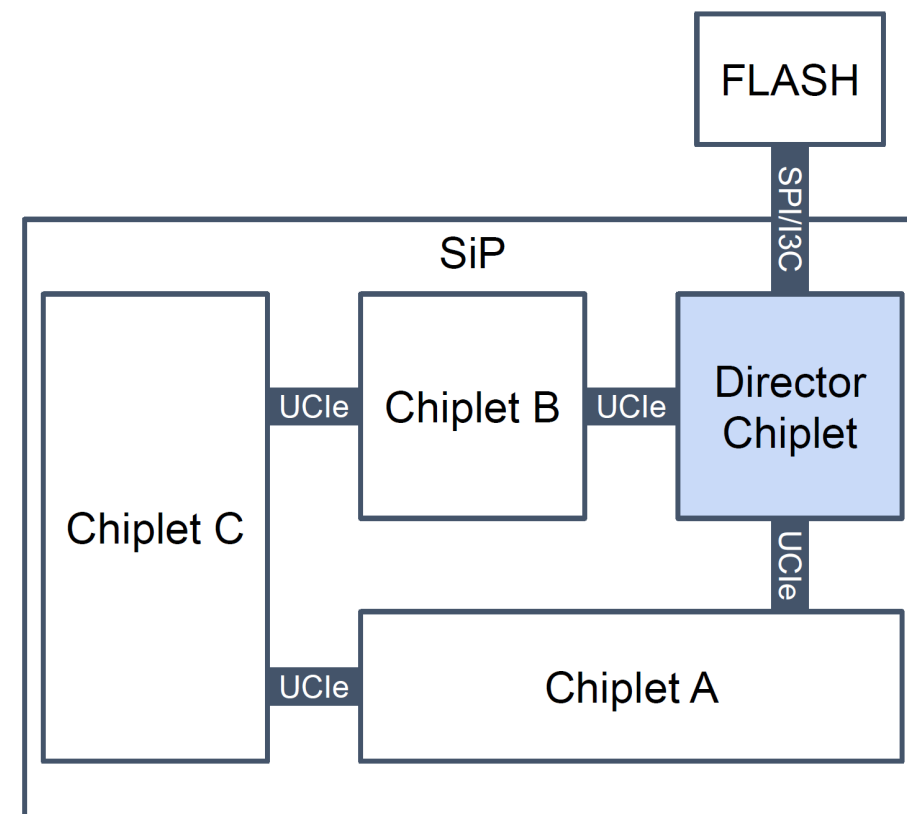
- **Function:** Open Drain Pins enable low latency, bi-directional events
- **Benefit:** Simultaneous SiP wide broadcast to all chiplets

## ■ Fast Throttle and Emergency Shutdown

- **Function:** Setup open drain IO and map critical notification events as potential broadcast in System-in-Package (SiP)
- **Benefit:** Provides a standard approach across chiplet vendors to ensure critical function interoperability at the SiP level


# Early Firmware download Flow overview

- Director Chiplet
  - Can load mutable firmware from outside SiP
  - Initialize side-band management network
  - Download first mutable firmware to chiplets
- Chiplet waits for director to download firmware
  - ⇒ Chiplet boot first mutable
  - ⇒ Chiplet can request further firmware update
    - Through side-band
    - Through main-band
    - Using MCTP / PLDM / ...
- Data structures like circular buffer are also defined for interoperability between chiplets



# Priority sideband packets

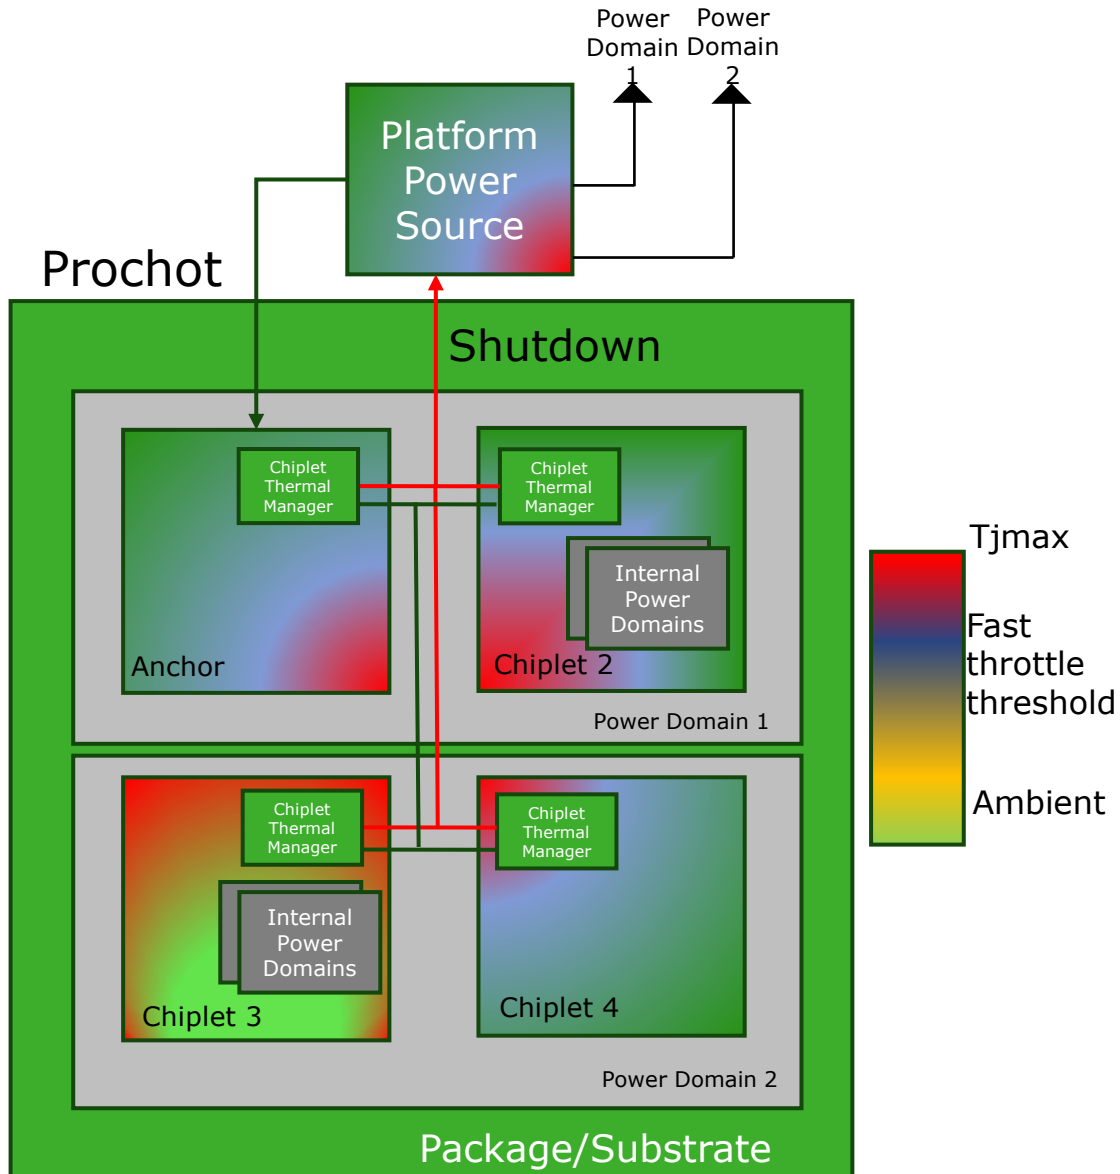
- Motivation: Several events need high priority notification over others
  - Example: Power down, Power wake up, low-latency telemetry data, power supply switch to redundant supply – needs 1-10 us latency. Do not want these to be stuck behind say a FW download/ debug dump
- Approach: Create a mechanism to interrupt sideband packets (“normal traffic”) at an 8UI interval to insert the priority vector (“priority traffic”) that is to be transported to the remote Link partner
- Mechanism:
  - A trigger from the transmitter to indicate it is switching to “priority traffic” from “normal traffic”
    - This trigger is in the form of the clock remaining 0b for 8UI before beginning the priority transfer.
    - The receiver detects this (implementation specific means), and it expects a priority vector next.
  - The priority vector is sent from transmitter to receiver – total 32UI, with 23 bits for the vector, 5b opcode, 3 reserved bits and 1 bit even parity.
  - After this 32UI has completed, the packet from “normal traffic” is resumed or another priority packet can be sent (based on opcode) without any gaps in clock
  - Max time to transfer priority packet: 8UI (boundary) + 8UI (switch) + 32 UI = 48 UI = 60 ns



# Extended reach sideband (UCIe-S only)

- Permit 100mm sideband channel to minimize hops/daisy chaining in SiP (enables practical usage of Star topology in SiP with sideband)
  - Connect the Director Chiplet directly to each chiplet for better manageability and security using a UCIe-S SB only link
- UCIe 2.0 Specification: Sideband like main-band reach specified at 25 mm, even for sideband-only link
- Given that the operating frequency is 800 MHz, we can easily extend that reach to 100mm for sideband
- Details: Provide appropriate guidelines for the extended reach
  - $V_{ih} = 70\%$  of  $V_{CCAON}$ ,  $V_{il} = 30\%$  of  $V_{CCAON}$
  - At Longer Channel Lengths, Slope is not a meaningful measurement. Waveforms take a long time to reach  $0.8 \cdot V_{CCAON}$ , and sometimes do not reach 80% of  $V_{CCAON}$ . So,  $T_{xron}$  is a more meaningful measurement for measuring Eye Height and Eye Width.
  - For longer reach, Driver Ron needs to be limited to 60 Ohm worst case (including worst case errors and supply variations)

# Throttle and Shutdown




## ■ Fast throttle

- Introduce a common dedicated open drain bidirectional pin on the chiplets
- Wires from all chiplets participating in a thermal zone are tied together, such that any combination of chiplets can pull down the pin
- Chiplet(s) assert the pin when the respective internal fail-safe temperature limit is hit or if signaled from the external platform running hot.
- When asserted, all participating chiplets throttle to the pre-negotiated level and at a defined rate

## ■ Emergency Shutdown

- Introduce a dedicated open drain bidirectional pin on the chiplets and an off-package driver
- Wires from all chiplets participating in a thermal zone are tied together, such that either one can pull down the pin
- Chiplet(s) assert the pin when their max temp limit is hit.
- SiP exposes the shutdown to internal (through the directional nature of this communication) and external power source i.e., off-package driver for shutdown

# Conclusion





# Summary

- The UCIE Consortium is committed to establishing an open chiplet ecosystem and a ubiquitous interconnect at the package level.
- The UCIE specification is continuing to evolve, based on end-user feedback, to meet the new usage models.
  - Tremendous support across the industry with several companies announcing IP/VIP availability
  - Evolving as the interconnect of SoCs just as PCIe and CXL at the board level
  - UCIE 3.0 Specification is available to the public at [uciexpress.org/specification](https://uciexpress.org/specification)
- Get involved! Learn more by visiting [UCIexpress.org](https://UCIexpress.org)

# Thank You

[www.UCIexpress.org](http://www.UCIexpress.org)