



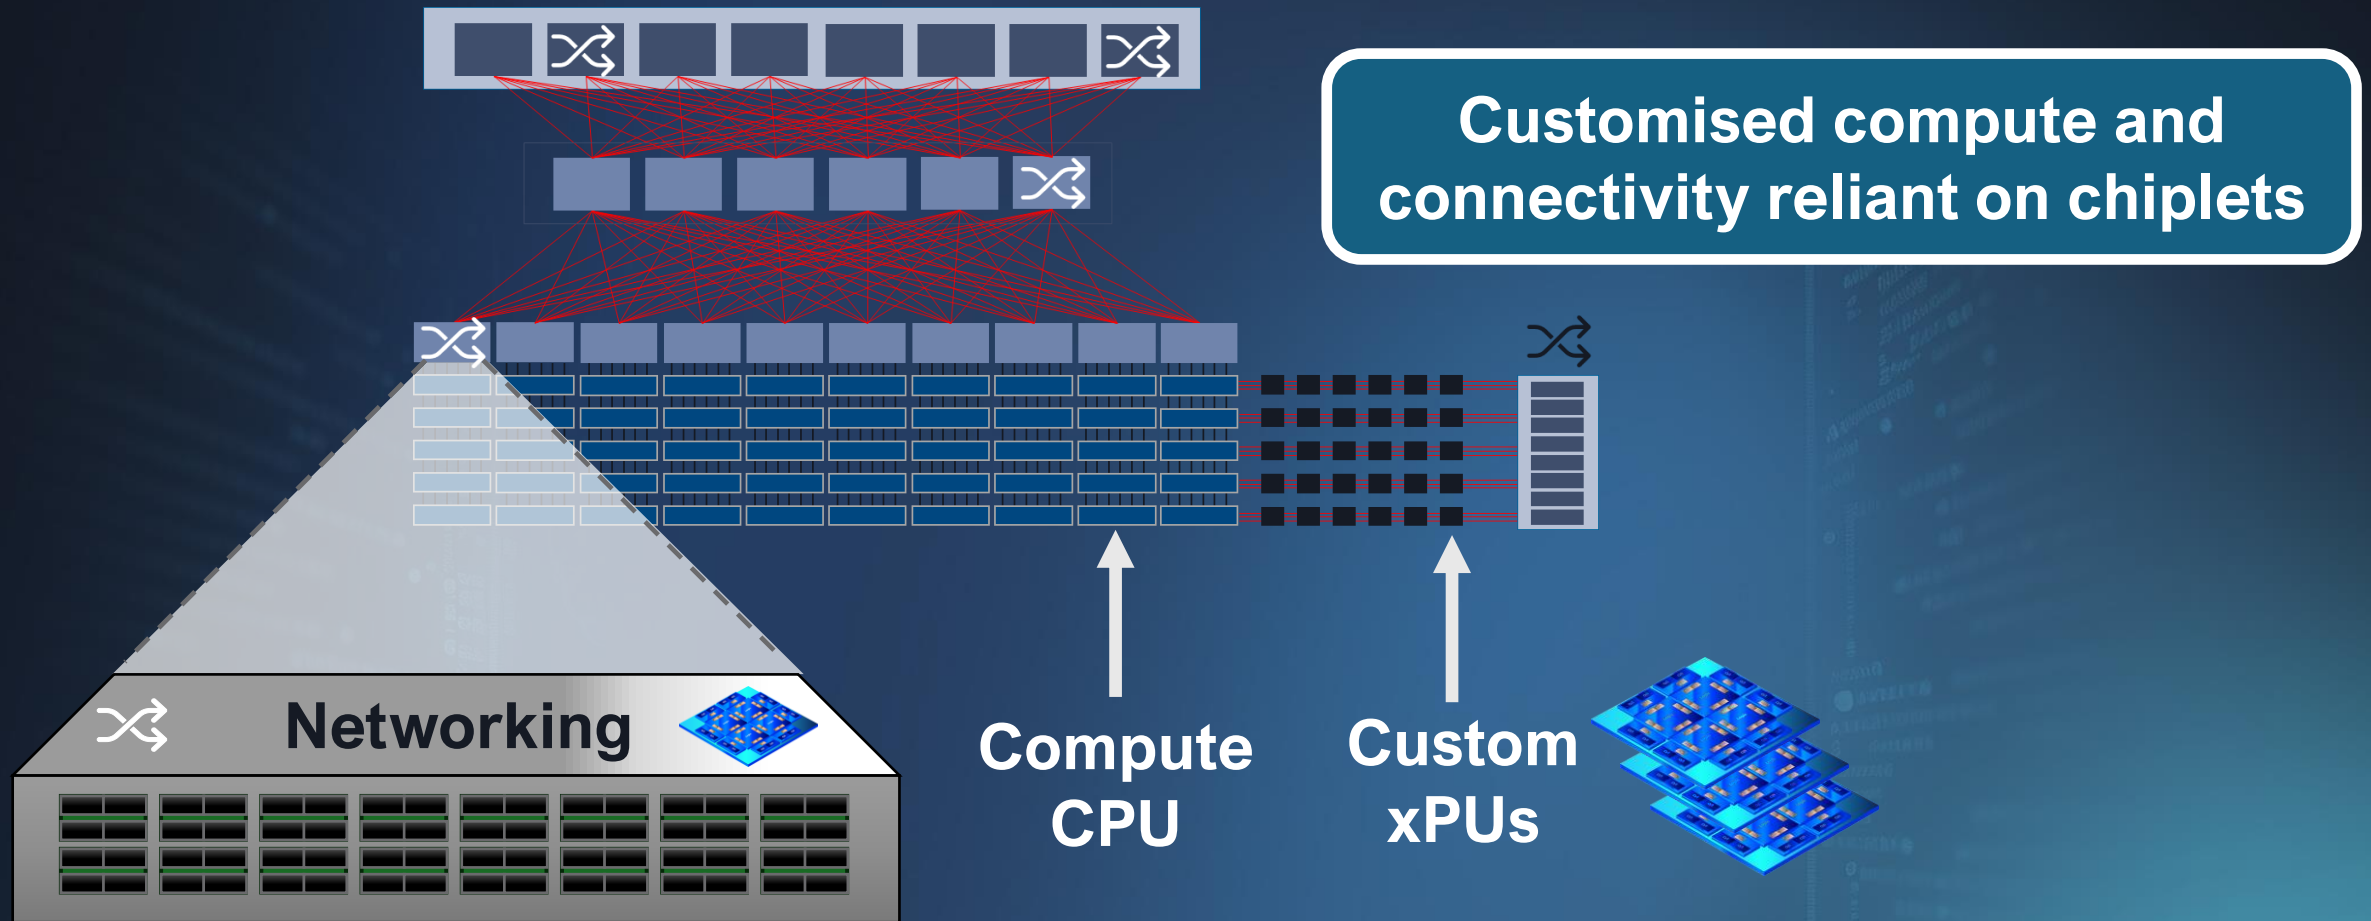
Redefining Connectivity: Charting Next-Gen Pathways in Chiplet Interconnects

Tony Chan Carusone, CTO


AI Hardware and Edge AI Summit

September 11, 2024


Delivering Custom Silicon in the Data Centre



Outline



- Scaling AI with Connectivity




- The Role of Chiplets for AI Compute



- Enabling a Chiplet Ecosystem

Scaling AI with Connectivity


Evolving Data Centre Connectivity Landscape



- Optical and electrical links
- Flexible and redundant networking

This evolution is accelerating and diversifying with AI deployment in the data centre


AI in the Data Centre – Proliferating Connectivity



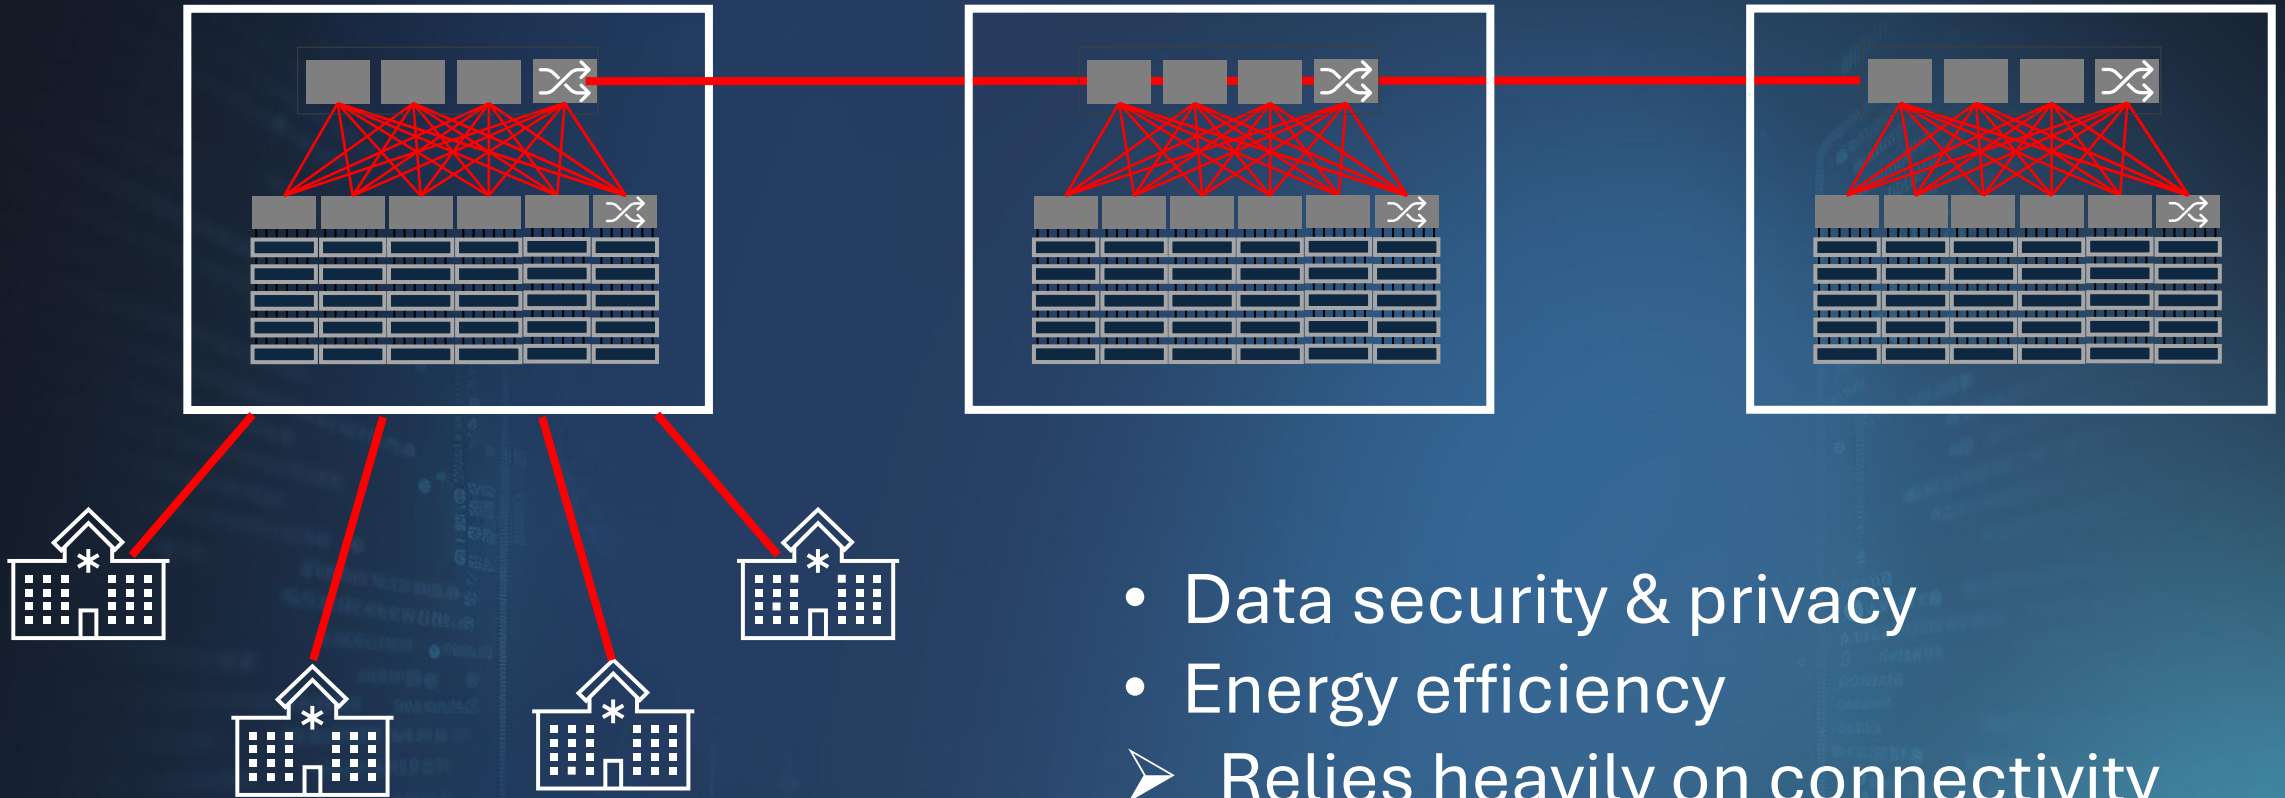
Back-End ML Network

- Low-latency & high-speed

Scaling Up and Out








Distributed AI



- Data security & privacy
- Energy efficiency
- Relies heavily on connectivity

Connectivity Demands for AI

-  Higher per-lane data rates
-  Scale up and out to larger clusters
-  Low latency; Low power; Low failure rate
-  Proliferation of new optical connectivity technologies

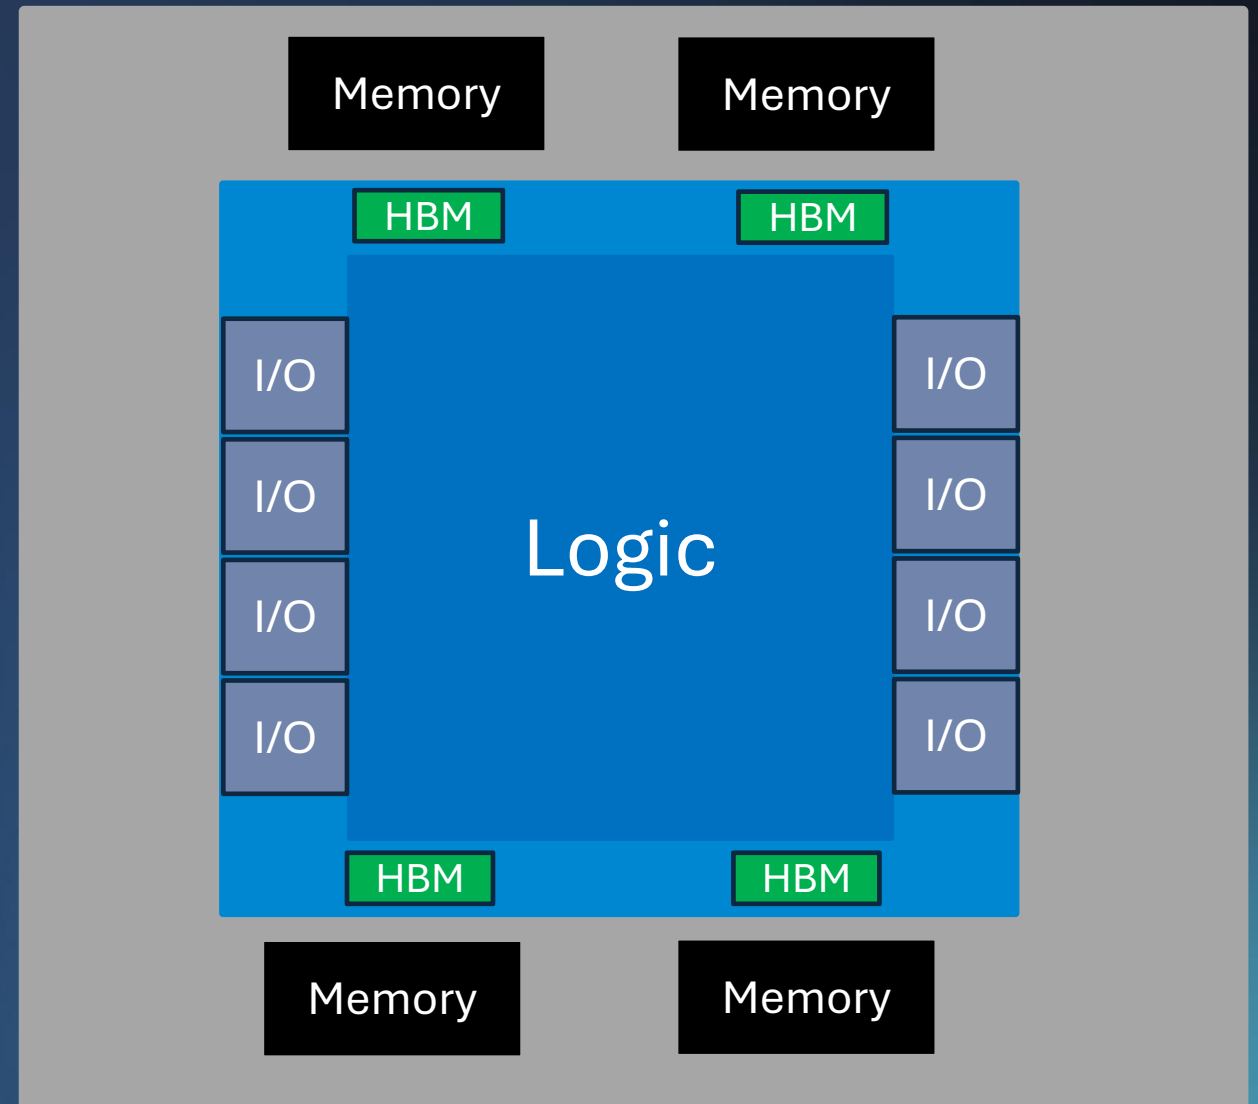


The Role of Chiplets for AI Compute

Monolithic Solutions

Benefits and Drawbacks

- Requires integration of external IP
 - Design and verification time and risk
 - Licensing cost
- Maximum size is limited by reticle
- Yield reduced by large die size
- I/O must be in the most advanced technology
- Every I/O must be capable of driving any interconnect




Bandwidth Density

- Reticle-limited die sizes afford roughly 50mm die edge on the east + west
- 50Tbps aggregate I/O bandwidth requires 1 Tbps/mm

Adv. Logic Node

Chiplet-Optimized Node

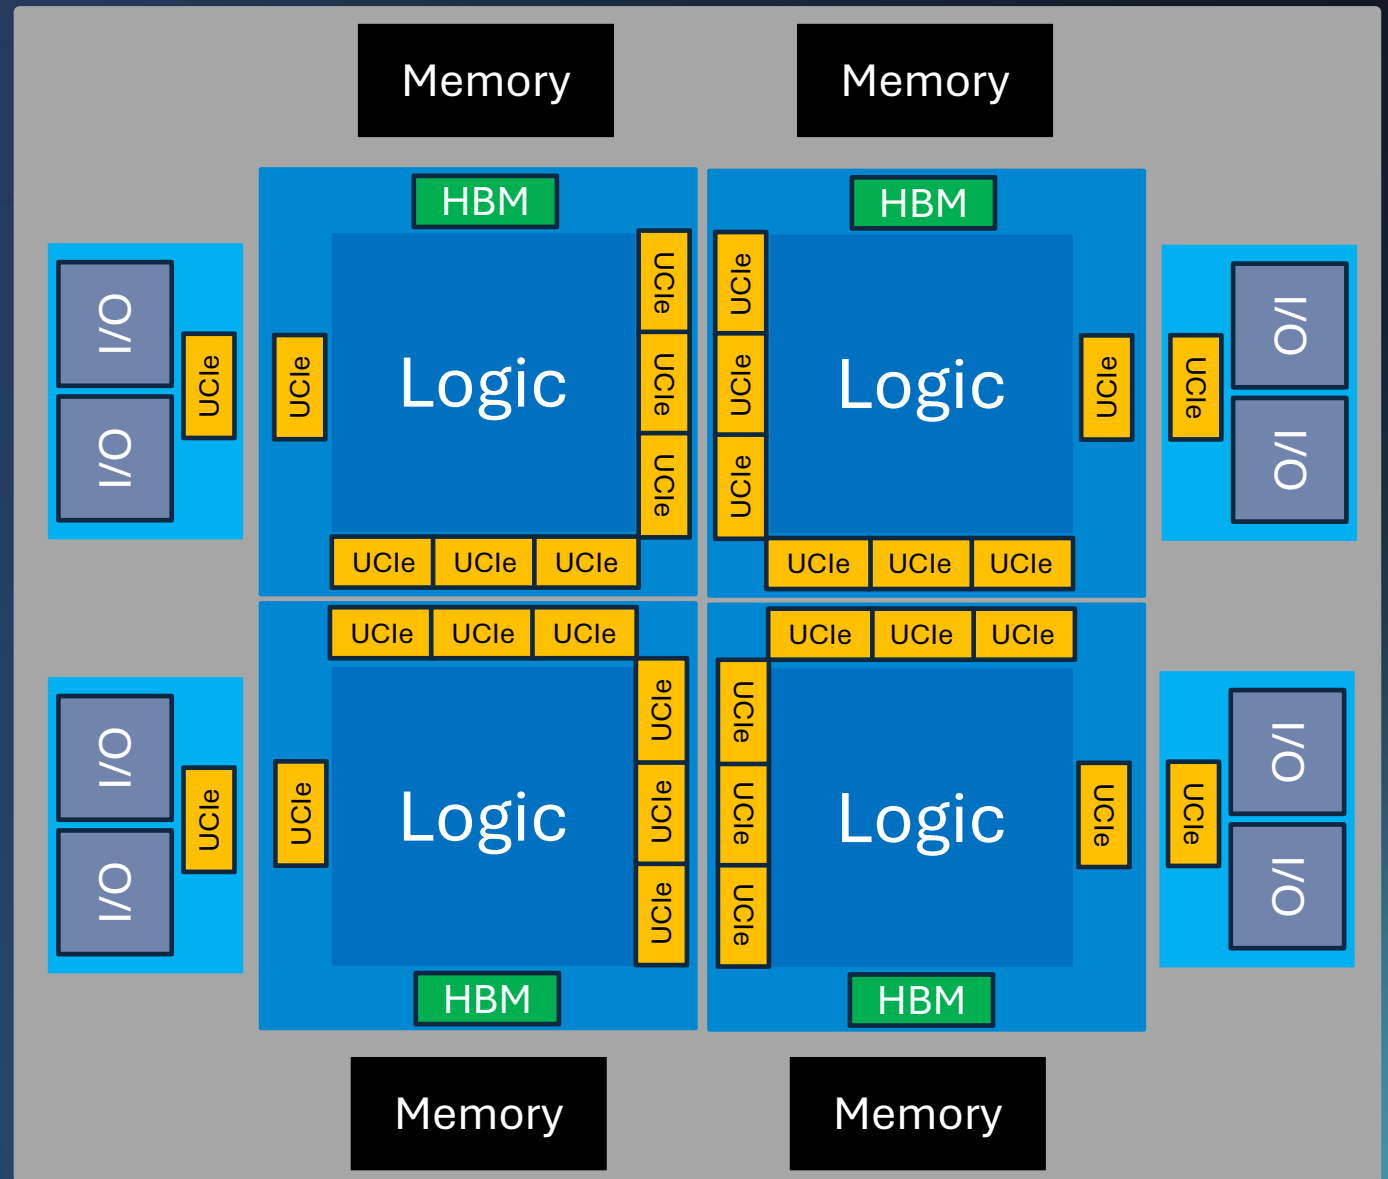


Chipllets

- Reduced design time, risk, NRE and silicon costs
- Composability

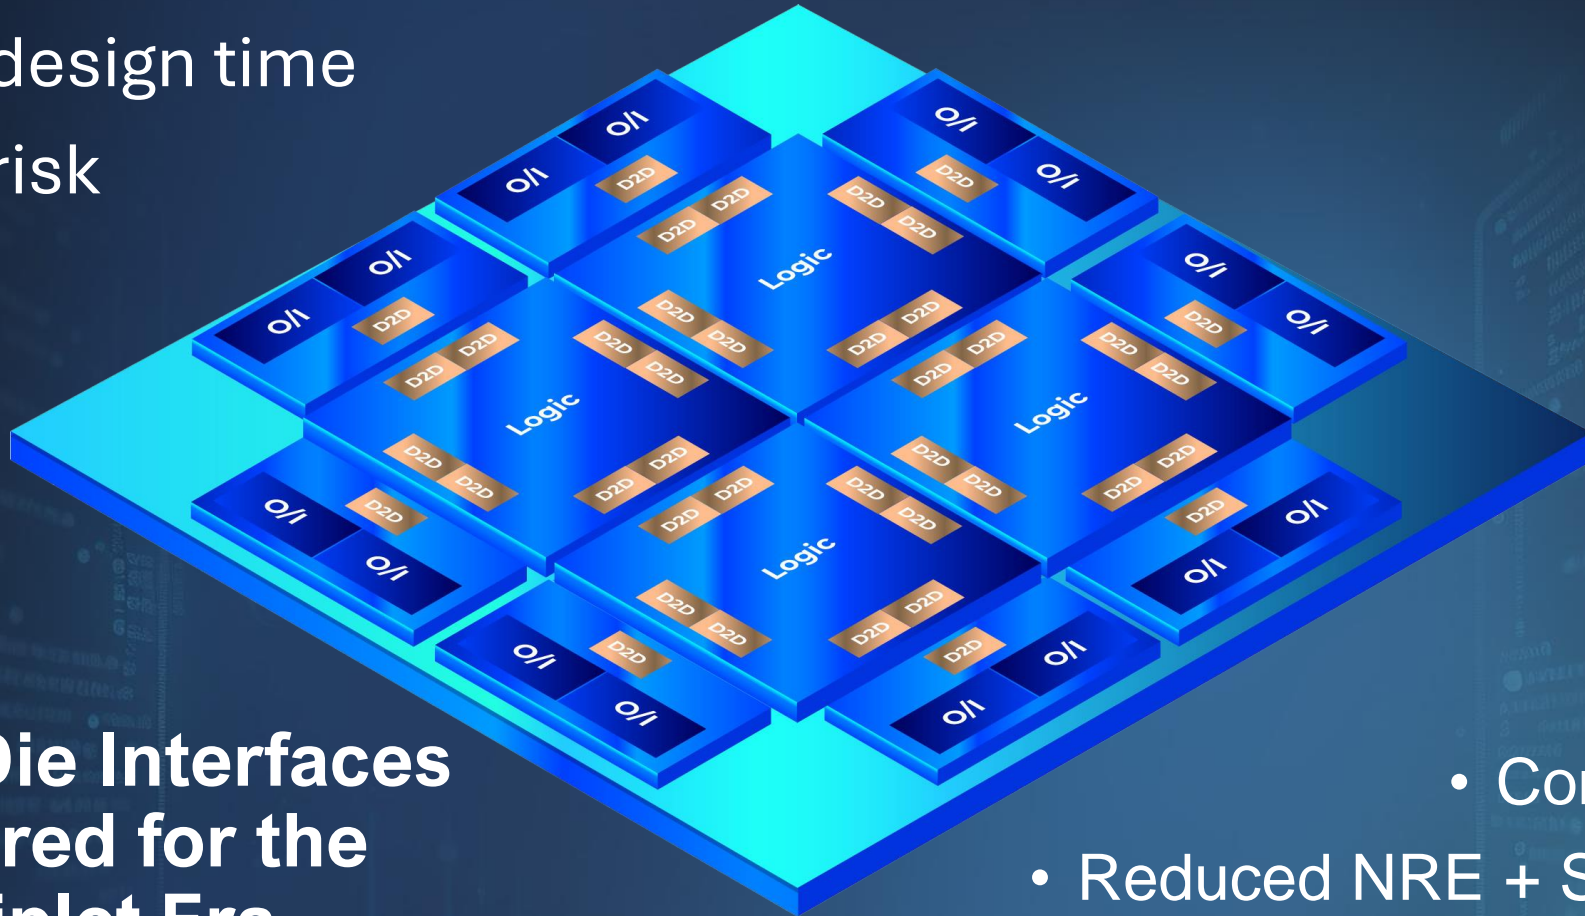
Adv. Logic Node

Chiplet-Optimized Node



Accelerating Hardware Upgrades

- Reduced design time
- Reduced risk






Die-to-Die Interfaces Required for the Chiplet Era

- Composability
- Reduced NRE + Silicon Cost


Connectivity Hardware in the Datacenter

Pluggable
Modules or
Transceivers



Optical Module Anatomy with Chiplets

Chiplet devices on a low-cost substrate

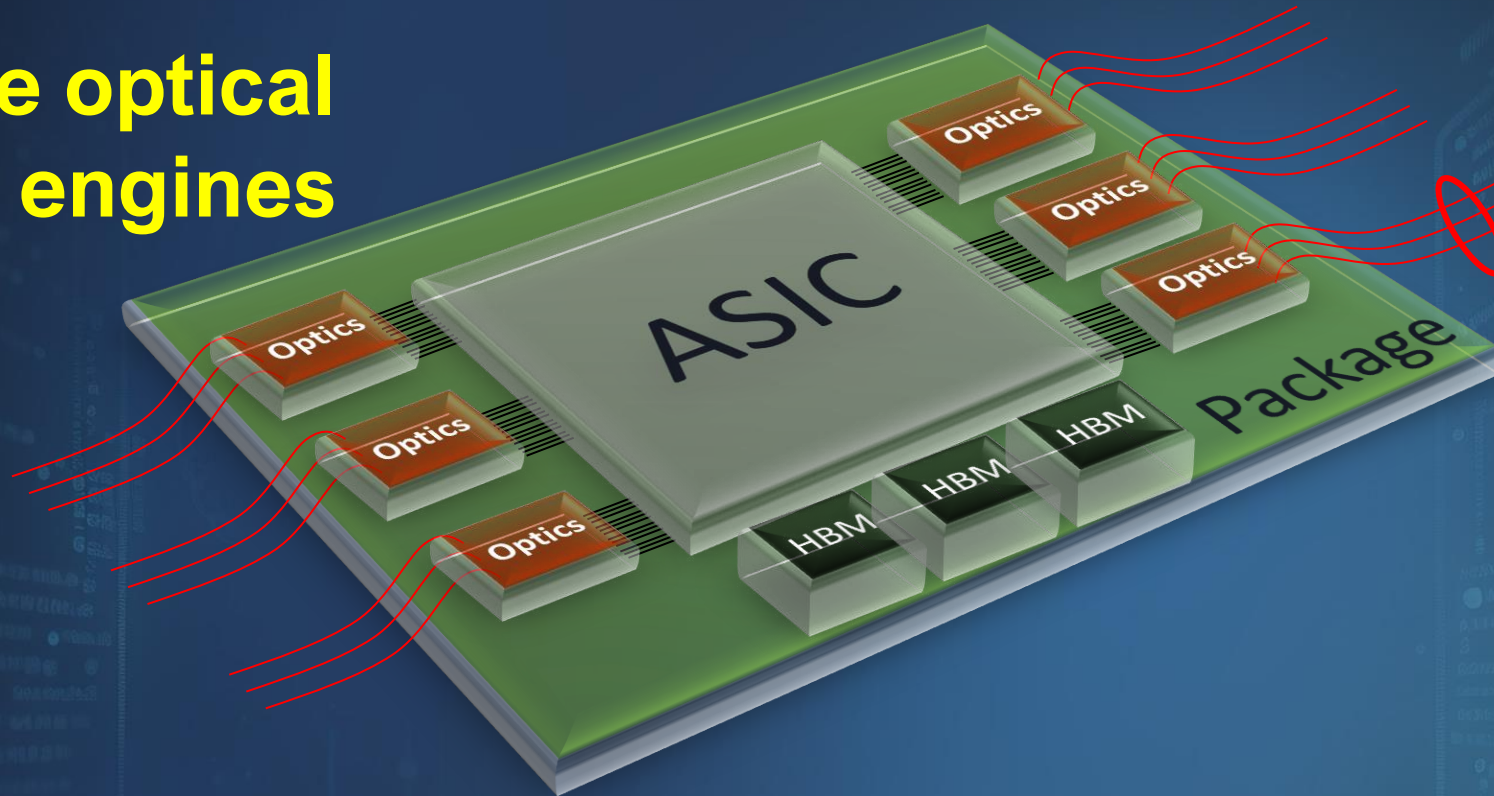


Benefits:

- ✓ Power
- ✓ Performance
- ✓ Cost

Co-Packaged Optics (CPO)

Miniature optical engines

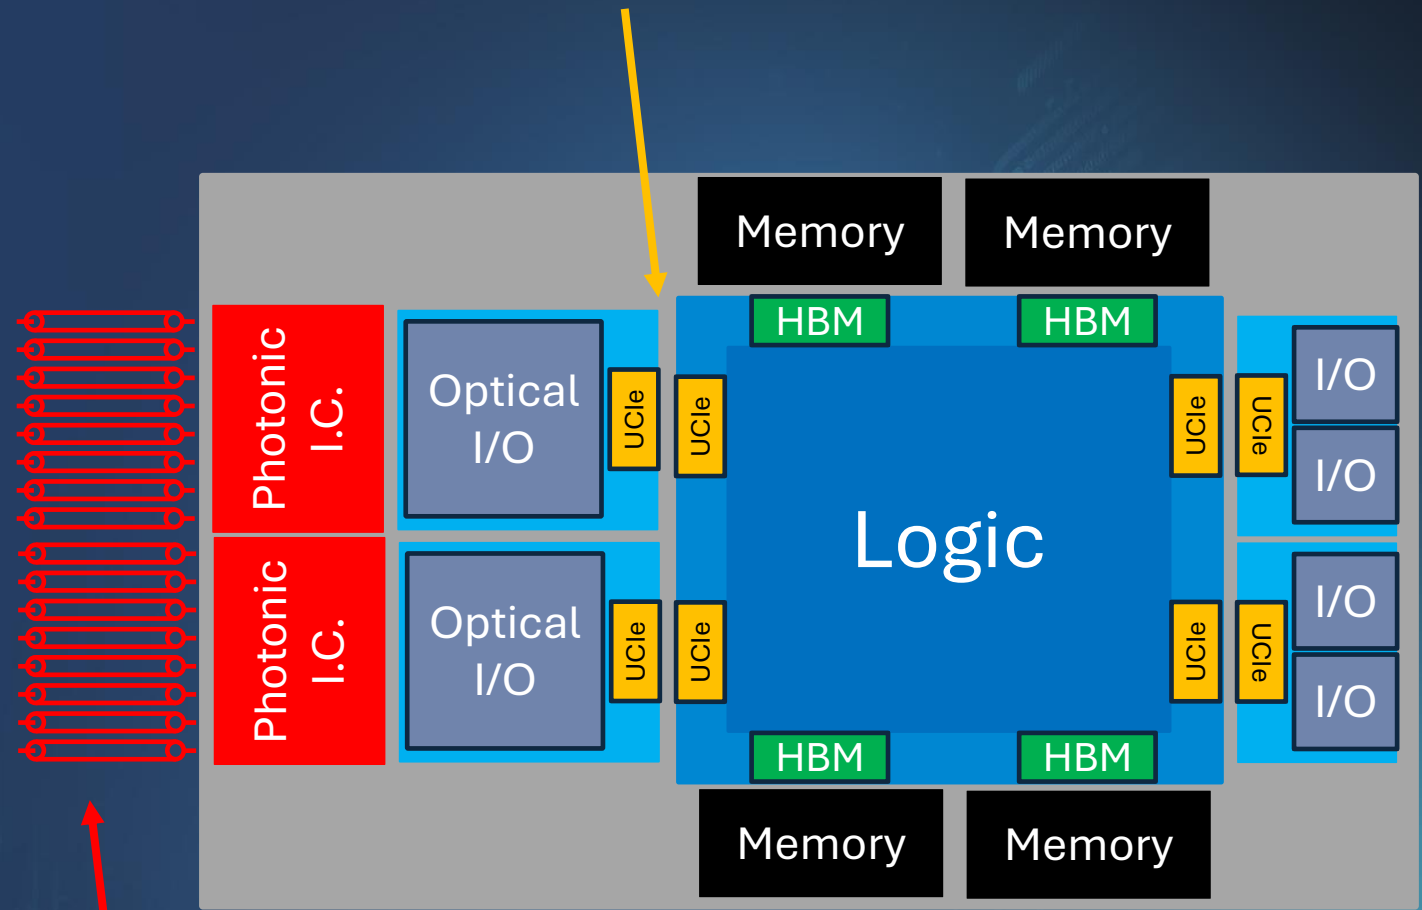


Optical fiber

CPO Chiplets


- Mix-and-match with electrical I/O chiplets for different applications
- Optical I/O density being pushed by new solutions:
 - Multiple wavelengths
 - Dense fiber arrays
 - Package fan-out

Bandwidth (“Beachfront”)
Density Up to 10 Tbps/mm



Distributed Data Centers Driving New Connectivity

Geographically-distributed compute requires broadband connectivity



Future solutions are needed to address these new requirements

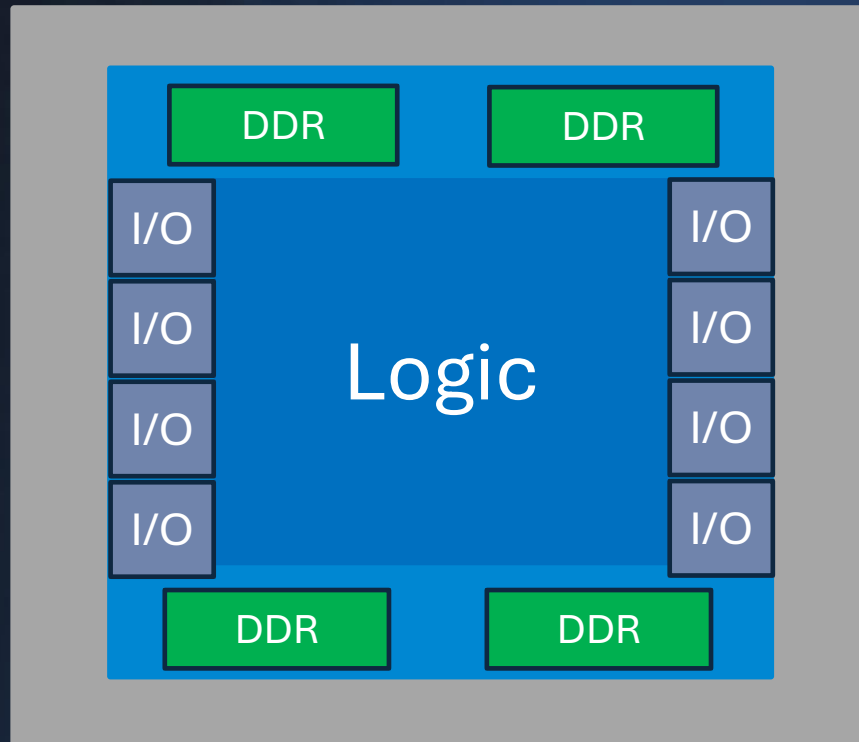


Enabling the Chiplet Ecosystem

Design Paradigm Transformation

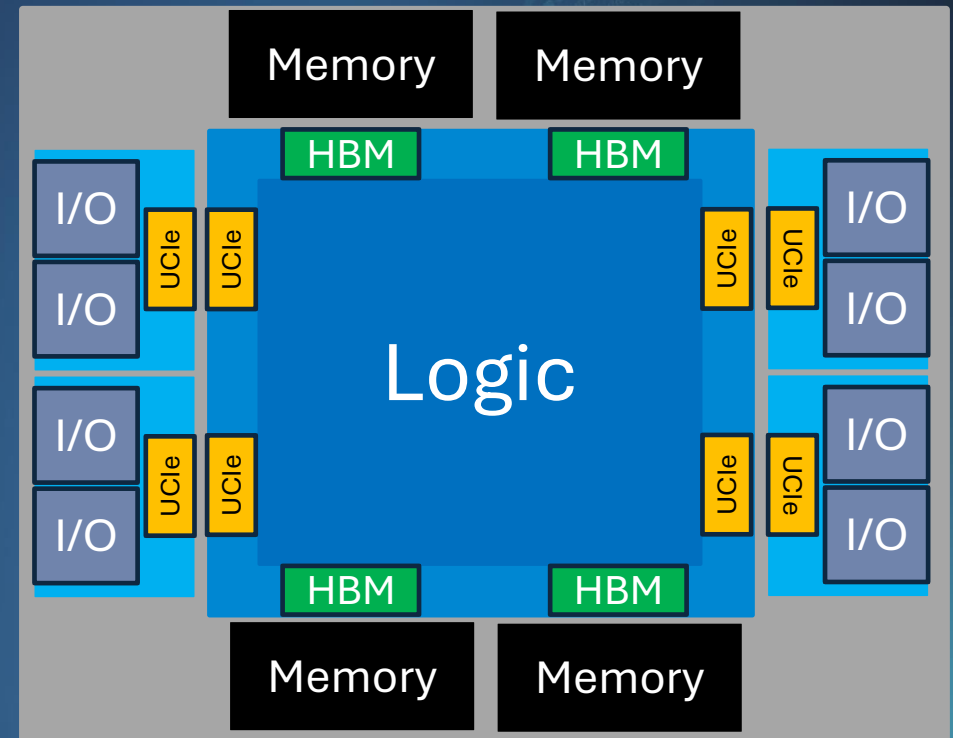
Monolithic Design with Foundry-Enabled Foundation IP

➤ Accelerated design



Heterogeneous Design with Foundry-Enabled Foundation **Chiplets**

➤ Evolution of Hardware Design



Die-to-Die Interfaces Supporting a Chiplet Ecosystem



Bandwidth density



Energy efficiency



Latency



Reach



Availability, Interoperability, Reliability

Key Technologies in Die-to-Die Interfaces



Signal Integrity

- Managing Crosstalk
- Interconnect loss compensation



Power Integrity


- Limited area for supply decoupling
- Variation in packaging technology and dense signal routing complicate power delivery




Clocking

- Low-power and low-jitter clock distribution
- Clock/Data alignment


A Chiplet Portfolio



- I/O Chiplet
 - **Multi-Standard SerDes IO with Integrated Protocol Controllers: PCIe Gen6 / CXL 3.0 / 112Gbps Ethernet**




- Compute Chiplet
 - **High Performance, Arm-Based Compute**



- Memory Expansion Chiplet
 - **Low Latency DDR/HBM**

➤ **UCle Die-to-Die Interfaces**

Takeaways



- **AI has redefined datacenter infrastructure**
- Connectivity technologies are the key to scaling AI clusters and geographically distributed datacenters effectively
- **Chiplets enable custom silicon solutions optimized for AI workloads**
- Essential to affordably scale performance, achieve lower power, and faster time to market
- **A chiplet ecosystem is emerging, enabled by die-to-die interfaces** and allowing for a wide variety of chiplet use cases



ALPHAWAVE SEMI

Accelerating the Connected World

Thank You!