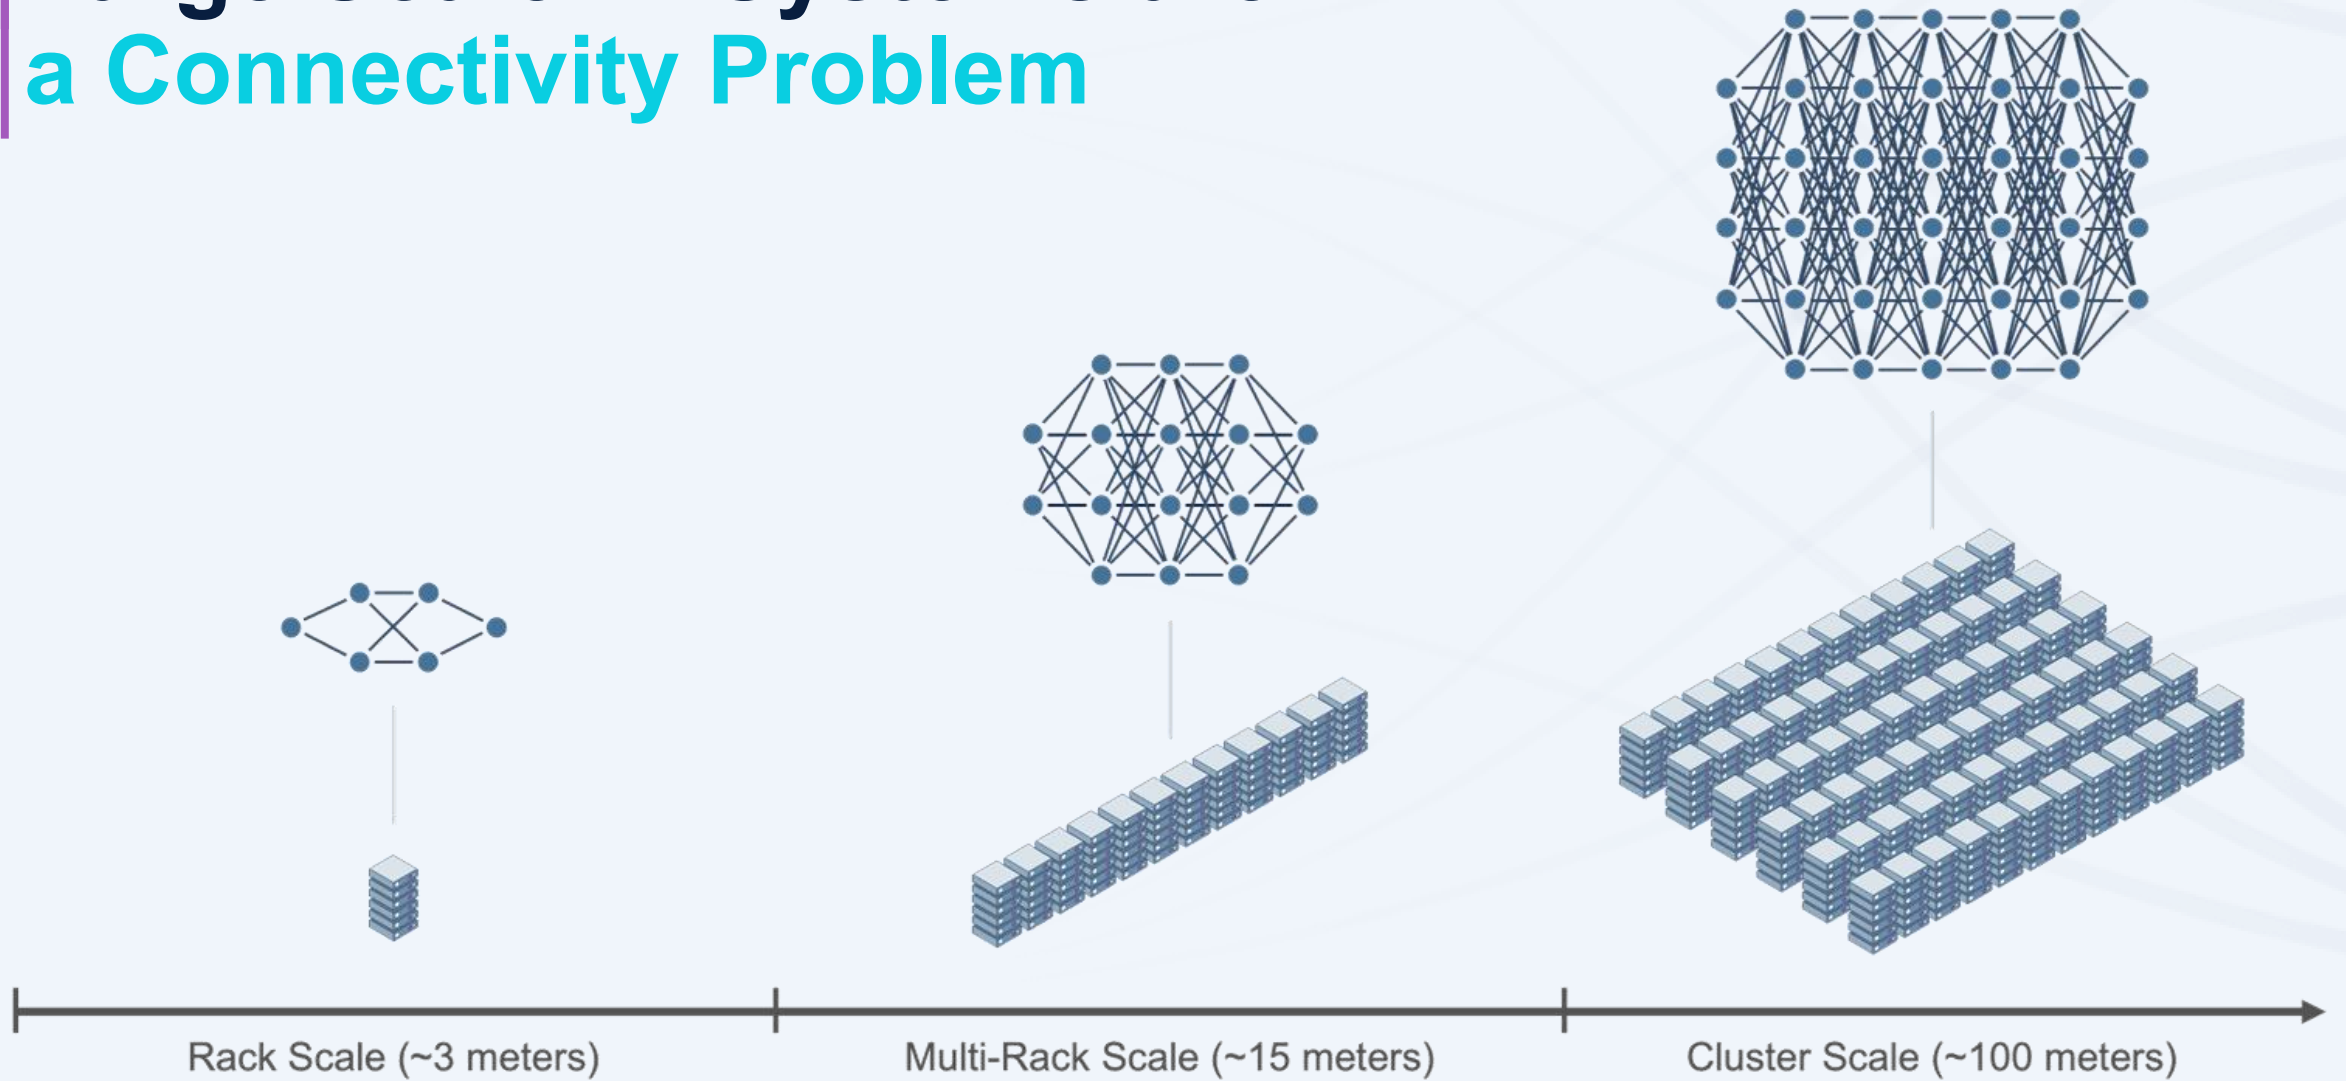




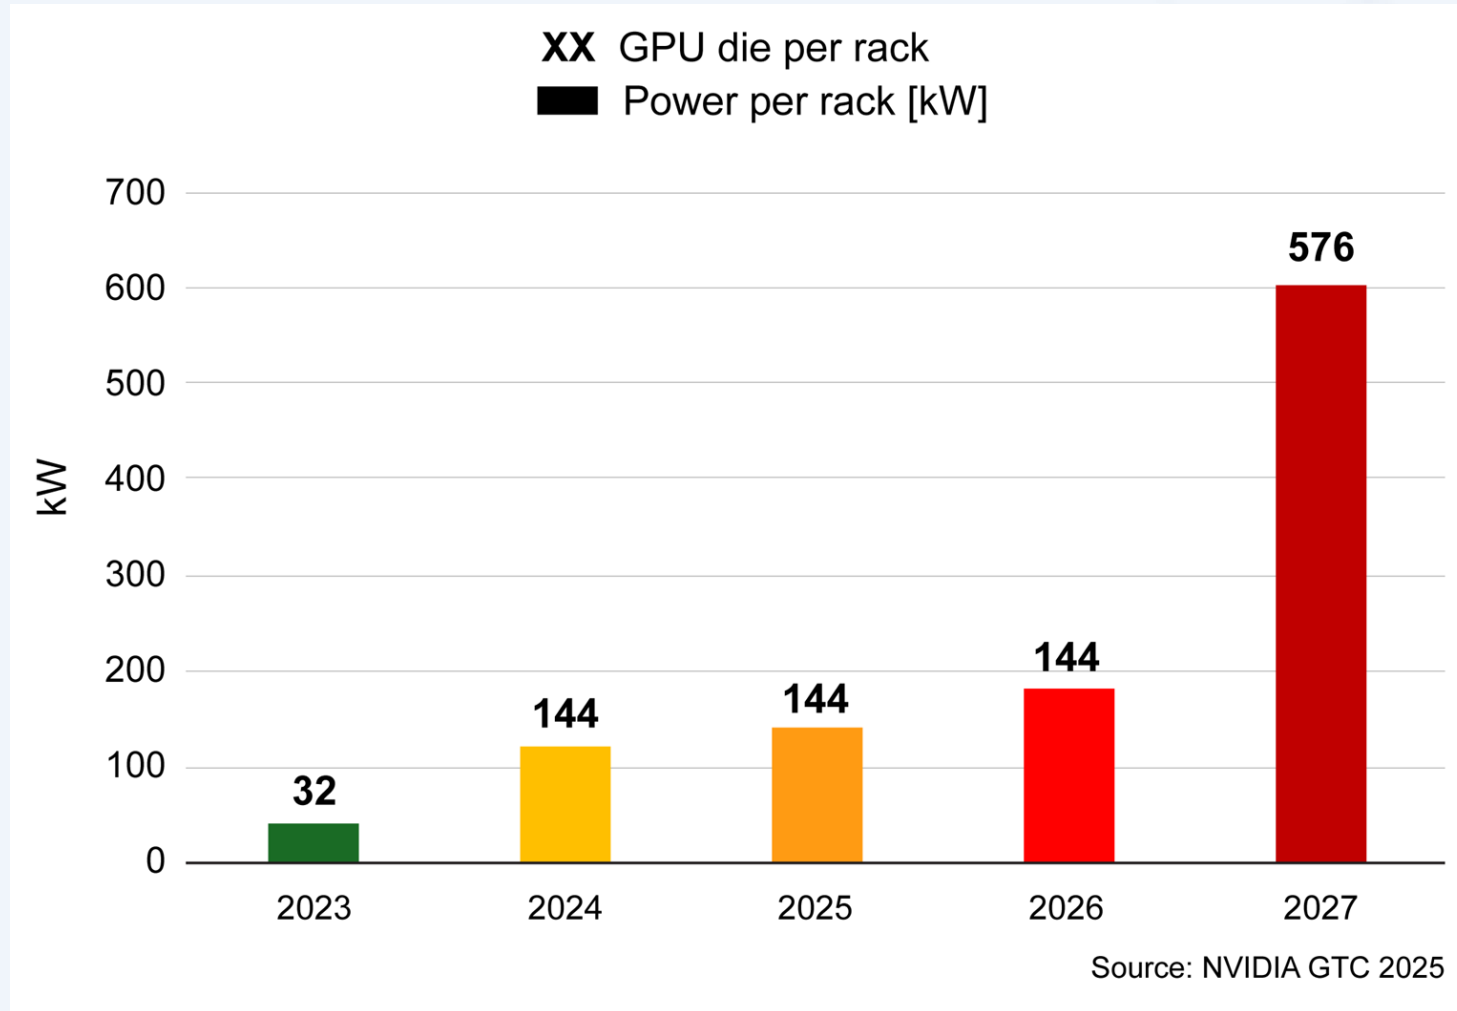
## A UCle Optical I/O Retimer Chiplet for AI Scale-up

Vladimir Stojanovic  
CTO and Co-Founder

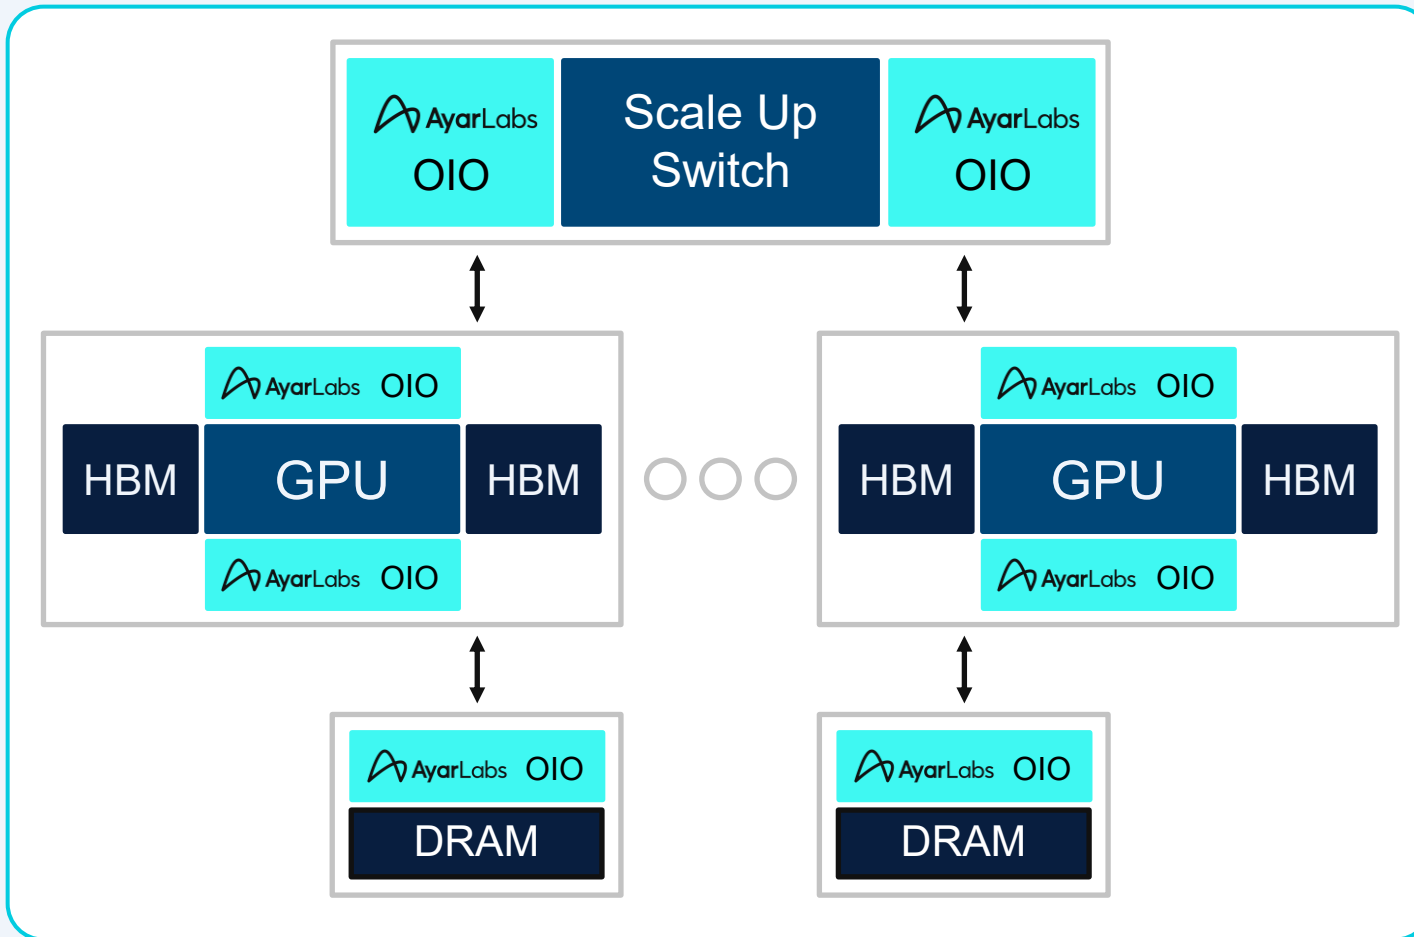
# Large Scale AI Systems are a Connectivity Problem



# Challenge: Power Per Rack with Electrical I/O




# Scalable, Flexible, and Composable AI Infrastructure



## Scale-Up Network

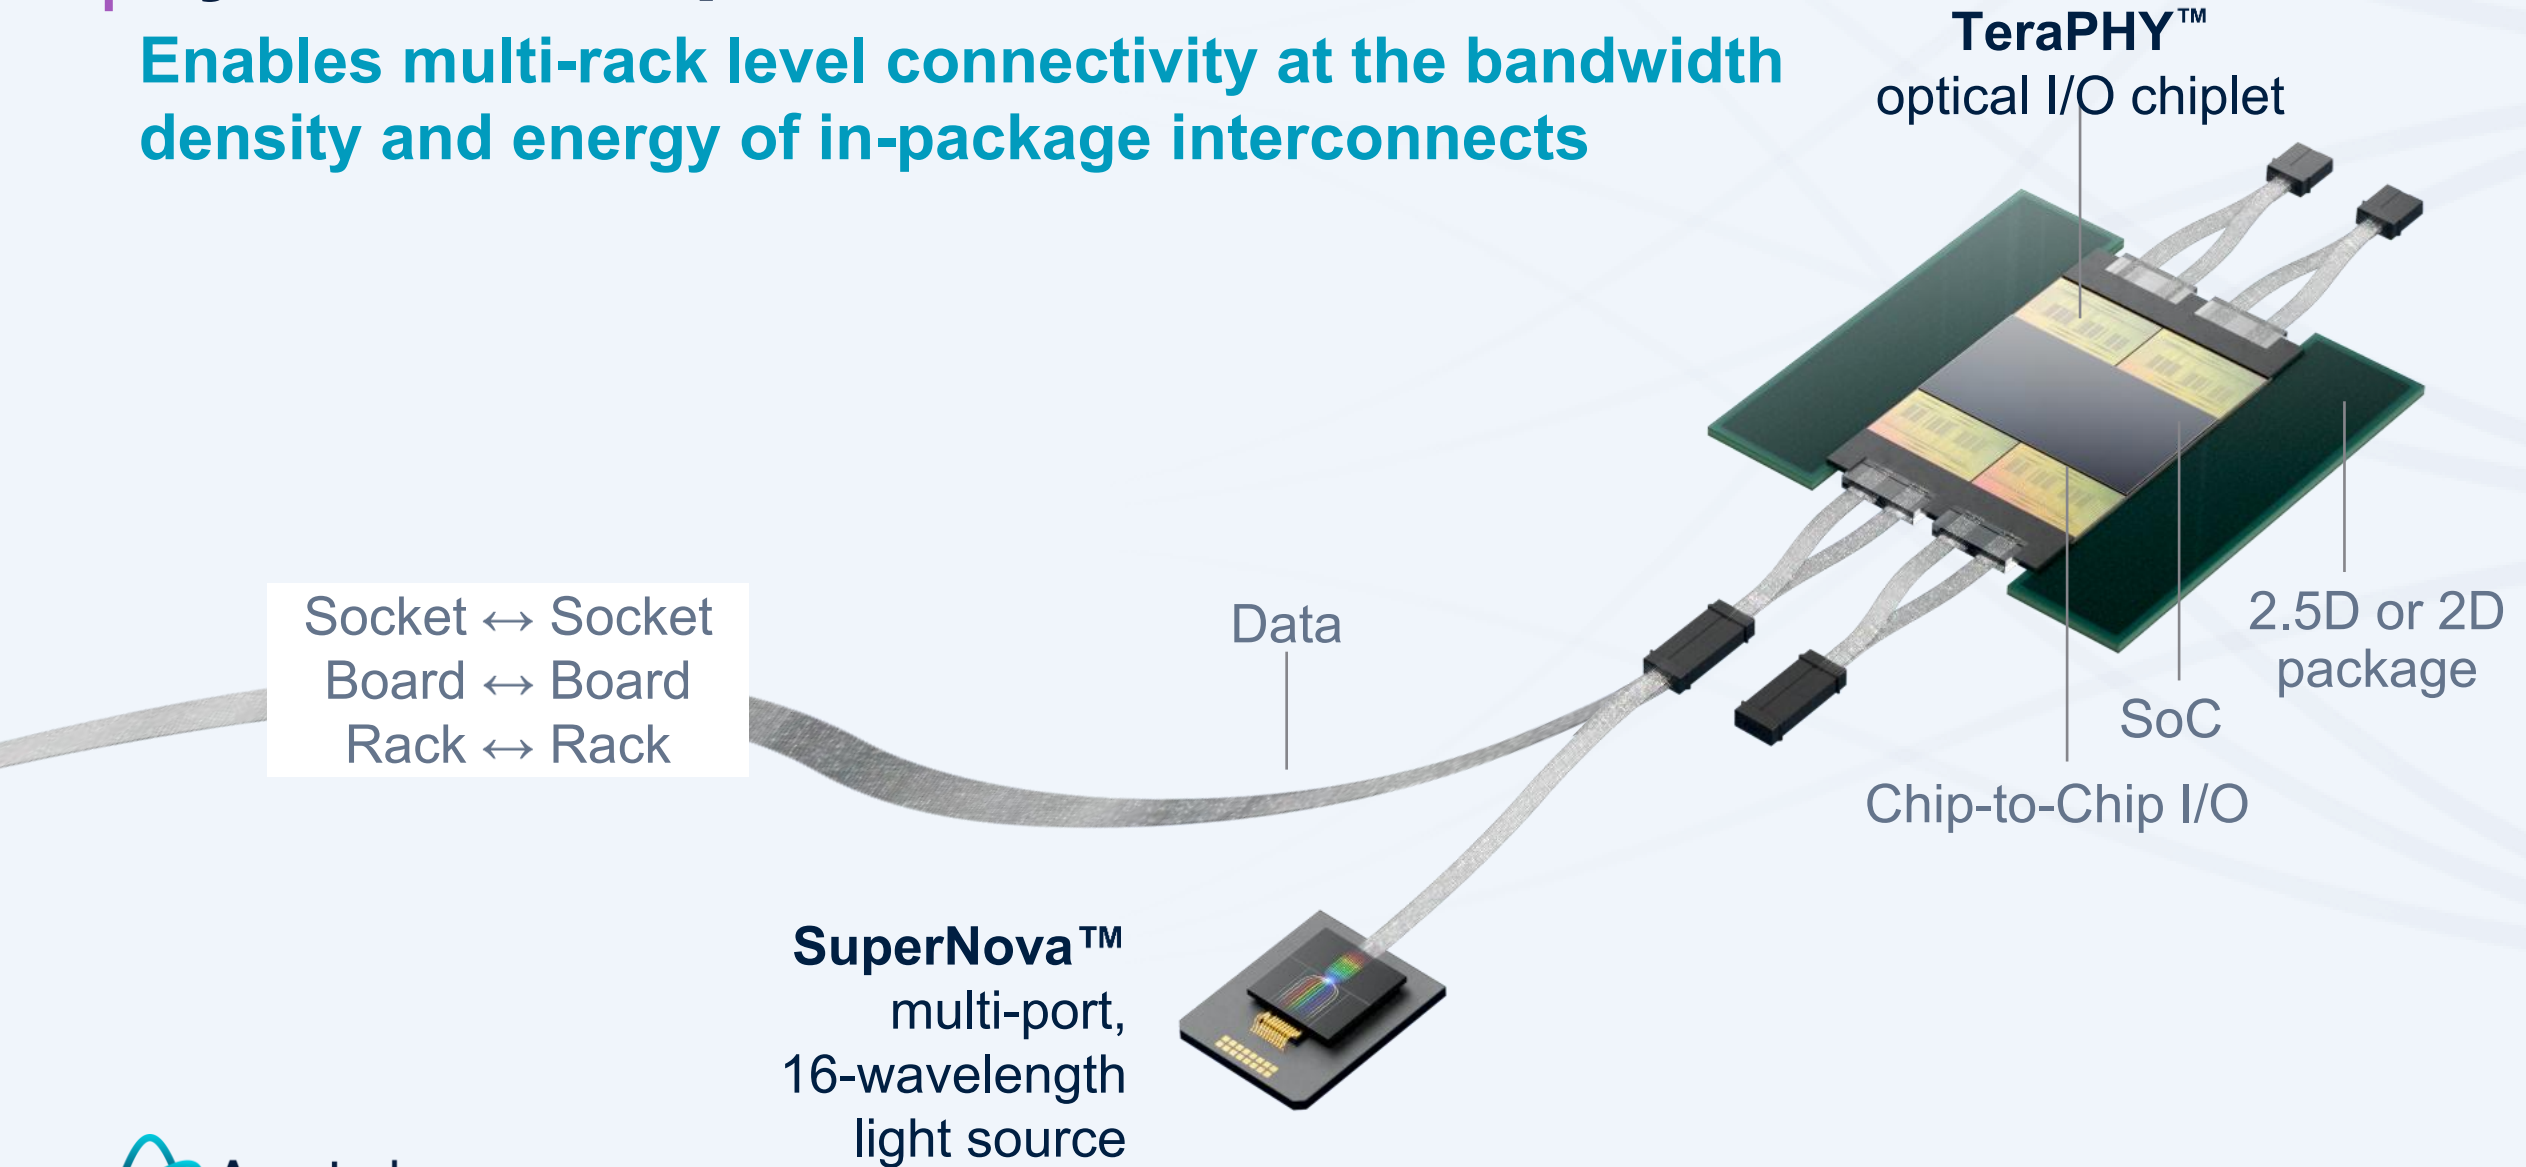
- GPU-to-GPU
- GPU-to-Switch
- Switch-to-Switch

## Extended Memory

 = Ayar Labs Optical I/O

# Ayar Labs Optical I/O Solution


Enables multi-rack level connectivity at the bandwidth density and energy of in-package interconnects



# UCle 8 Tbps TeraPHY™ Optical I/O Chiplet

## Key Innovations:

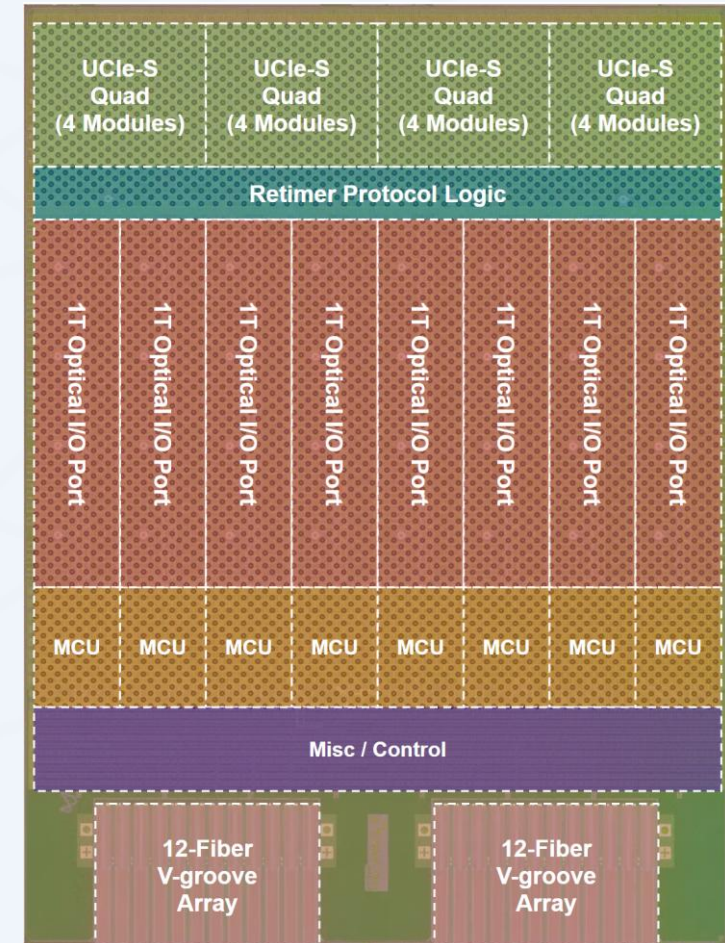
- First implementation of a UCle optical retimer
- First demonstration of 16-wavelength microring based links
- DWDM w/ SMF fiber between TX and RX
- Support of up to 8.192 Tbps of bidirectional bandwidth



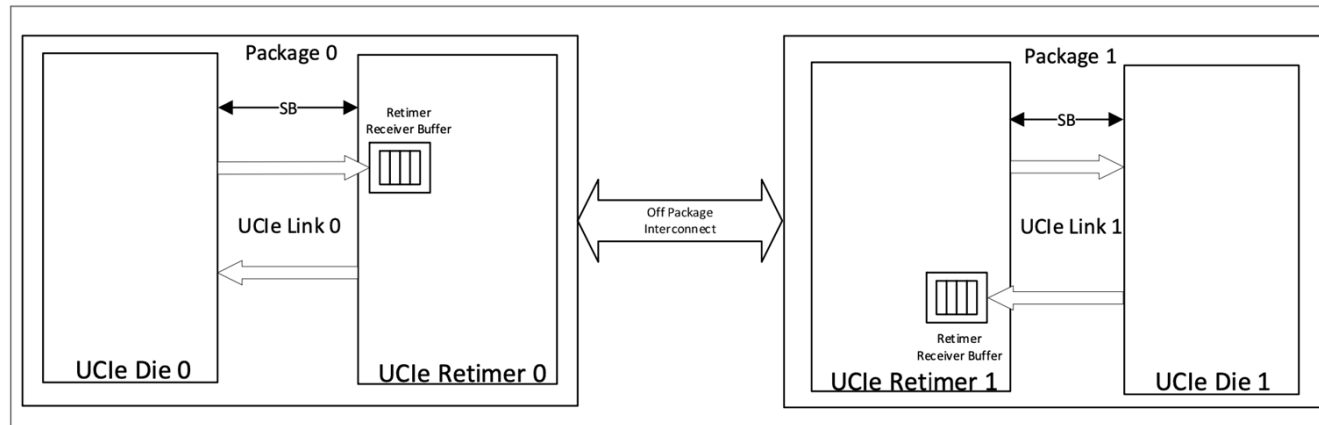
# UCle Optical Chiplet Organization

## Key Features:

- Fully retimed chiplet architecture
- Streaming Raw mode
- Supports various host UCle-S interfaces (stacked, unstacked)
- 8 TeraPHY™ optical I/O ports
- Embedded MCUs



# UCle Architecture

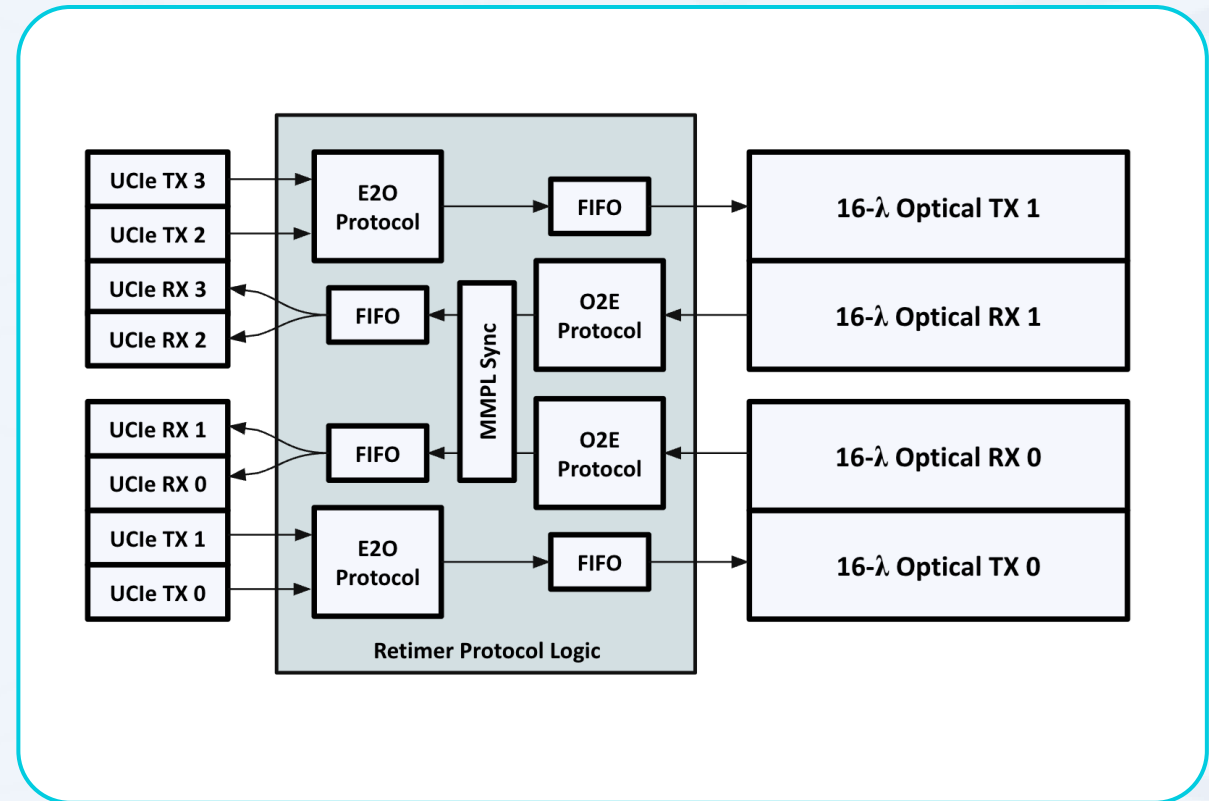


[UCle Specification 2.0]

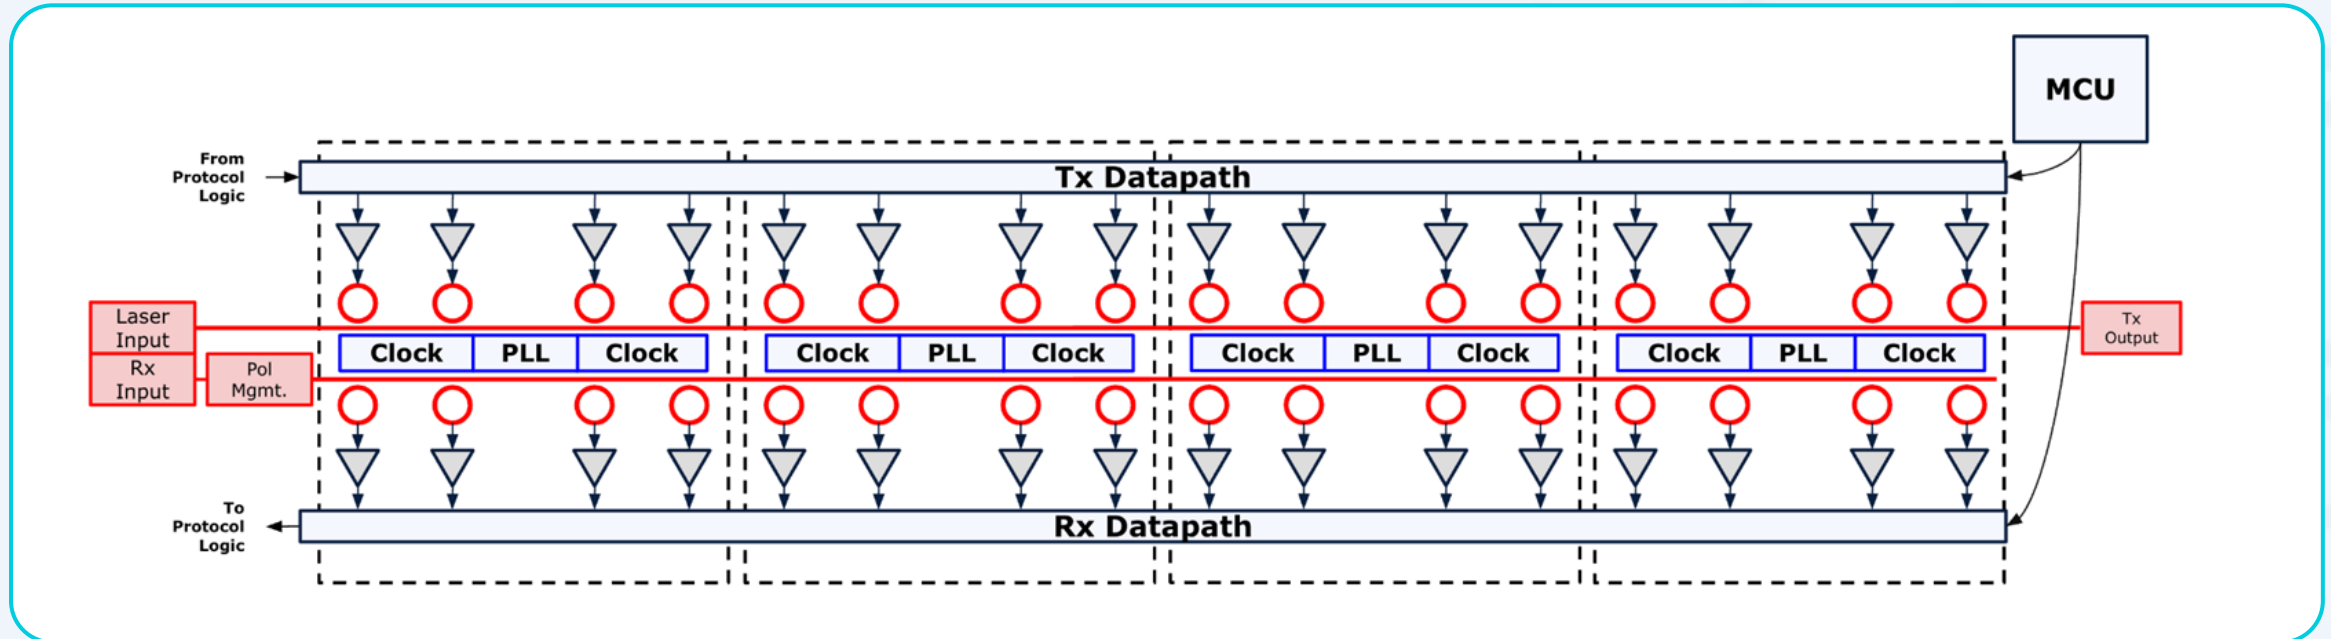
- UCle standard package @ 16 Gbps per lane w/ termination
  - 512Gbps bidirectional BW per module
- 16 modules per chip, total 8.192 Tbps bidirectional BW
- 123um bump pitch, bandwidth density of 896 Gbps/mm
  - Limited by UCle Standard bump map
- Includes retimer protocol logic to transfer UCle traffic over optical I/O

# Retimer Protocol Adapter

- 2:1 gearbox
- Performs encoding of UCle valid signal over optics
- Synchronizes optical lanes and compensates PPM offset between optical chiplets
- Enables UCle x4 MMPL by synchronizing between two fiber ports
- Supports message transport over optics for UCle sideband or other management protocols
- ~1.5% bandwidth overhead


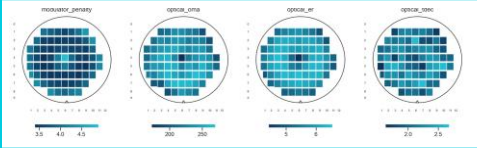
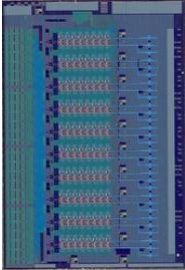
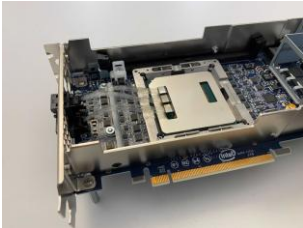
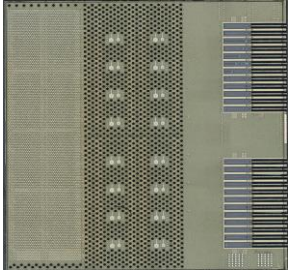

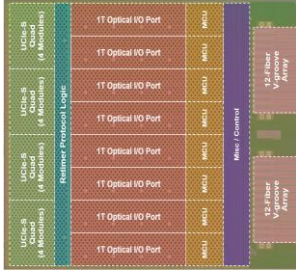
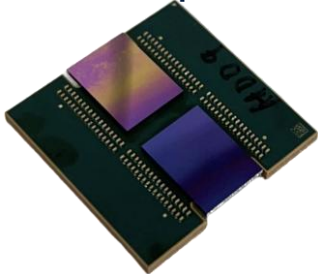


# Optical I/O Port Architecture

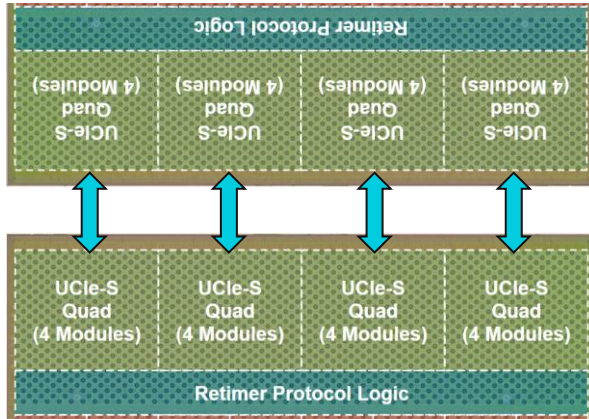


- 16 TX and 16 RX lanes @ 32Gbps NRZ per lane, 1024Gbps bi-directional BW
- Shared PLL and clocking for groups of 4 lanes
- Includes built-in polarization management to support SMF fiber on RX input
- On-chip MCU enables firmware control of initialization and management


# Maturing the TeraPHY™ Chiplet Development Toward HVM

2019	2021	2023	2025	2027
	<ul style="list-style-type: none"> <li>✓ EWS</li> <li>✓ Bring-up</li> <li>✓ MCP Integration</li> <li>✓ MCP Demo</li> </ul>	<ul style="list-style-type: none"> <li>✓ EWS, OWS, EOWS</li> <li>✓ Bring-up, EVT, DVT</li> <li>✓ Reliability</li> <li>✓ MCP/System Integration</li> </ul> 	<ul style="list-style-type: none"> <li>✓ EWS, OWS, EOWS</li> <li>✓ Bring-up, EVT, DVT</li> <li>▪ Reliability/Qual</li> <li>▪ MCP/System Integration</li> </ul> 	
	 <p>2Tbps</p>  <p>Advanced Packaging/MCP</p>	 <p>4Tbps</p>  <p>MCP/System Integration</p>	 <p>8Tbps</p> 	

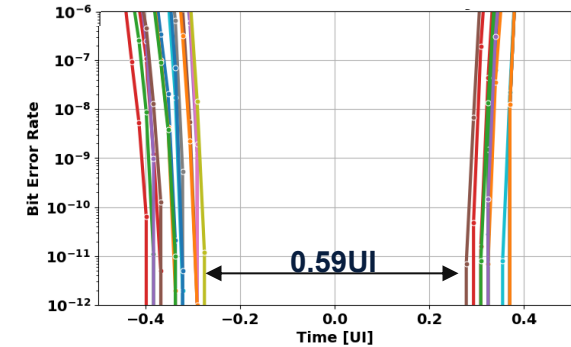
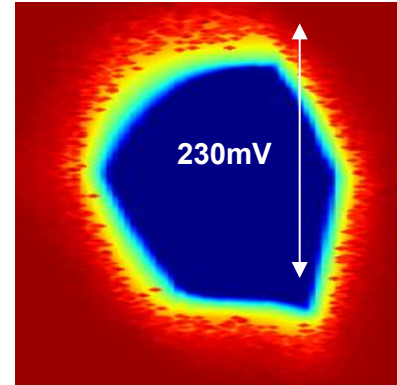
# UCle Link Test Results



Dual-Die Package for UCle Testing



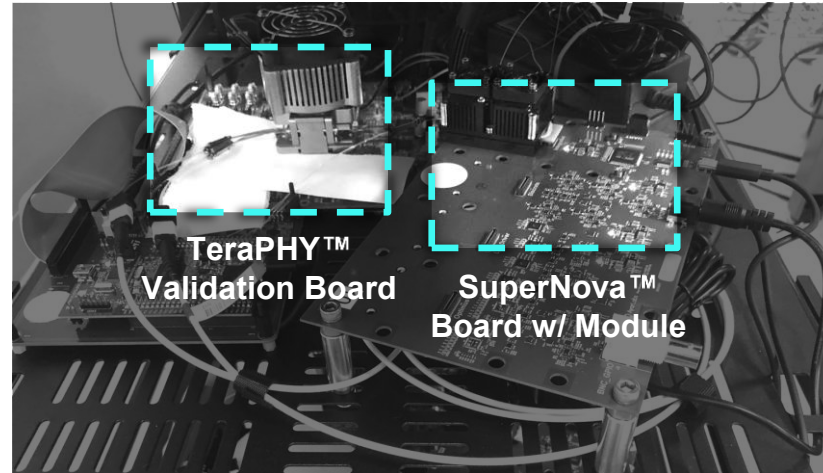

## UCle Eye Diagram + BER Bathtub



## ~5 Day Test, Error-Free Across All 16 UCle Modules

Accumulated Bits = 1.8002e+18;	BER = 0.0000e+00
Accumulated Bits = 1.8004e+18;	BER = 0.0000e+00
Accumulated Bits = 1.8007e+18;	BER = 0.0000e+00
Accumulated Bits = 1.8009e+18;	BER = 0.0000e+00
Accumulated Bits = 1.8012e+18;	BER = 0.0000e+00
Accumulated Bits = 1.8014e+18;	BER = 0.0000e+00
Accumulated Bits = 1.8017e+18;	BER = 0.0000e+00
Accumulated Bits = 1.8019e+18;	BER = 0.0000e+00

# EVT Link Margin Validation



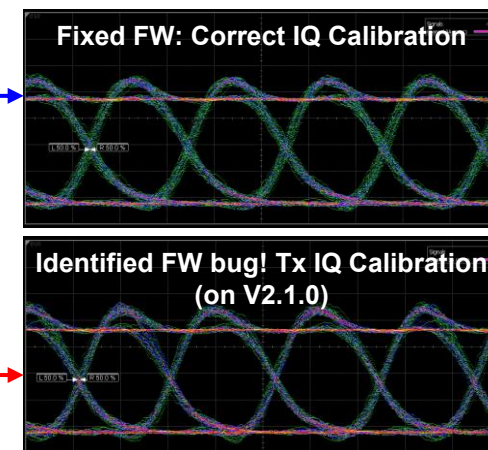
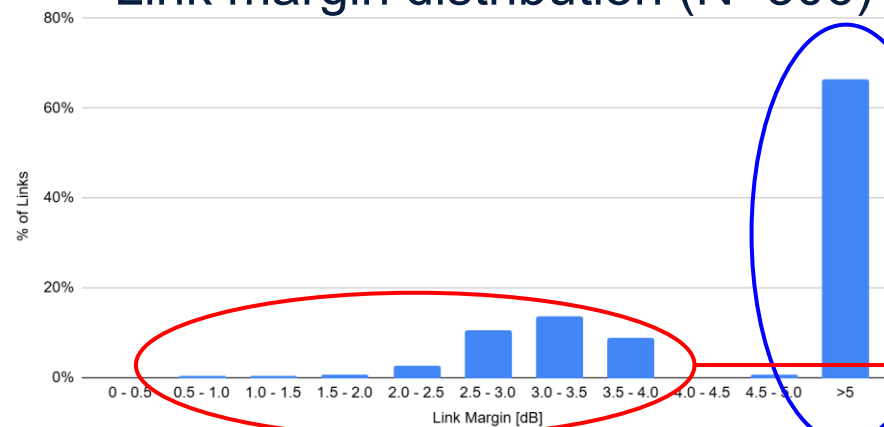
Links setup for cross-port loopback

## Using SuperNova™ for 16λ laser source

- Lowered laser power by 2.5dB (to 2.5dBm) for EVT link test
- Using on-chip VOA to sweep input power
- PRBS31 in all link tests

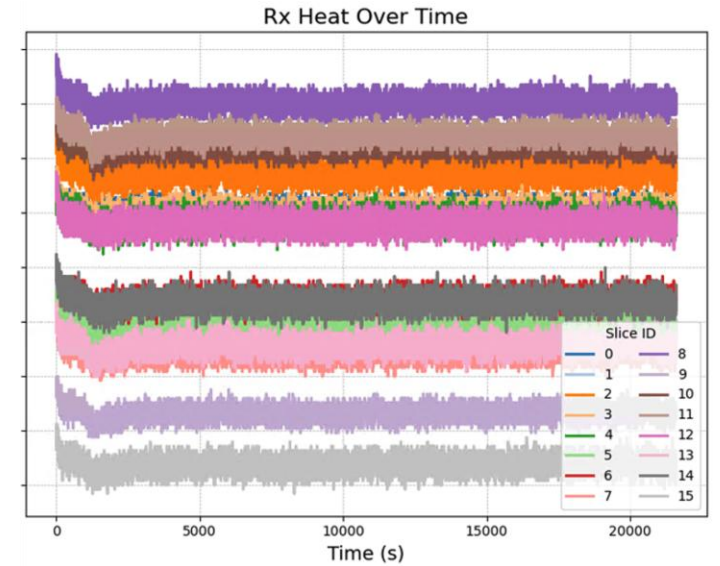
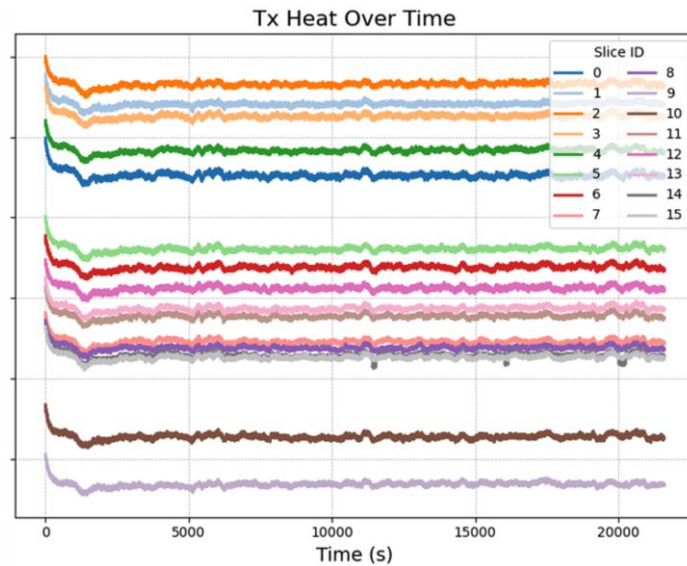
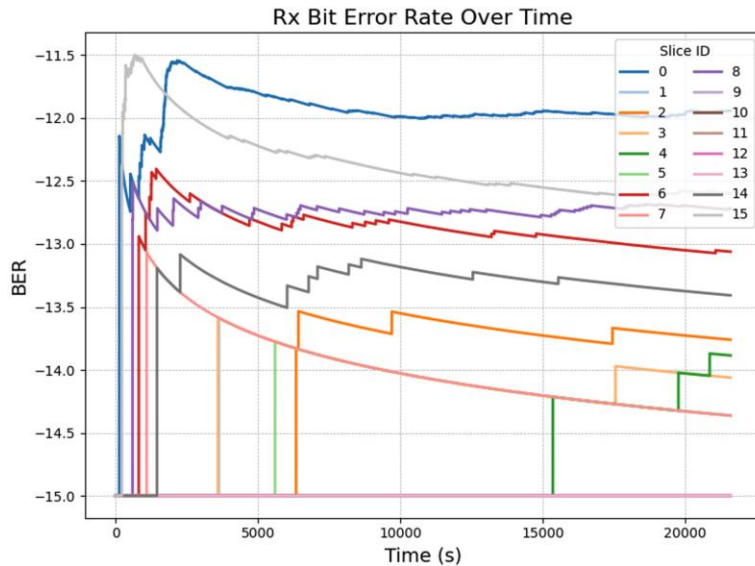
Links have 5+dB of margin

## Link margin distribution (N=306)



# System EVT: Long-Term Link Stability

Long-term link stability tests @5dB of attenuation (to force bit-errors)

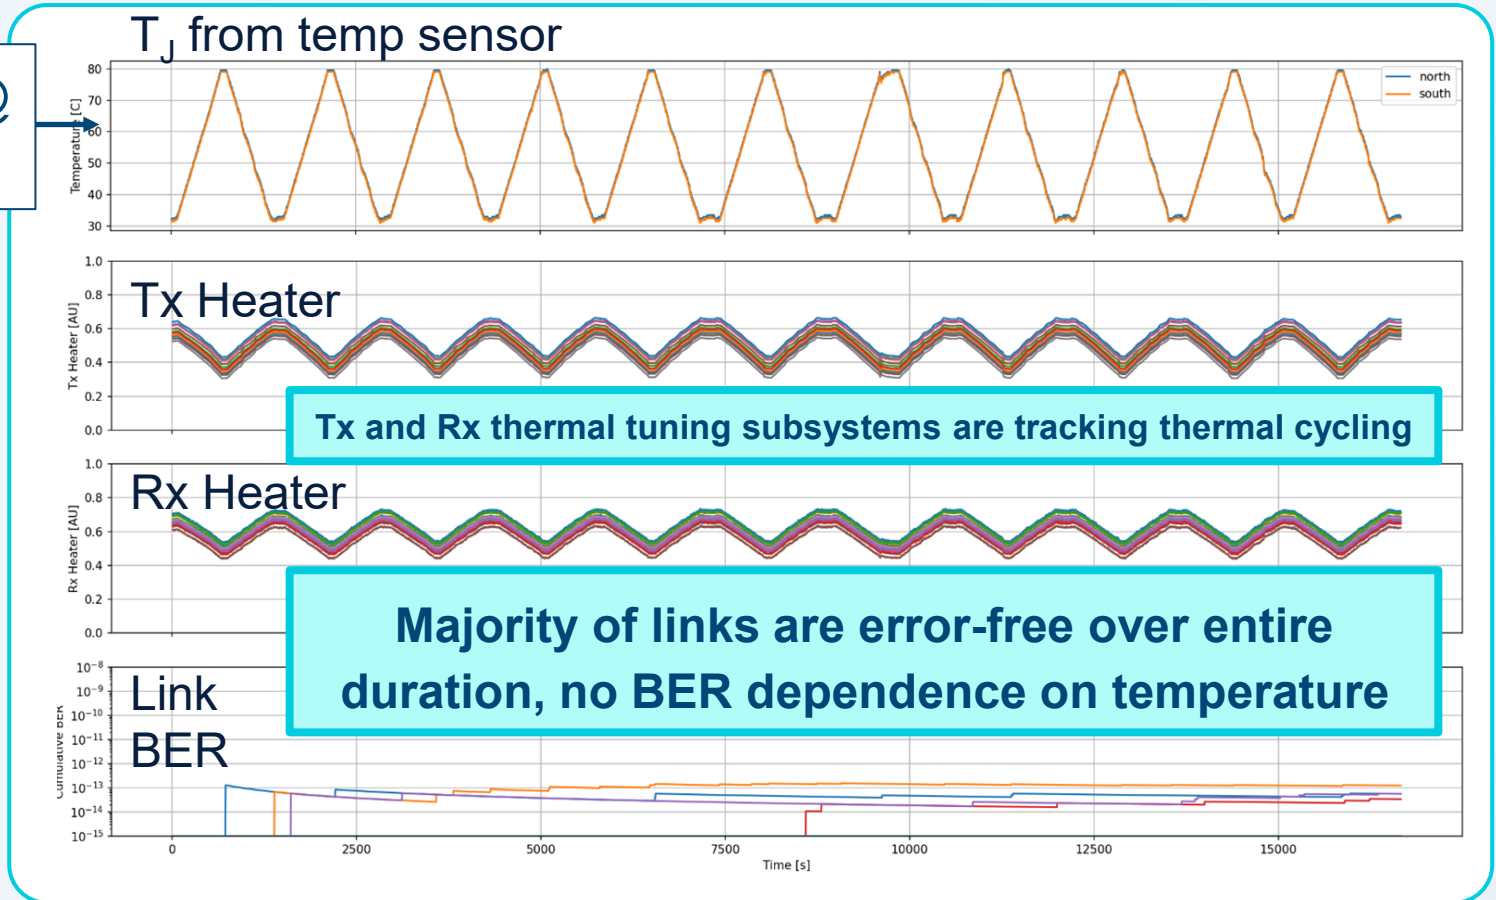

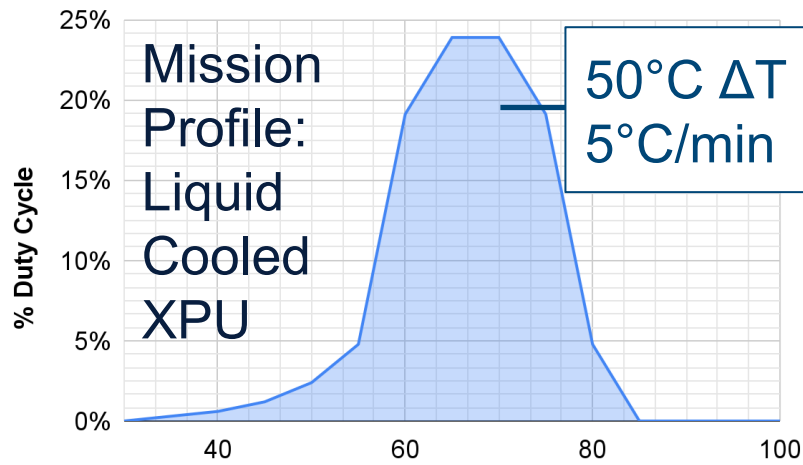


**Applies 5dB of attenuation to force bit-errors on links**

- Checks consistent BER over time
- Note: some links remain error-free during this test
- Baseline test for other longer-running system tests

**All links maintained stable BER over >6 hours**



# System EVT: Thermal Cycling Link Test



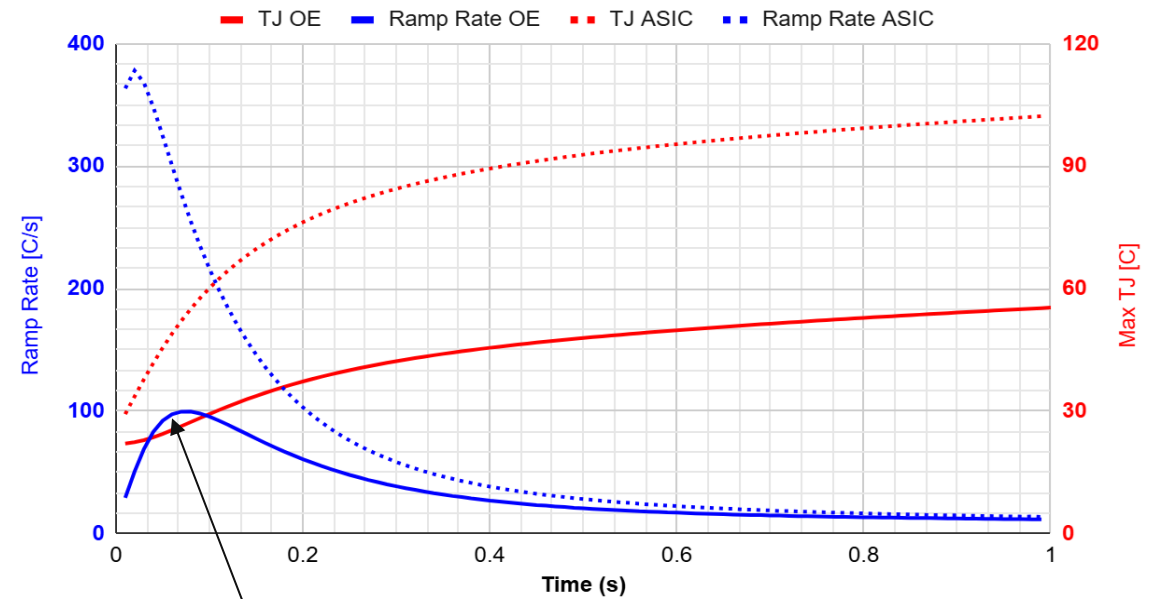
Thermal cycling test using FlexTC temperature controller mimicking system thermal cycling behavior of liquid cooled XPU

- Slow ramp test over extended duration – 5°C/min

# Application Study: Co-packaged 500W ASIC Link Thermal Transient



ASIC has  
500W power  
step at  $t=0s$



Link on OE can briefly see a thermal transient of  $\sim 100^{\circ}\text{C/s}$  over a few degrees

# Application Study: Links up to 800°C/s (Emulated)



Emulate fast temperature ramps by sweeping laser  $\lambda$

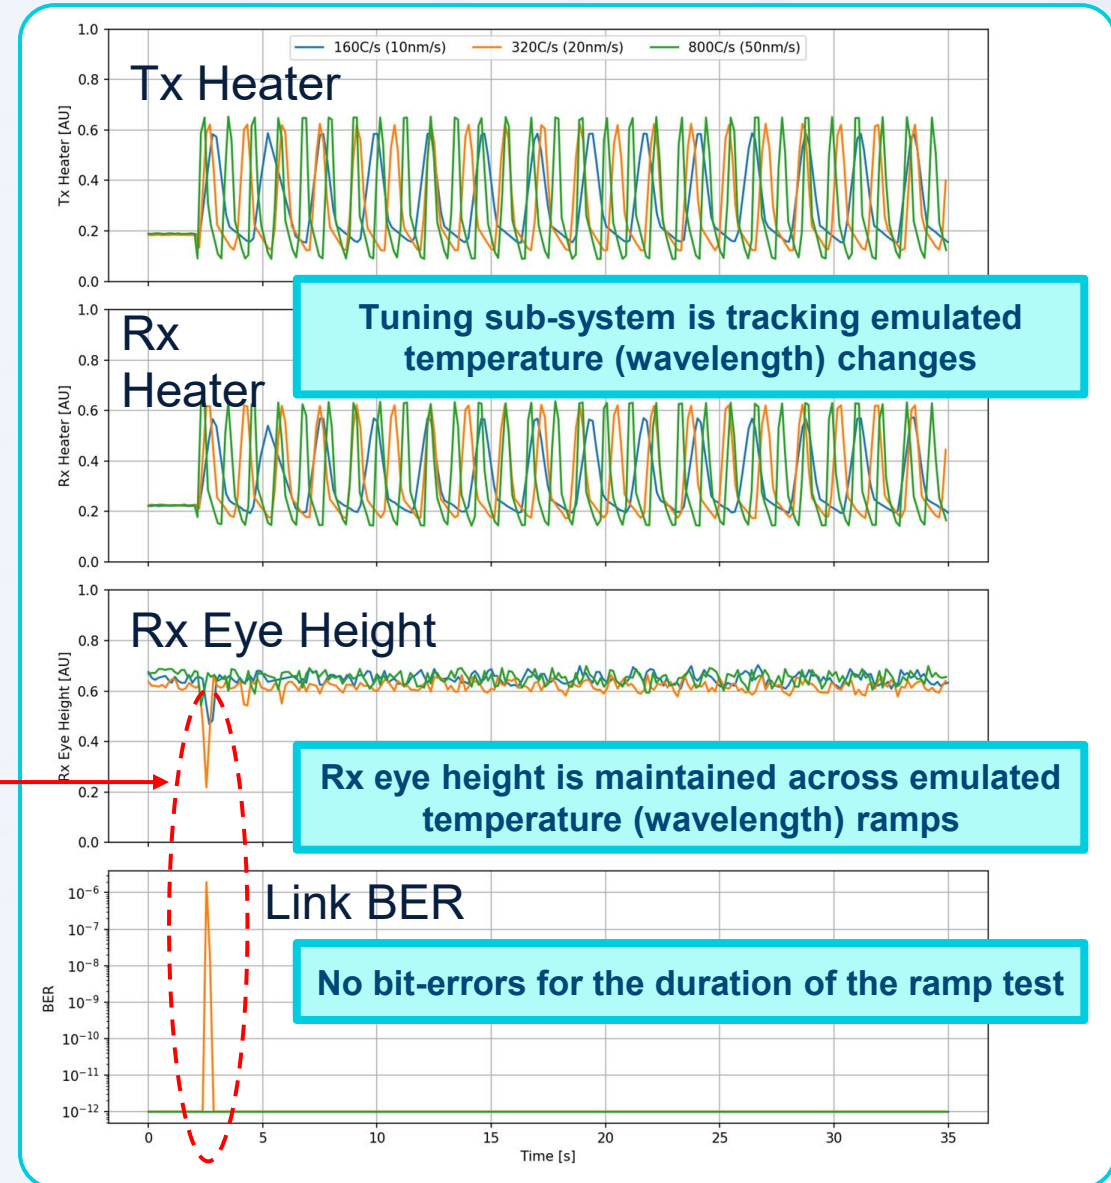
- Sweep over 4 nm (64°C) in a sawtooth pattern

Sweep Rate	Sweep Range	Equivalent * $\Delta T$	Sweep Time	Equivalent $\Delta T/s$
10 nm/s	4 nm	64°C	0.4 s	160 °C/s
20 nm/s	4 nm	64°C	0.2 s	320 °C/s
50 nm/s	4 nm	64°C	0.08 s	800 °C/s

\*1 nm change = 16°C change ( $\Delta T$ )


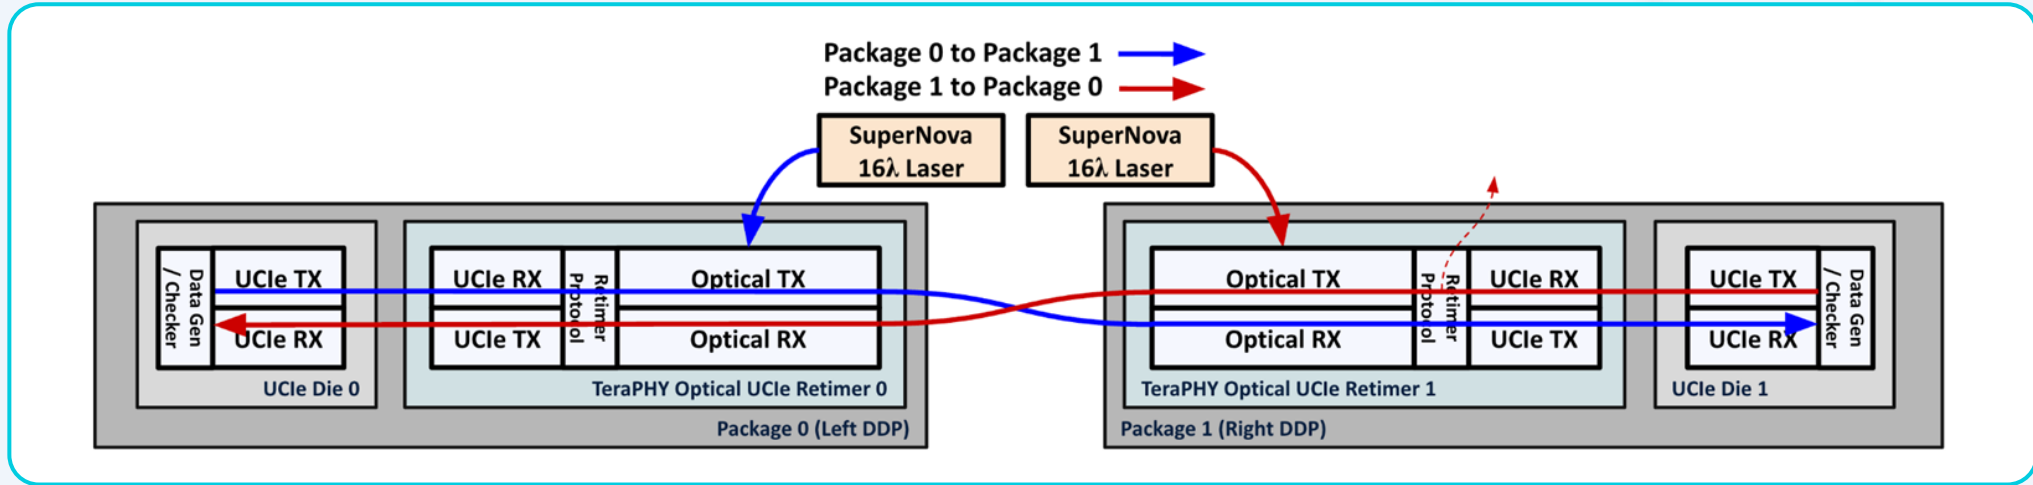
50 nm/s max laser sweep rate

Power glitch from tunable laser at start of sweep

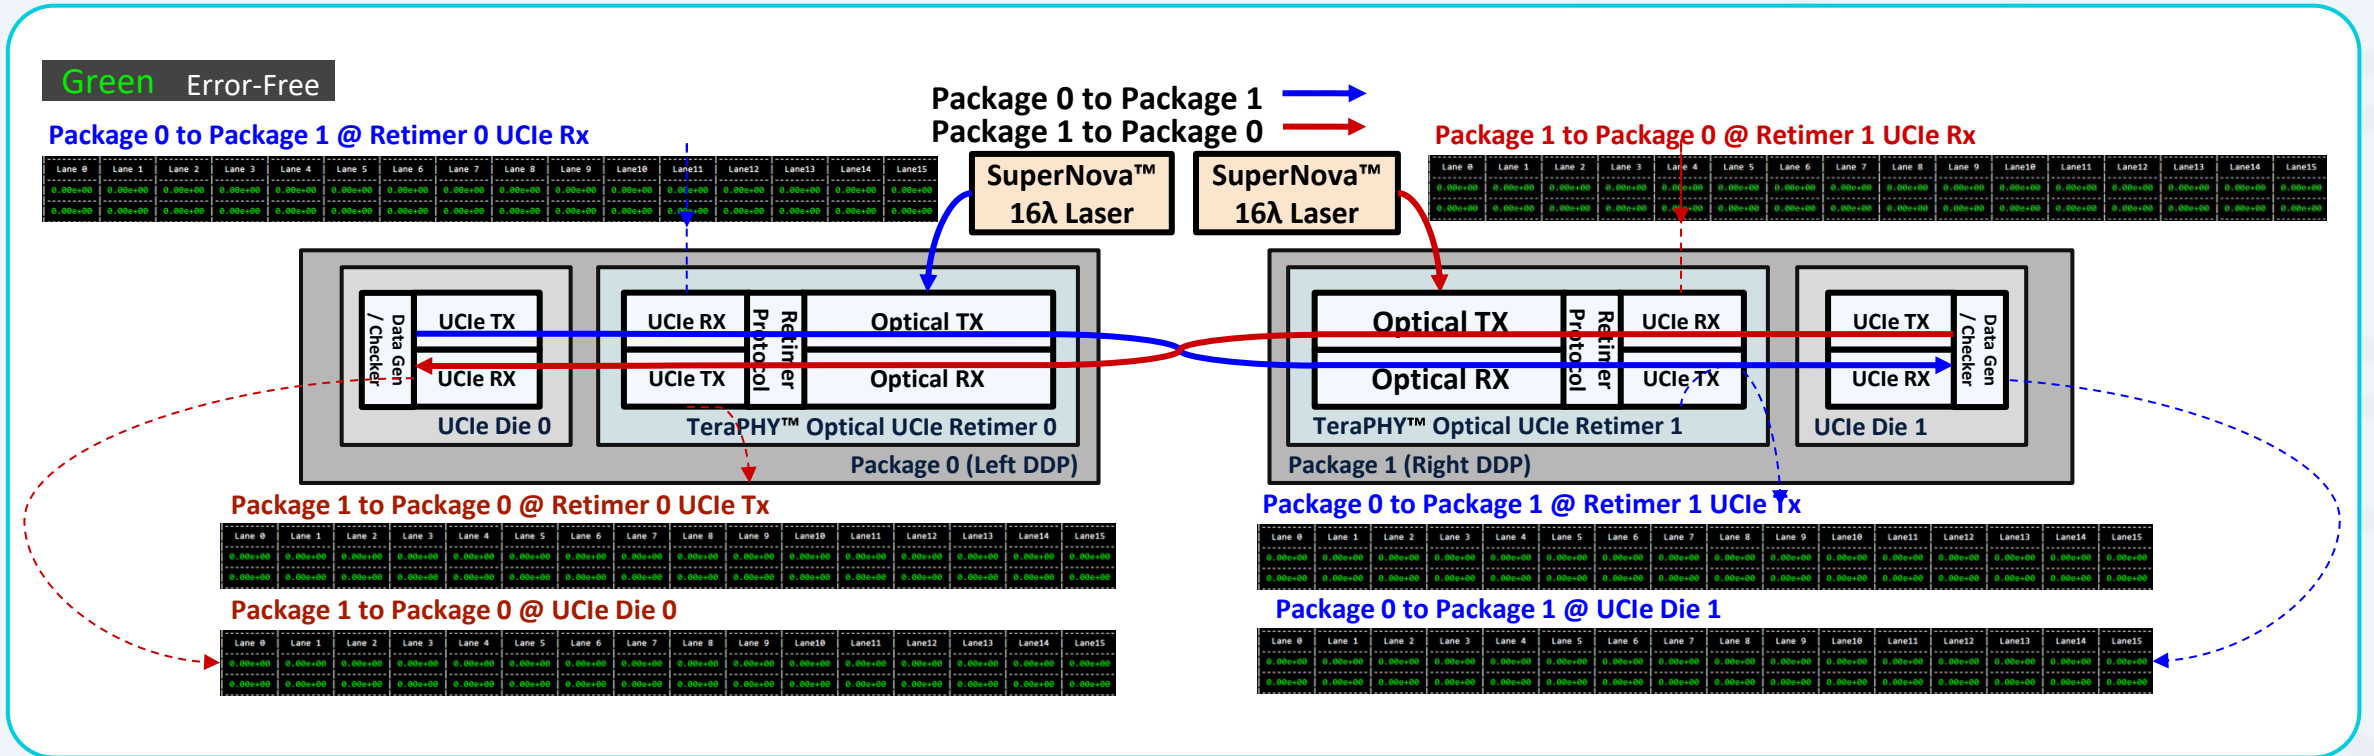


**No bit-errors at up to 800°C/s emulated  $\Delta T$ . Tuning subsystem tracks shifts, maintains Rx eye height**

# UCle over Optics End-To-End Link Test Setup

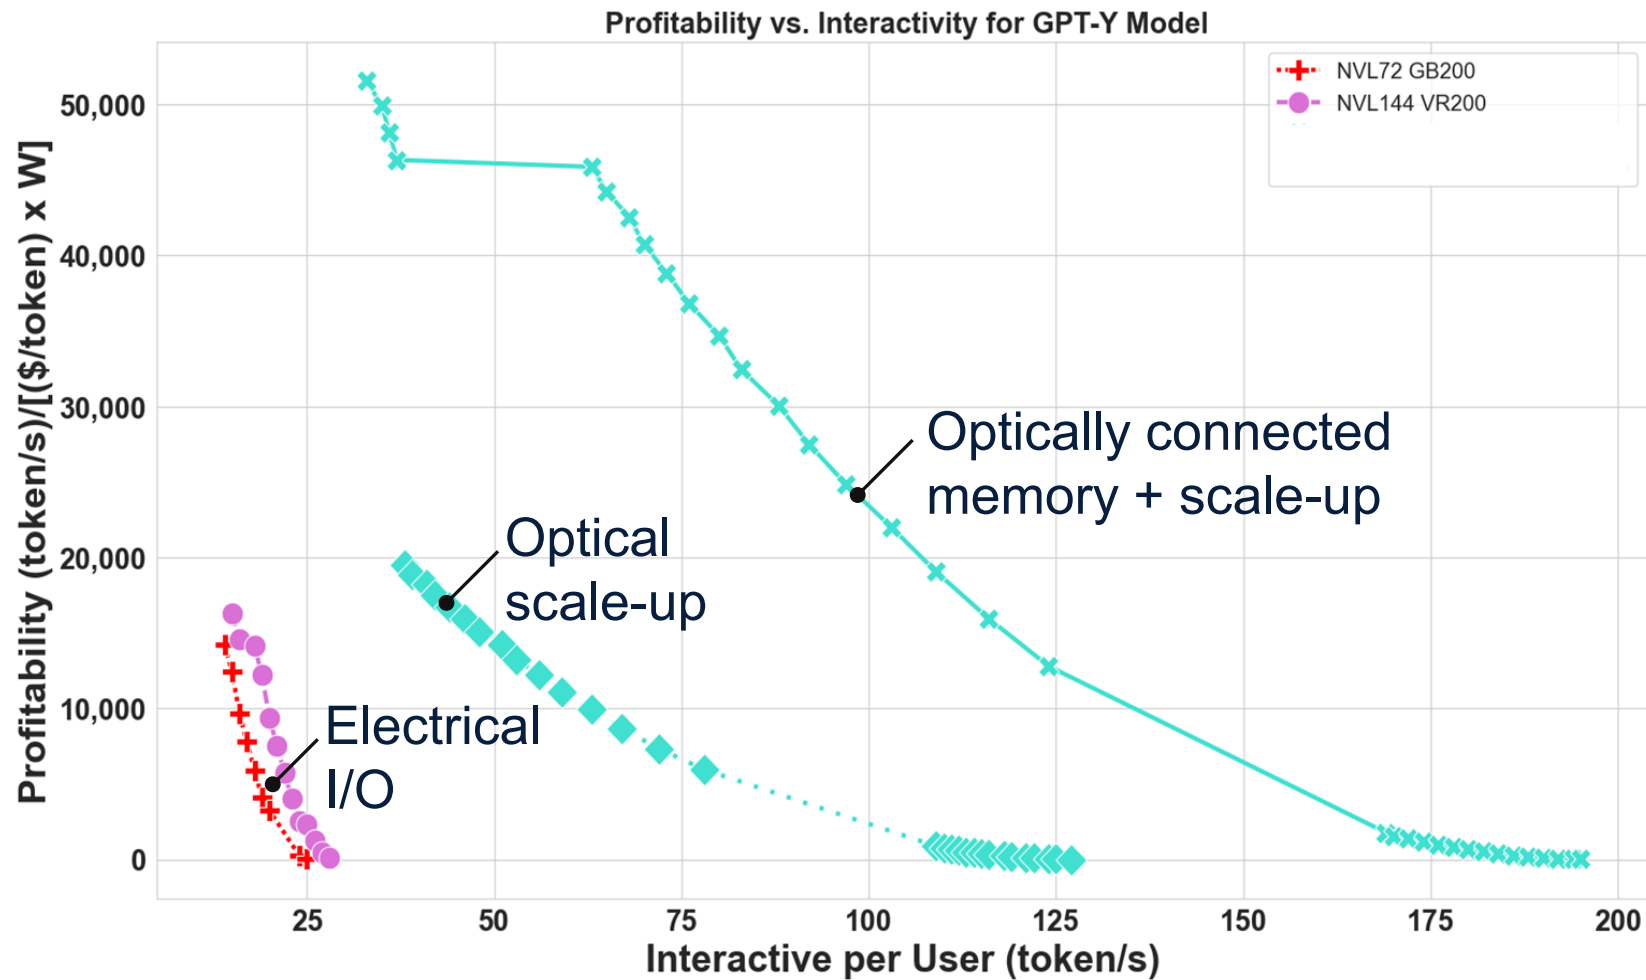


# UCle over Optics End-To-End Link Test Result



- Full-duplex error-free BER observed at every test point over >10 hours
- End-to-end latency < 25ns

# Application-level Impact of Optical I/O



# Conclusions

- Large-scale AI systems are a connectivity problem
  - Limited by copper reach and high-power/high-cost of pluggable optics
  - Exponential growth in rack power density
- Ayar Labs' 8Tbps TeraPHY™ chiplet is the first demonstration of UCle traffic over optics
  - Fully integrated optical retimer chiplet enables co-packaging with AI compute hardware, breaking the scale-up interconnect bottleneck
- Microrings work in the TeraPHY™ chiplet production validation and are ready for the HVM ramp

