Addressing Better Optical Connections in AI Networks

为AI网络提供更优的光互连方案

Dec. 8, 2023

Confidential

AI Model Sizes are Growing 10x Annually
大模型规模每年扩增10倍

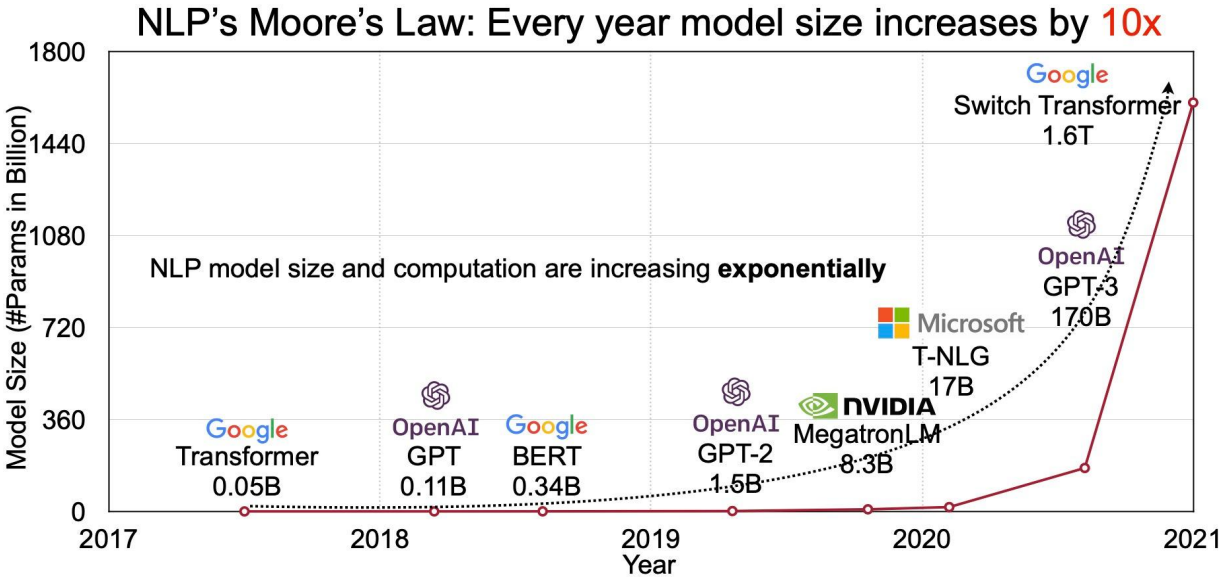

GPT3 Example

- 50,257 word vocabulary
- 2,048 word sequence length
- 175B Parameters : >1TB to store model
- 300B tokens in training data set
- Training required 10,000 NVIDIA V100 GPUs for 1 month in 2021
- Power ~4.6MW

Microsoft's Zaid Khan, GM Cloud AI, Apr, 2023 "…we're now training models on 75MW"

CREDO

Introducing the Backend Network
AI后端网络

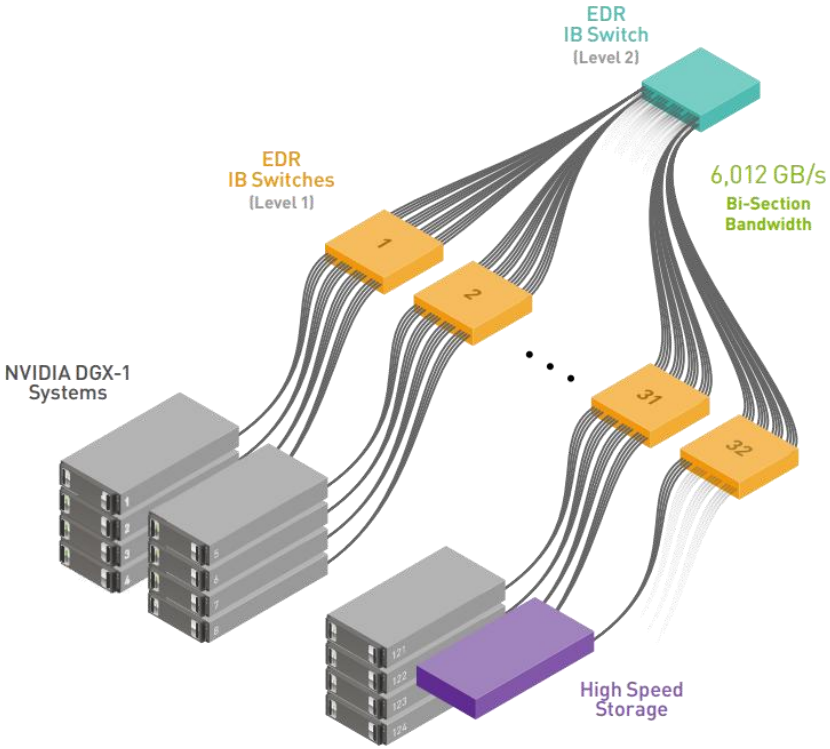

- Unlike traditional general compute networks, AI clusters are built with two separate networks
  - The front-end network is used for data I/O to the cluster
  - The back-end network creates a communication fabric between all GPUs
- The back-end network can be 10-20x more dense than the front end network
- The result is a huge number of new optical connections at 400G or 800G
- Given the huge volume of these optical components, improvements in energy efficiency are essential as cluster sizes continue to expand

CREDO

Credo: Addressing Every High-Speed Connection in the HSDC

Credo致力于高速互连解决方案

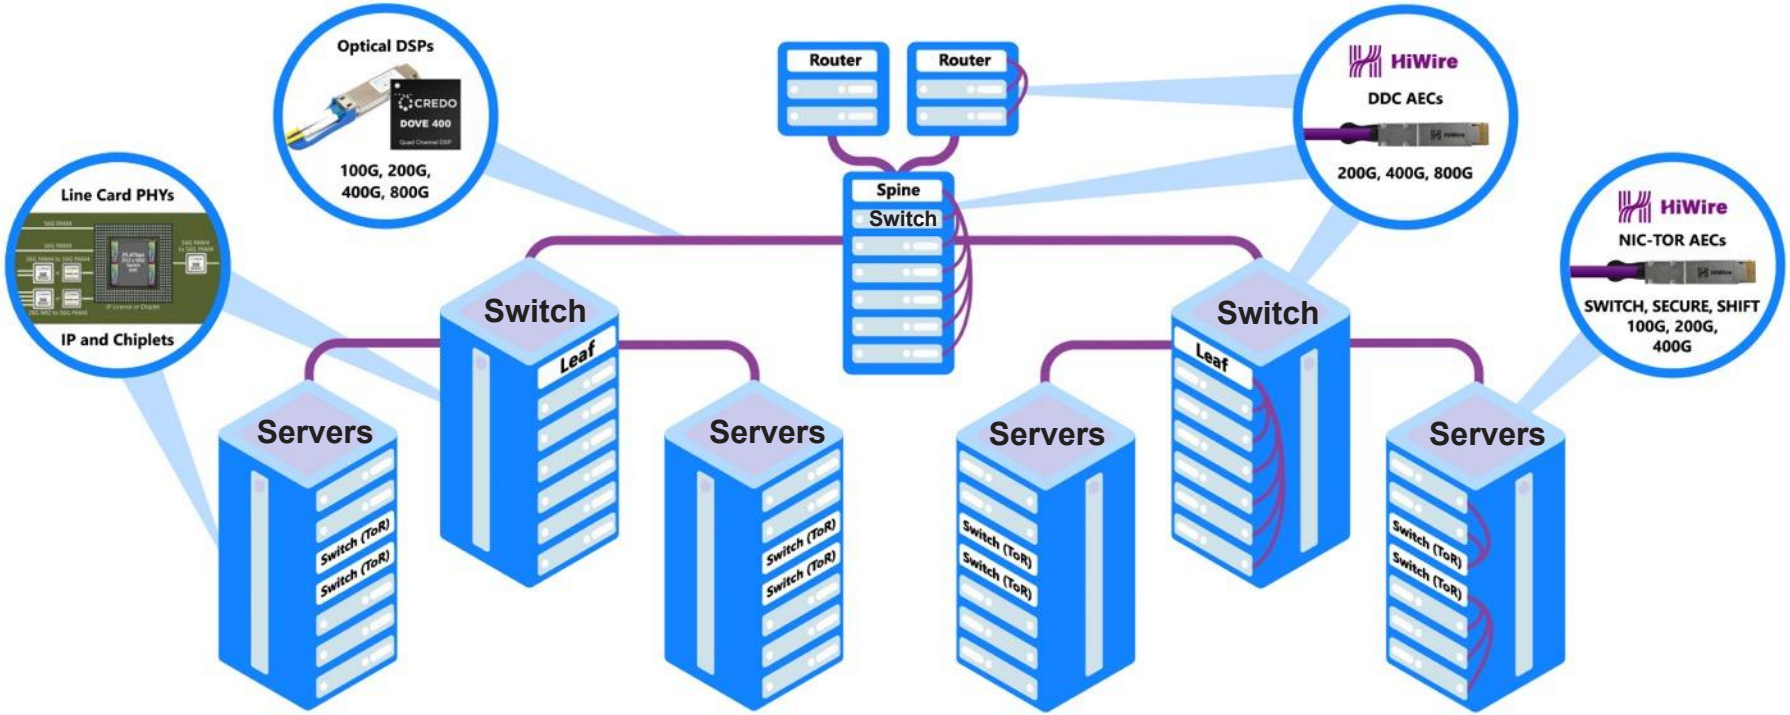

Confidential

CREDO

Credo Solutions
Credo解决方案

Line Card PHYs

Optical DSPs

DDC and TOR-to-NIC HiWire AECs

Product Solutions

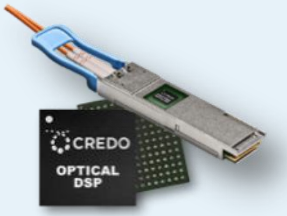

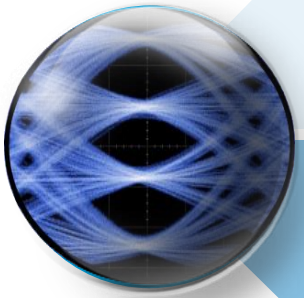

IP and Chiplet Solutions

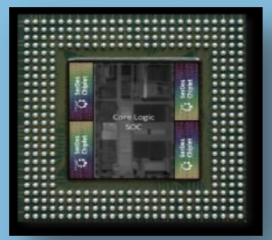

SerDes IP Licensing

2.5D and MCM SerDes Chiplets

CREDO

Core Technology Drives Competitive Advantage
核心技术优势

CREDO

LPO Challenges and a Unique Solution
LPO的挑战以及新方案

- The LPO story is appealing

……but there are many challenges

A Tx DSP solves the performance problems
1. Optical standard compliance
2. Eliminates manual tuning
3. Bounds the range of signals at the far end Rx

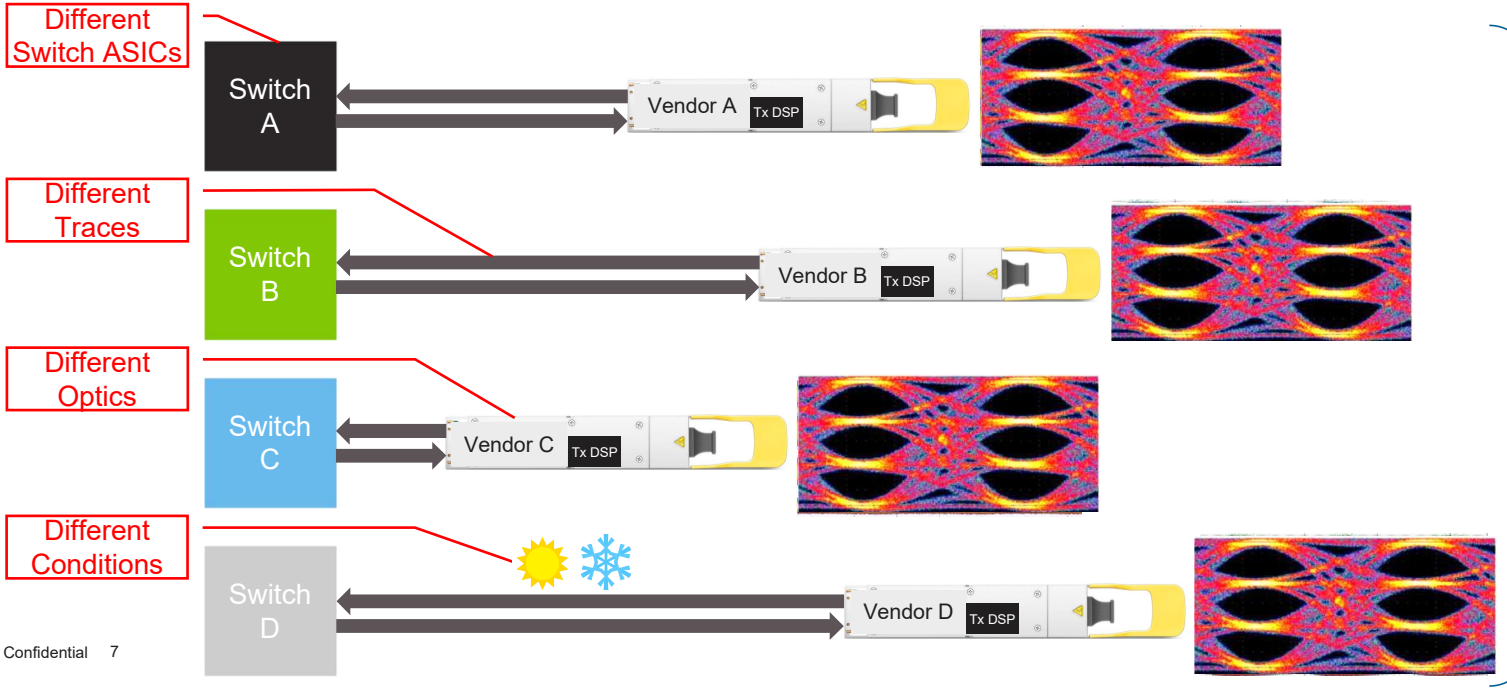

Introducing Linear Receive Optics (LRO)
介绍LRO（线性接收光连接）

- **LRO** = Linear Receive Optics
- Optimized DSP Tx lanes only (no Rx)
- Suitable for MMF and SMF designs
- LRO advantages
  - Decouples optical performance from the host
  - IEEE compliance at TP2 and TP3
  - Factory optical calibration (no field tuning)
  - Adaptive temperature compensation
  - Additional diagnostic capabilities
  - Allows for any type of laser (EML, VCSEL, etc)
- Facilitates multi-vendor interop and high volume deployment

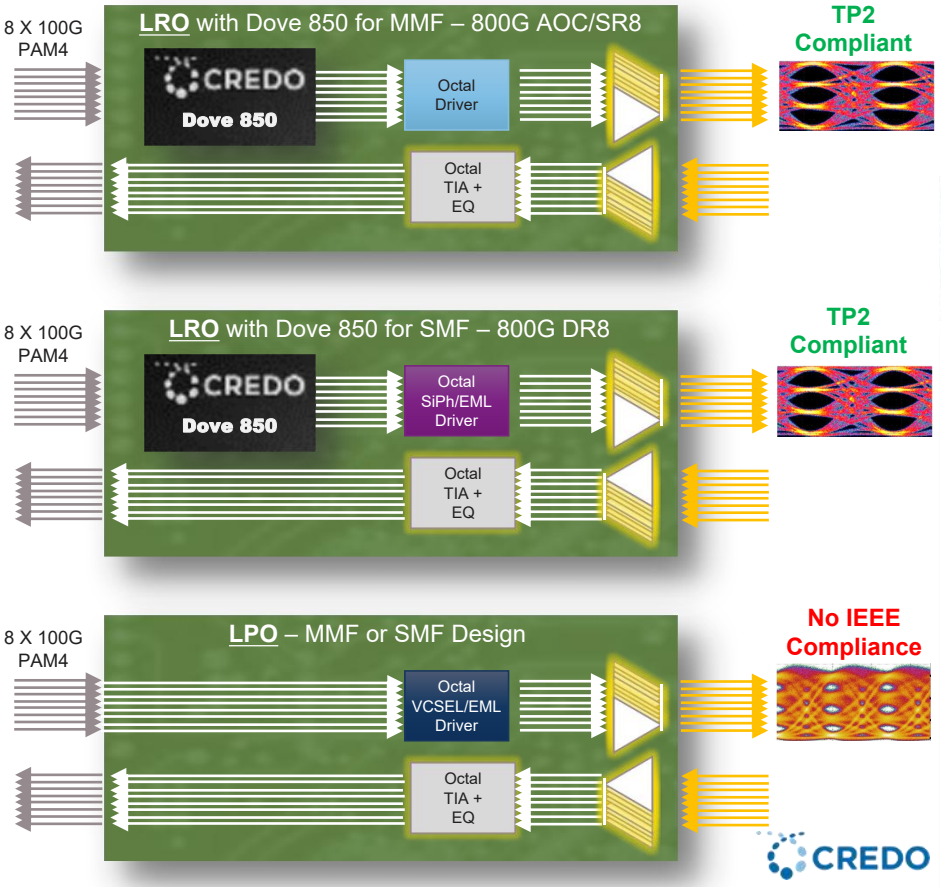

Confidential

LRO Compared with Full 800G DSP Solutions

LRO方案与 DSP方案的对比

- Competitor's 7nm and 5nm full 800G DSPs
- Dove 850 LRO DSP Design versus 7nm DSP
  - **50% power savings**
  - **50%+ cost savings**
- Dove 850 LRO DSP Design versus 5nm DSP
  - **30% power savings**
  - **65%+ cost savings**
- Compelling power and cost savings with LRO

800G Application	Full 7nm DSP	Full 5nm DSP	Dove 850 LRO DSP
Octal TIA	☑	☑	☑
Octal Driver	☑	☑	☑
MMF or SMF Optics	☑	☑	☑
DSP Power	100%	75%	**50%**
DSP Price	X	1.2X	**0.5X**

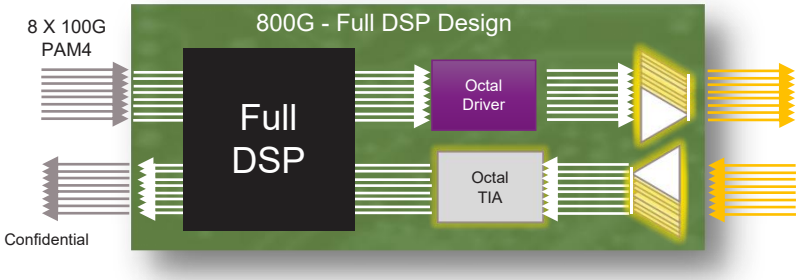

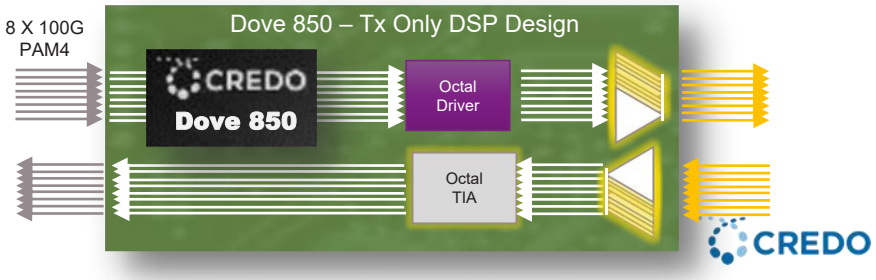

Confidential

LRO versus LPO Performance Evaluation Setup
LRO和LPO 性能对比

- All modules are 800G DR8 using Silicon Photonics
- Host PCB traces are varied to simulate different switch ports

- **Test Setup #1**: LPO on the Tx and LPO on the Rx

- **Test Setup #2**: Full DSP on the Tx and LPO on the Rx (simulates LRO case)

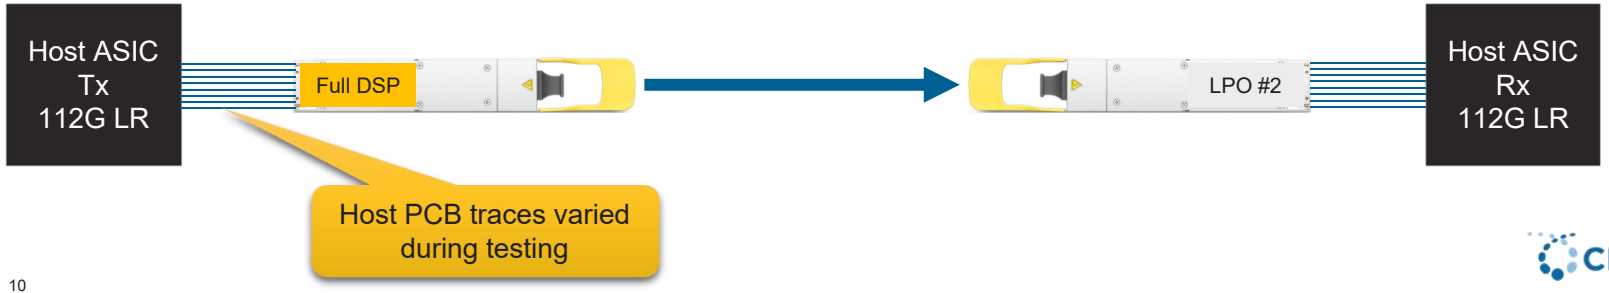

LRO versus LPO Performance Results
LRO和LPO 性能对比

Host Trace	Pre FEC BER
6.1dB	2.0e-9
9.3 dB	2.0e-7
10.1 dB	1.2e-6
12.0 dB	1.5e-3
14.1 dB	N/A (no link)
16.2 dB	N/A (no link)

- **Test Setup #1**: LPO on the Tx and LPO on the Rx

Host Trace	Pre FEC BER
6.1dB	2.6e-10
9.3 dB	1.4e-9
10.1 dB	7.3e-9
12.0 dB	9.0e-9
14.1 dB	9.6e-9
16.2 dB	2.0e-8

- **Test Setup #2**: Full DSP on the Tx and LPO on the Rx (LRO case)

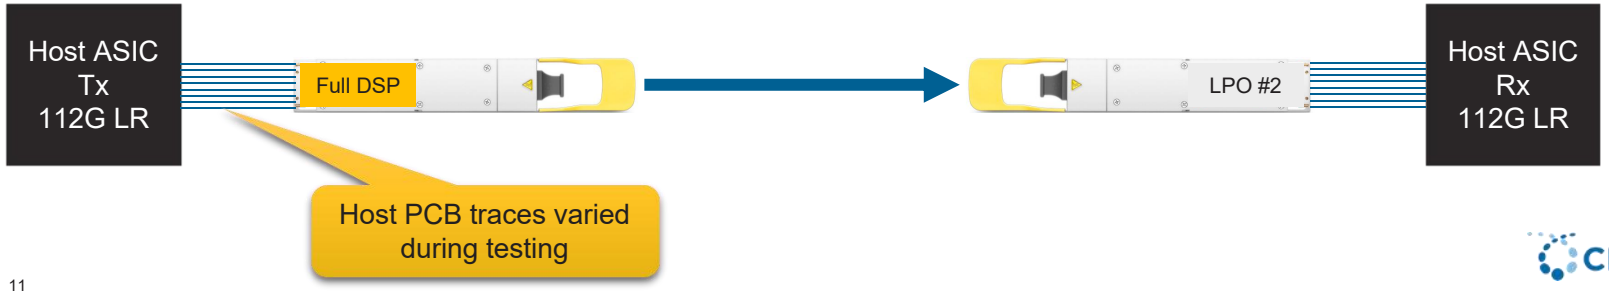

The LPO Latency Myth
LPO 低时延“神话”

- Supporters of LPO claim a major reduction in system latency
- The fact is, it makes very little overall difference
- Every 400G/800G Ethernet or Infiniband system uses FEC
- Fiber length + FEC are the dominant drivers of system latency
- The latency difference for a 100m link is only 3%
  - The longer the fiber, the smaller the transceiver impact
  - Looking at the LPO latency on its own is meaningless!

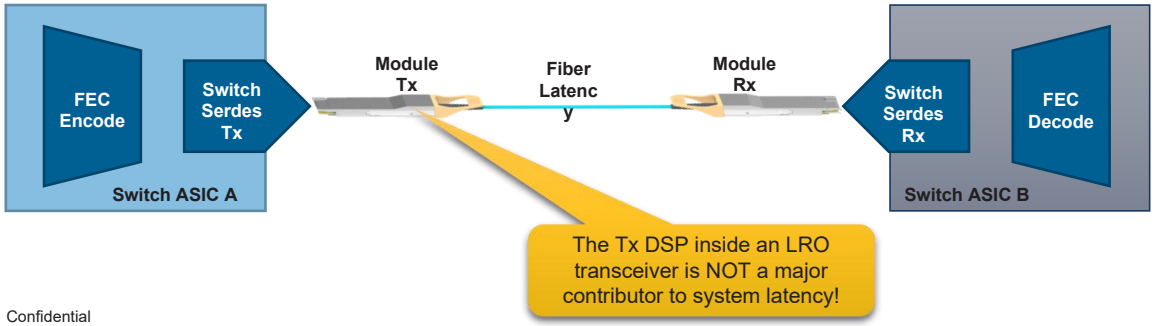

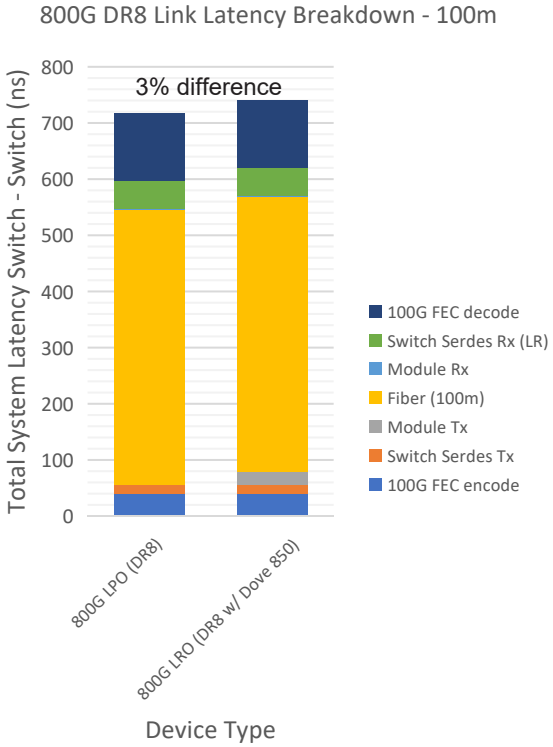

Confidential

CREDO

LRO versus LPO Summary
LRO和LPO 对比总结

- Linear Pluggable Optics (LPO) were proposed to reduce the power and cost of pluggable optics
- Eliminating a DSP means no adaptive compensation for:
  - Various host ASICs
  - Host PCB trace lengths
  - Different optical components
  - Different link partners
  - Environmental conditions
- Interop and volume deployments are infeasible
- Lower power comes with a performance and OPEX cost

- Linear Receive Optics (LRO) are an alternative approach addressing the short-comings of LPO
- Maintaining a DSP in the transmit path only:
  - Decouples optical performance from the host
  - Ensures IEEE compliant optical signals
  - Enables factory optical calibration (no field tuning)
  - Provides automatic temperature compensation
  - Allows for any type of laser (EML, VCSEL, etc)
- The Dove 850 is the world's first DSP for LRO.
- Save power without sacrificing performance and OPEX

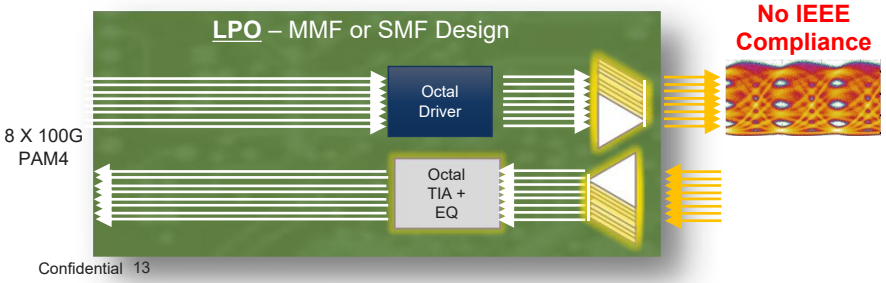

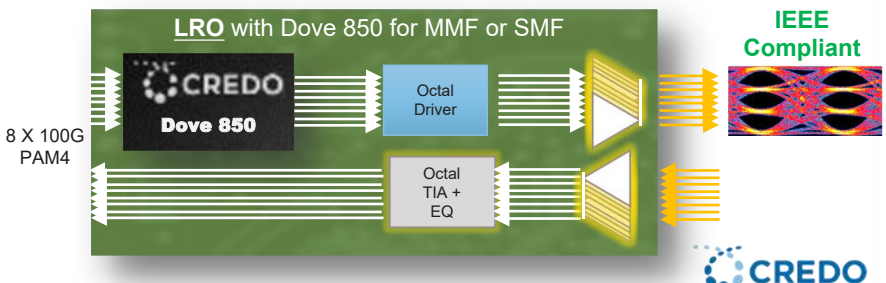

The Path Forward Better Connections in AI Networks
更优的AI网络互连

- As the demands for AI increase, the industry call for lower power networking must be addressed
- LRO will answer this call by saving 50% of the DSP power in every module
- High performance makes LRO a suitable choice for demanding, high-reliability applications
- Cost savings are an attractive side benefit
- LRO is suitable for 400 Gb/s and 800 Gb/s connections today
- LRO has a very clear path to 1.6 Tb/s (with 224G/lane) tomorrow
- The Credo Dove 850 is the first available DSP completely optimized for LRO applications

Confidential