Tecnologias avançadas de interconexão

Redefinindo a escalabilidade e eficiência nos data centers

Décio Coraça – Systems Engineering Director

March 2026

The Real Question: What problem are we solving?

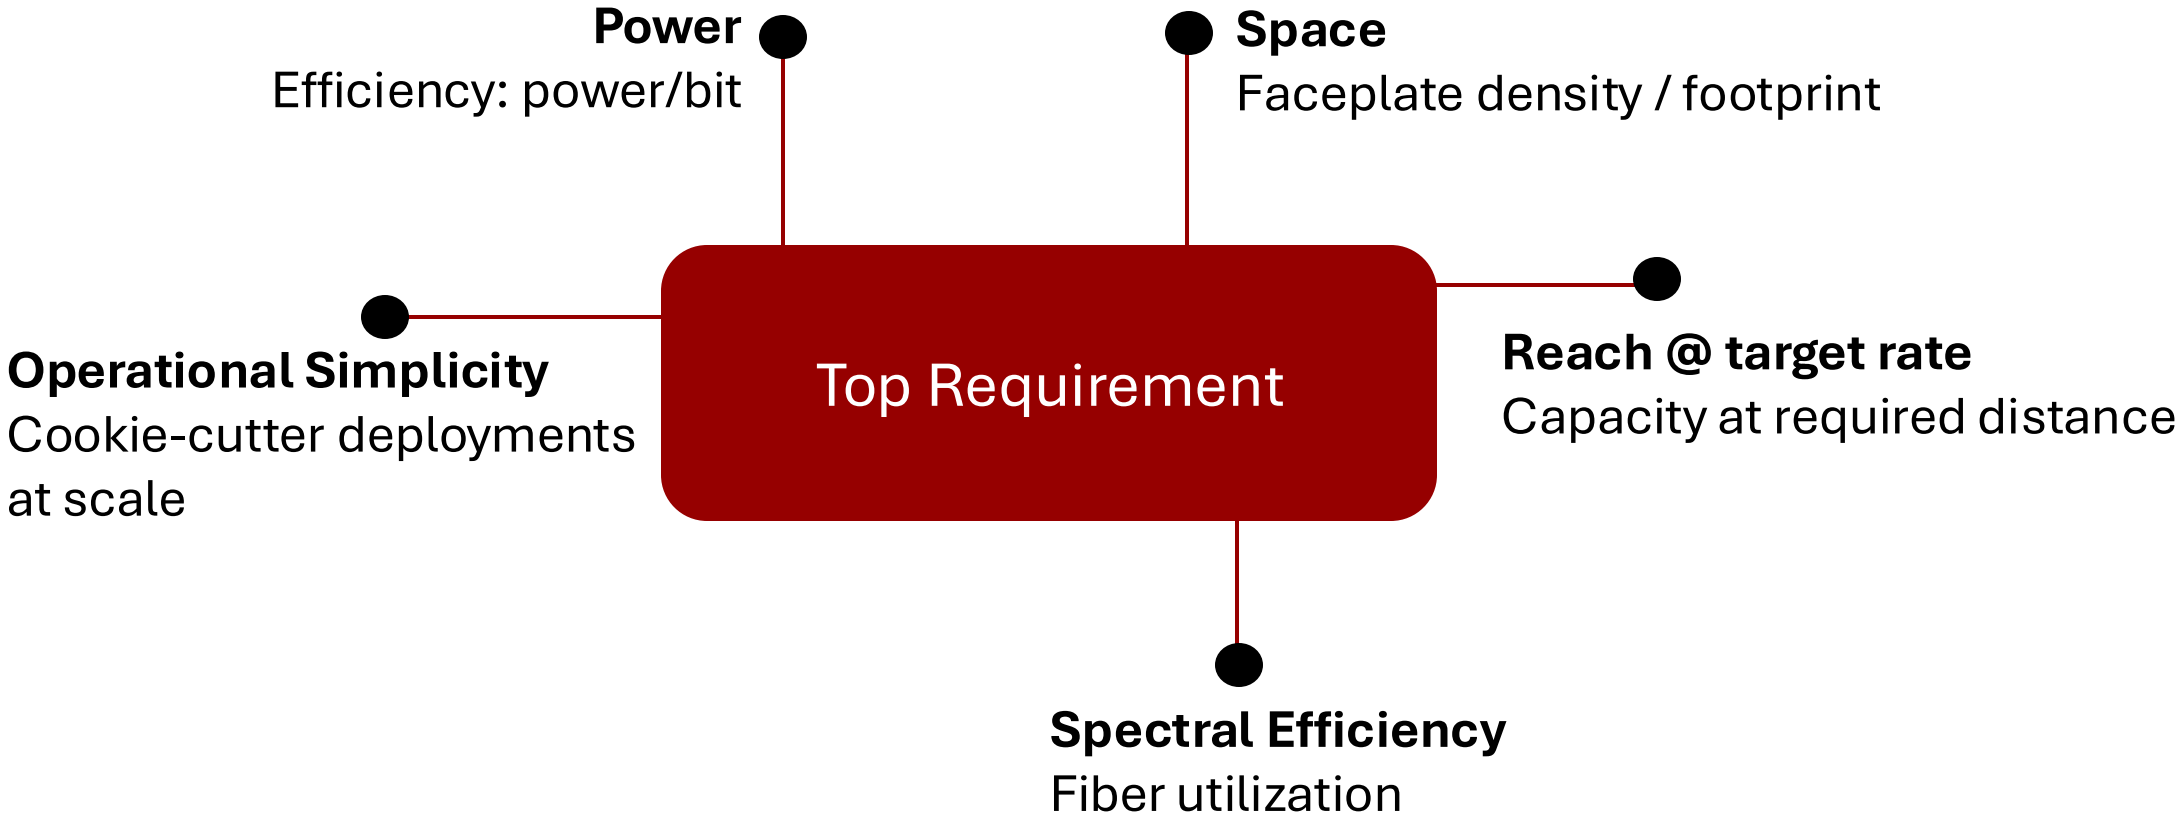

Explosive compute demand → explosive demand for power

"Deep learning worked, got predictably better with scale, and we dedicated increasing resources to it." **Sam Altman,** 09/23/2024

Massive implications on connectivity

DeepSeek R1
- 5.2 e24 FLOP
- 670 billion parameters
- $6.8M

Gemini 1.0 Ultra:
- 5e+25 FLOP
- 100 days
- $30M

Llama3.1:
- 16 trillion datapoints
- 3.8e+25 FLOP
- 16,000 H100 GPUs
- **27 MW power***
- 89 days

GPT 4:
- 4.9 trillion datapoints
- 2.1e+25 FLOP
- 25,000 NVIDIA A100 GPUs
- **10 MW POWER[2]**
- 95 days
- $41M

Epoch AI, 'Data on Notable AI Models'. Published online at epochai.org. Retrieved from 'https://epochai.org/data/notable-ai-models' [online resource]. Accessed Feb 2025.

*equivalent to 23,000 US residences
[2] Estimate based on 25000 GPUs, 400w/GPU

AI compute growth will require massive power

Projected power growth for frontier AI training

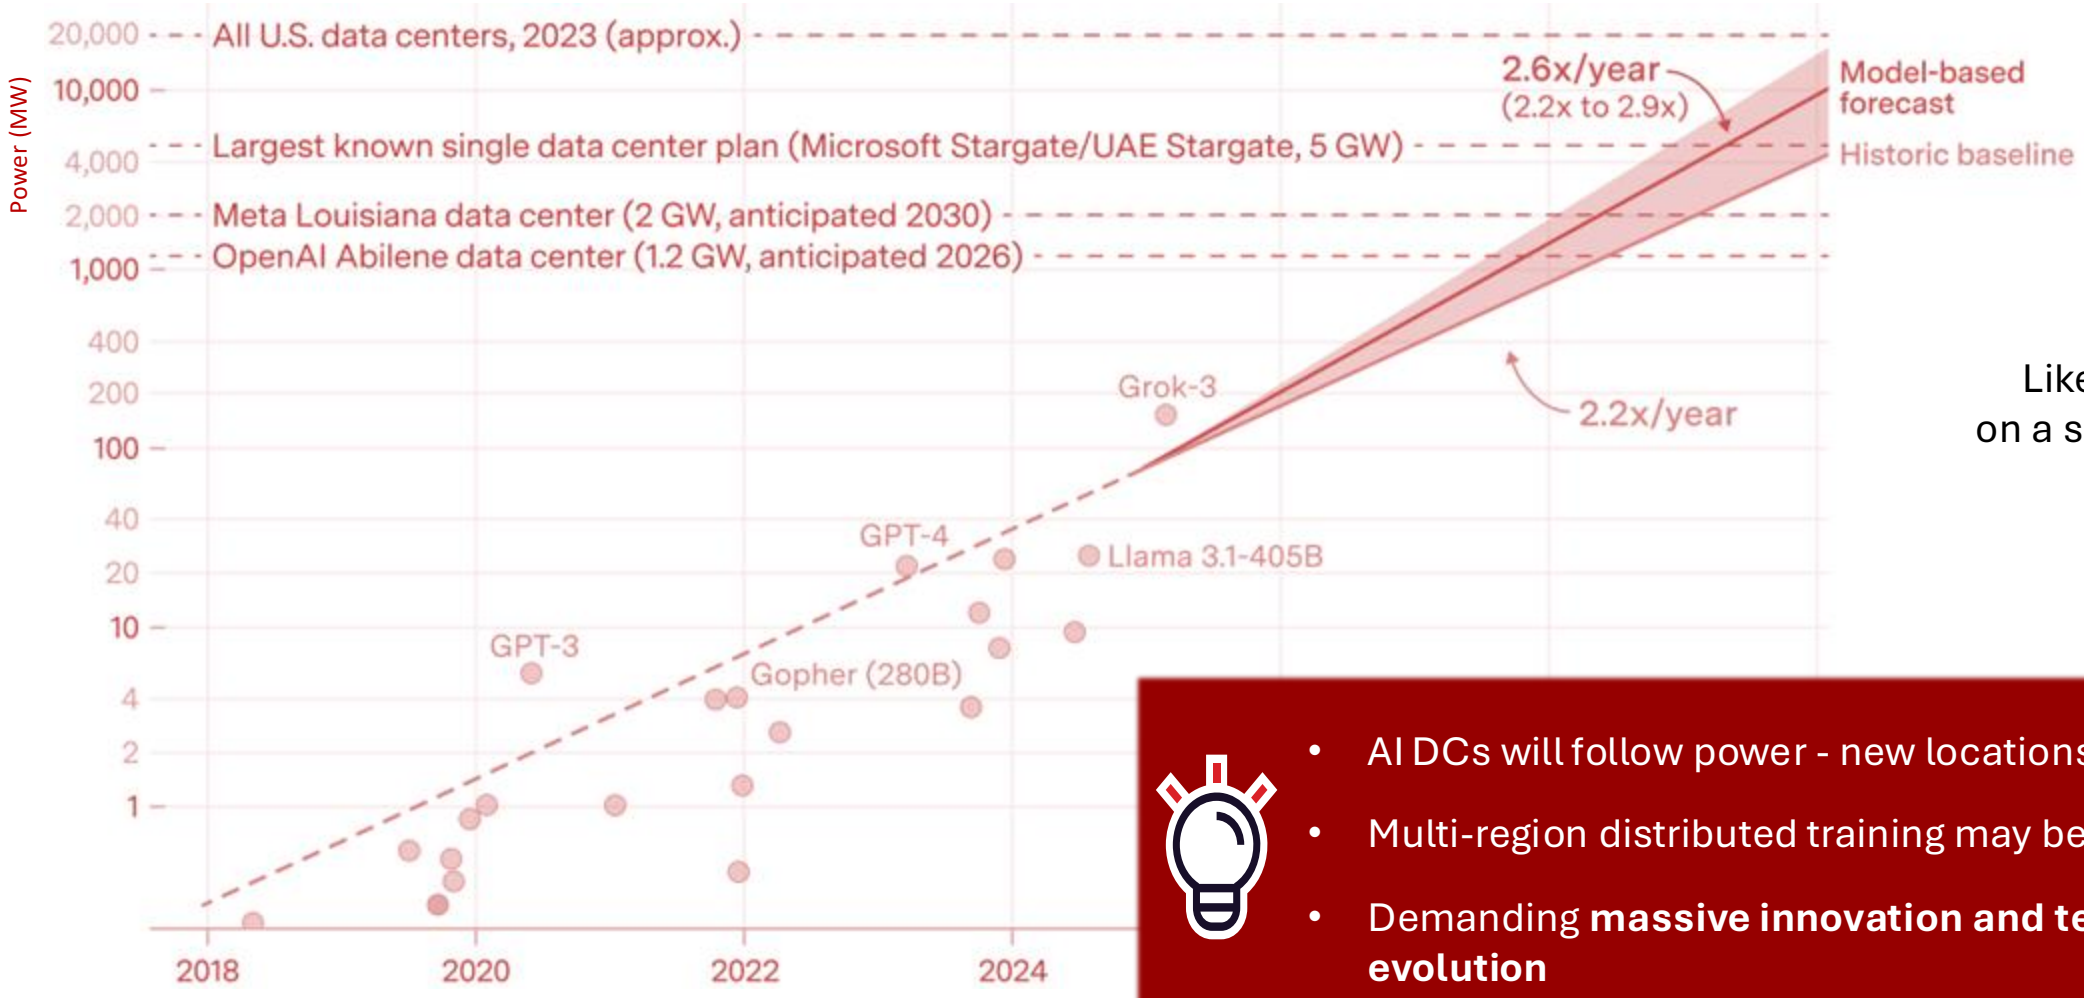

2030:
5-15GW

Likely not feasible
on a single DC campus

- AI DCs will follow power - new locations to connect
- Multi-region distributed training may be necessary
- Demanding **massive innovation and technological evolution**

Source: Epoch AI (https://epoch.ai/blog/power-demands-of-frontier-ai-training). Joshua You, David Owen, David Porter, and Tom Wilson. 'Scaling Intelligence: The Exponential Growth of AI's Power Needs'. EPRI, 2025. https://www.epri.com/research/products/000000003002033669

Data Center technology evolution

Eixos complementares

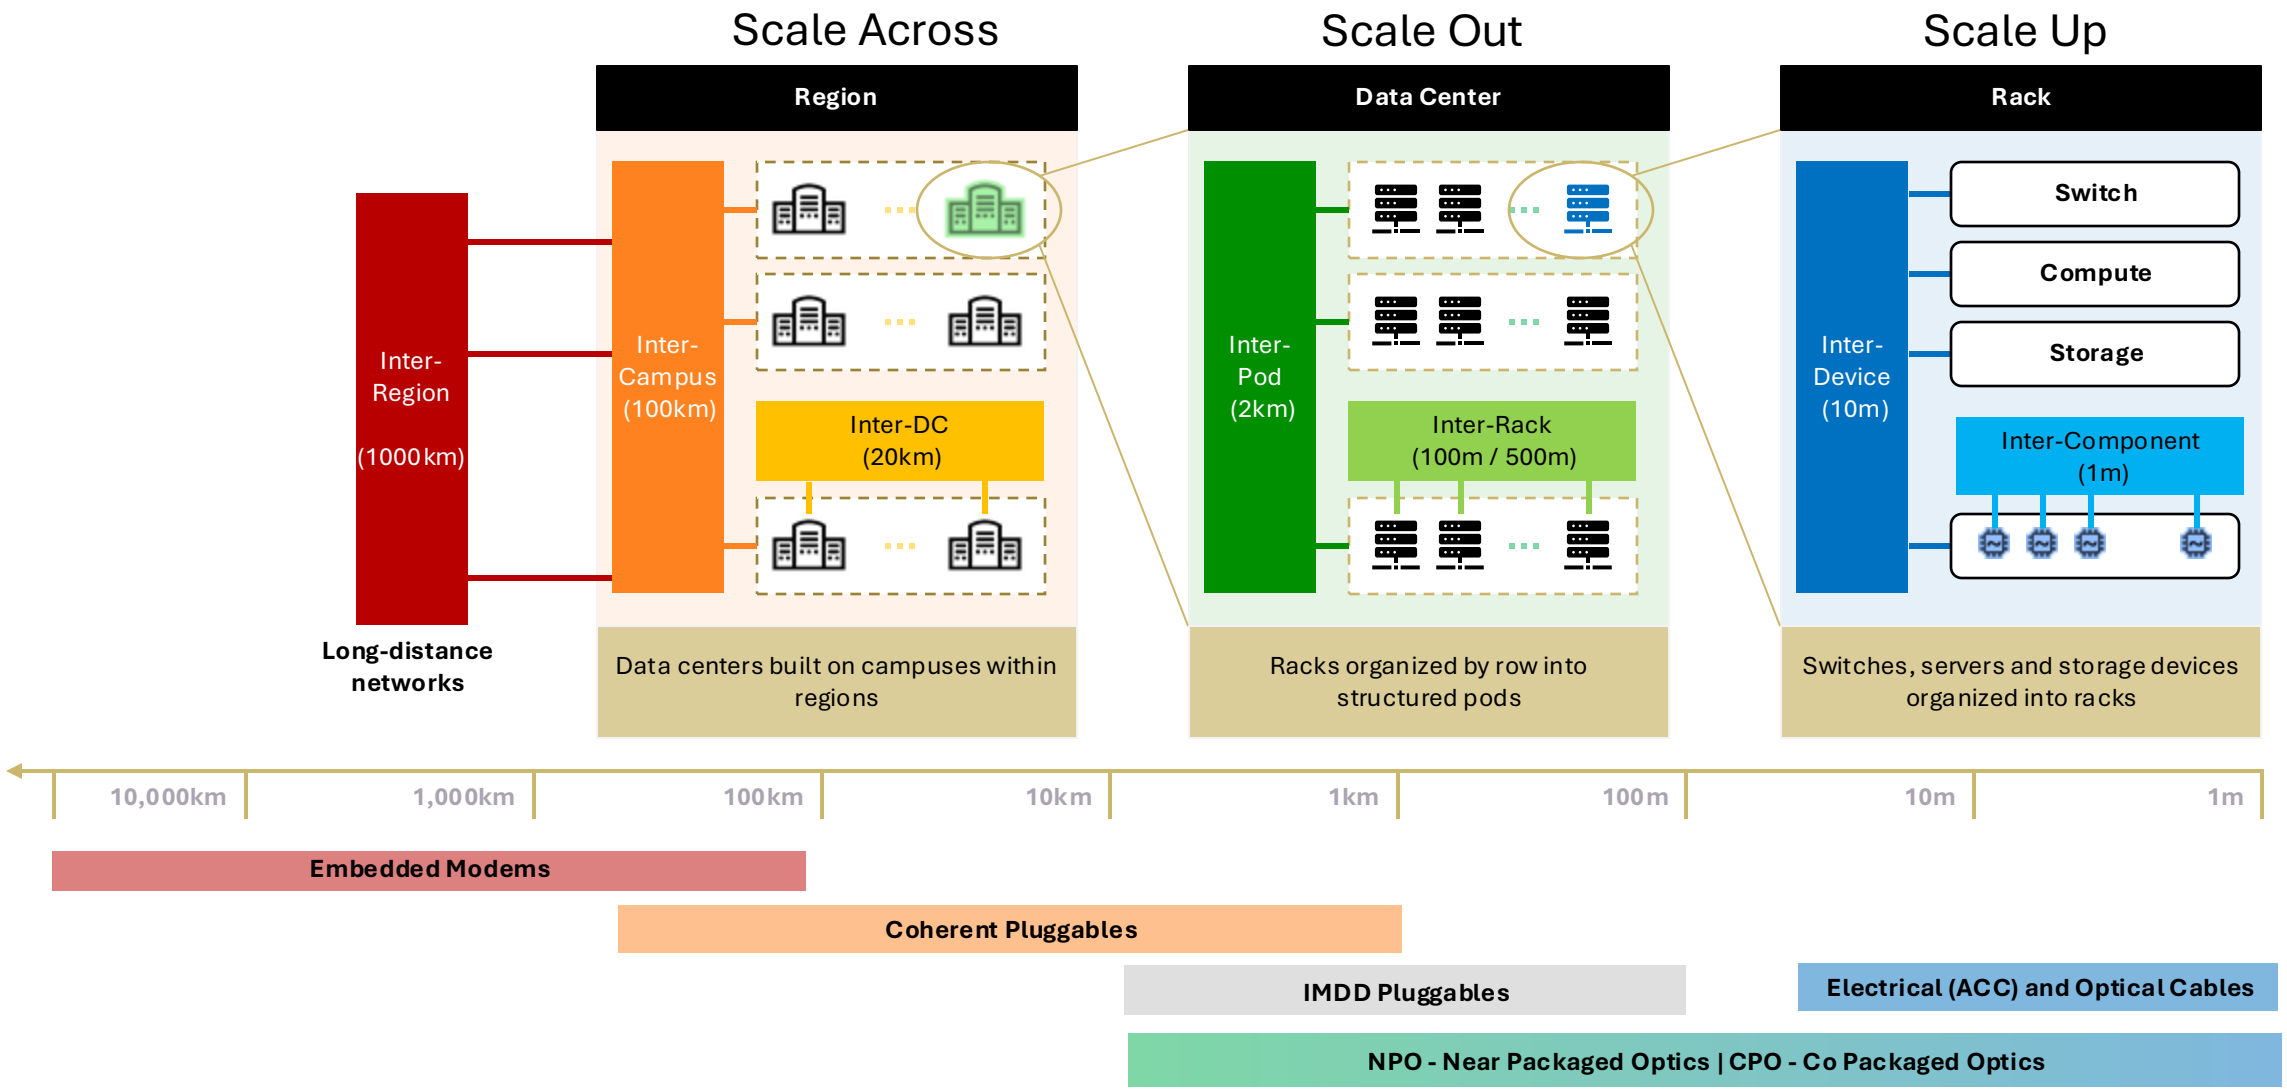

(*) IMDD = Intense Modulation Direct Detect

Same Fiber, Different Optimization

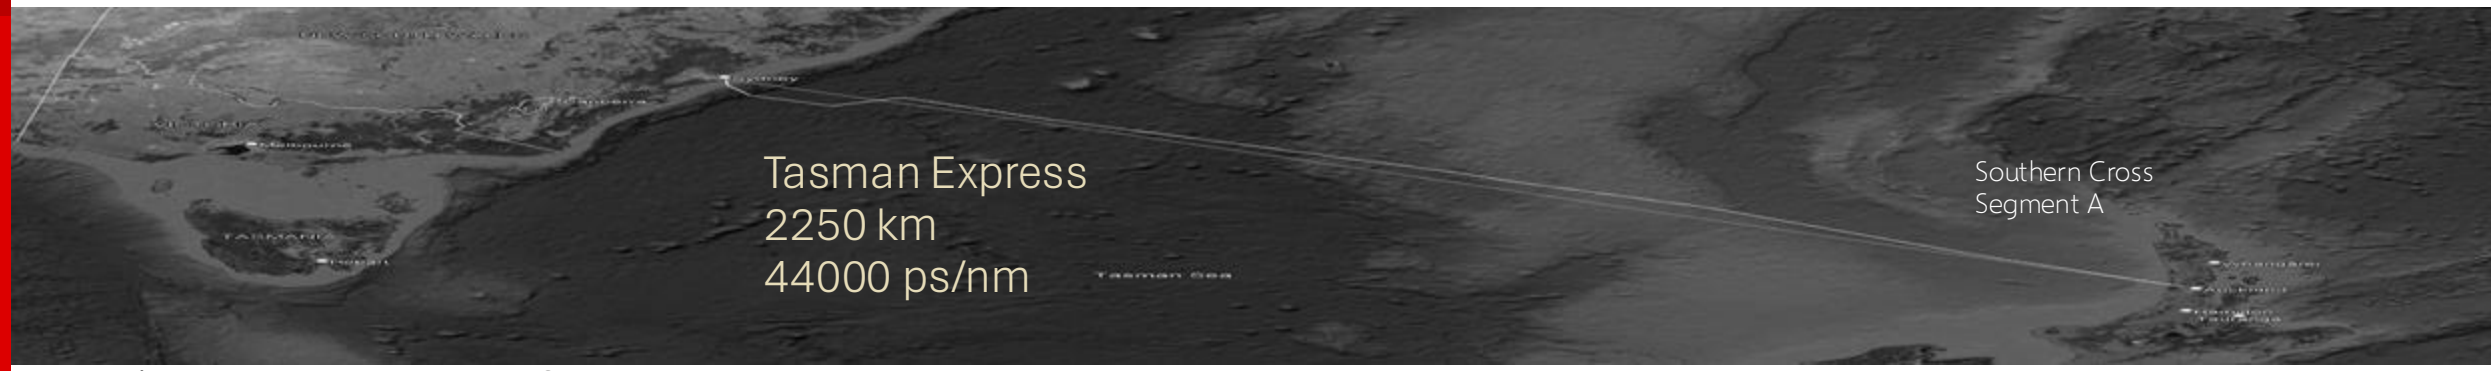

Australia - New Zeland (Tasman Sea)

	Coherent Pluggable Transceiver	Performance Transponder
Platform	Coherent plugs	High performance transponders
Capacity	23.2 Tb/s	27 Tb/s (+16%)
Channels	29 x 800 Gb/s (150GHz)	27 x 1 Tb/s (162GHz)

Power & space optimization → Pluggables

Spectral Efficiency Optimization → Transponders

ciena

What coherent pluggables have enabled

Best Power and Space Efficiency	Simpler Network Architecture	Multivendor Ecosystem

Expanded the Market for Coherent Transceivers

What's Next: 1600ZR+ pluggables

- 200G/lane hosts
- Longer reach at target rate
- Same spectral efficiency as 800G pluggables
- Leveraging DSP innovations from performance transponders

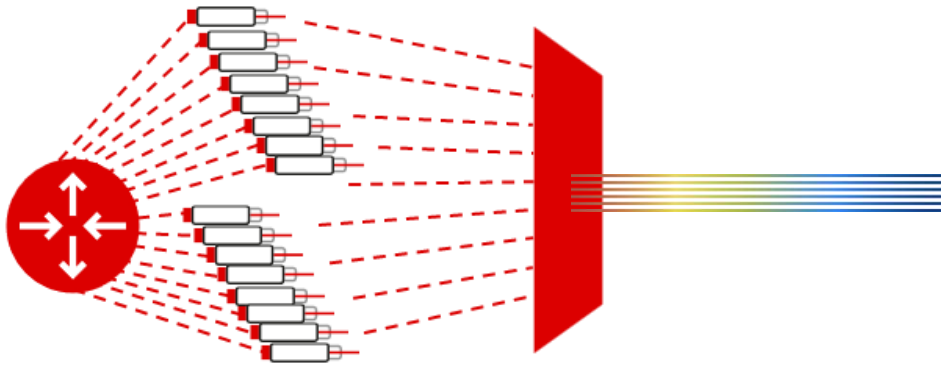

Pluggables enable simpler network designs, have accelerated growth of coherent optics

The New Limits of Optical Scaling: Density and Power

Router capacity scaling

- 51.2 → 102.4 → 204.8 Tb/s
- Faceplate density limited to **32 OSFP/OU**

Physical limits emerging

- Pluggable power approaching **air-cooling limits**

Power and faceplate density are the new scaling limits.

Co-packaged optics (CPO)

- Higher density
- Lower electrical power

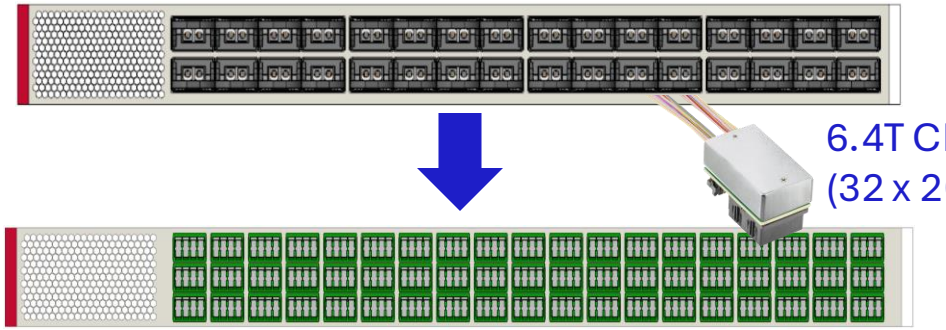

6.4T CPO
(32 x 200G/lane)

Return to grey optics and transponders

Direct-to-plug liquid cooling

- Enables higher-power pluggables
- Maintains the benefits of pluggables

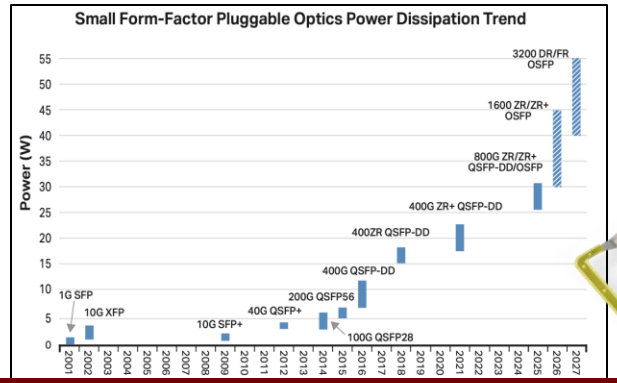

12.8T XPO
(8 x 1.6T)

Requires liquid cooling infrastructure

(*) OU = Operational Unit

Cracking the XPU beachfront bottleneck *(example with CPC)*

The *"beachfront":* chip perimeter where I/O pins reside. It doesn't scale proportionally with chip area.
As GPUs got bigger (lots of cores), **I/O per chip became a bottleneck**.

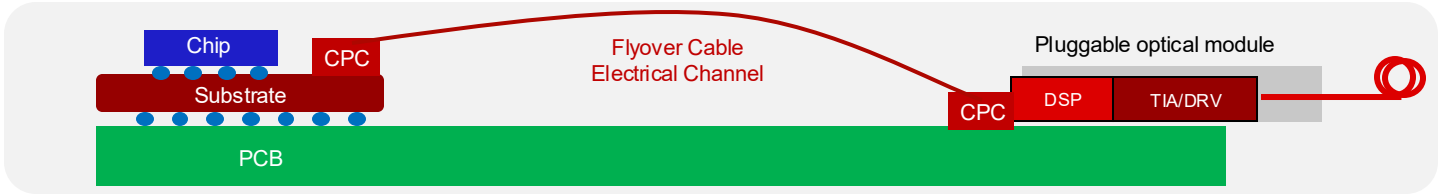

Tradditional setup: GPU → long PCB traces → front-panel optical module (often with DSP).

PAM6-PAM4

Fully Retimed (DSP)
Power: 30w (18.5 pJ/bit)

Optical module with no DSP, relying on host ASIC's analog signal.

PAM4-PAM4
PAM6-PAM6

Linear Pluggable (LPO)
Power: 10w (6.5 pJ/bit)

PAM4-PAM4
PAM6-PAM6

Co-Package Optics (CPO)
Power: 5w (3 pJ/bit)

30-60% power saving but shorter reach

70-85% power saving but shorter reach

Scale-up networks must go both faster and wider: higher lane speeds (toward 448 Gbps) combined with increased I/O width through co-packaging.

(*) CPC = Co-packaged copper

From Lighting Wavelengths to Lighting Full Fibers

- **Today: wavelength-by-wavelength deployment**
  - 128 x 400G x 2 pluggables per fiber pair
  - 64 x 800G x 2 pluggables per fiber pair
- **Scaling this model for 10s-100s fiber deployments becomes operationally complex.**
- **Full spectrum transponder**
  - Lights the entire fiber at once
  - Returns to grey optics and external transponders
  - Reduces reliance on DWDM expertise
  - Enables cookie-cutter deployments
  - Can be higher performance without pluggable power constraints

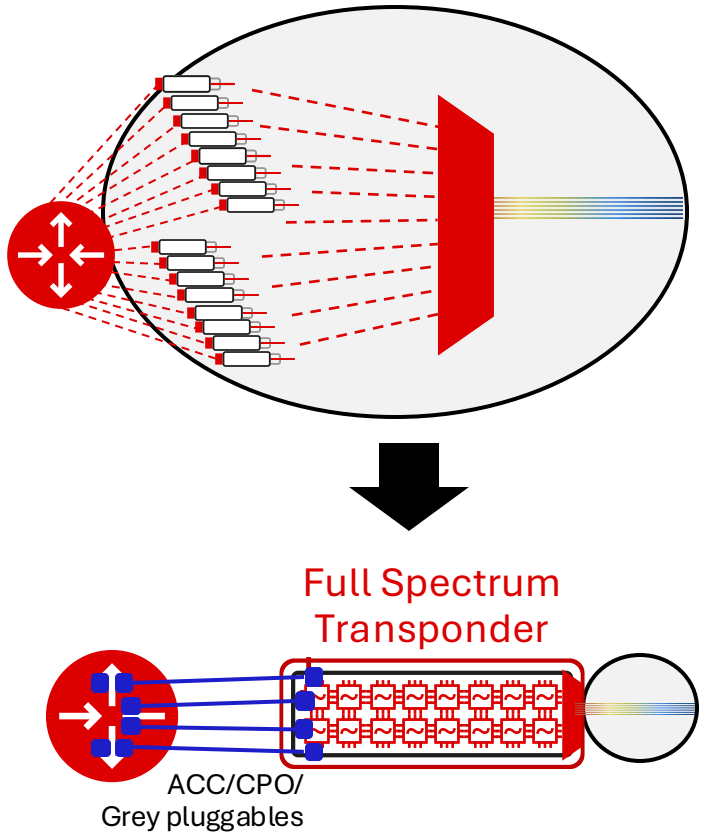

Optimizing for Deployment Velocity Favors Full-Fiber Architectures

Technology consumption models are widening

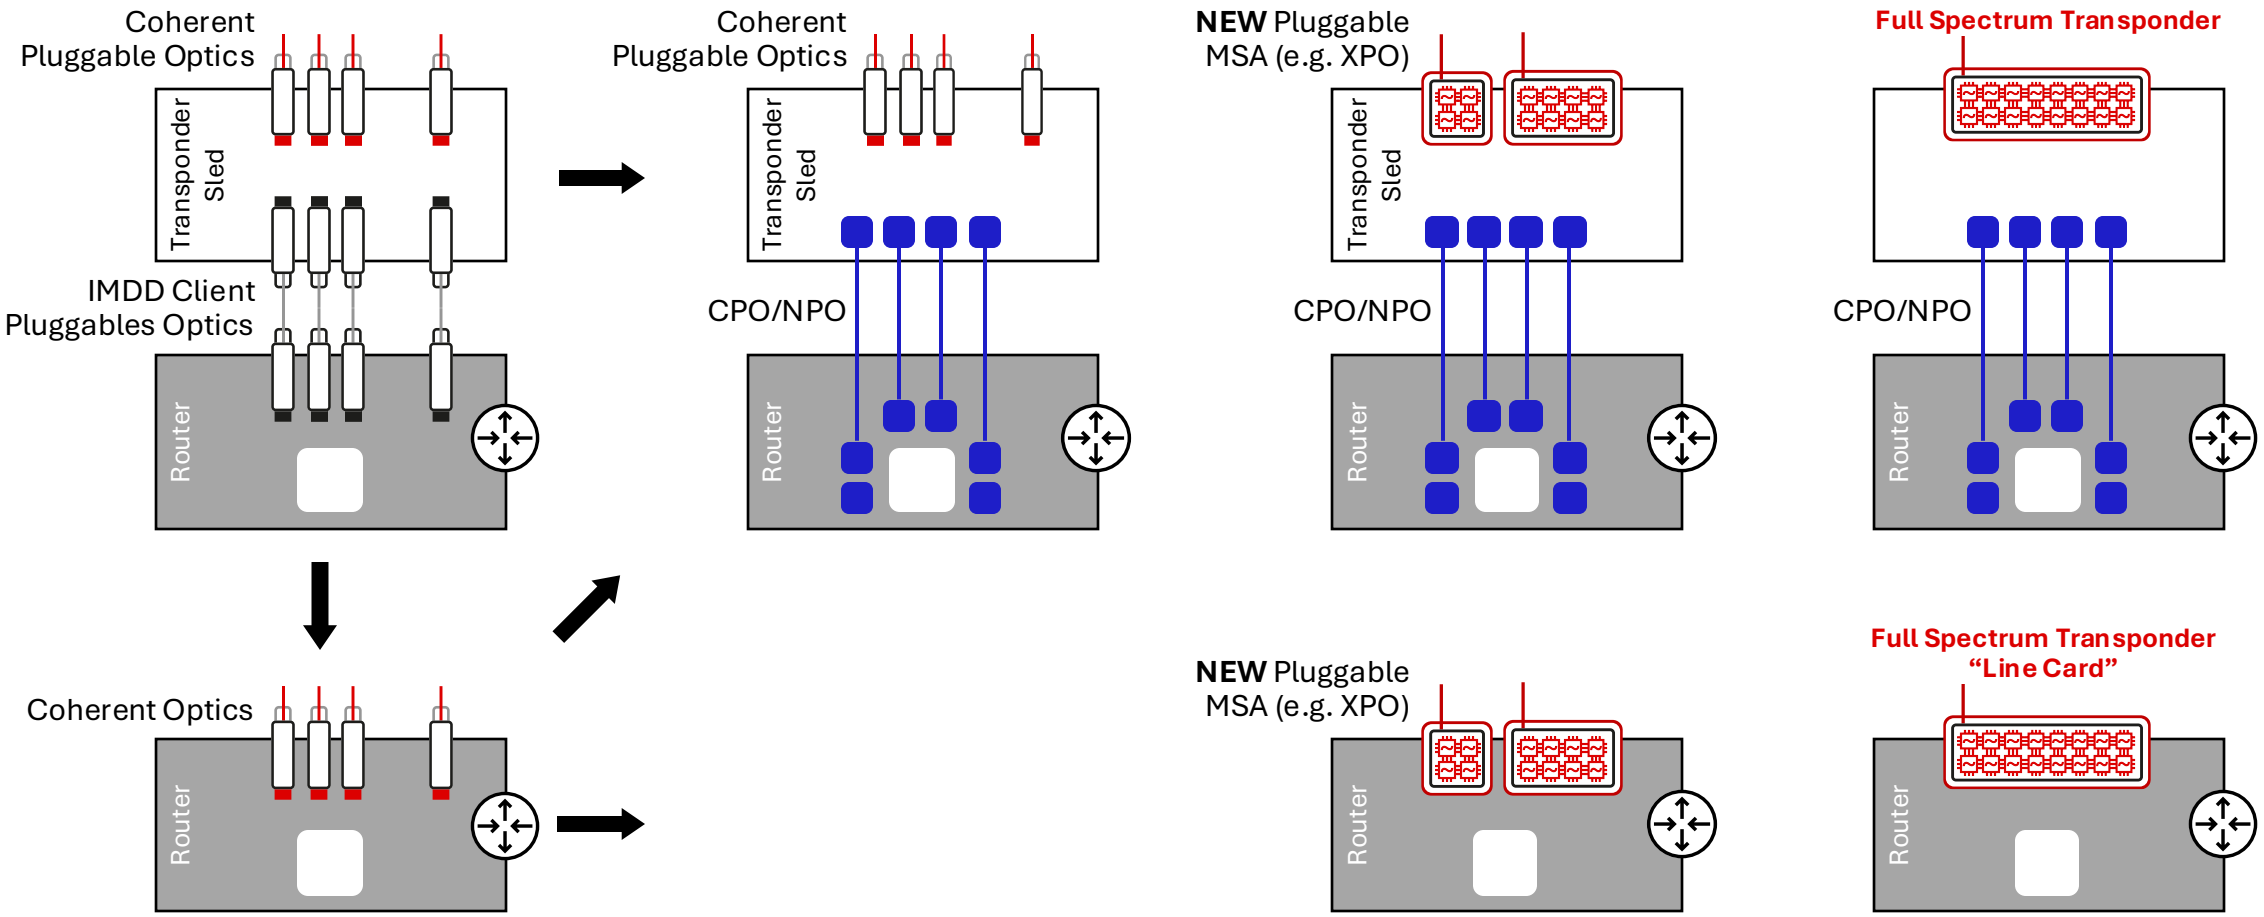

Coherent Optics Moving Beyond Conventional Pluggable Form Factors

Evolution of coherent optics consumption models

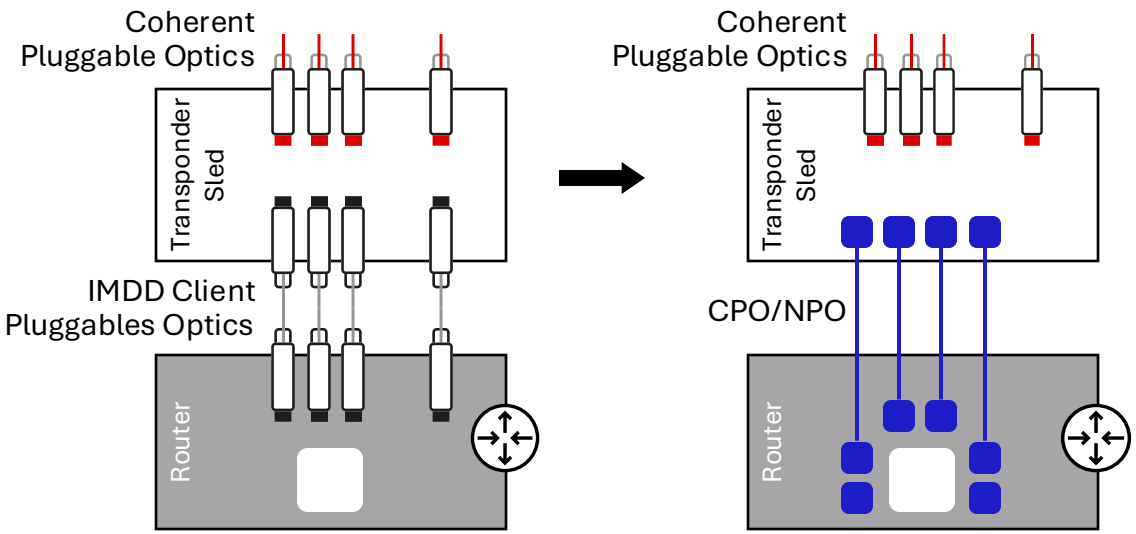

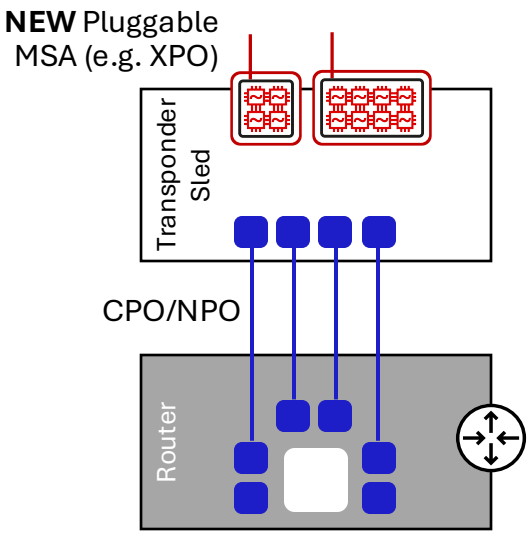

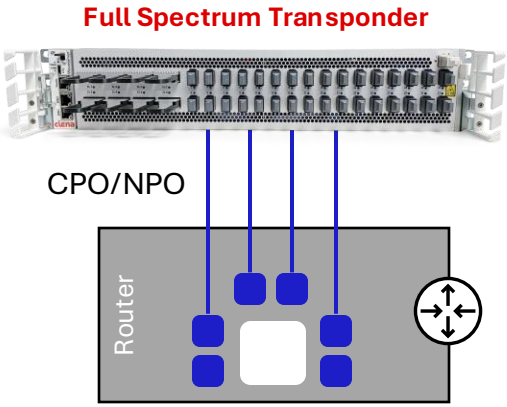

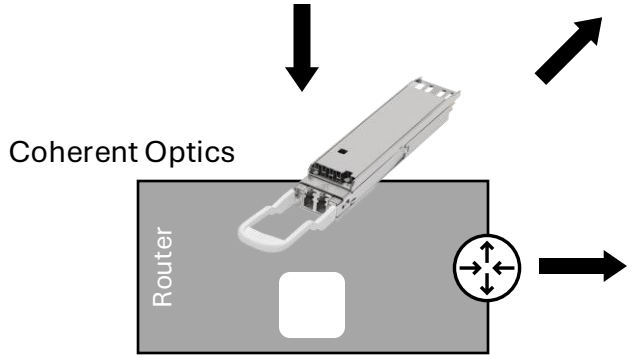

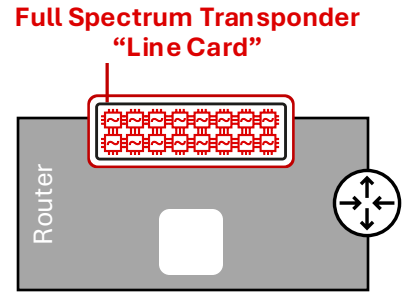

Coherent Optics Moving Beyond Conventional Pluggable Form Factors

Hyper-rail provides massive densification of amplifier sites for space, power, capacity, and cost

Line amp sites **TODAY**
4 rails/rack, 16 rails/hut
All capacity within one hut

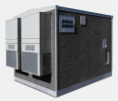

4 rails / rack
1 hut

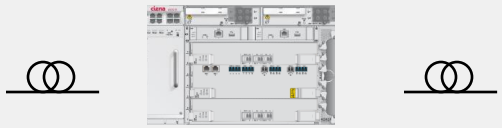

20 Petabit/sec **AI example**
With today's technology
22 huts!

4 rails / rack
22 hut

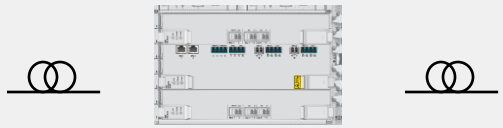

New hyper-rail design

32x
density improvement!

75%
Power savings

New hyper-rail innovation:
20 Petabit/sec
1 hut!

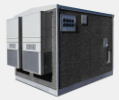

128 rails / rack
1 hut

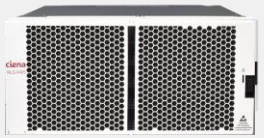

AI-scale optical transport physically and economically viable

Drivers for Full Spectrum Transponder and Hyper Rail

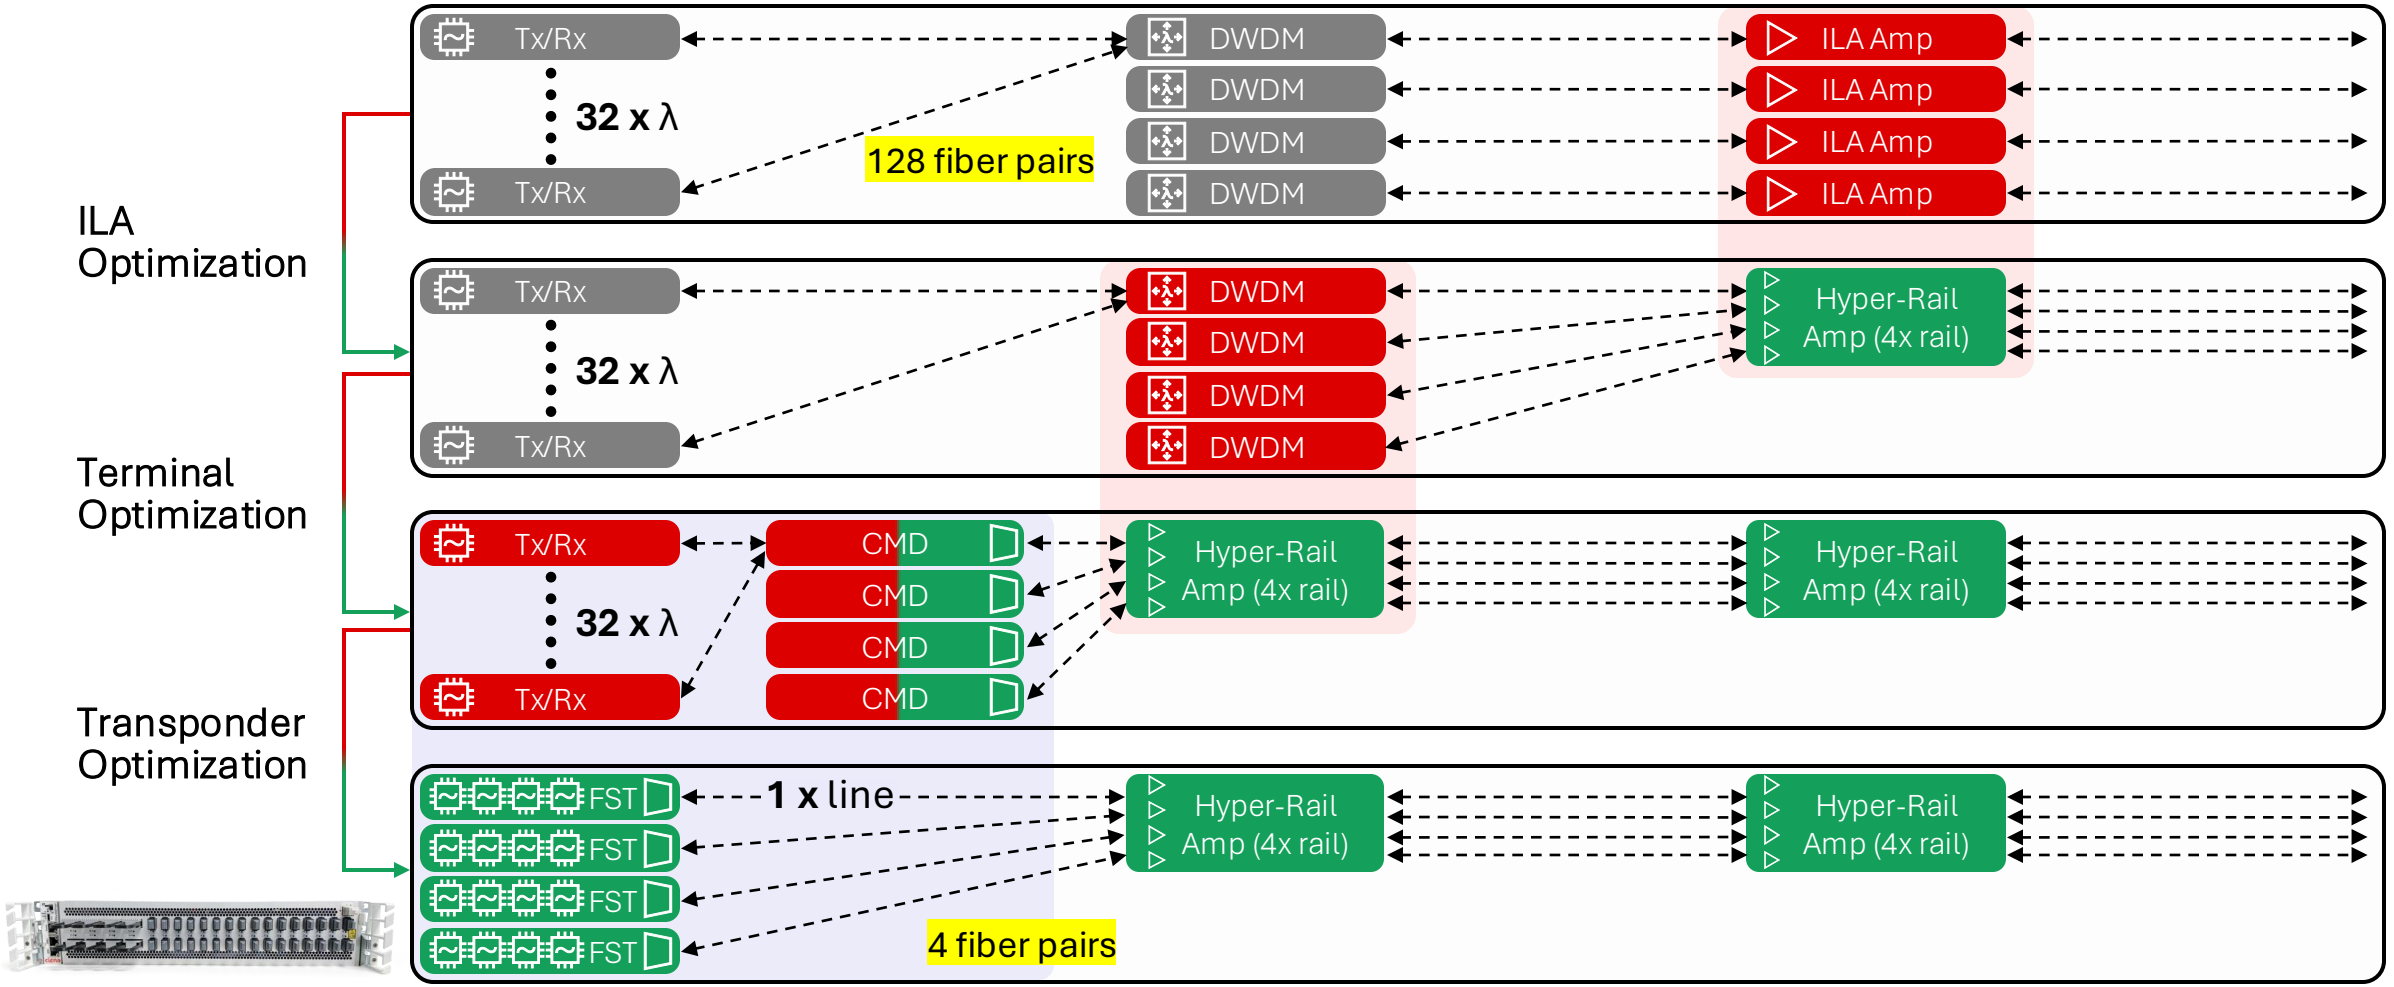

Enables significant improvement in deployment velocity with cookie-cutter configurations

ciena

Choosing the Right Architecture Depends on the Requirement

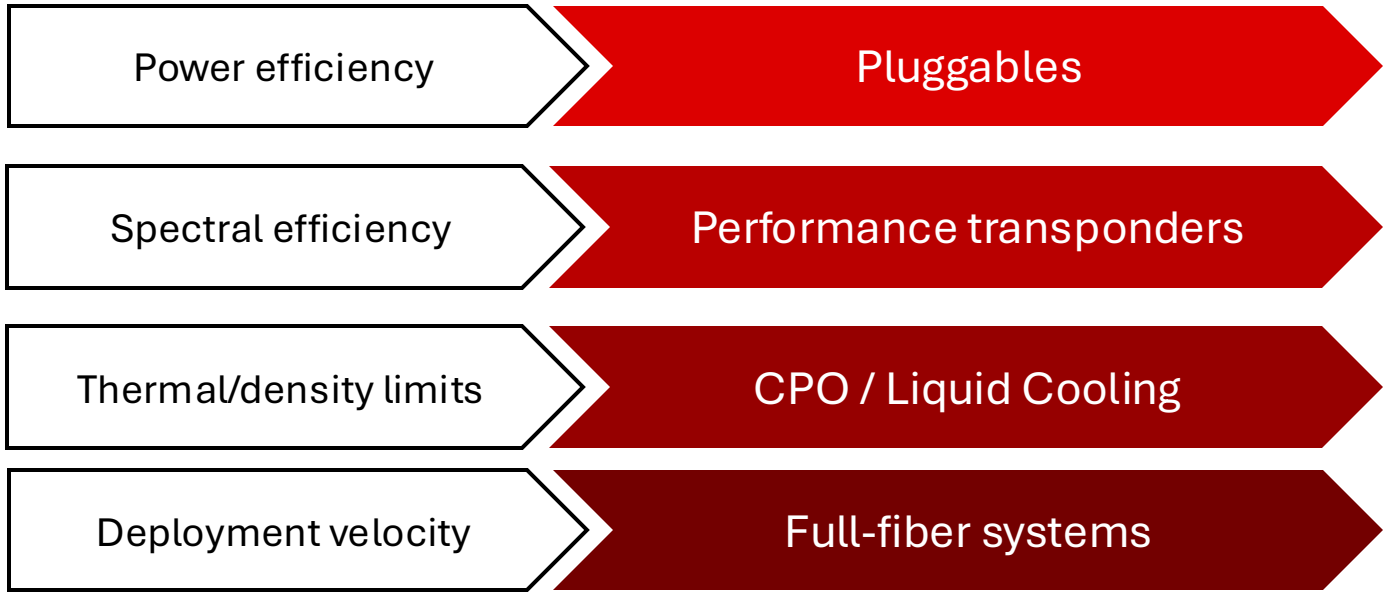

Multiple architectures will coexist as networks are designed for different constraints

Thank you

Décio Coraça

dcoraca@ciena.com