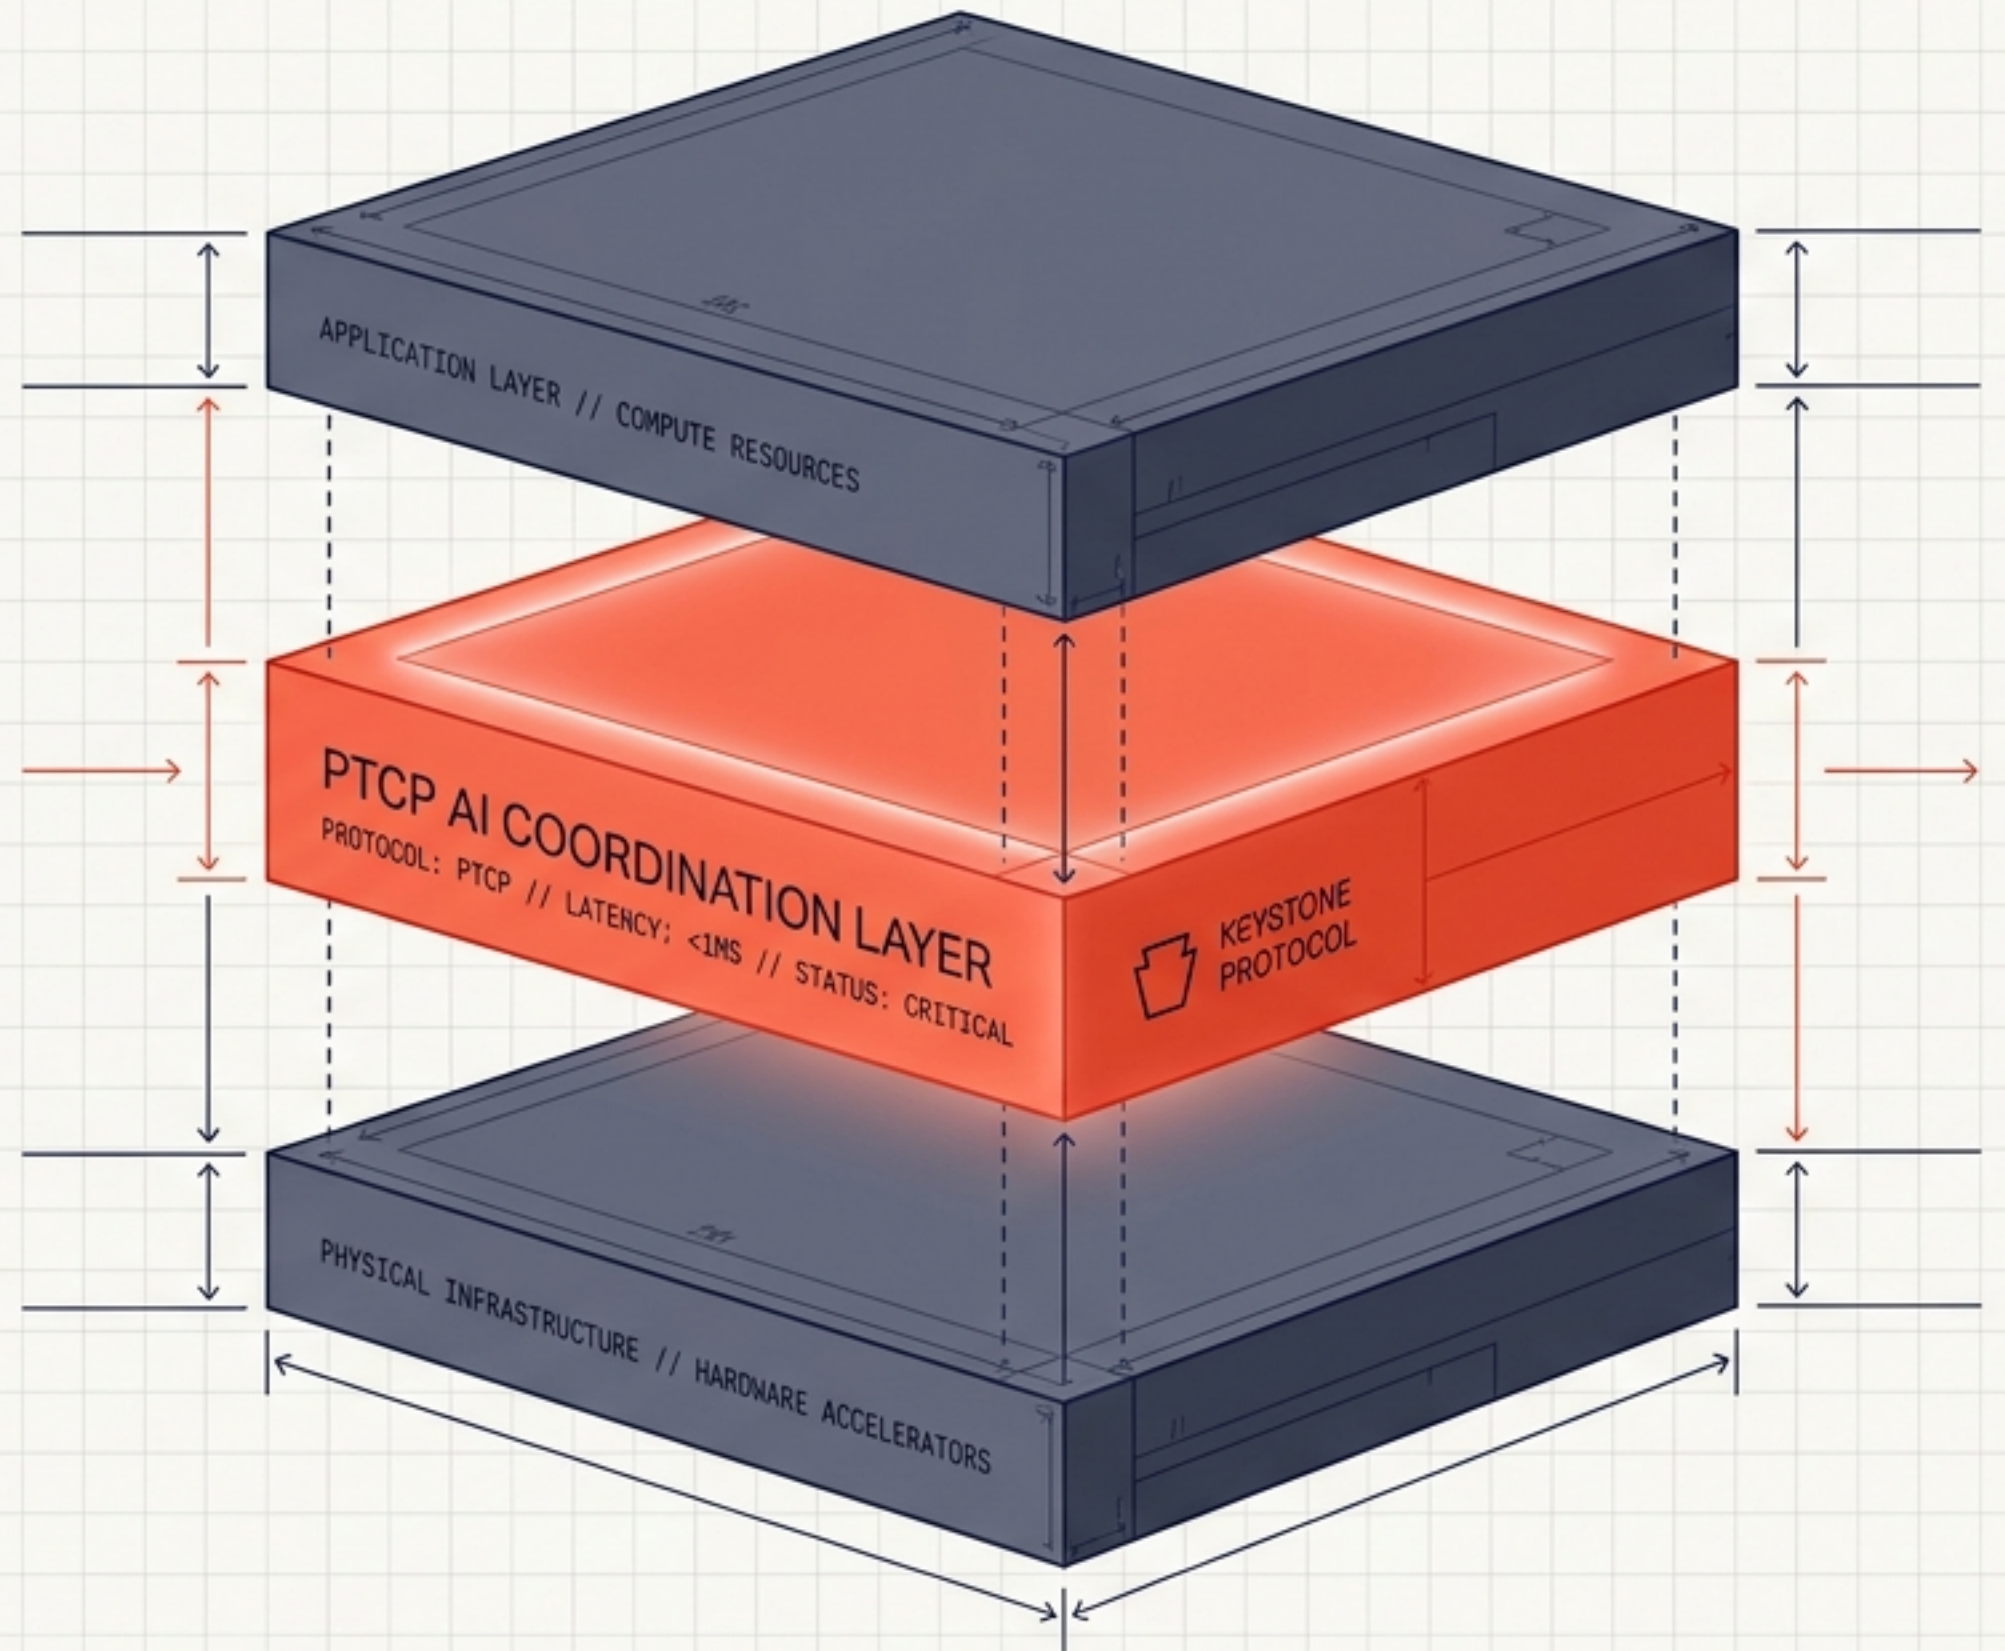


The Missing Protocol

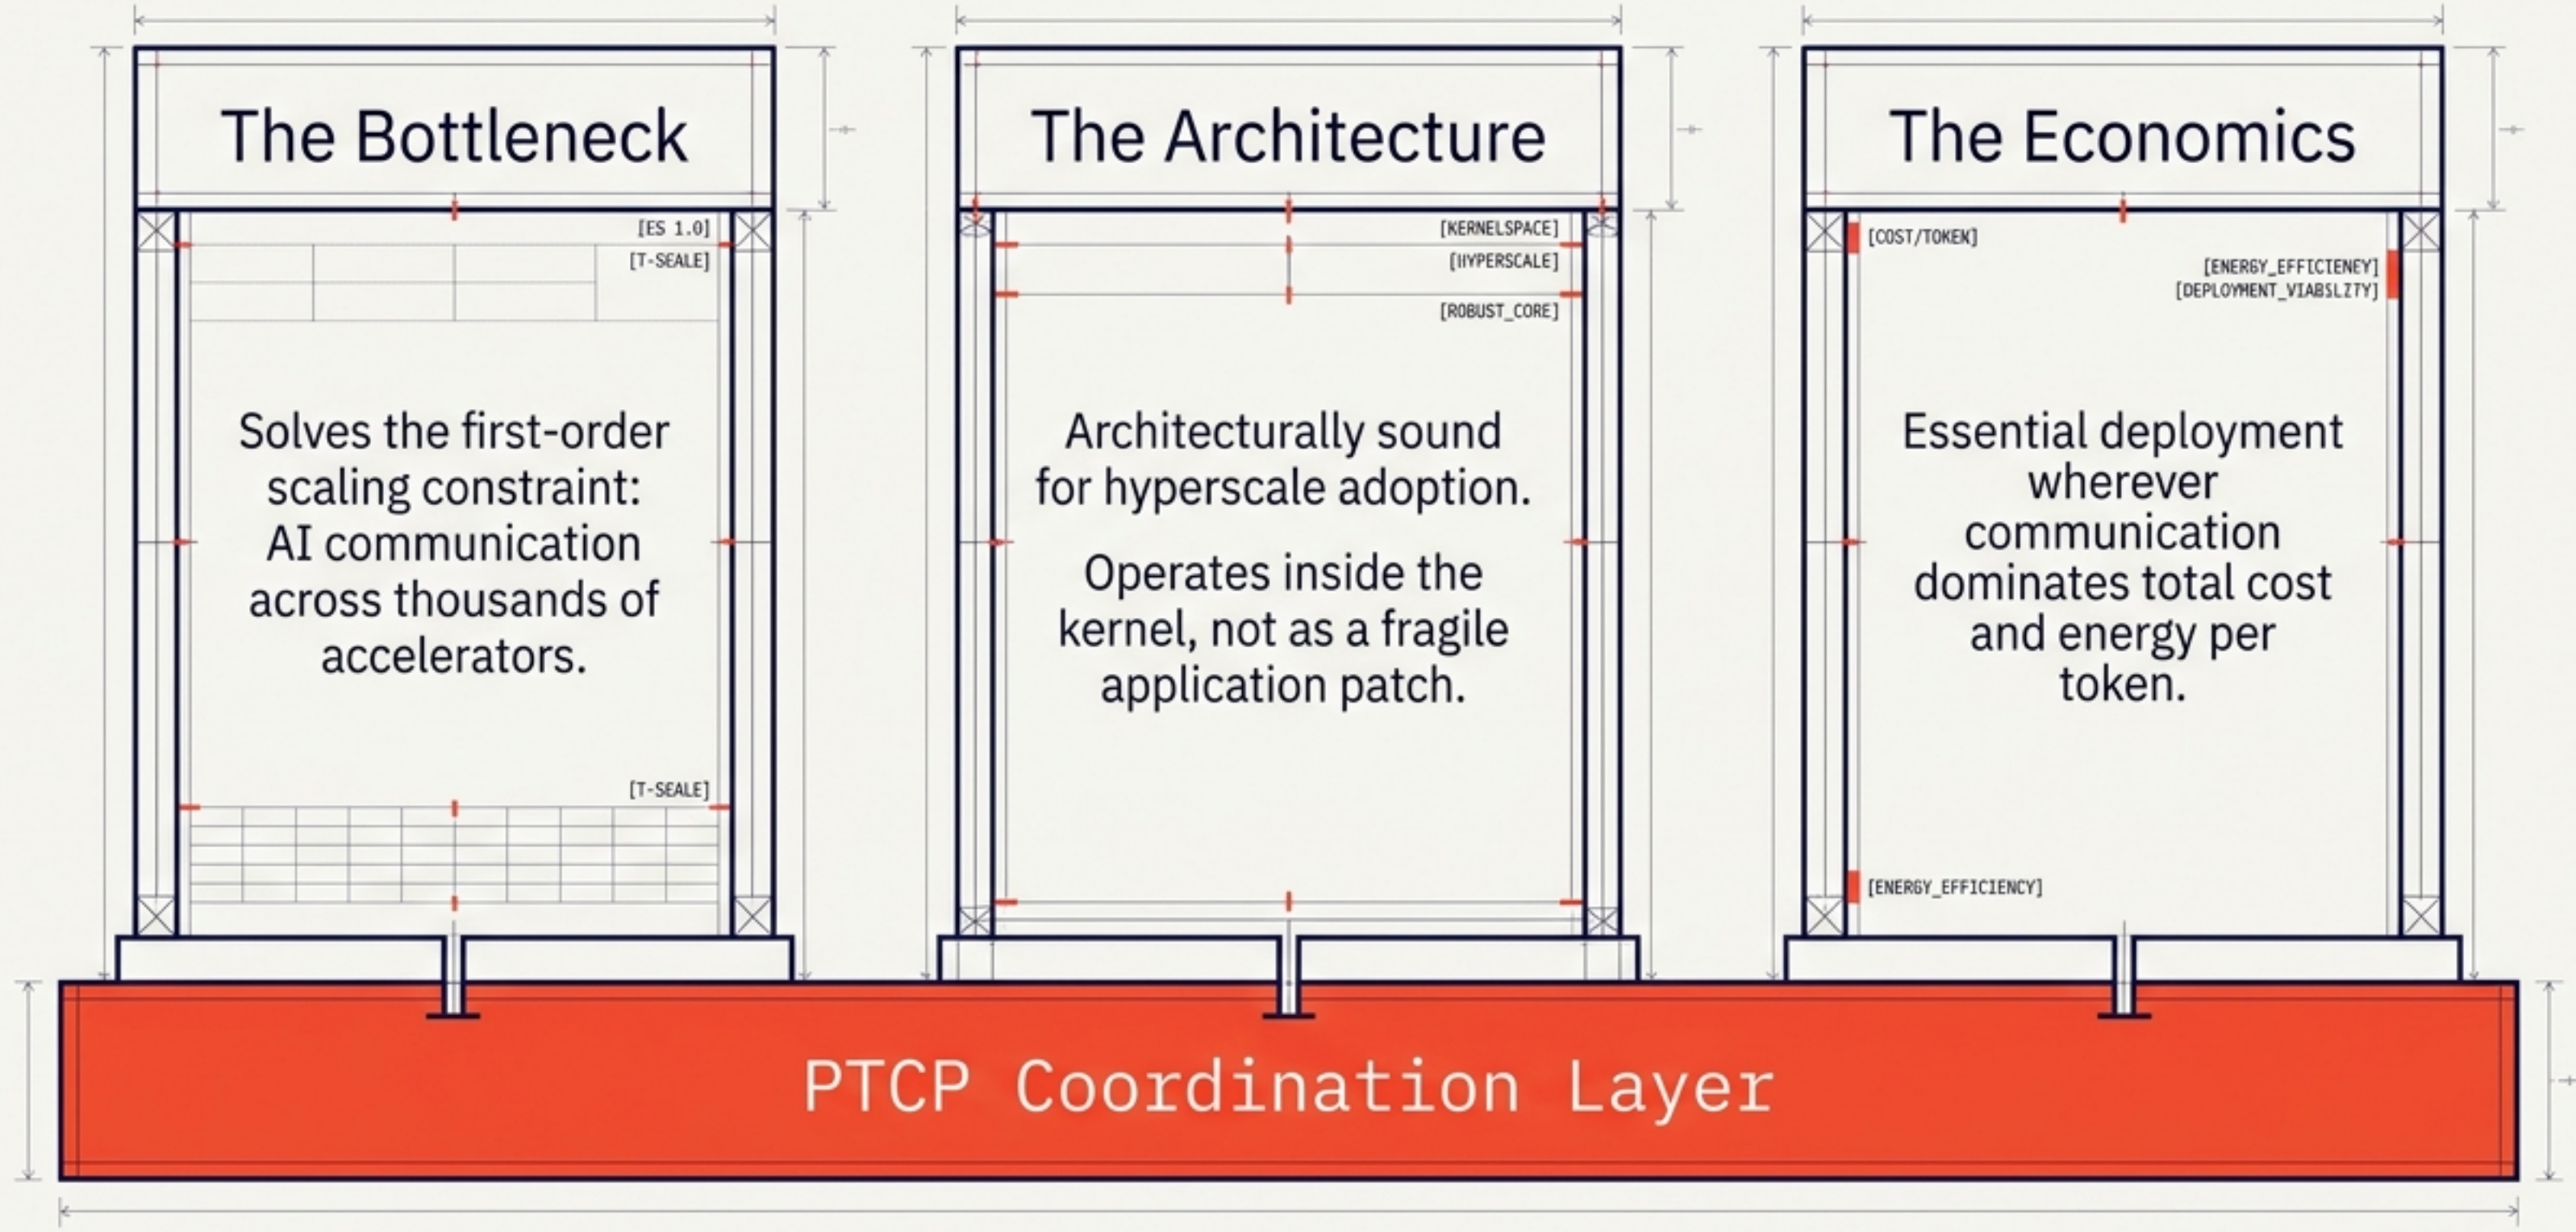
PTCP AI as the new coordination layer for hyperscale AI systems.




R1.	1	REVISION #1	DATE
S3	1	BY155-ABRAT18.9ET	2024-01-01

SCALE DESIGN

Infrastructure Coordination, Not Optimization Polish



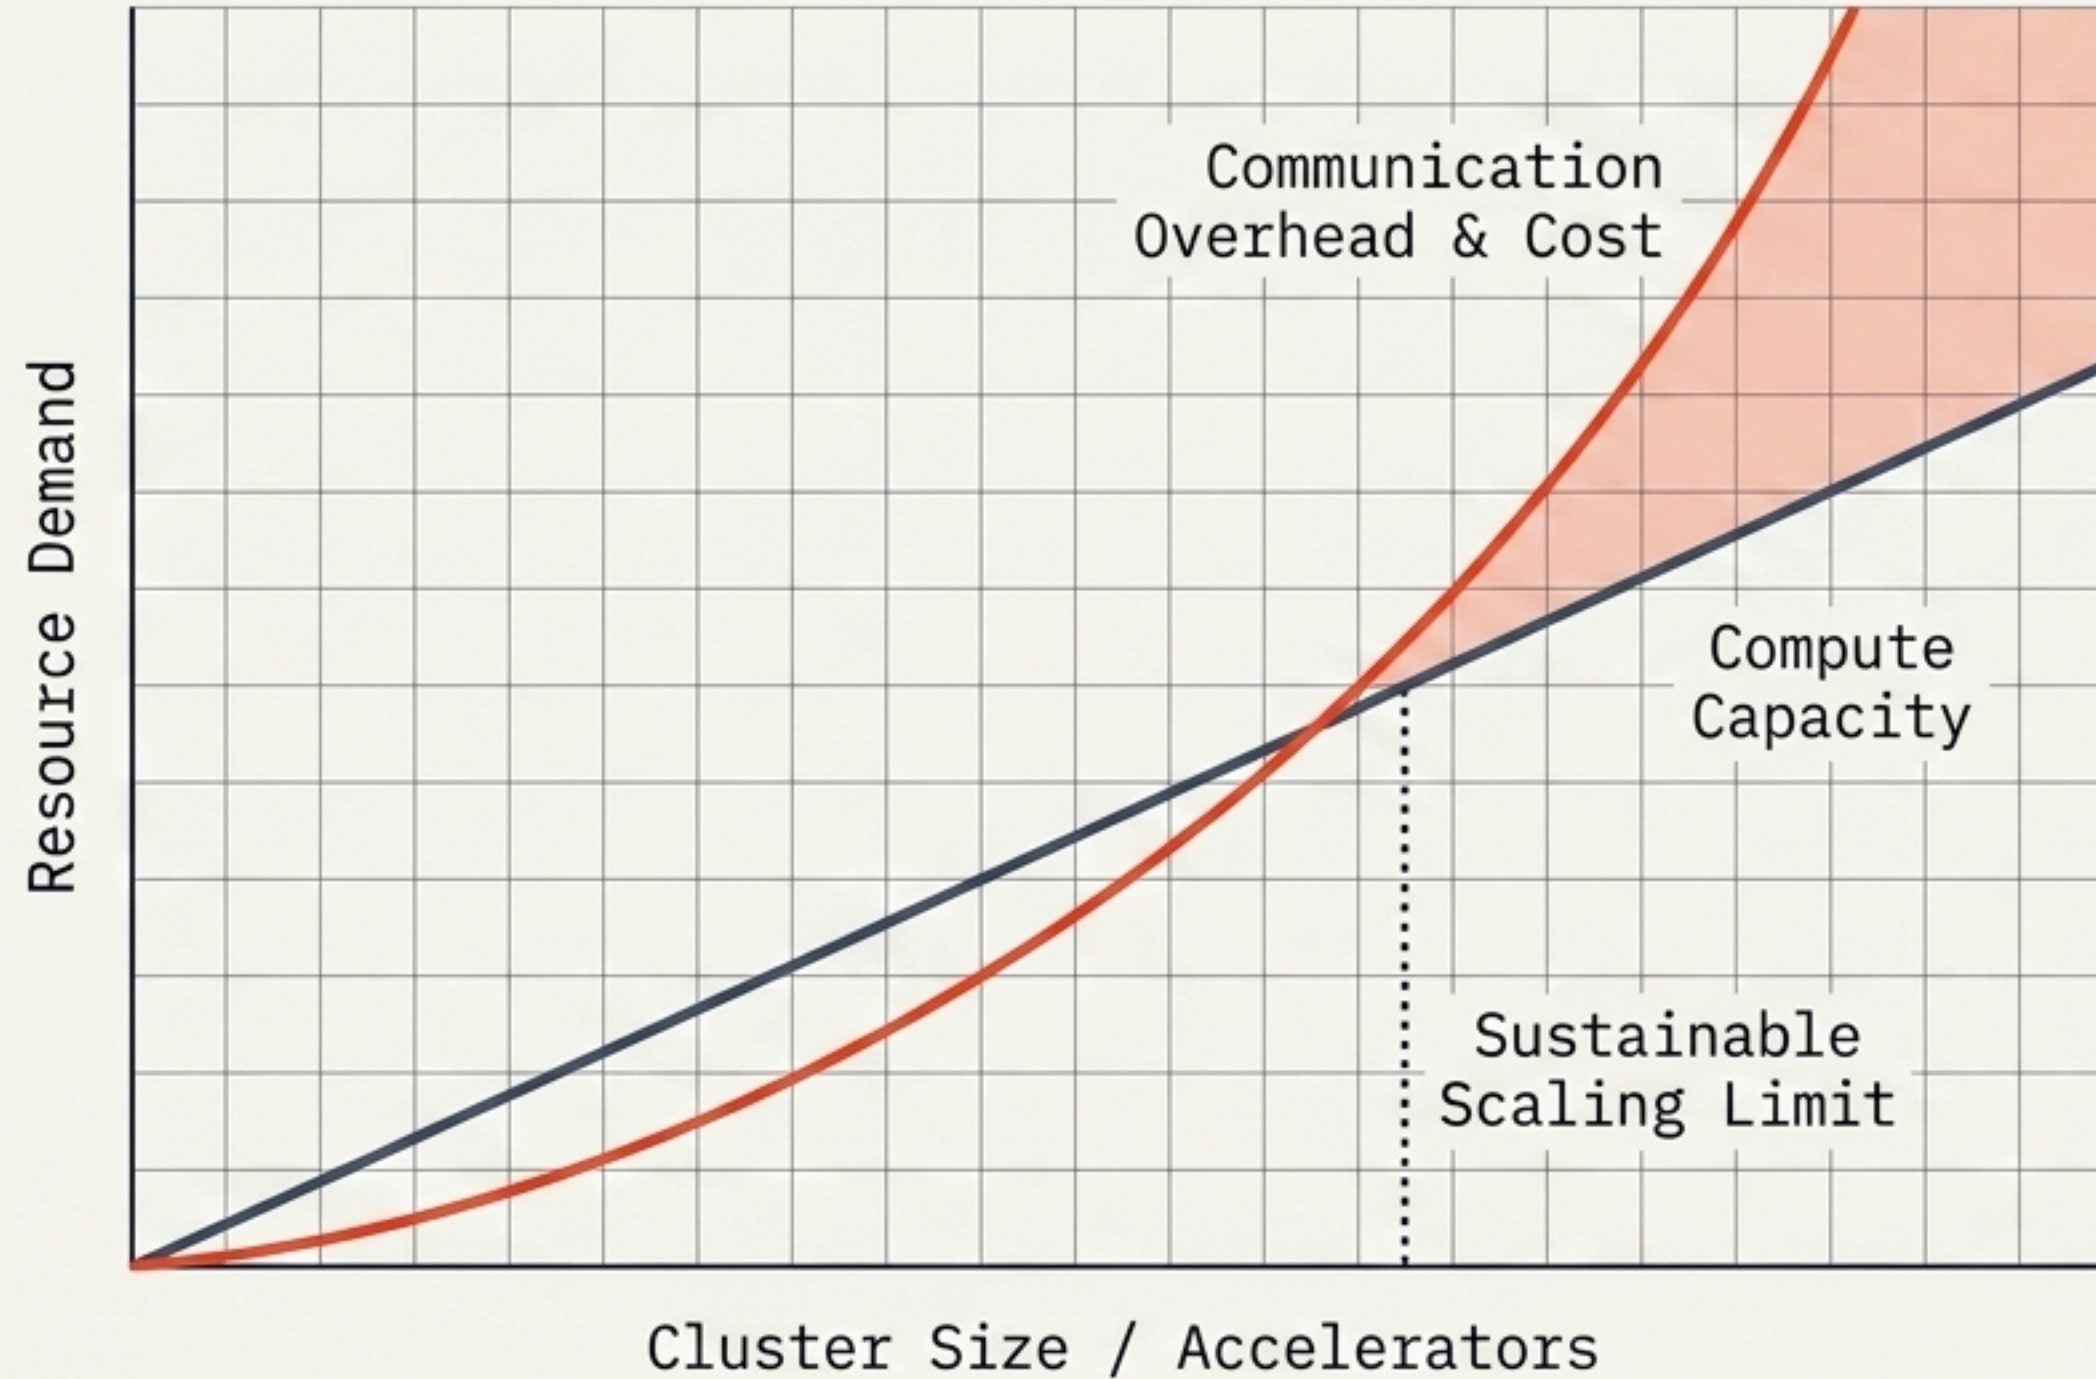
REVISION. 2 2022-09-15



SCALEZENI, RA6T DBT. J32

RI.	1	REVISION #1	BATE
S3	1	BF155-ABRAT18.9ET	2624-01-01


The AI Scaling Wall



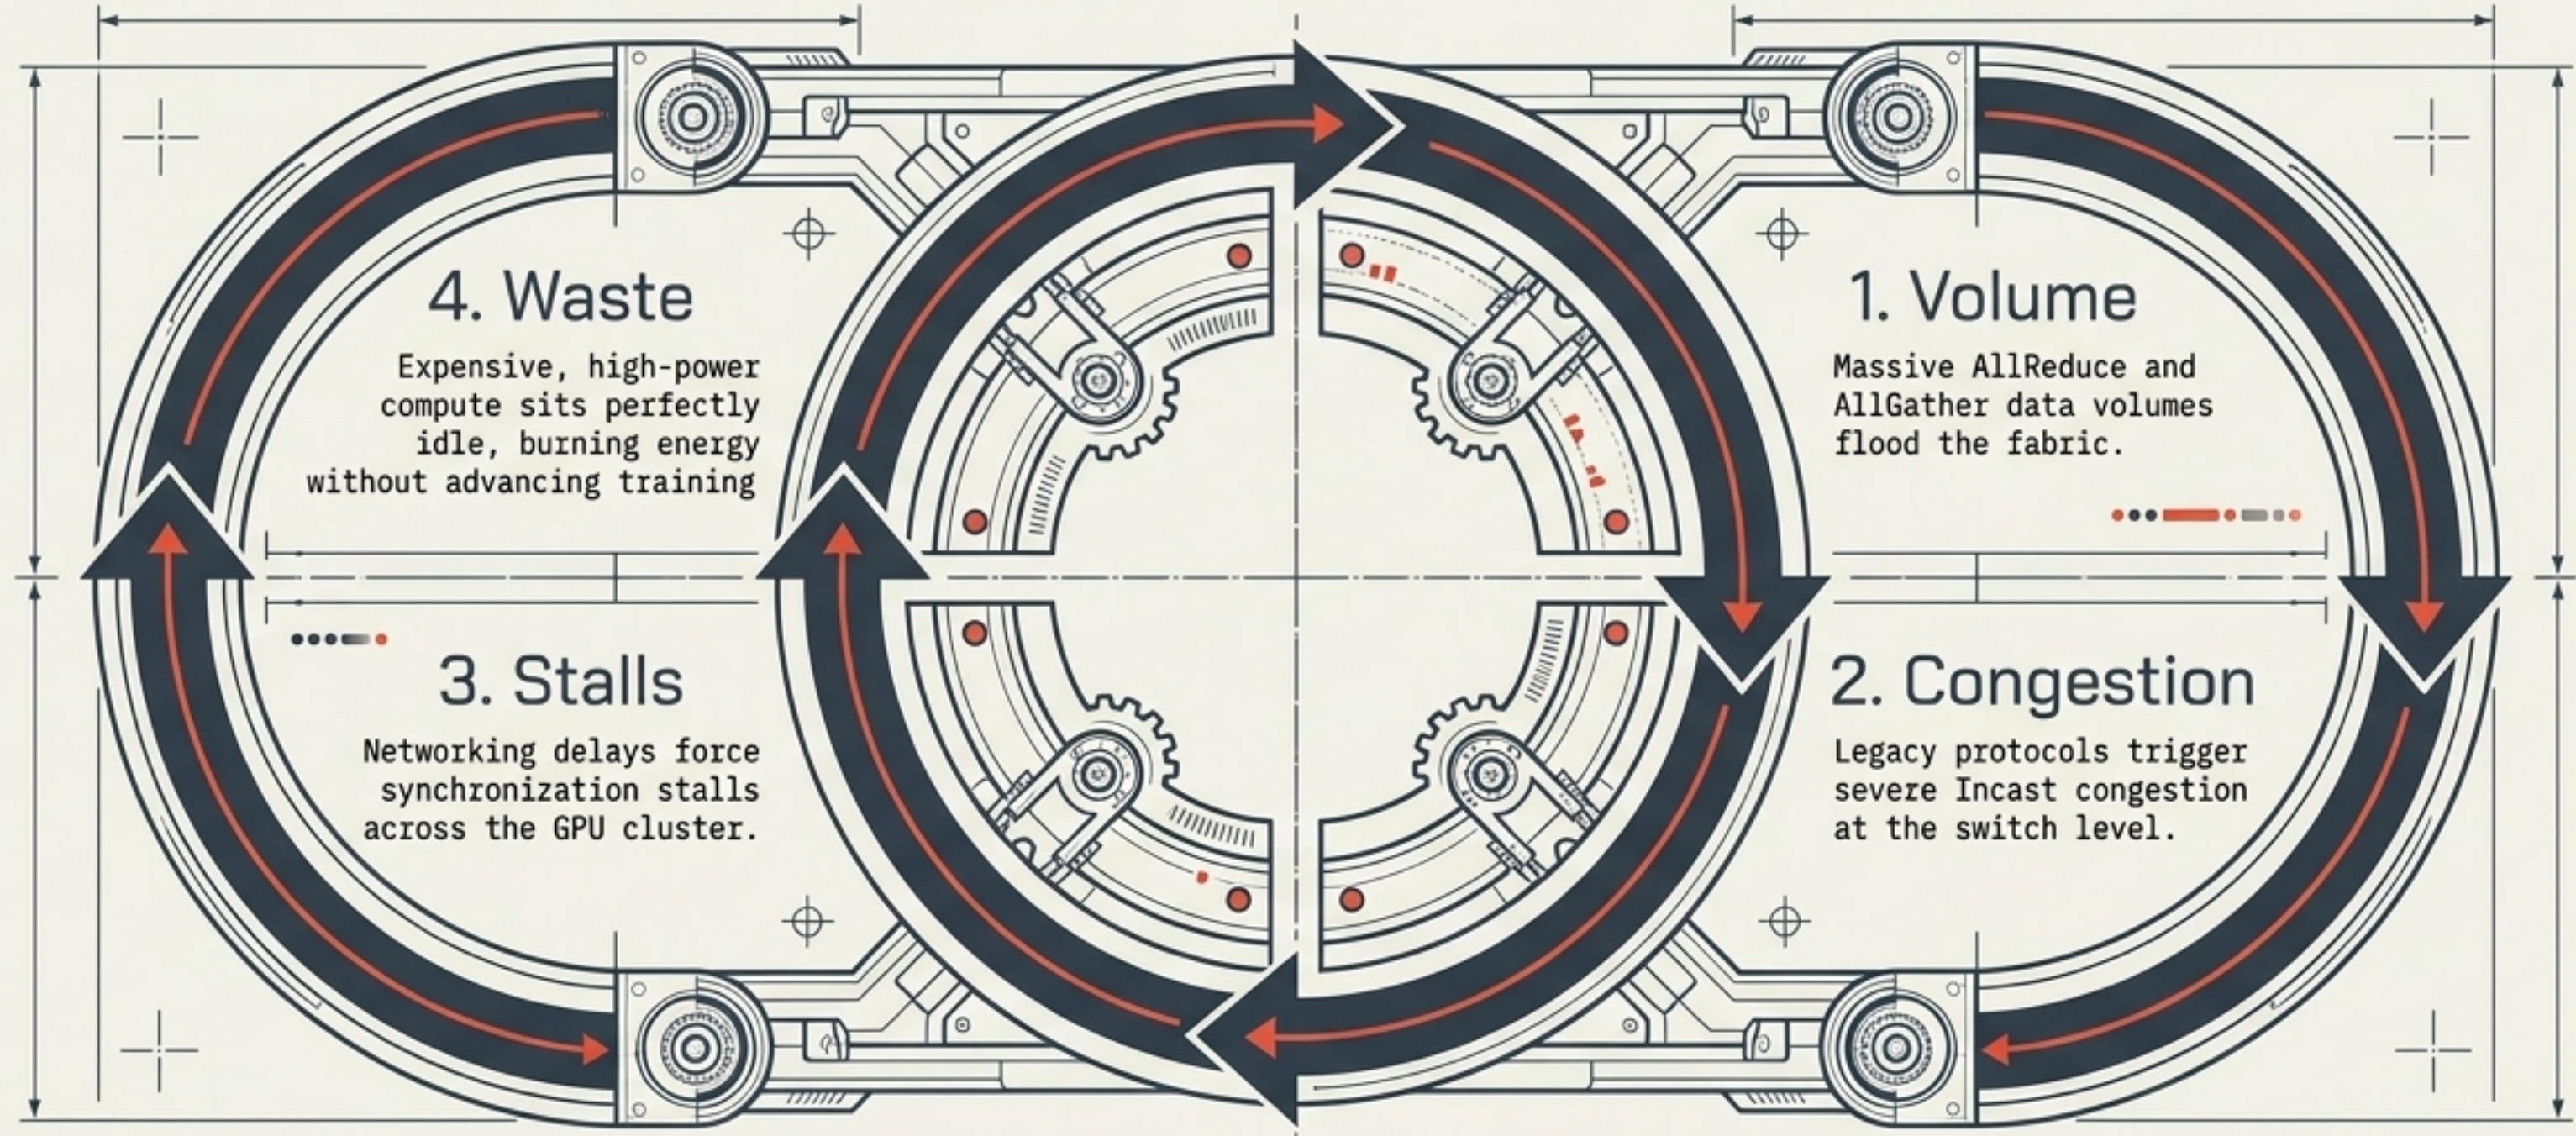
Key Insight:

The value proposition of PTCP improves specifically as cluster size grows.

In modern training, communication demand grows faster than compute capacity, rendering hardware-centric scaling mathematically and economically unsustainable.



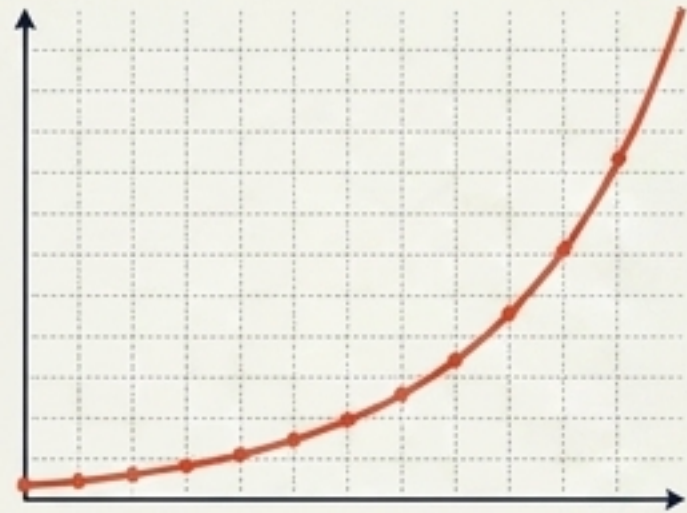
Diagnosing the Bottleneck: The Cascade Failure



Modern training speedups are strictly bounded by this cascade.

The Hardware Trap: Why You Cannot Buy Your Way Out


GPU roadmaps are permanently outpacing interconnect evolution.



InfiniBand

Status: Marginal costs are rising super-linearly.

Flaw: Roadmaps physically cannot match the I/O demands of next-gen GPUs.



Optical Interconnects

Status: High bandwidth, but introduces fatal coordination discontinuities.

Flaw: Creates gaps in the control plane that hardware alone cannot bridge.

The Missing Protocol: Historical Precedent

Era	Ad-Hoc Solution	Protocol Paradigm
Early Internet Scaling	Bespoke vendor networking	TCP
Cloud Containerization	Bespoke bash scripts	Kubernetes
Data Center Topology	Hard-coded switch logic	SDN
Hyperscale Model Training	Reactive Congestion Control / Over-provisioning	PTCP


Multiplication of network congestion

Reactive Congestion Control zetes


Protocol actions target mapped

The Core Mechanism: Predictive vs. Reactive

Reactive Congestion Control



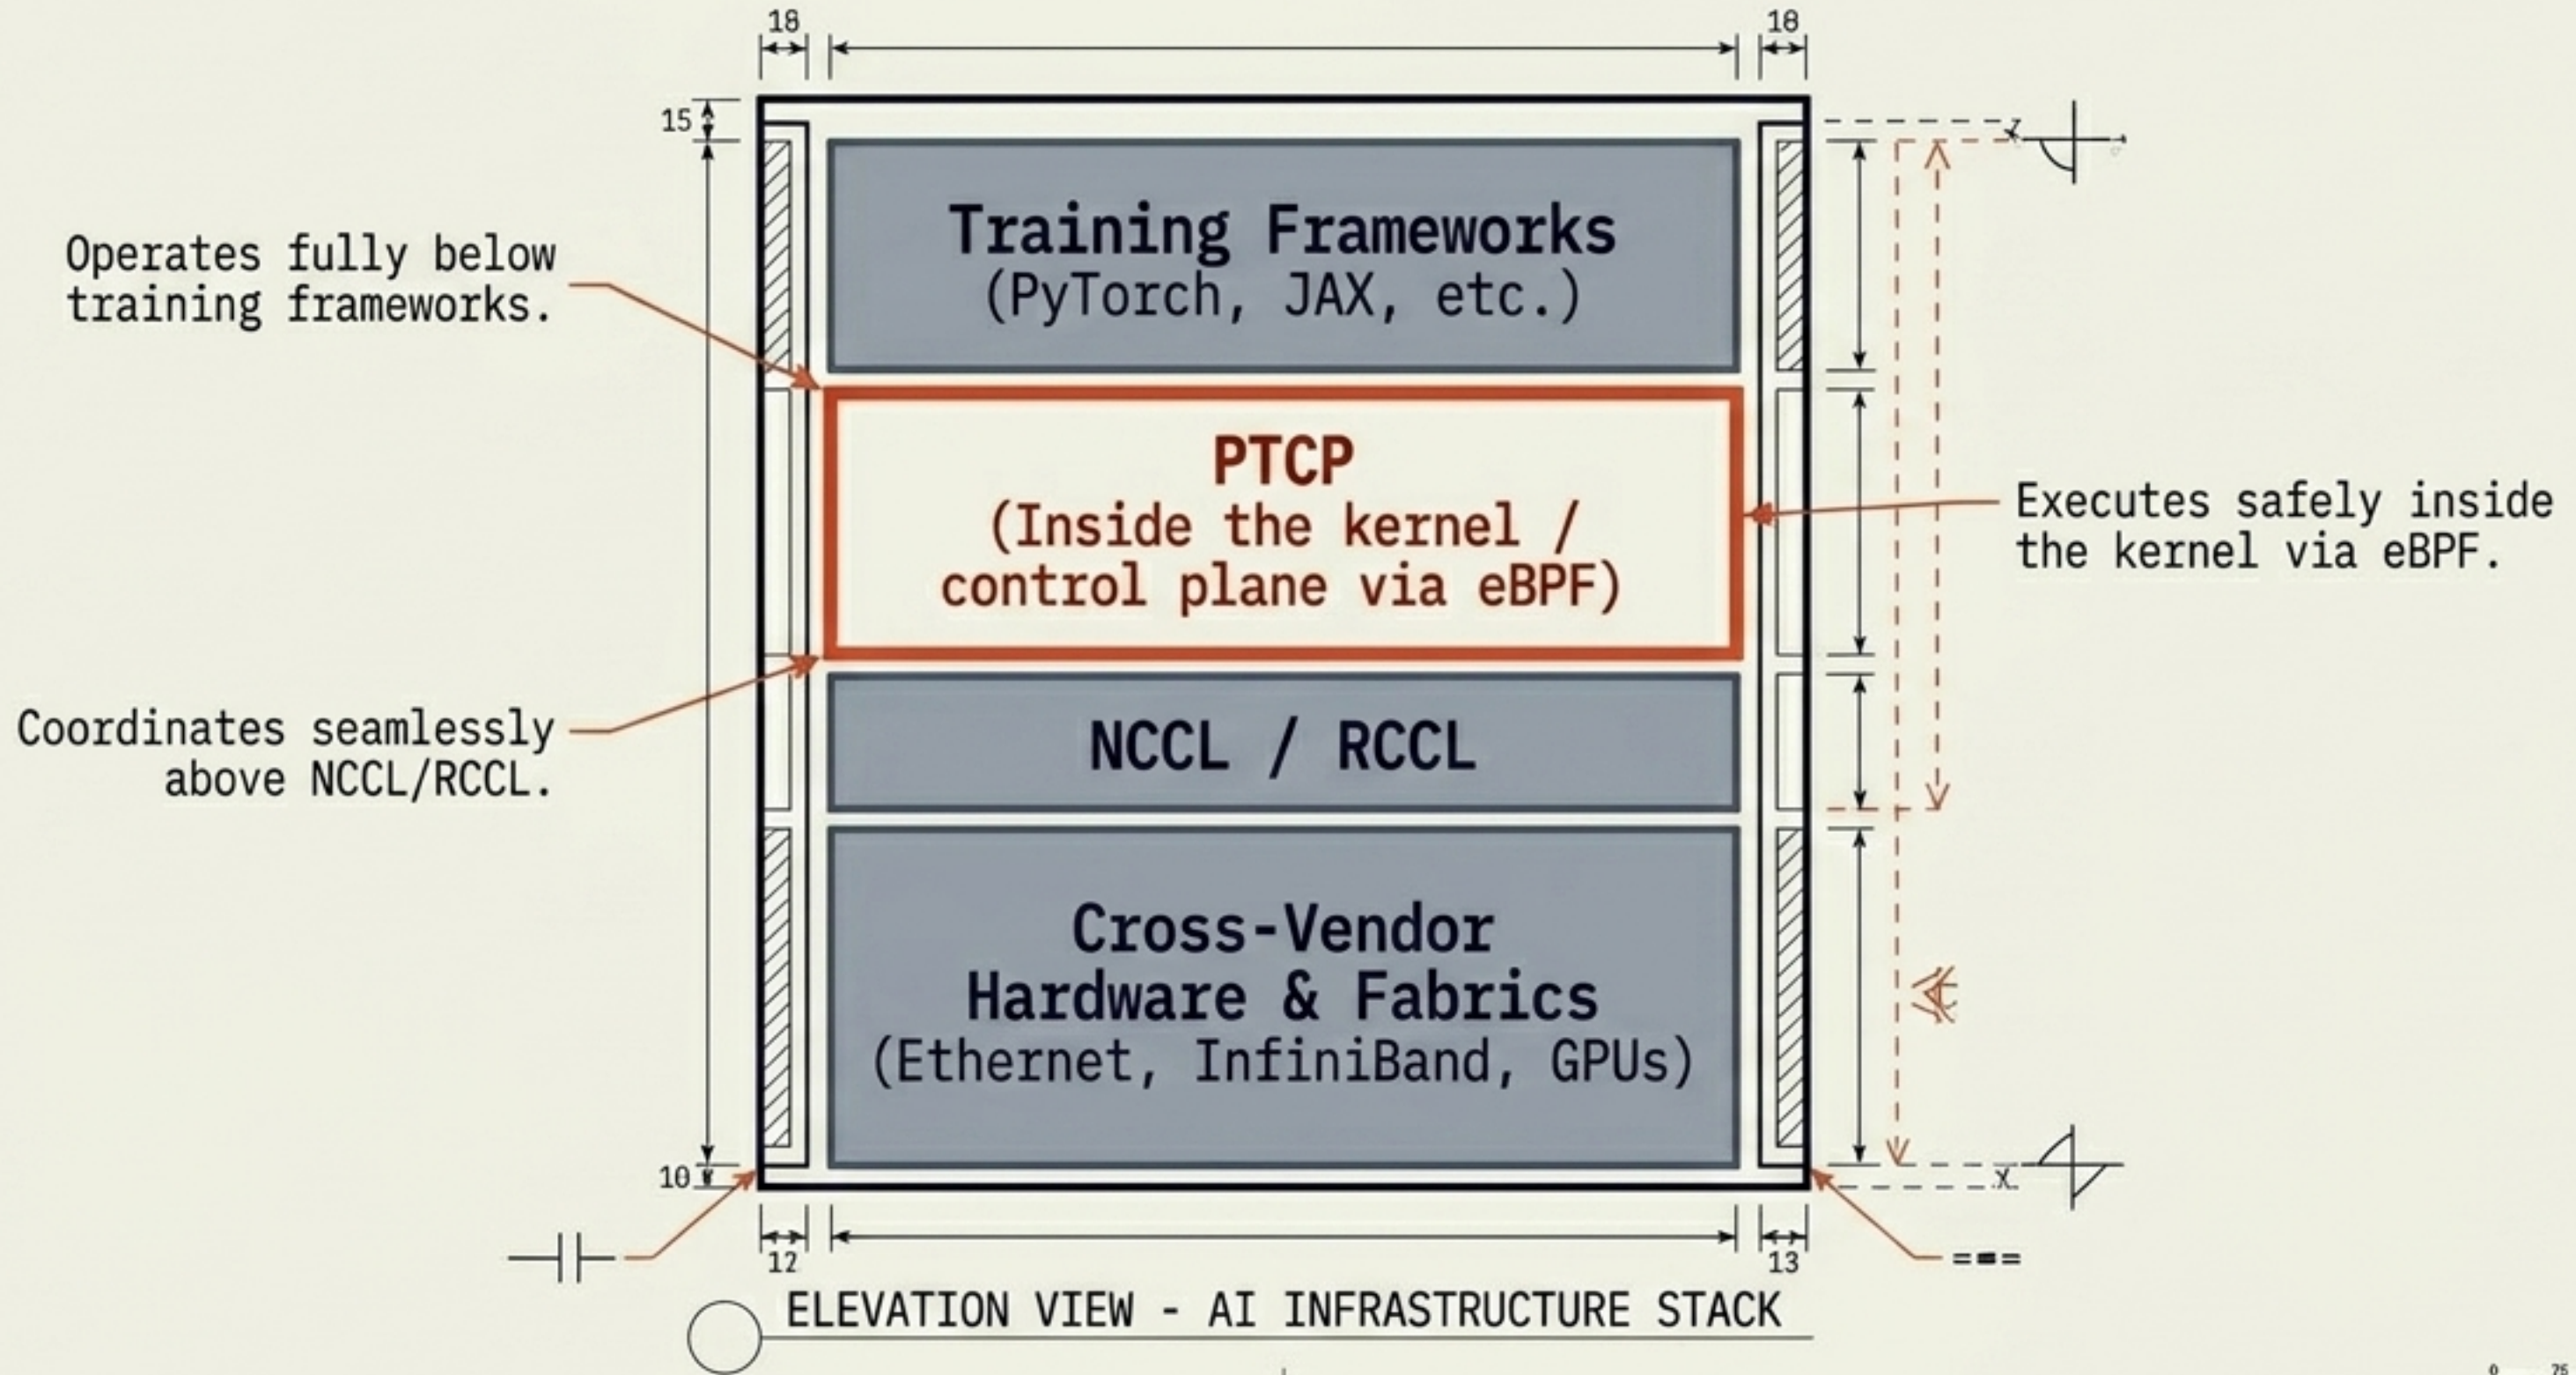
Predictive PTCP



PTCP targets the dominant constraint—model-parallel and data-parallel communication. Instead of relying on legacy reactive congestion control (waiting for a bottleneck to occur), **PTCP** directly reduces bytes moved using **predictive tensor-network compression and scheduling**.

R1.	3	REVISION #1	BATE
S3	3	BPTSS-ABPAT18, 9ET	2624-01-01

The Architectural Stack



ELEVATION VIEW - AI INFRASTRUCTURE STACK

SCALE SECTION

REVISION. 2 2022-05-15

SCALE SECTION. MUST LIST. 13R

R1.	3	REVISION #1	BATE
S8	3	BPTSS-ABPAT18, 9ET	2624-01-01

Zero-Friction Integration




NO Hardware Replacement.

Operates on your existing physical fabric.




NO Framework Rewrites.

Models and training scripts require zero modification.



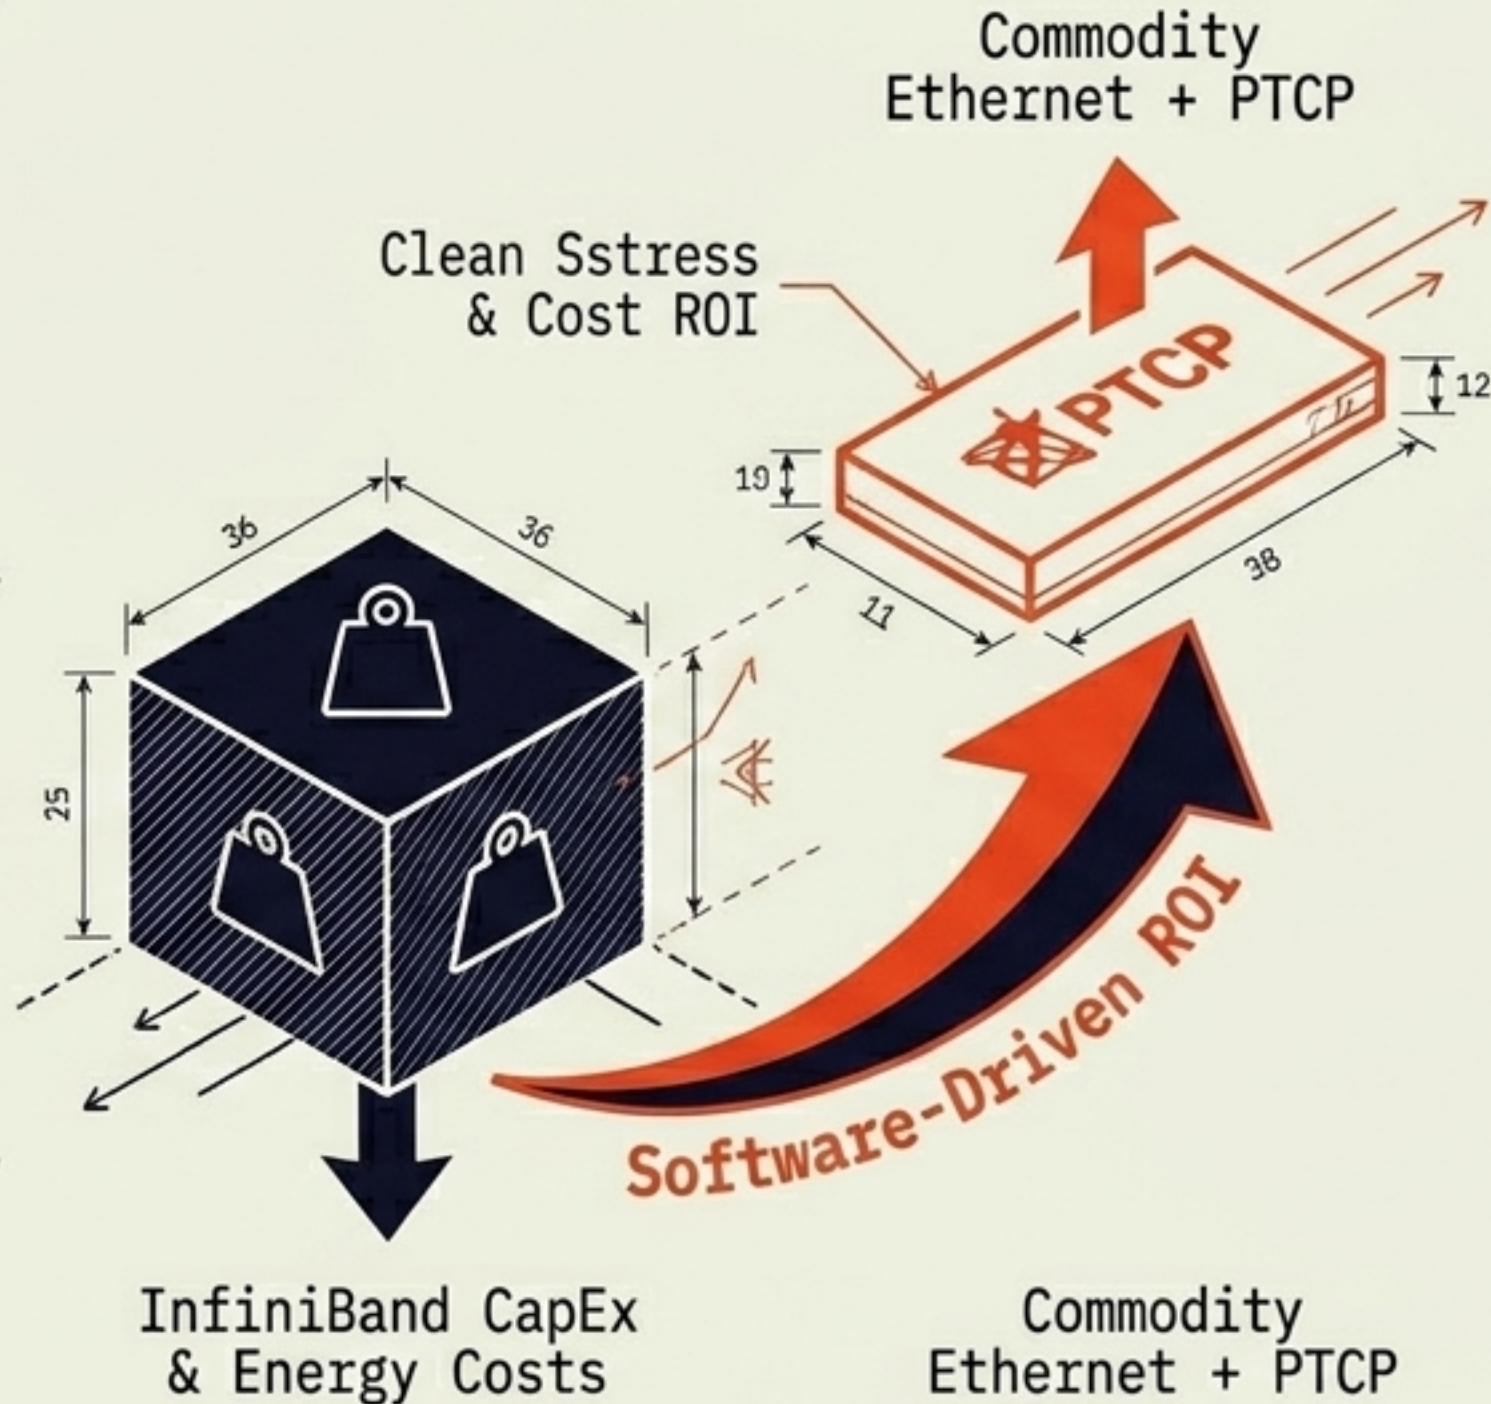
NO Vendor Lock-in.

Agnostic across different accelerator vendors and network fabrics.



RI.	3	REVISION #1	RATE
S8	3	BPTSS-ABPAT18, 9ET	2624-D1-01

The New Economics of AI Infrastructure



Commodity Foundations

Runs optimally on commodity Ethernet, breaking dependence on premium proprietary fabrics.

Asset Extension

Extends the operational lifespan and viability of existing GPU fleets.

Financial Transformation

Converts crushing hardware CapEx pressure into scalable software ROI.


Energy Efficiency

Mitigates the dominant operational cost: energy per token.

R1.	3	KEVISION #1	DATE
S8	3	BPTSS-ABPAT18, 9EY	2024-D1-01

Quantitative Protocol Impact

15-40%+

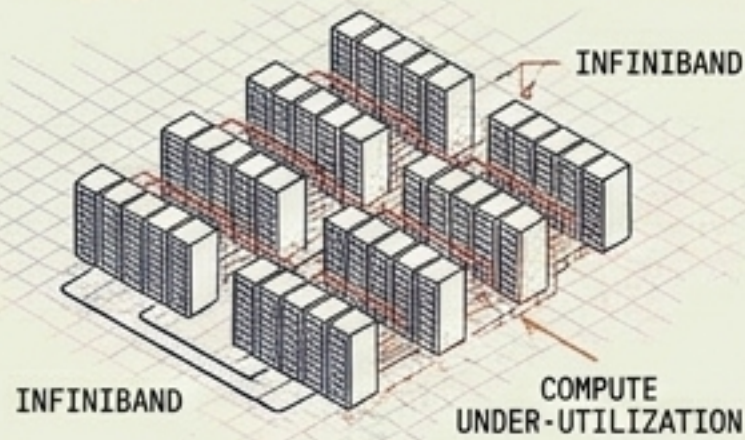


Moving 40% fewer bytes over the fabric translates directly to faster training iterations, higher GPU utilization, and drastically lower power consumption at the exact same mathematical accuracy.

0 15 100 200

Target Environments: Where to Deploy Today


Hyperscale Clusters



Organizations fighting compute under-utilization where **InfiniBand costs** have become unsustainable.

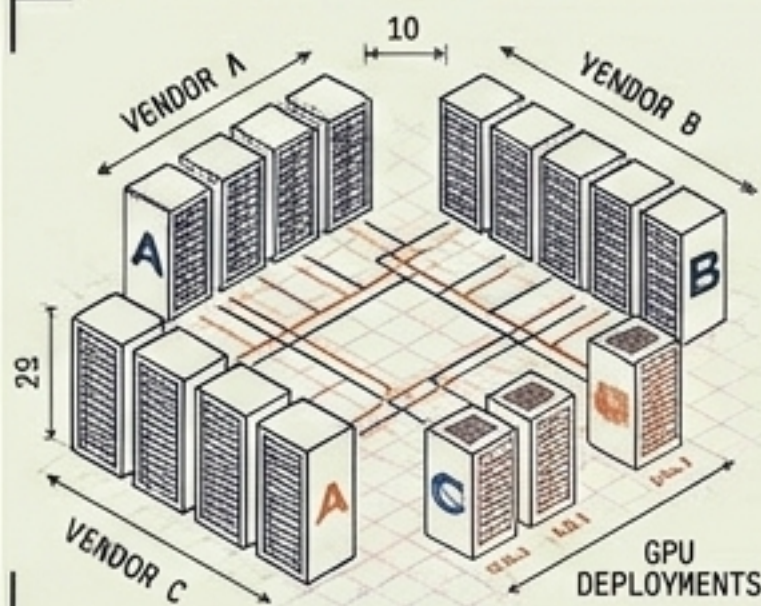
Constrained Power DCs

Facilities hitting physical power and cooling limits, where **energy cost per token** is the primary bottleneck.



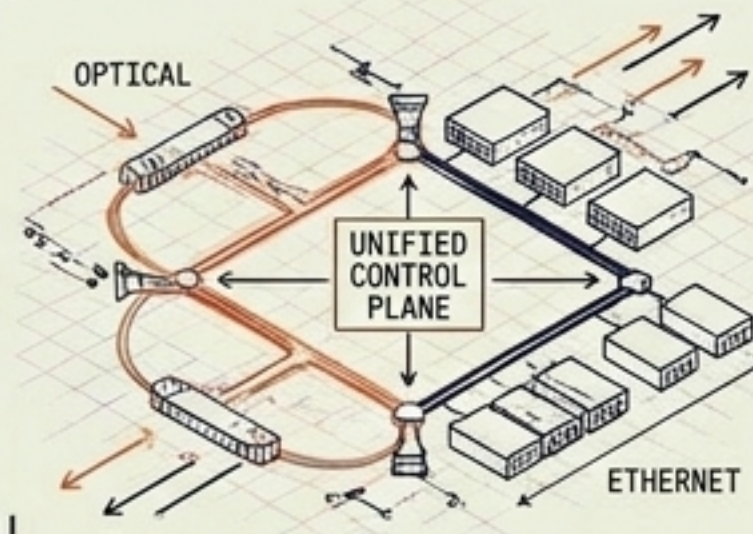
Multi-Vendor Fleets

Entities actively seeking **hardware independence** and avoiding vendor lock-in across their GPU deployments.

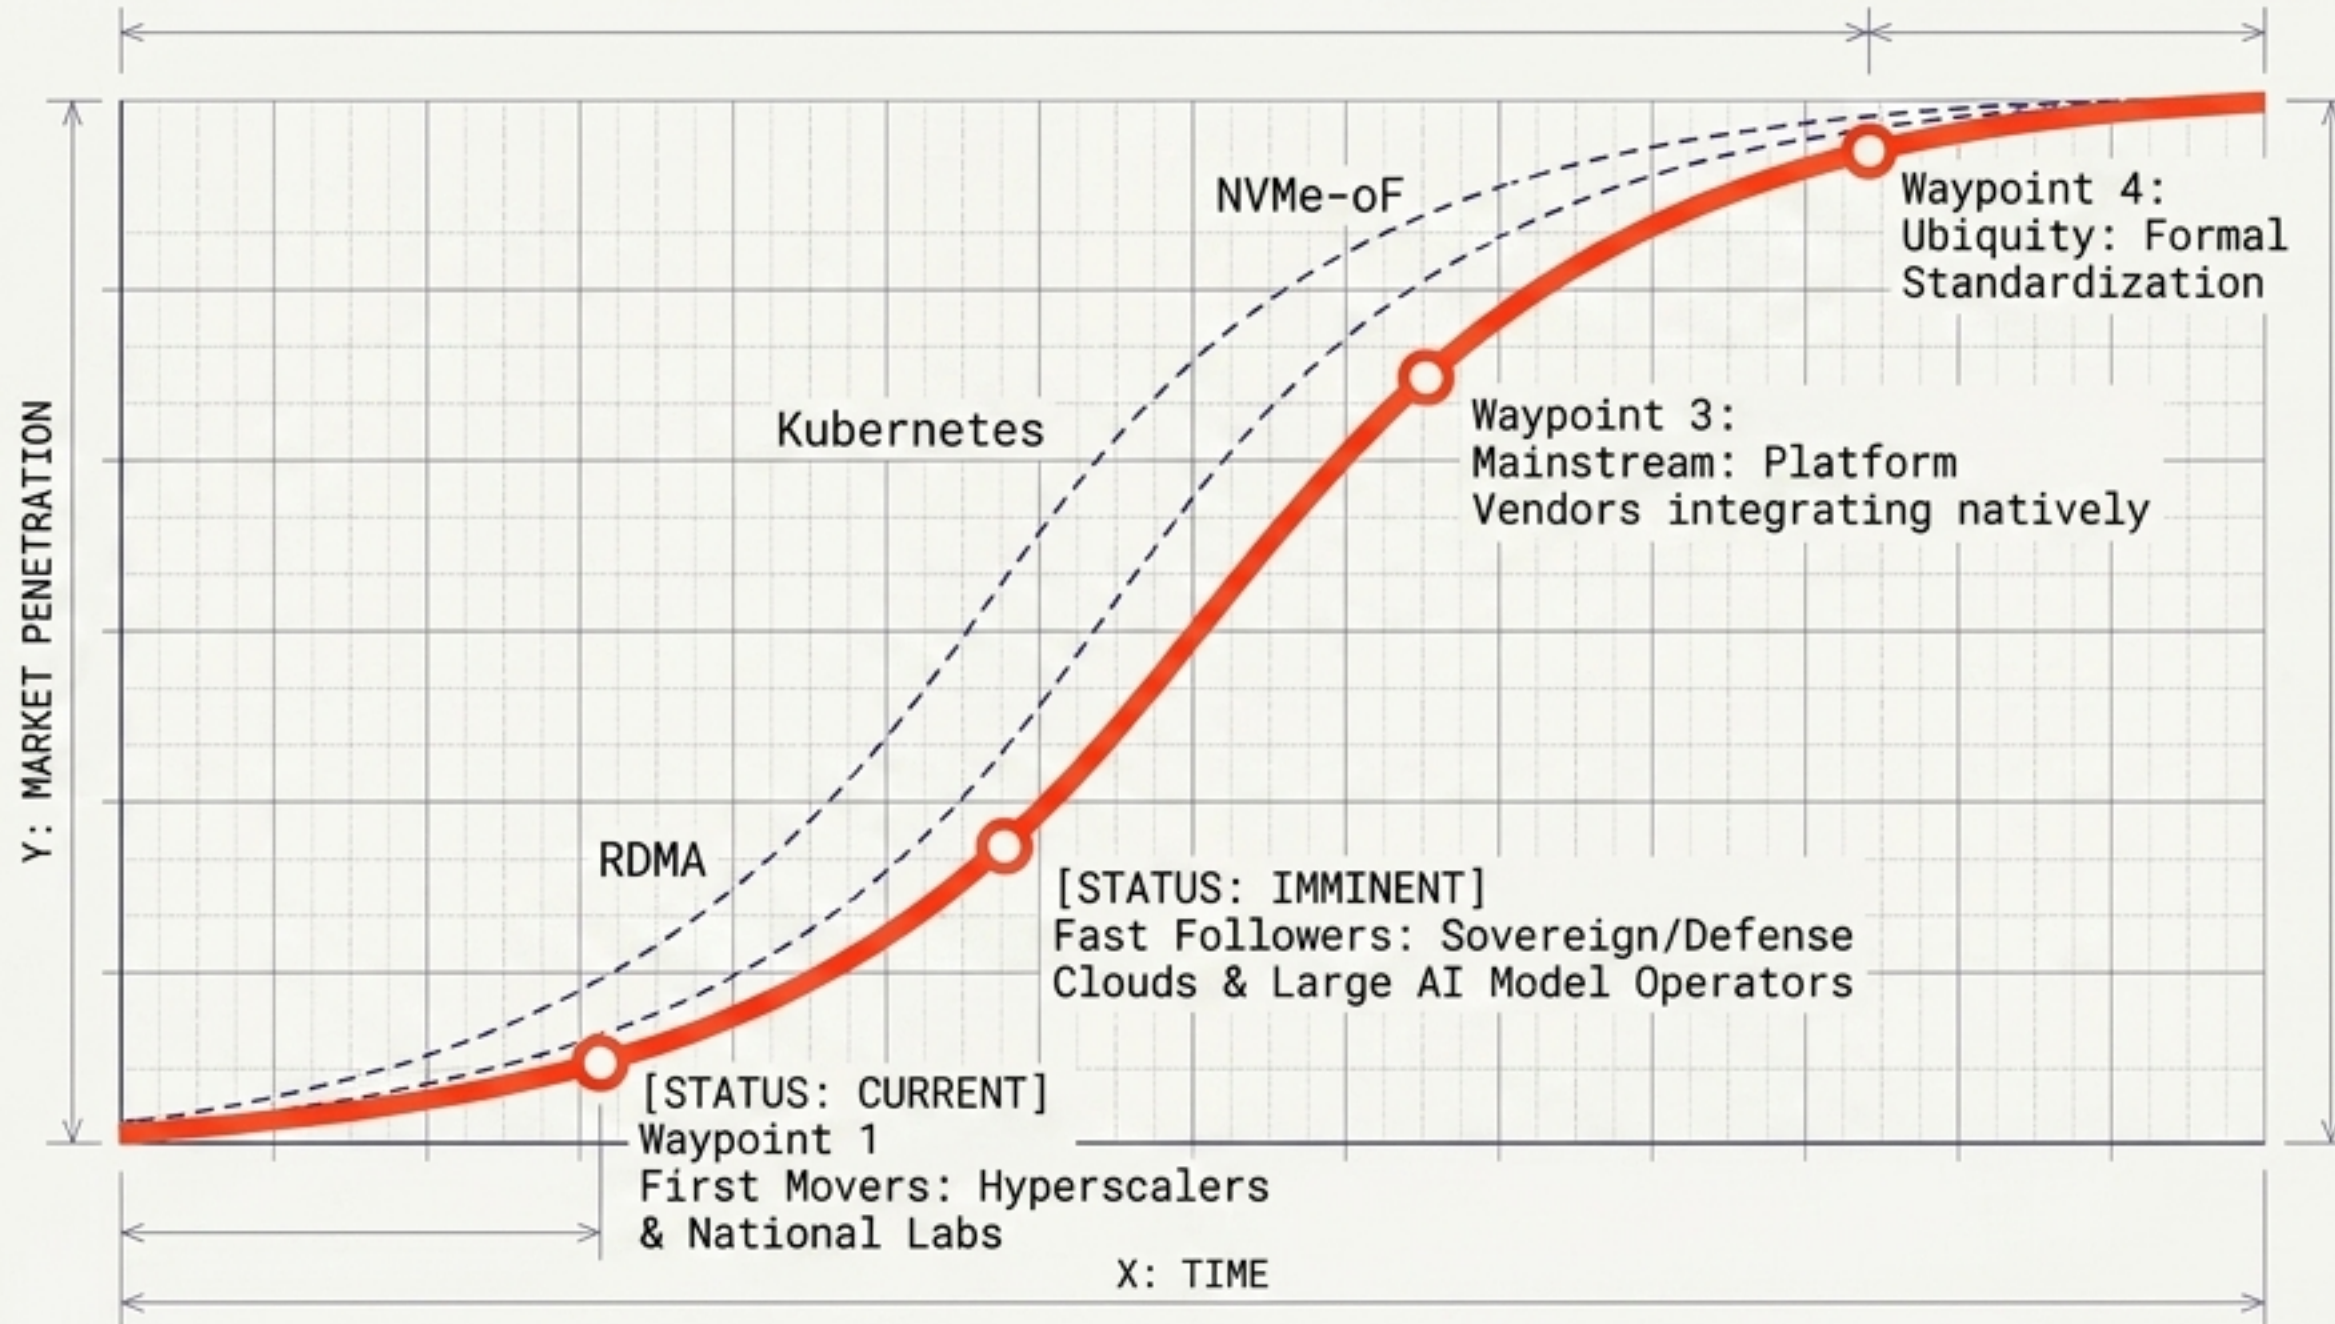


Hybrid Fabrics

Cutting-edge **optical-Ethernet hybrid fabrics** that require a unified control plane to smooth out coordination discontinuities.

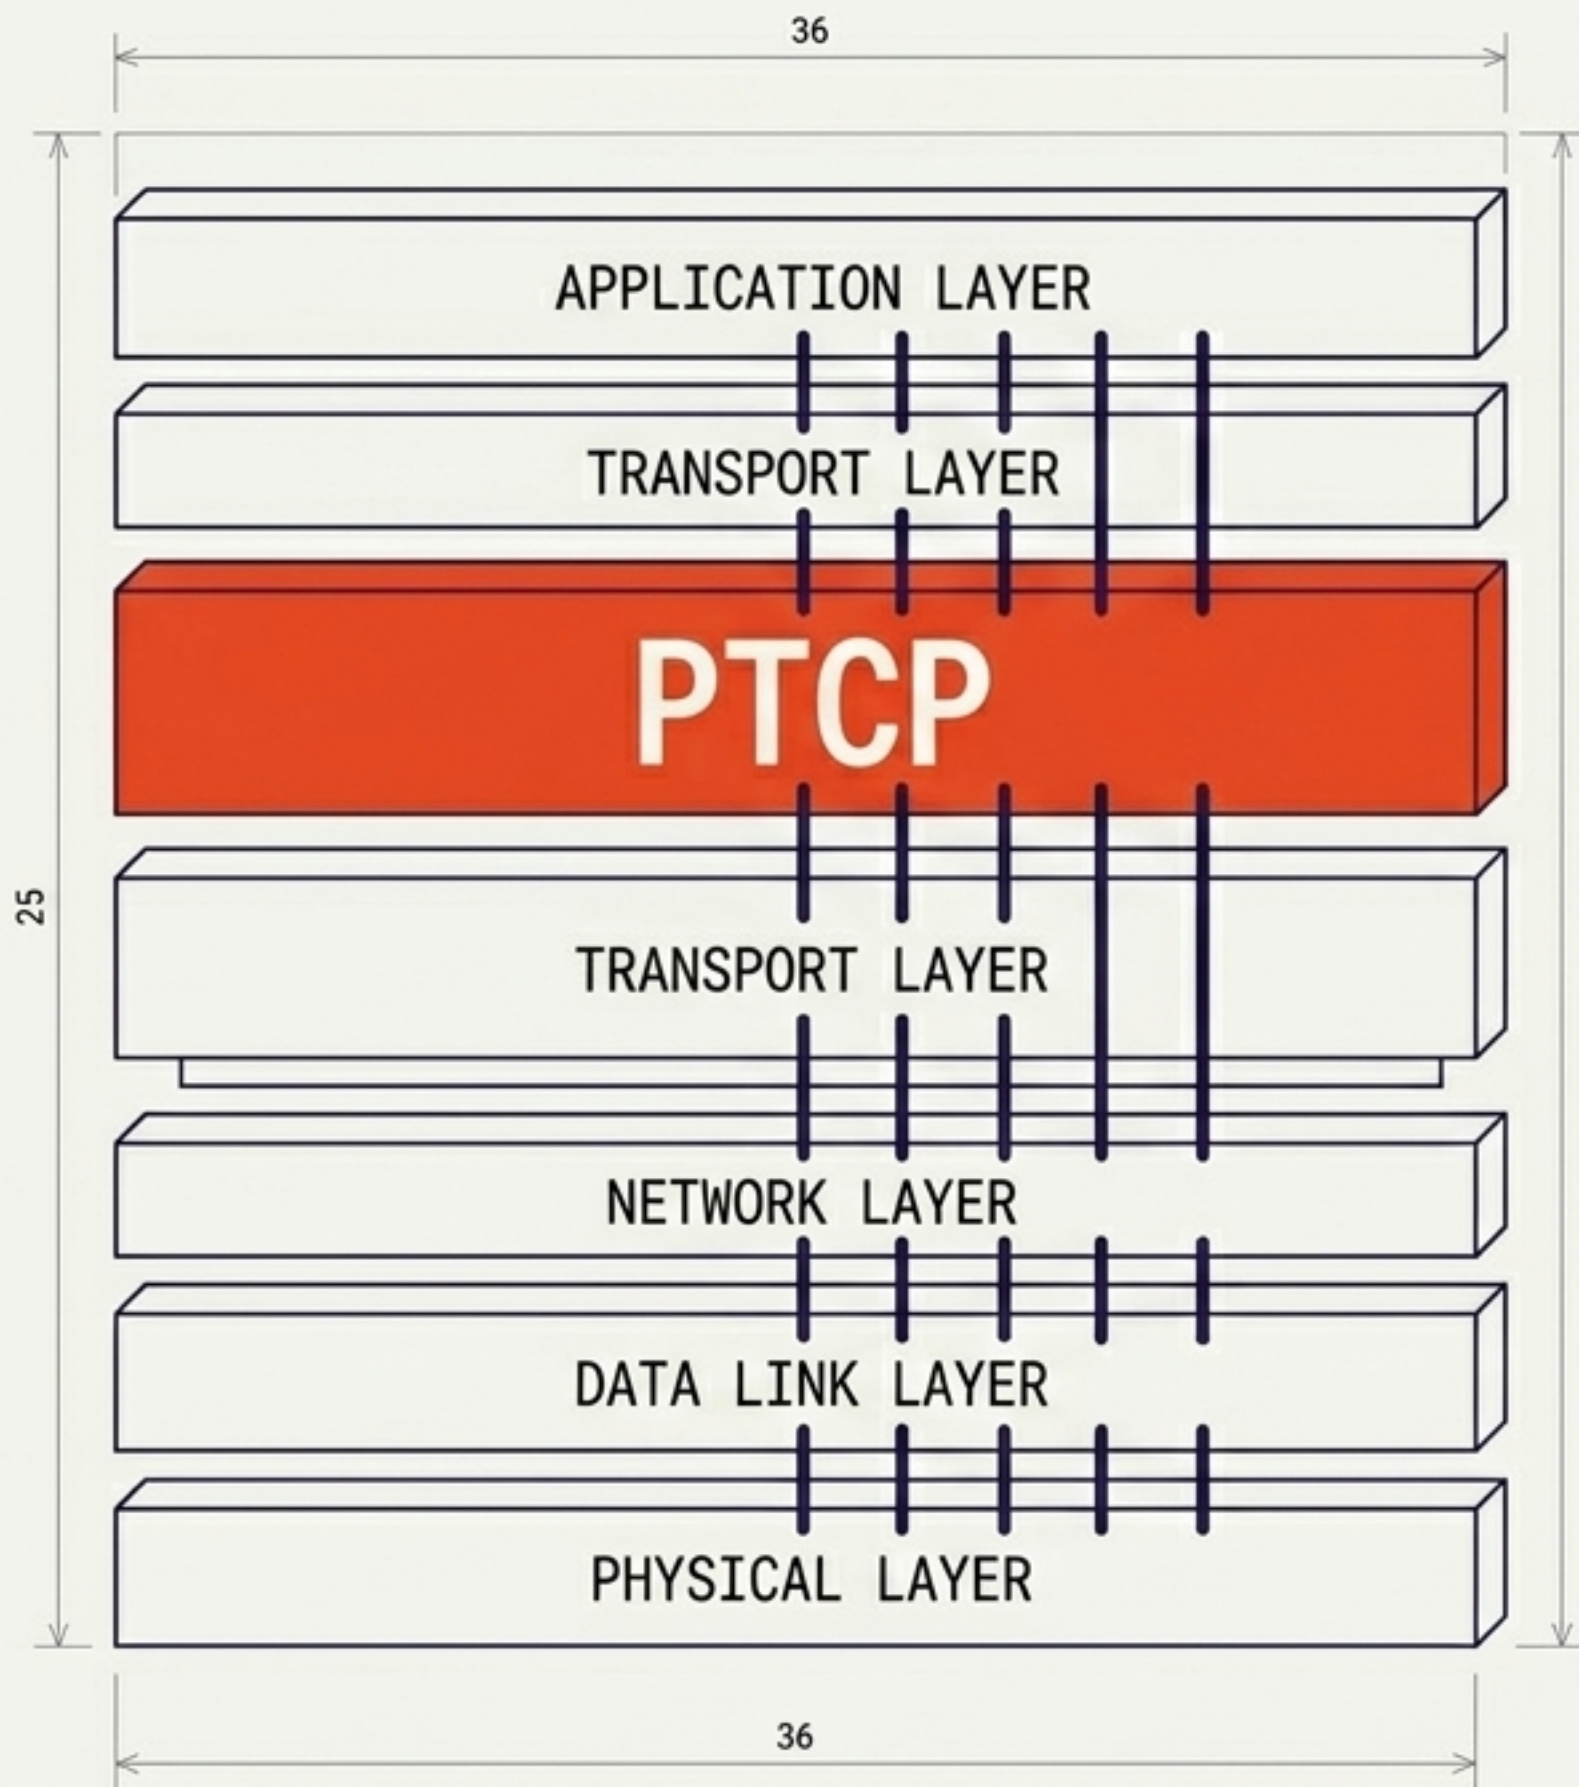


The Adoption Precedent



Infrastructure owners dictate the protocol layer. The sequence is already in motion.

SCALE	SCALE	REV-S.
0 06	1	



The New Coordination Layer

AI scaling has reached its physical and economic communication limits.

Hardware alone cannot bridge the gap.

PTCP is the inevitable protocol-level layer that coordinates hyperscale AI.

Runs on commodity Ethernet to unlock the next generation of model scaling.

SYSTEM STATUS: COORDINATED.