


Scale 100



Scale 10

Scale 10

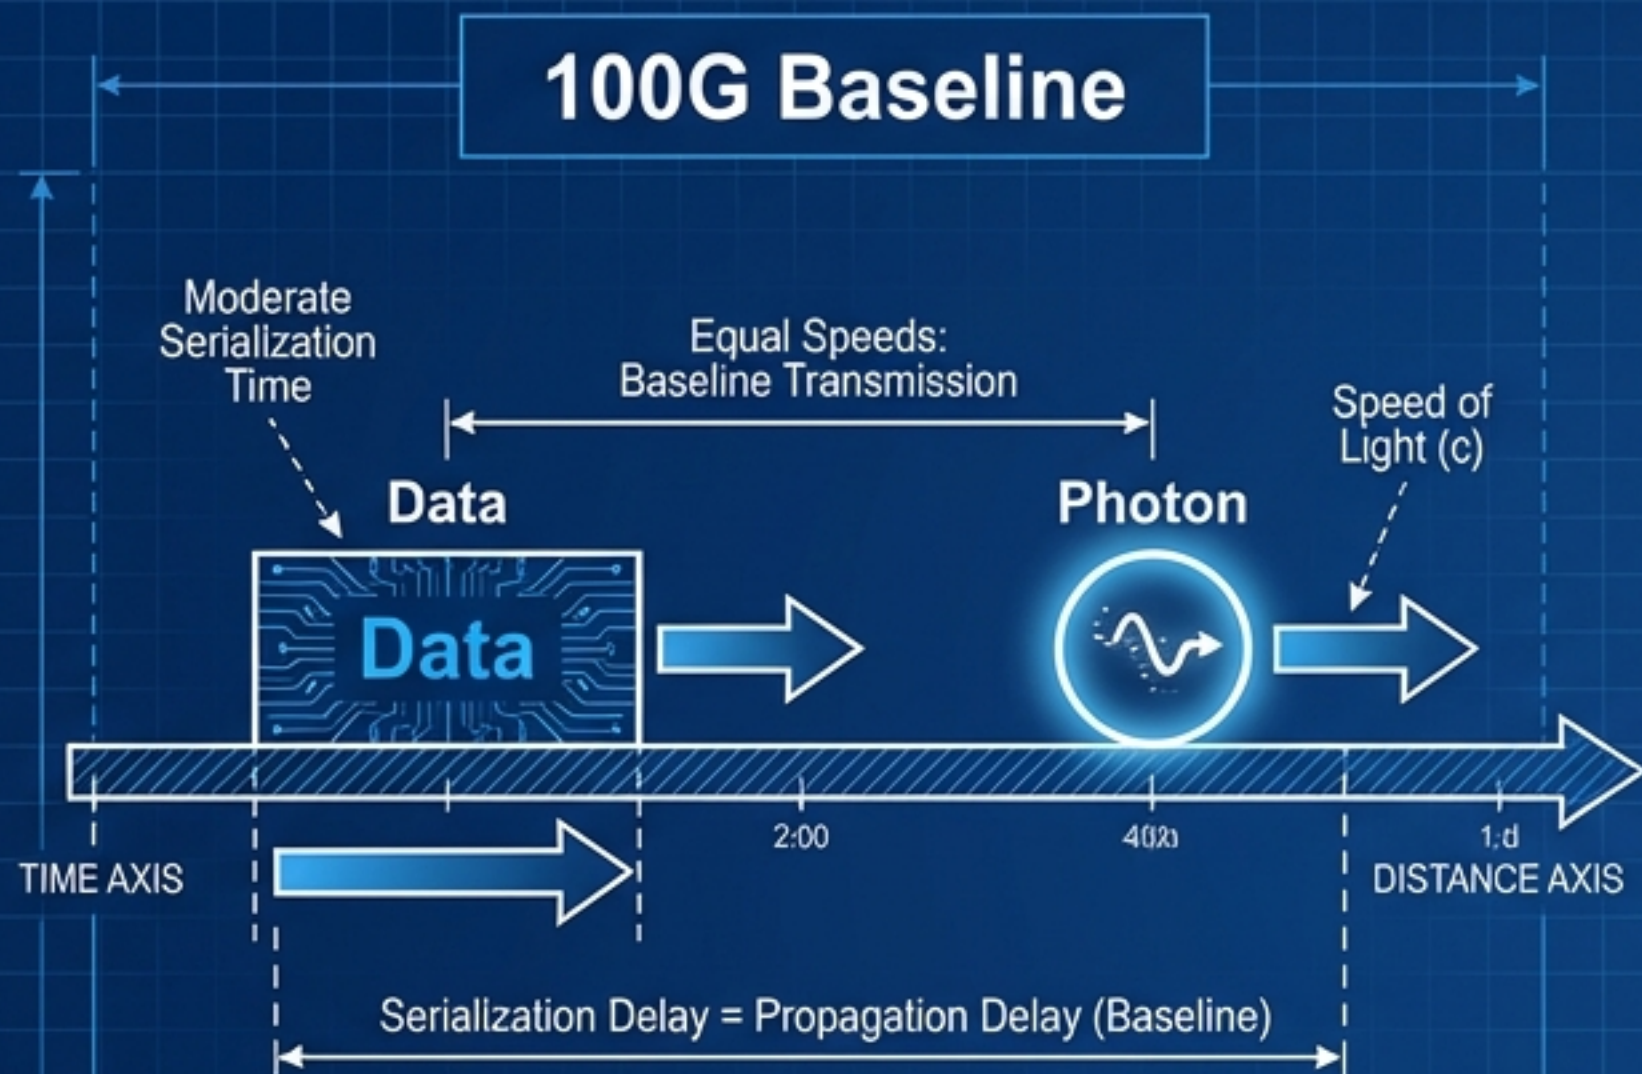
$$N_c = E \left( \frac{3/2}{P/LG} - \frac{10^{-7} N^2}{10^{-7} N^2} \right)$$

# The 800G Speed of Light Wall

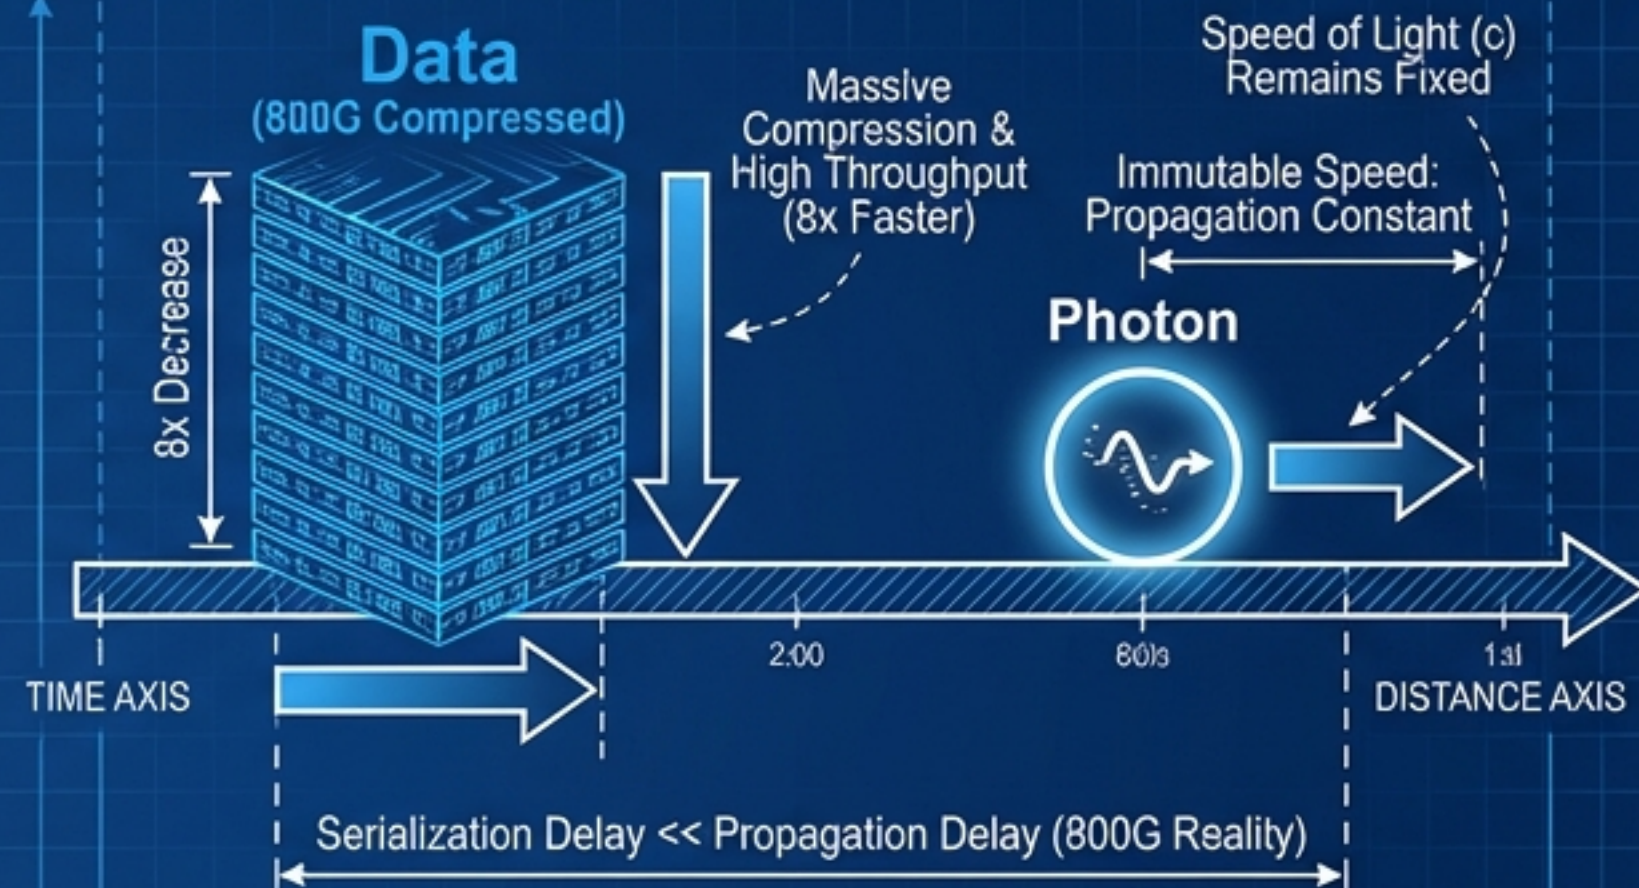
The physics of exascale AI, the Bandwidth-Delay Product crisis, and the transition to predictive network determinism.

# Bandwidth scales. Physics does not.

## 100G Baseline




## 800G Reality




① **Serialization Delay:** Decreased by 8x (Data hits the wire at ~100GB/s).



② **Propagation Delay:** Constant (Fixed by the speed of light in optical fiber).



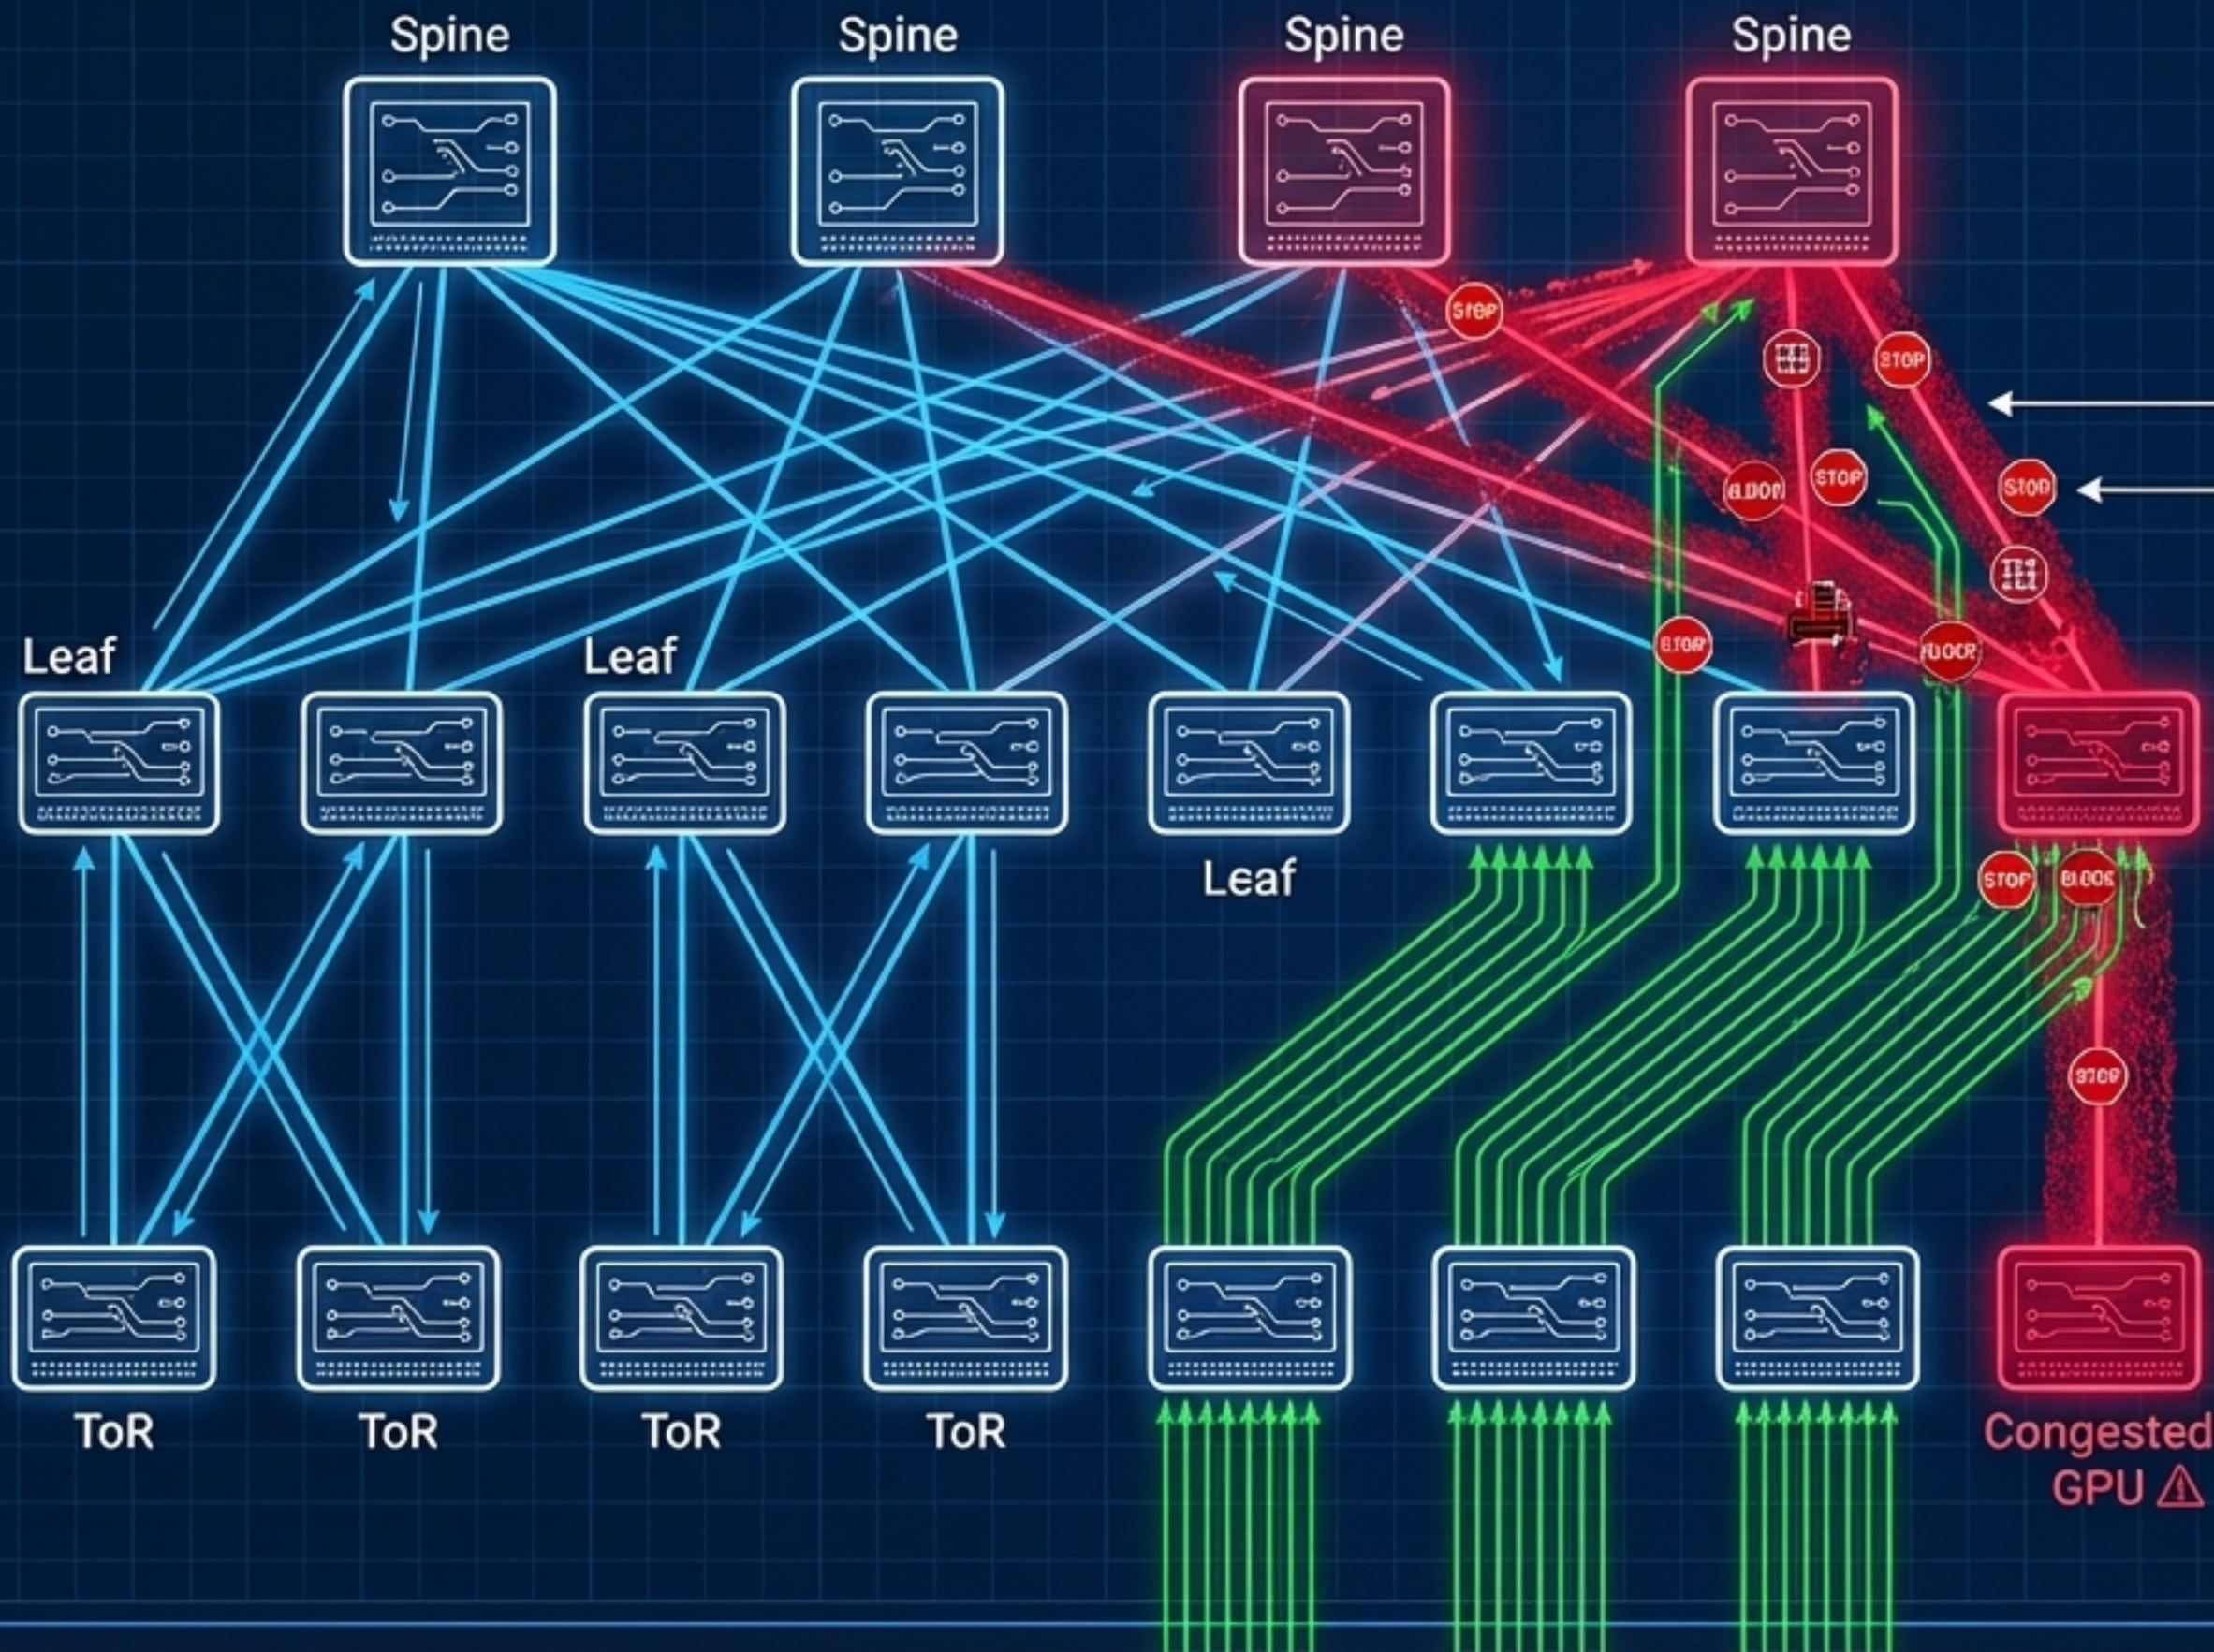
# The Bandwidth-Delay Product (BDP) Explosion



# The Failure of Reactive Telemetry (RoCEv2)



# InfiniBand's Exascale Bottleneck

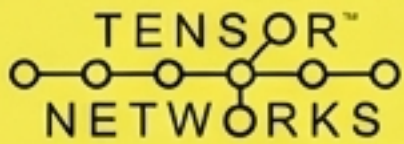


1 **Hardware Cryptographic Credits** prevent packet loss but introduce Congestion Spreading.

2 **Head-of-Line (HoL) Blocking:** One slow GPU stalls unrelated "victim flows" across a \$1 billion supercomputer.

3 **Credit RTT Penalty:** Waiting for hardware credits introduces severe jitter and tail-latency.

# The Interconnect Diagnostic Scorecard

	<b>NVIDIA InfiniBand (Quantum-X800)</b>	<b>Standard RoCEv2 Ethernet (800G)</b>	 <b>Ethernet + PTCP (800G)</b>
Flow Control	Credit-based (Hardware)	Reactive (PFC/ECN)	<b>Predictive</b> (O(1) eBPF Pacing)
Packet Loss	Zero	High	<b>Zero</b> (Mathematically Prevented)
Tail Latency	Moderate	Extreme	<b>Absolute Minimum</b>
Topology	Strict Fat-Tree	Flexible	<b>Seamless Multi-pathing</b>
AI Throughput Yield	~85-90%	~60-75%	<b>&gt;95% (Max Theoretical)</b>


PROJECT:  
INTERCONNECT DIAGNOSTIC SCORECARD

DESIGN NO:  
DWG-IC-SCORECARD-001

SCALE:  
1:1 (SCHEMATIC)


# Redefining the Control Plane

## Legacy Approach




- Post-Congestion Timing
- Hardware Buffer Thresholds
- Constrained by the past

## PTCP Approach



- Pre-Congestion Timing
- $O(1)$  eBPF Kernel Logic
- Constrained by the mathematical future

# The Math of Foreknowledge



**Pattern-of-Life:** AI workloads utilize static computation graphs, making traffic bursts 100% predictable.

**PoL-TT Mathematics:** Tensor Train math maps the synchronization schedule of the GPUs globally.


**Bounding Policy Envelope:** The exact deterministic path and timing are calculated before a single packet is serialized.

PROJECT: DETERMINISTIC FOREKNOWLEDGE MAPPING

DRAWING NO: DWG-DFM-003

SCALE: 1:1 (SCHEMATIC)

# Bypassing the Switch via eBPF



**1. Intercept:** Agents catch 800G bursts at the host kernel layer.

**2. Assign:** The Earliest Departure Time (EDT) is applied based on the orchestrator's map.

**3. Pace:**  $O(1)$  kernel logic completely removes the need for hardware-level telemetry.




BLUEPRINT

PROJECT:  
INTERCONNECT CONTROL PLANE EVOLUTION

DRAWING NO:  
DWG-IC-CONTROL-PLANE-004

SCALE: 1:1 (SCHEMATIC)

# The 'Micro-Gap' Interleaving Engine



**Nanosecond** pacing creates **Zero Queueing Delay**. Switches operate as pure transit, mathematically preventing buffer utilization.



PROJECT:  
INTERCONNECT CONTROL PLANE EVOLUTION  
DRAWING NO:  
DWG-IC-CONTROL-PLANE-005  
SCALE: 1-1 (SCHEMATIC)

# Eradicating the **InfiniBand Tax**

## The InfiniBand Tax




### Proprietary Hardware & Single-Vendor Supply Chains

Decouples high-performance determinism from proprietary hardware dependencies.

Enables an open, multi-vendor ecosystem (Broadcom, Cisco, Marvell).

## Predictive Ethernet



### Vendor-Neutral Nodes

(Broadcom Tomahawk 5, Cisco Silicon One)

Drives an immediate 30% to 40% reduction in network CapEx.


## BLUEPRINT

PROJECT:  
INTERCONNECT CONTROL PLANE EVOLUTION

DRAWING NO:  
DWG-IC CONTROL-PLANE-006

AVVE:

# Future-Proofed for Exascale



As wire speed doubles to 1.6T, the complexity of PTCP's  $O(1)$  kernel logic remains completely constant.

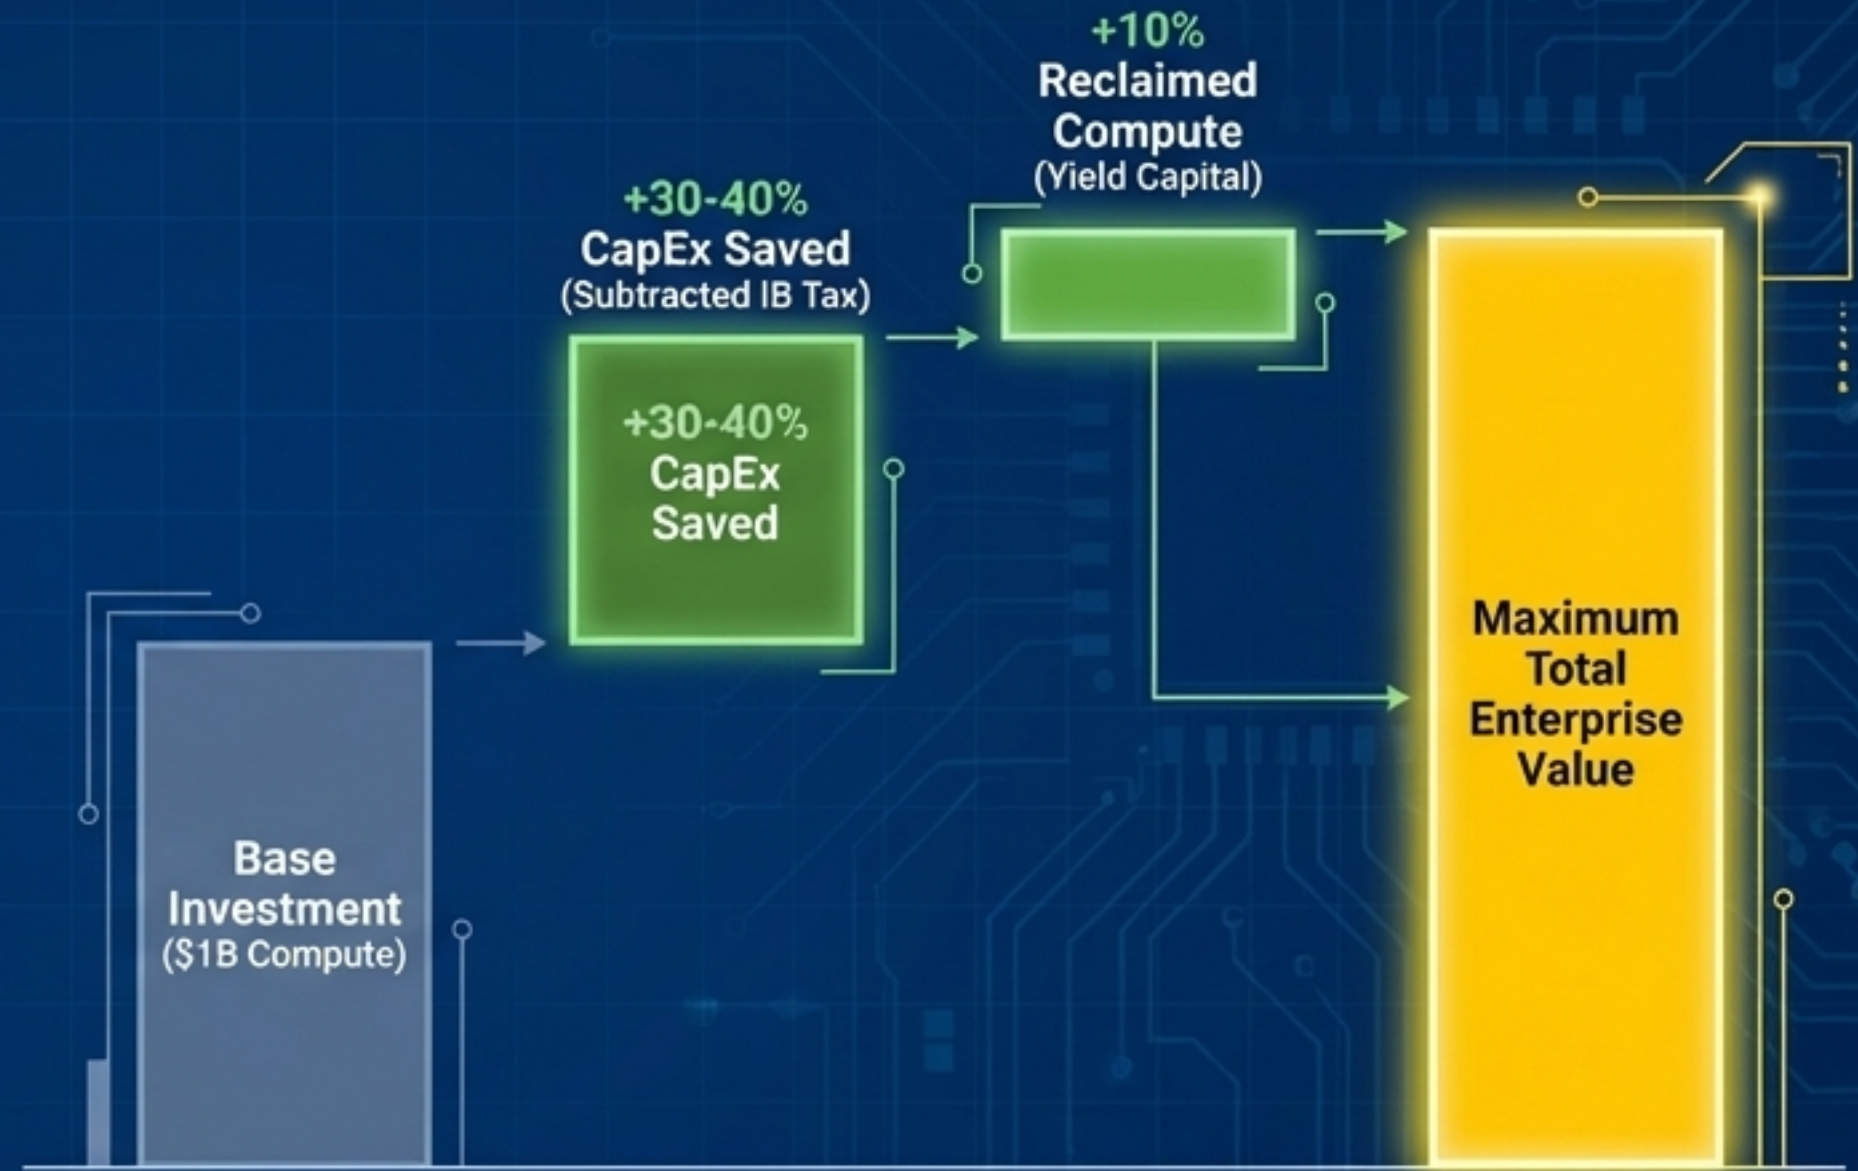
The architecture **scales infinitely** against the limits of the physical world.



PROJECT:  
INTERCONNECT CONTROL PLANE EVOLUTION  
DRAWING NO:  
DWG-IC-CONTROL-PLANE-007  
SCALE: 1:1 (SCHEMATIC)


# The Economics of Nanosecond Math

connective thread



**Reclaiming just 10% of idle GPU time yields \$100 Million in capital efficiency per \$1 Billion invested**

# At Exascale, Network Physics is Data Center Economics.



**The shift from reactive telemetry to predictive determinism is not an upgrade.  
It is the mandatory architectural foundation for planetary-scale AI.**