turba

Precision Timing as the Backbone of AI Hyper Scale

Ahmad Byagowi v1.0.3

May 2025

Content

- Advances in Precision
- Success
- Risk
- AI Workloads
- Next Steps
- Call to Action

Wait a Second

Standardization has made impacts. As you know, the present definition of the Second according to International System of Units (SI) is defined as the duration of 9,192,631,770 periods corresponding to the transition between two hyperfine levels of the ground state of the cesium-133 atom.

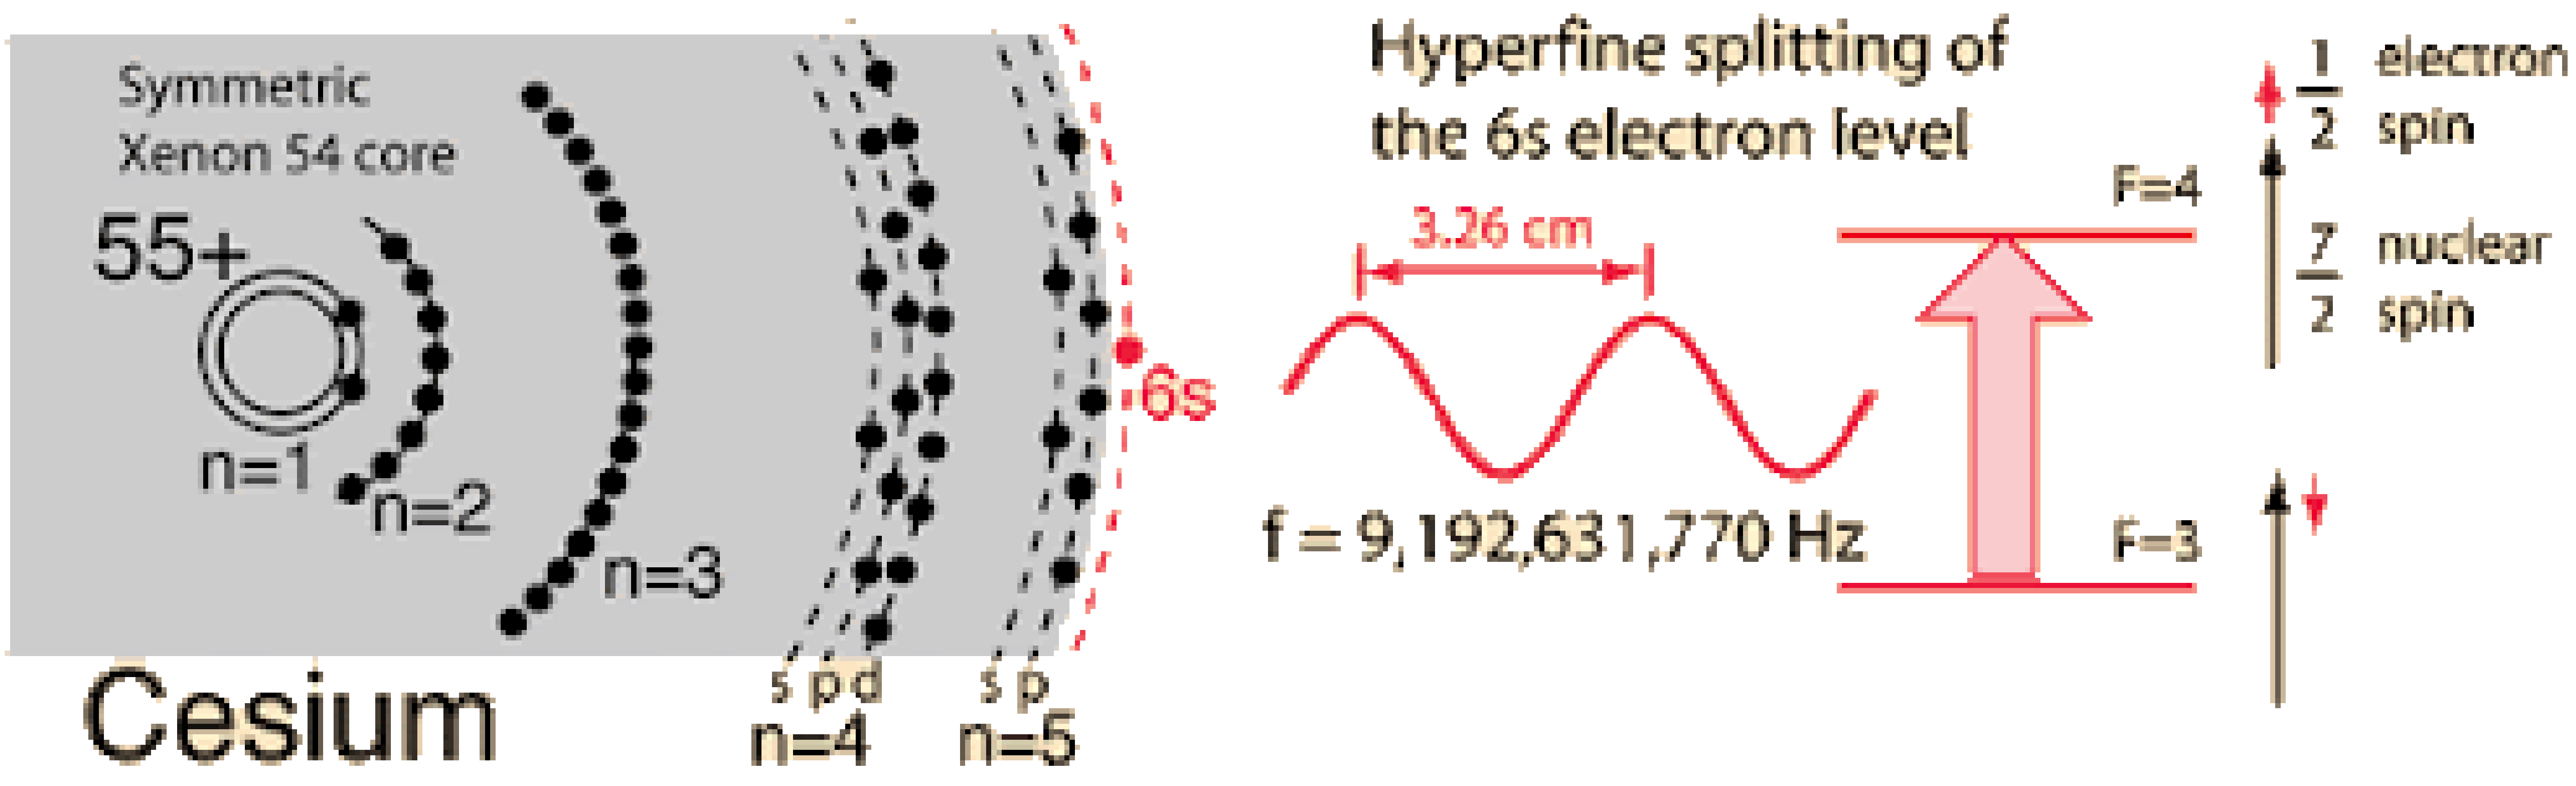

This level of precision has been a solid foundation to build today's technology.
It's just that Cesium clocks were not found on every playground, initially.

Synchronizing that Second

With the precise second, further standardization of accurate delivery followed:

- **Network Time Protocol (NTP)** followed by **Precision Time Protocol (PTP).**
- Further improvements such as **White Rabbit/SyncE, Precision Time Measurement (PTM)** are opening new doors.

This has led the way:

- A vibrant ecosystem of vendors to make these products **widely** available.
- Furthermore, open specifications for precision timing devices are available for all to tinker and build with. (Example: Open Compute Project, TAP)

Precision Timing, built on standardizations, enables markets.

State of Practice vs. State of Need

Review of Successes in Precision Timing: The Good

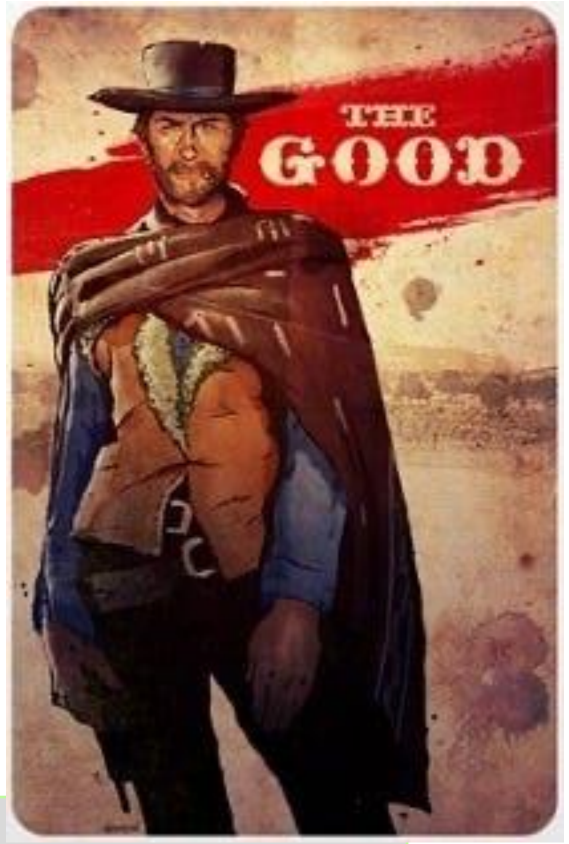

The existence of PTP (μs-level) and White Rabbit (sub-ns-level) precision **does not mean synchronization is a solved problem**; in practice, every time we ratchet precision up, new bottlenecks appear. Firstly, examples of the improvements!

Layer	What's improved with today's tech	What still breaks	Why it matters?
Physical / Network	Hardware timestamping in NICs plus transparent clocks gives sub-μs accuracy at light-to-moderate load.		
Scale / Topology	Open Time Server + PTP4U has shown **>1 M clients** on a single grand-master tree. Large AI Clusters of >100K accelerators.		
Security / Resilience	MACsec, ACLs, and PTP security TLVs help authenticate masters.		

The Bad

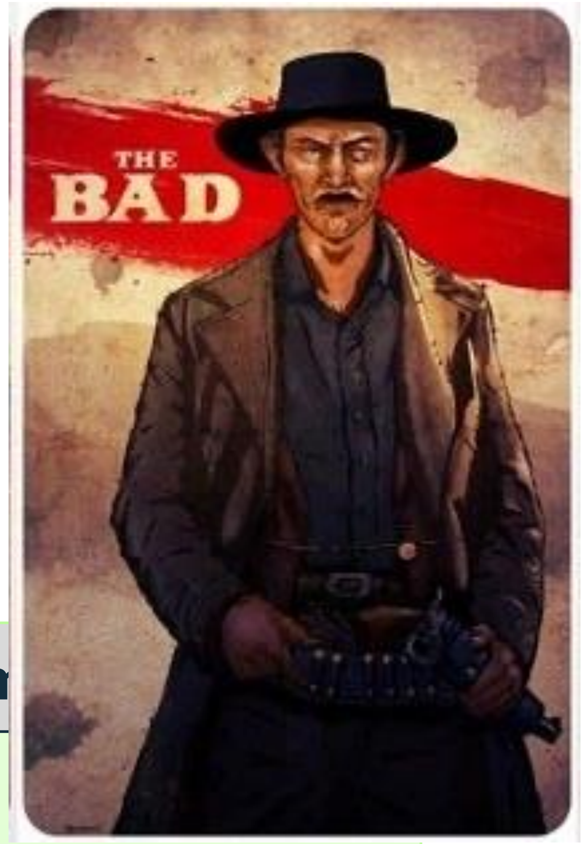

Layer	What's improved with today's tech	What still breaks	Why it m
Physical / Network	Hardware timestamping in NICs plus transparent clocks gives sub-µs accuracy at light-to-moderate load.	• Accuracy degrades by two–three orders of magnitude under micro-bursts or buffer contention. One study measured **360 µs–1 ms jitter** through commodity switches when queues build up • Fiber delay changes ±tens ns/km / °C; long White Rabbit links drift unless continuously temperature-trimmed	
Scale / Topology	Open Time Server + PTP4U has shown **>1 M clients** on a single grand-master tree. Large AI Clusters of >100K accelerators.	• **Boundary-clock hierarchies add noise** each hop; large brown-field networks can't enable BC mode safely • Today's designs assume homogeneous Ethernet; coming CXL/PCIe-over-fabric, optical circuit switches, and wireless edge tiers break that assumption.	
Security / Resilience	MACsec, ACLs, and PTP security TLVs help authenticate masters.	• **Delay-and-spoof attacks** can offset a 5G base-station or O-RAN cell by >2 • Insider spoofing can silently skew slaves by >100 µs without tripping alarms • If GNSS is jammed, quartz clocks walk off by >1 ms inside an hour; even large hyperscalers now bolt CSACs onto grandmasters for hold-over	

The Ugly

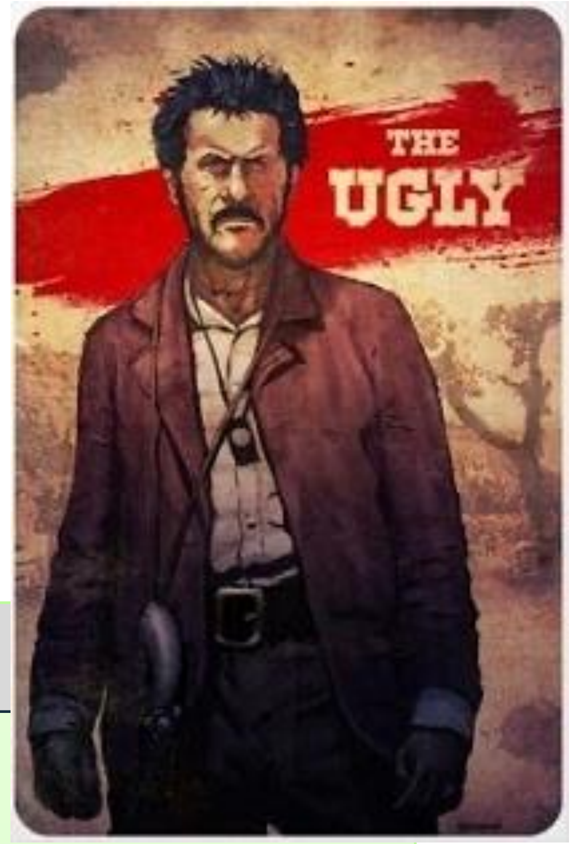

Layer	What's improved with today's tech	What still breaks	Why it matters?
Physical / Network	Hardware timestamping in NICs plus transparent clocks gives sub-µs accuracy at light-to-moderate load.	• Accuracy degrades by two–three orders of magnitude under micro-bursts or buffer contention. One study measured **360 µs–1 ms jitter** through commodity switches when queues build up • Fiber delay changes ±tens ns/km / °C; long White Rabbit links drift unless continuously temperature-trimmed	AI fabrics routinely push Tb/s across leaf–spine networks. Congestion episodes of a millisecond shove GPUs into idle-wait states, cutting utilization.
Scale / Topology	Open Time Server + PTP4U has shown **>1 M clients** on a single grand-master tree. Large AI Clusters of >100K accelerators.	• Boundary-clock hierarchies add noise each hop; large brown-field networks can't enable BC mode safely • Today's designs assume homogeneous Ethernet; coming CXL/PCIe-over-fabric, optical circuit switches, and wireless edge tiers break that assumption.	Next-gen clusters mix rack-scale optical fabrics with electrical backplanes and remote memory pools. They need picosecond-class alignment across heterogeneous media—well beyond current PTP profiles.
Security / Resilience	MACsec, ACLs, and PTP security TLVs help authenticate masters.	• **Delay-and-spoof attacks** can offset a 5G base-station or O-RAN cell by >2 • Insider spoofing can silently skew slaves by >100 µs without tripping alarms • If GNSS is jammed, quartz clocks walk off by >1 ms inside an hour; even large hyperscalers now bolt CSACs onto grandmasters for hold-over	AI governance and forensics depend on "ground truth" ordering. A malicious 100 µs skew can reorder trades, invalidate model audit trails, or mis-attribute sensor fusion events in autonomy scenarios.

turba

Today's Focus: AI

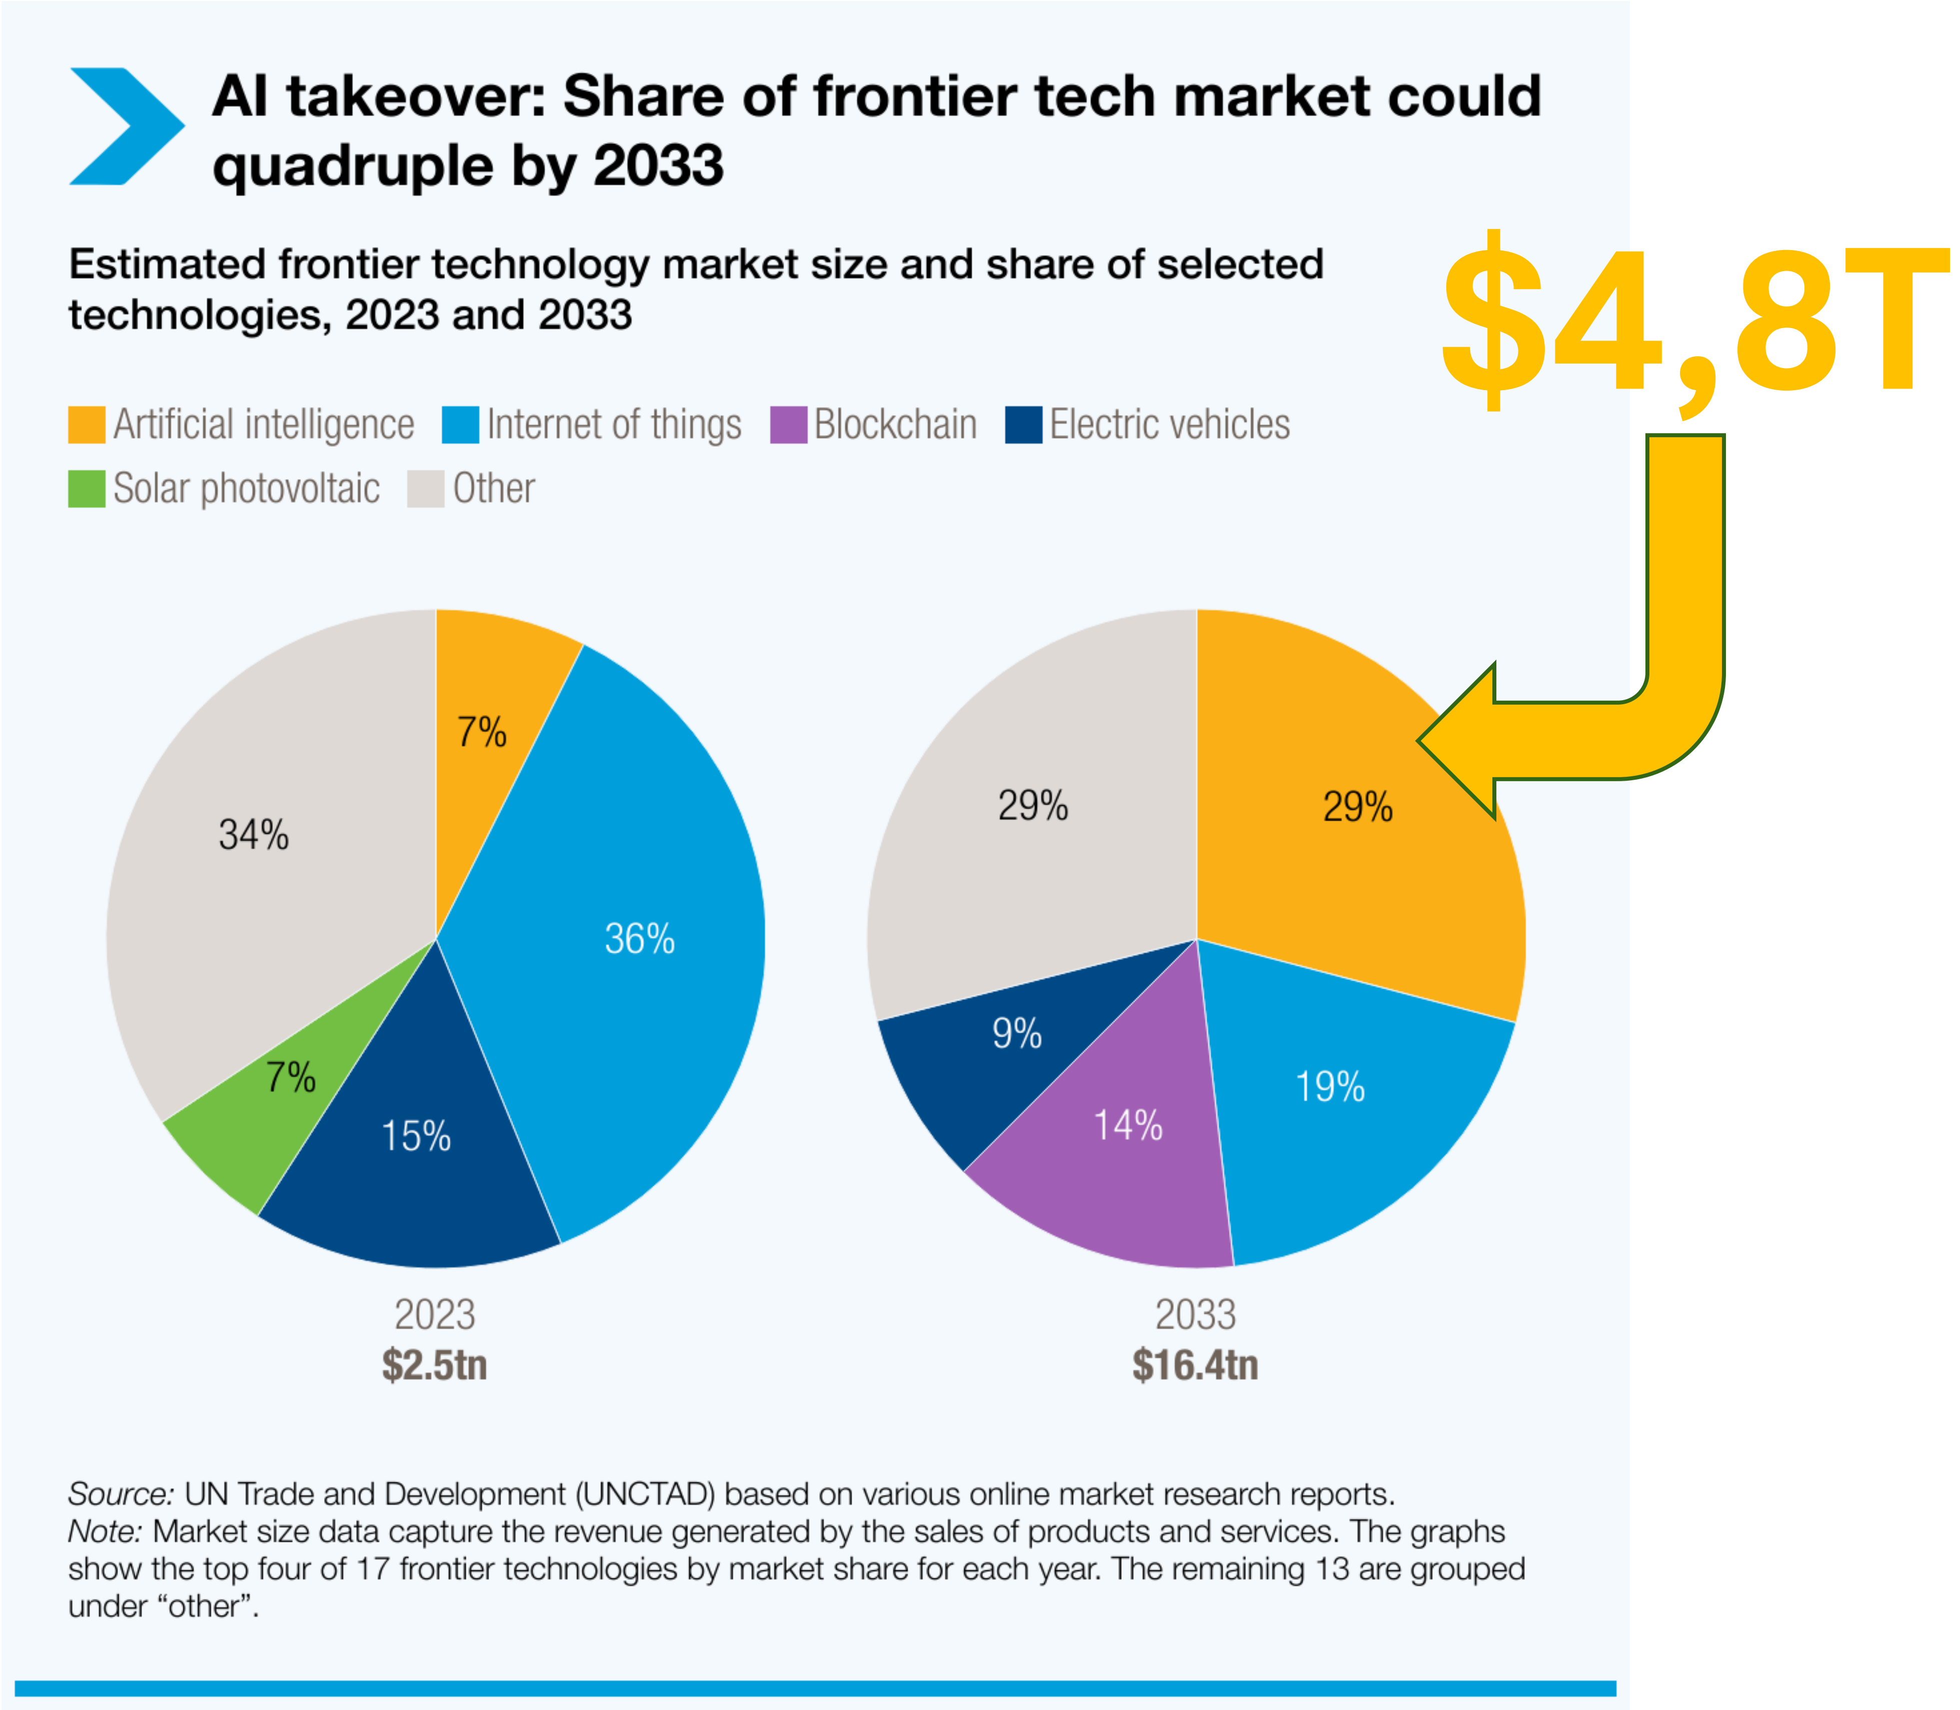

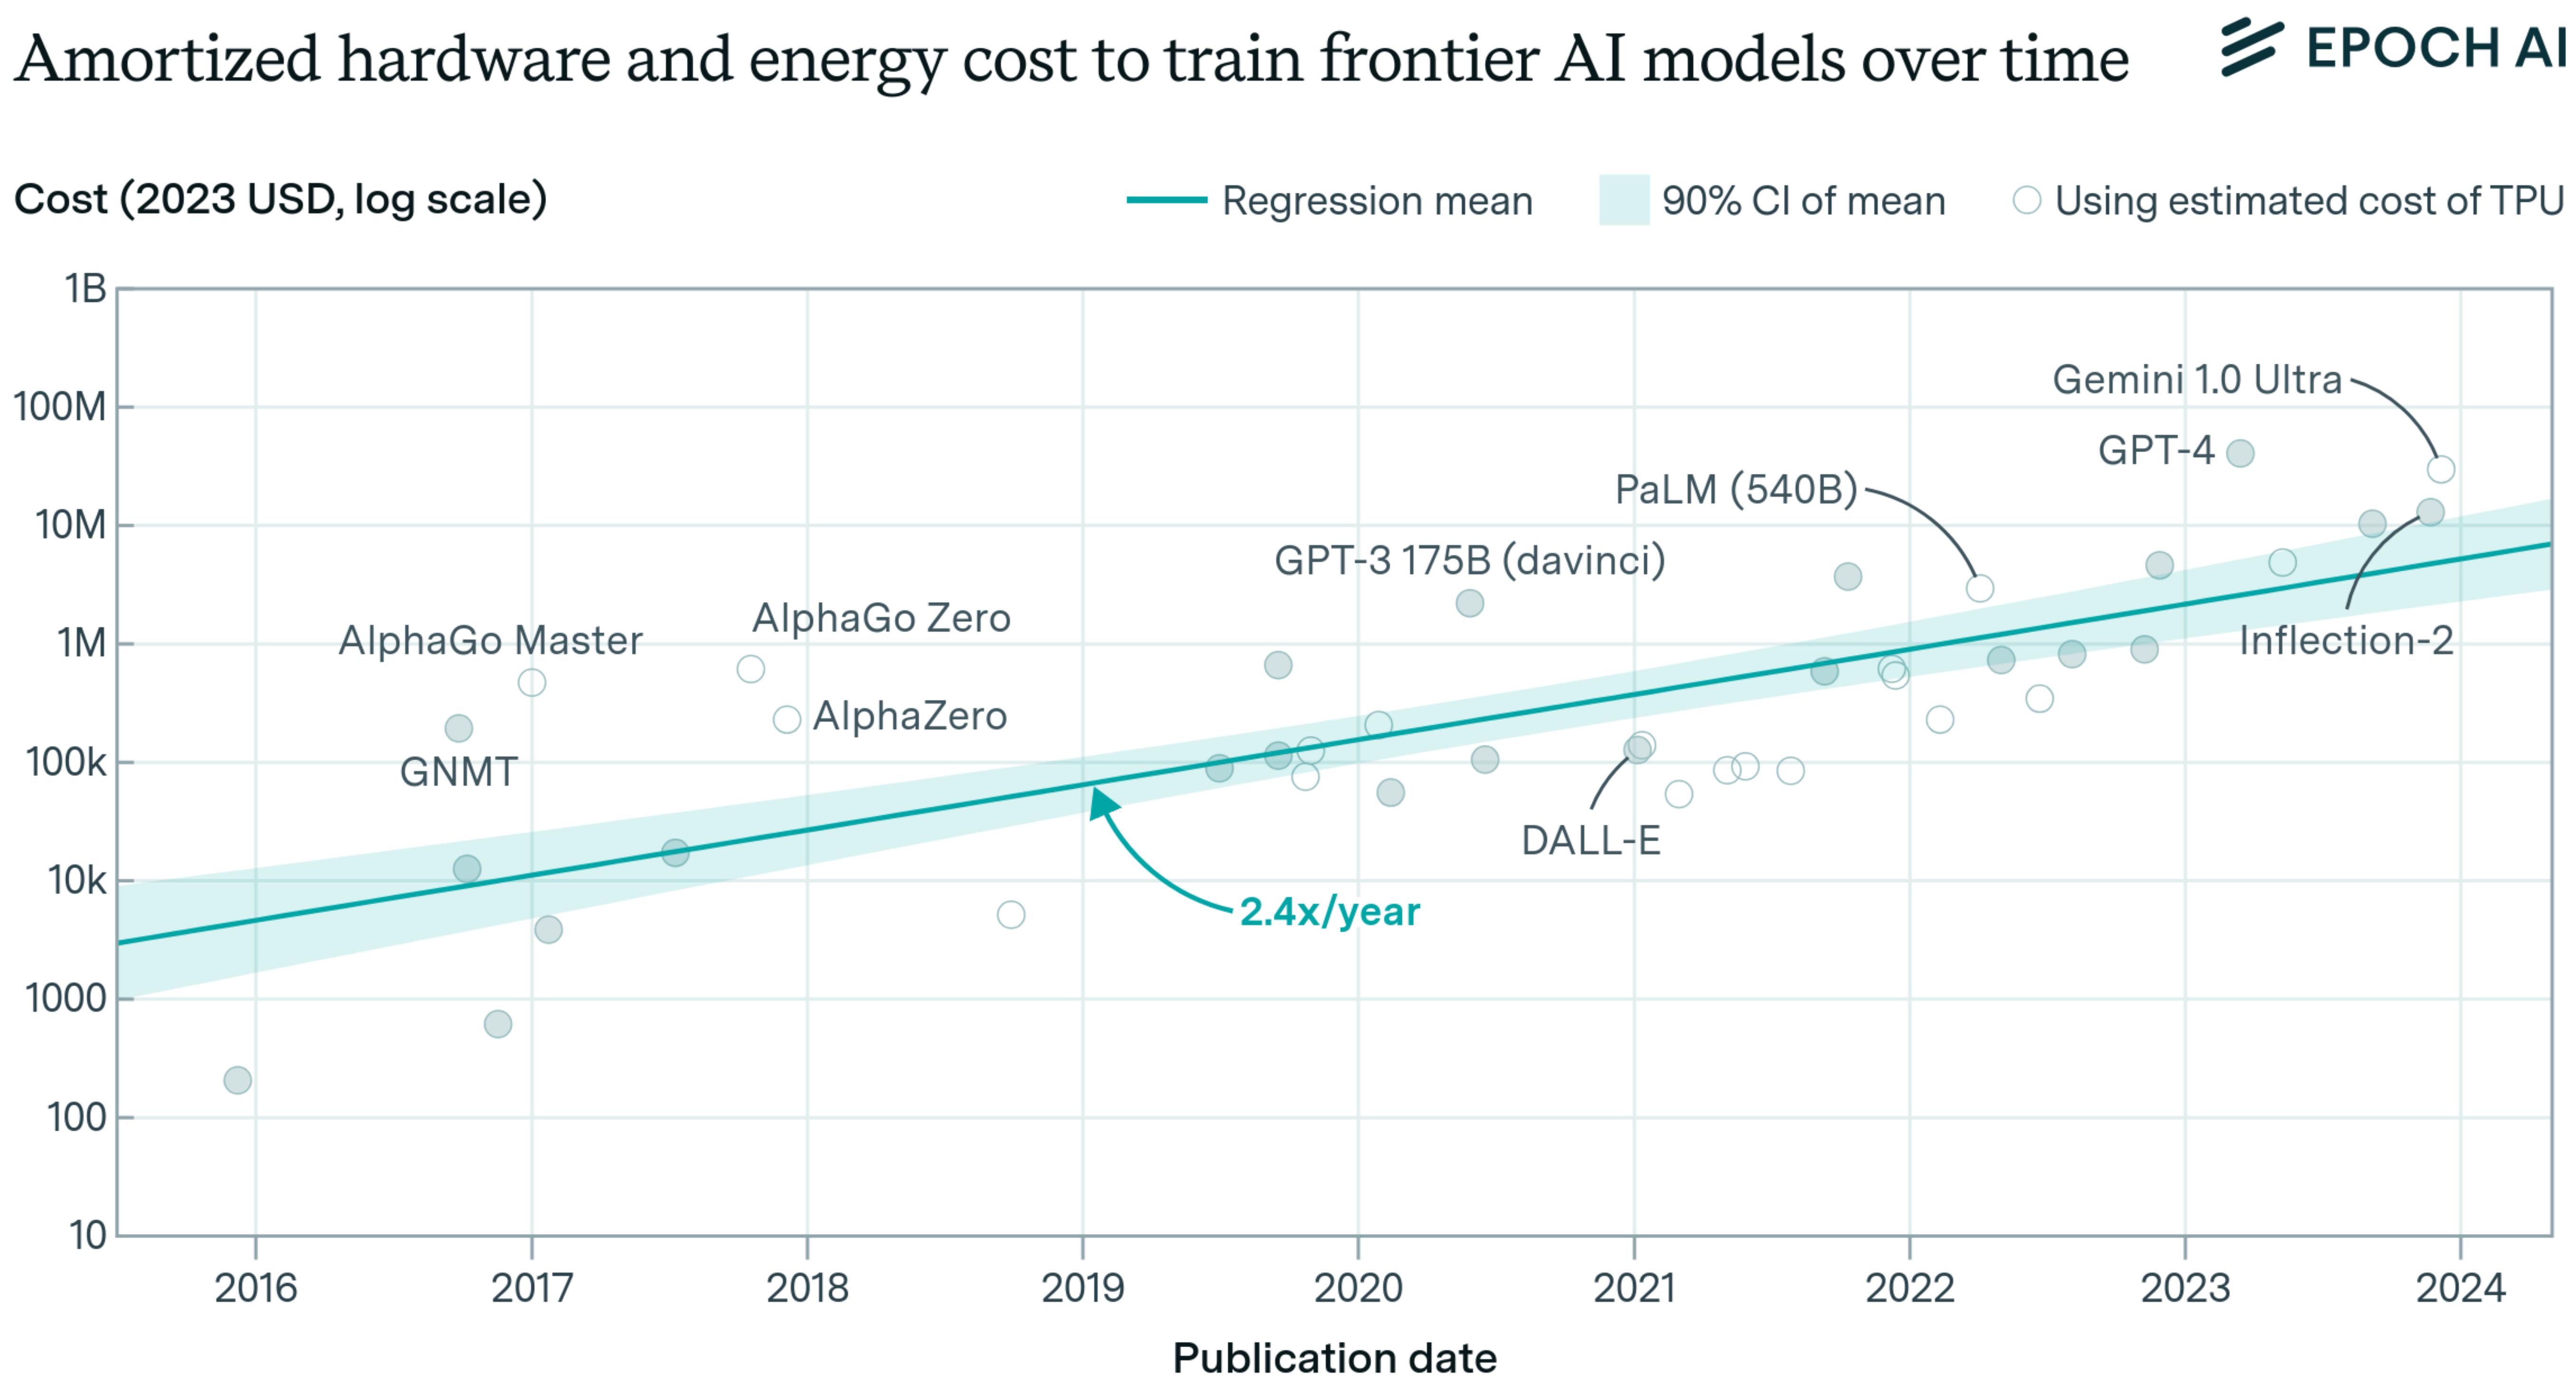

The Market is demanding more while the costs are approaching training runs of **$1B!**

October 2024

Hyperscale Cloud AI Growth

	Microsoft	aws	Google
Summary	Most new AI customers Most new GenAI customers	Most new traditional AI customers	Highest share of AI customers
# of new AI customer references*	274	207	102
Share by AI type (GenAI / Traditional AI)	46% / 54%	15% / 85%	37% / 63%
Selected new AI customers**	AXA, Schneider Electric, LEXMARK, Heineken, hp, KPMG	NTT docomo, SAMSUNG, RYANAIR, OMRON, AUTODESK, Nestlé	HDFC ERGO, MERCK, priceline, General Mills, Tchibo, University of Toronto

*=Customer wins/references refer to the published customer success stories on each vendor's website in the timeframe from June 2023 – June 2024; N=608. Customer references are considered AI once it includes the use of an AI service. The classification into GenAI also considers the text of the customer reference. **The listed companies were published as customer stories by the respective cloud vendor in the timeframe. For further information regarding the classification, please download a sample of the Global Cloud Projects Report and Database 2024 or contact the IoT Analytics team. **Source**: IoT Analytics Research - Global Cloud Projects Report and Database 2024. We welcome the republishing of images but ask for source citation with a link to the original post and company website.

Why Precise Timing in AI

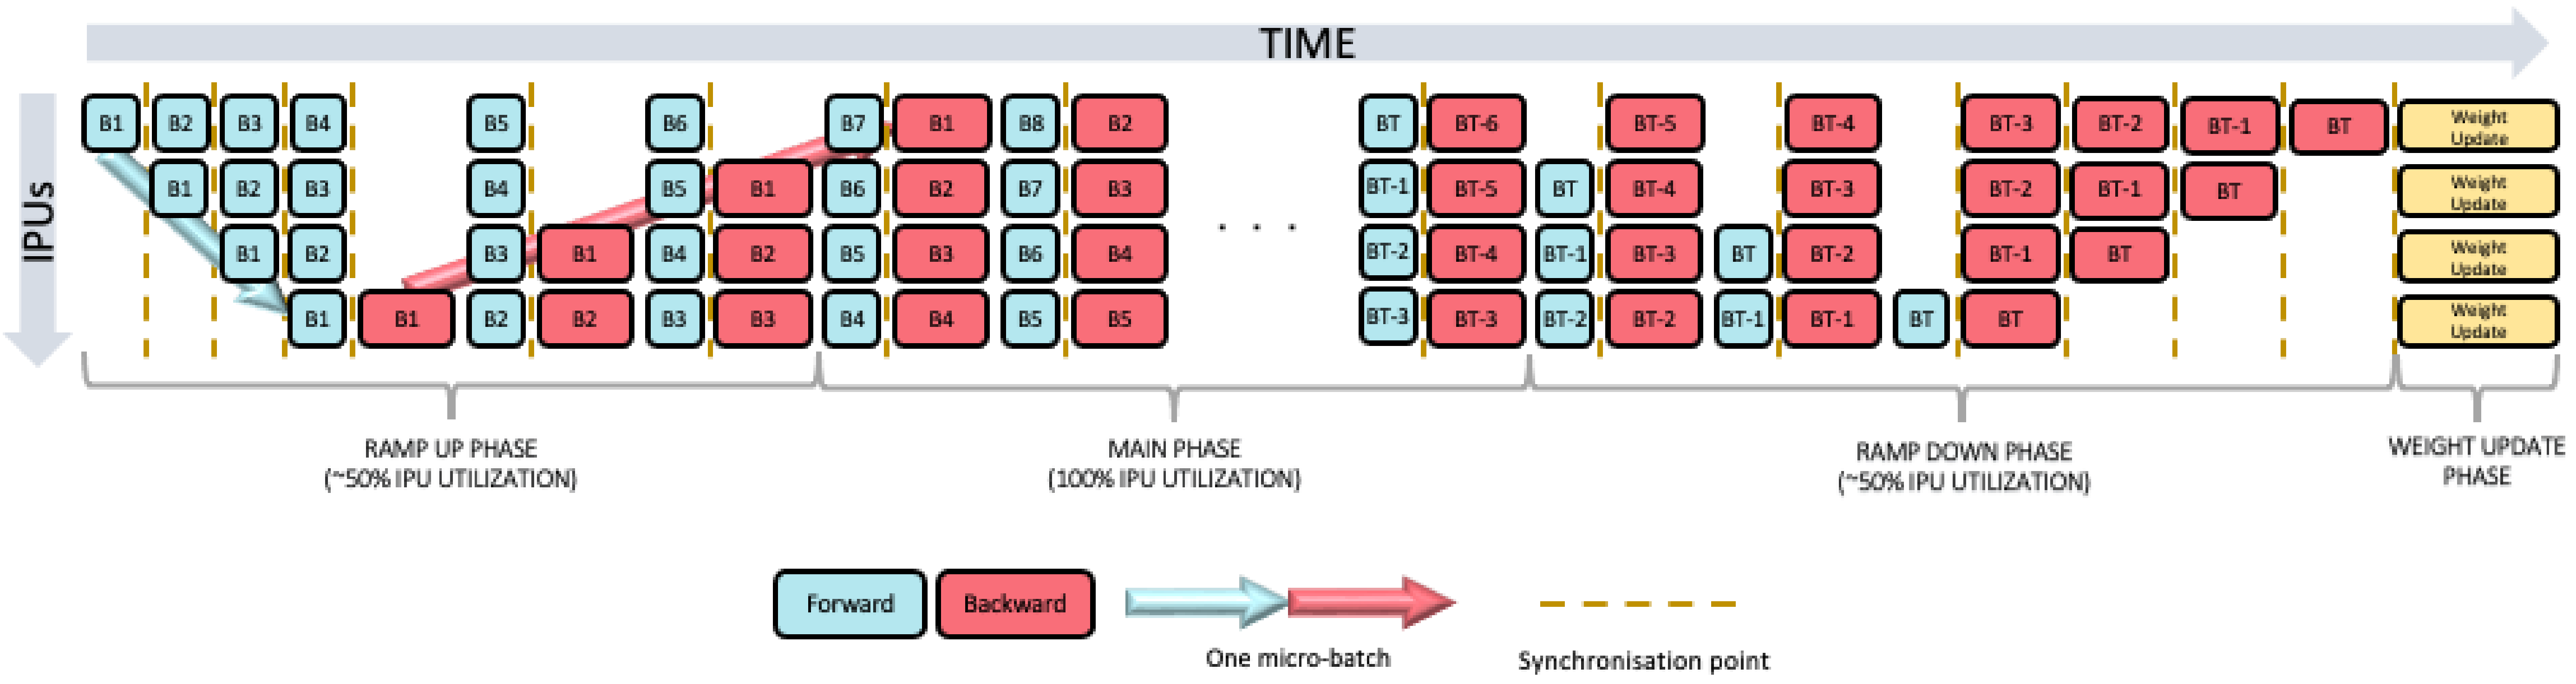

AI/ML workloads at hyperscale are highly distributed across clusters of GPUs, TPUs, and servers. Precise timing enables this.

- **Training** huge models (billions of parameters) involves splitting the model across multiple devices and nodes (pipeline, tensor and model parallelism) and splitting data across batches on different nodes (data parallelism).
- **Inference** is increasingly decentralized too – think of AI services running across cloud and edge, with multiple models across various topologies

In all cases, *precise time synchronization* is a critical requirement

https://docs.graphcore.ai/projects/tf-model-parallelism/en/latest/pipelining.html

What is Driving Precise Timing in AI

Model Parallelism- Synchronizing Multi-Node, Multi-GPU Execution

- Precision time enables synchronized operation scheduling, reducing tail latency and ensuring fair workload distribution

Model Checkpointing- Coordinated State Capture at Scale

- Accurate timing allows deterministic, synchronized checkpointing and ensures correctness in fault recovery

Reproducibility at Scale- Making AI Training Deterministic

- Precision time provides a shared reference for aligning execution order, logging events, and controlling nondeterminism

https://docs.graphcore.ai/projects/tf-model-parallelism/en/latest/pipelining.html

What if We Don't Get Precise Timing Right in AI?

- **Inefficiency and Wasted Hardware**
- **Undetected Errors and Monitoring Gaps**
- **Data Inconsistency and Corruption**
- **Security Attack Surface**

Precision timing errors can lead to slower training times, "mysterious" bugs, non-reproducible results, and even system failures.

https://docs.graphcore.ai/projects/tf-model-parallelism/en/latest/pipelining.html

Inefficiency and Wasted Hardware: GPU Example

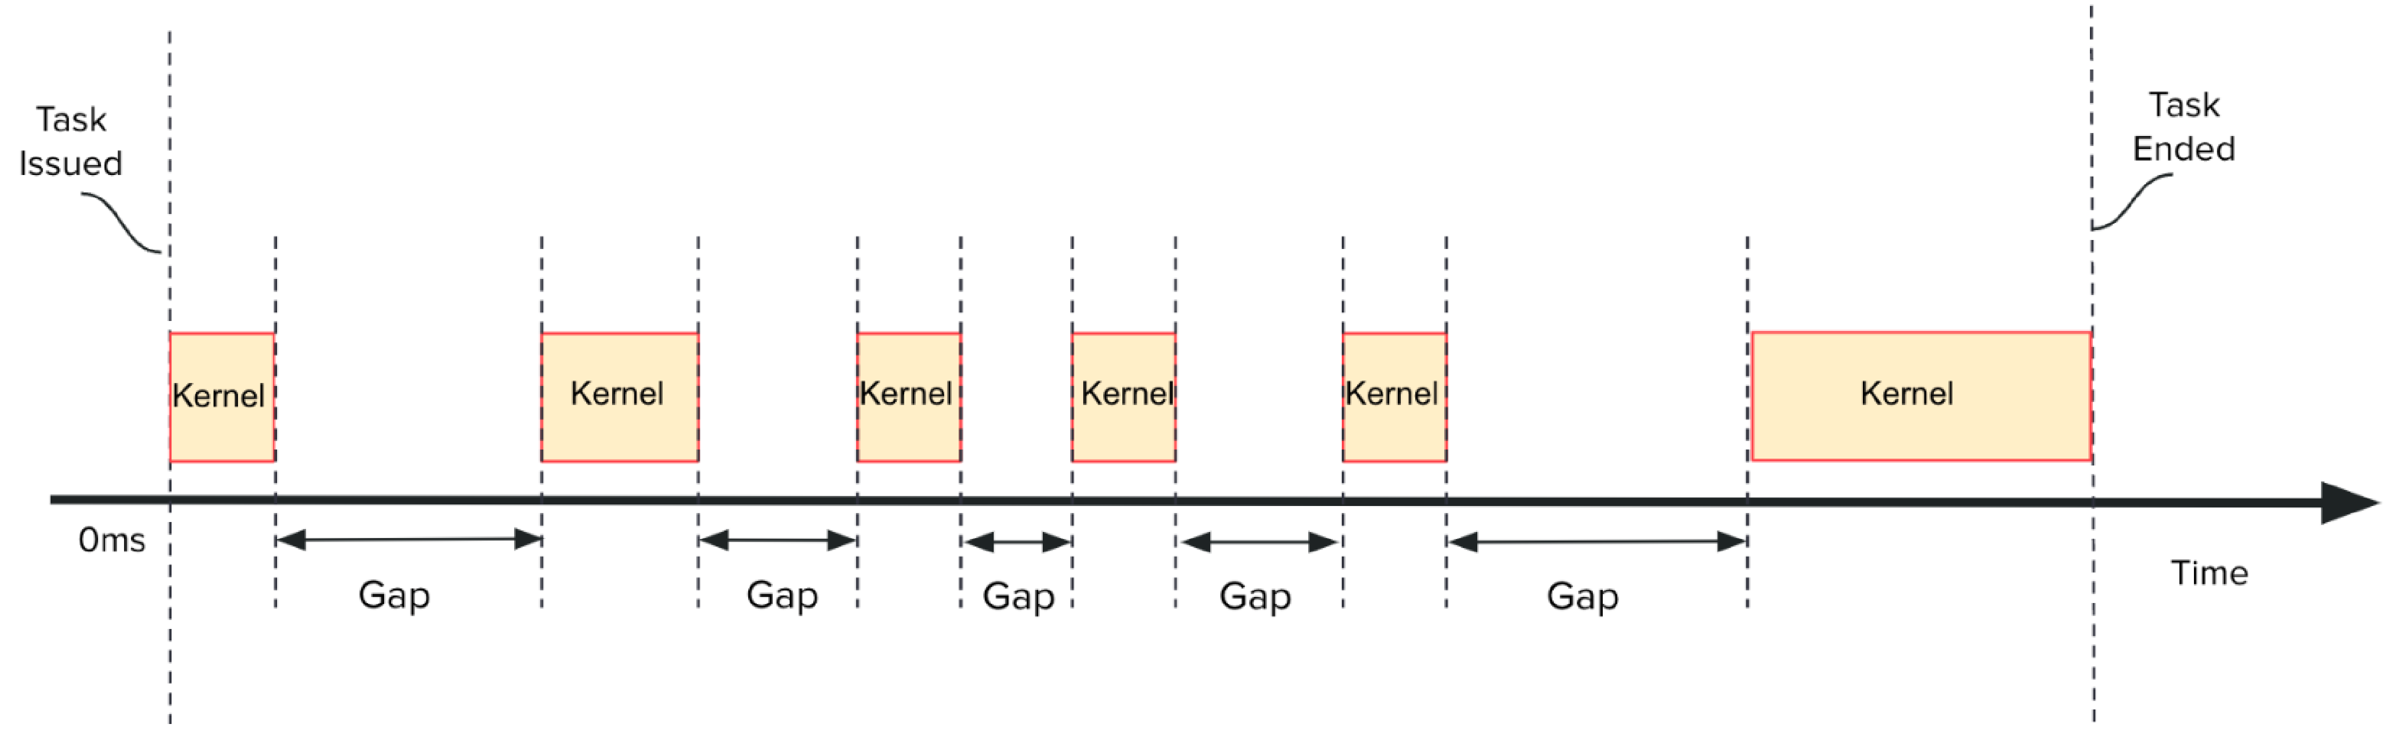

https://arxiv.org/html/2311.10359v5

- Typical task scheduling above
- +1 µs extra barrier wait = **0.1 % throughput loss/GPU**
- 1 ns-level sync, the waste reduces overall RAS

- *Millisecond* level example below of unavailable time during inference.

Unavailable time

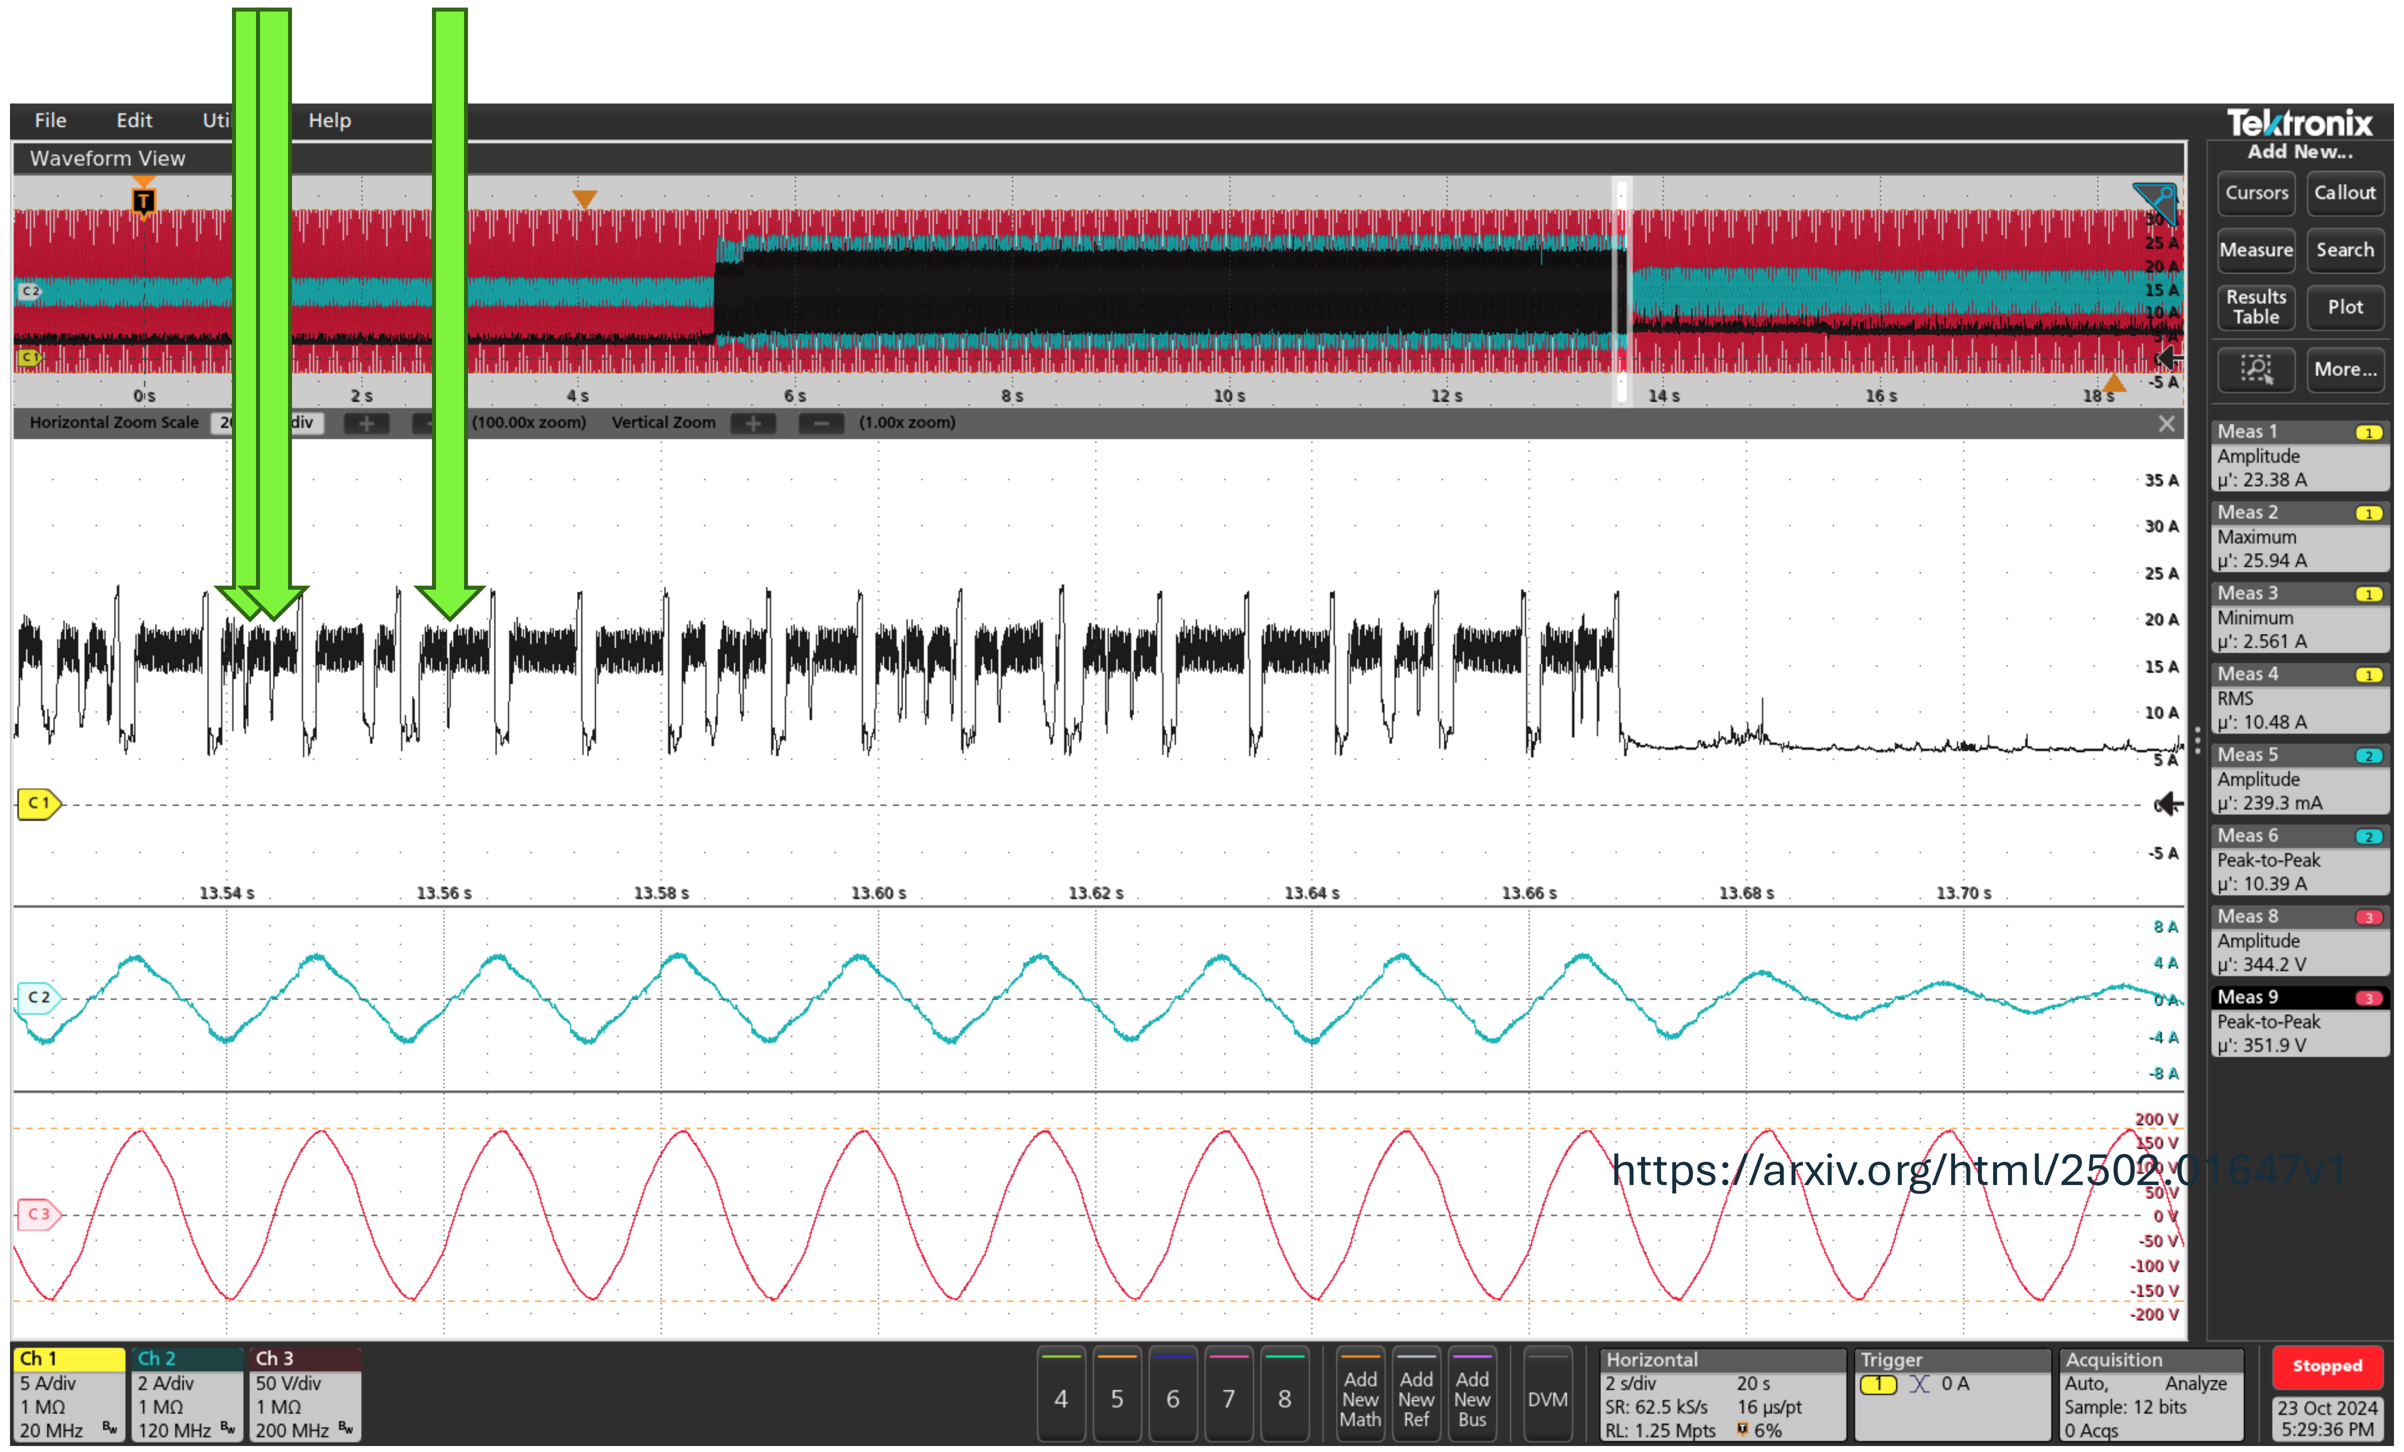

https://arxiv.org/html/2502.01647v1

Precision Timing is the Key to AI Performance

turba

Undetected Errors and Monitoring Gaps: RAS

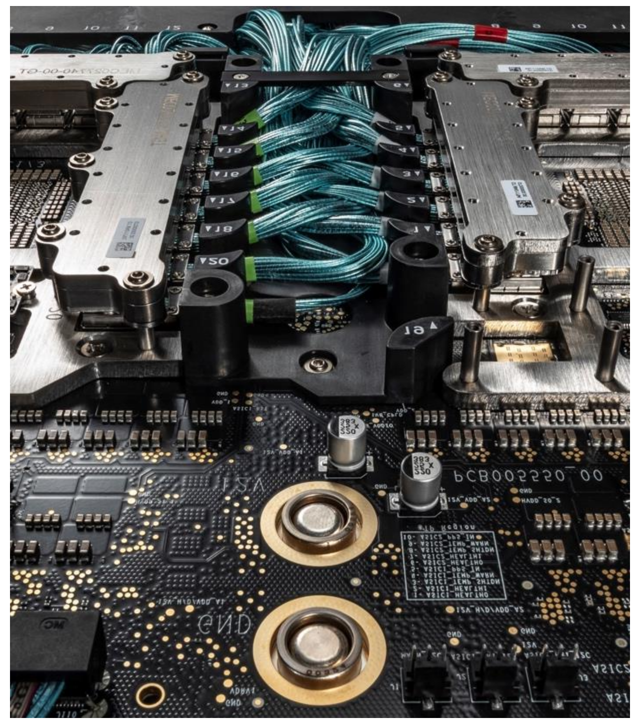

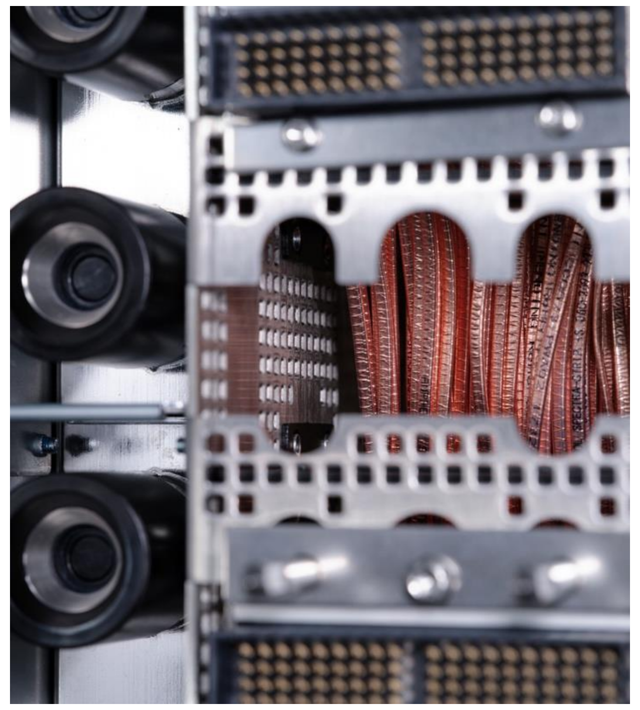

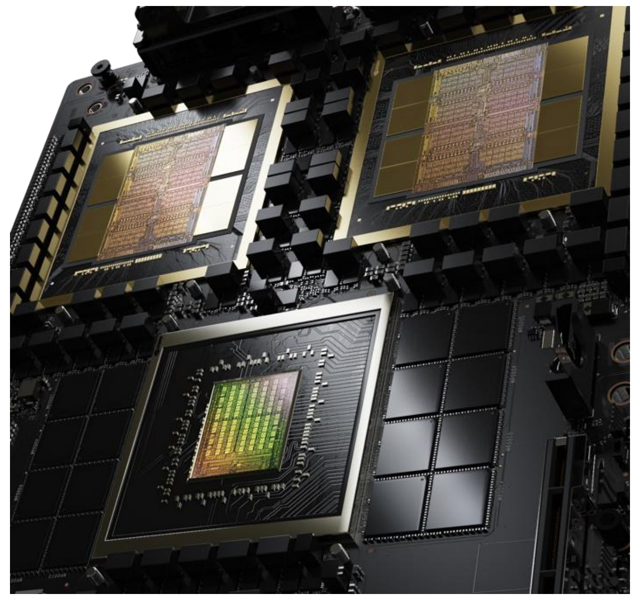

Data centers are rapidly scaling to accommodate the explosive growth of Big Data and the demands of emerging AI and HPC workloads.

This expansion involves deploying increasingly complex and highly valuable fleets of petaflop/petabyte scale AI/HPC pods.

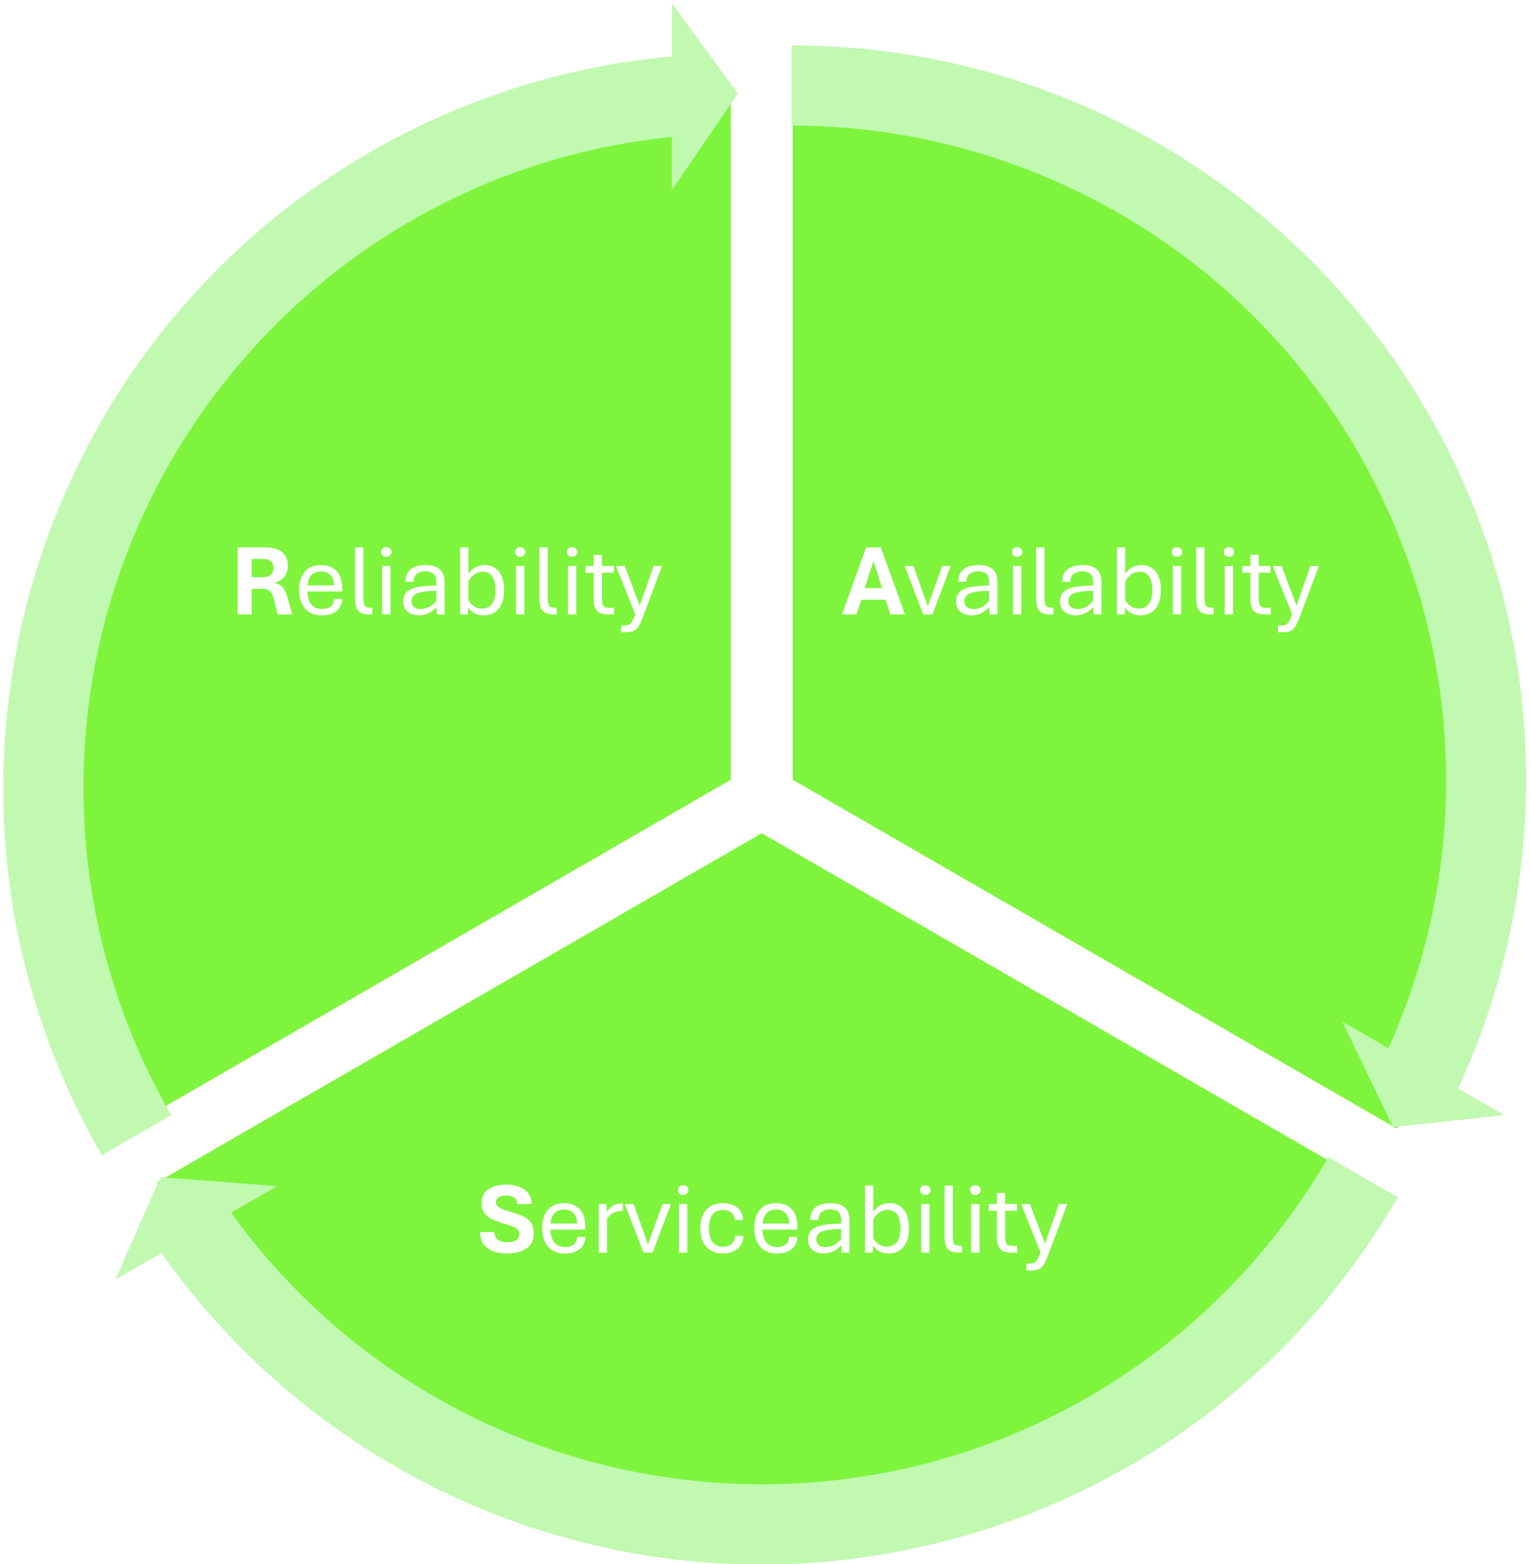

There is an **urgent** and **critical** need to develop robust data center and IT infrastructure that effectively **detects silent failures**, enhances **system reliability**, and **ensures** continuous **availability**.

RAS: Because Failure is Expensive!

An Example of RAS Analysis of AI Training Failures

Faults → **IT System faults occur** → Errors → **Error Rates grow with AI complexity and scale** → Failure

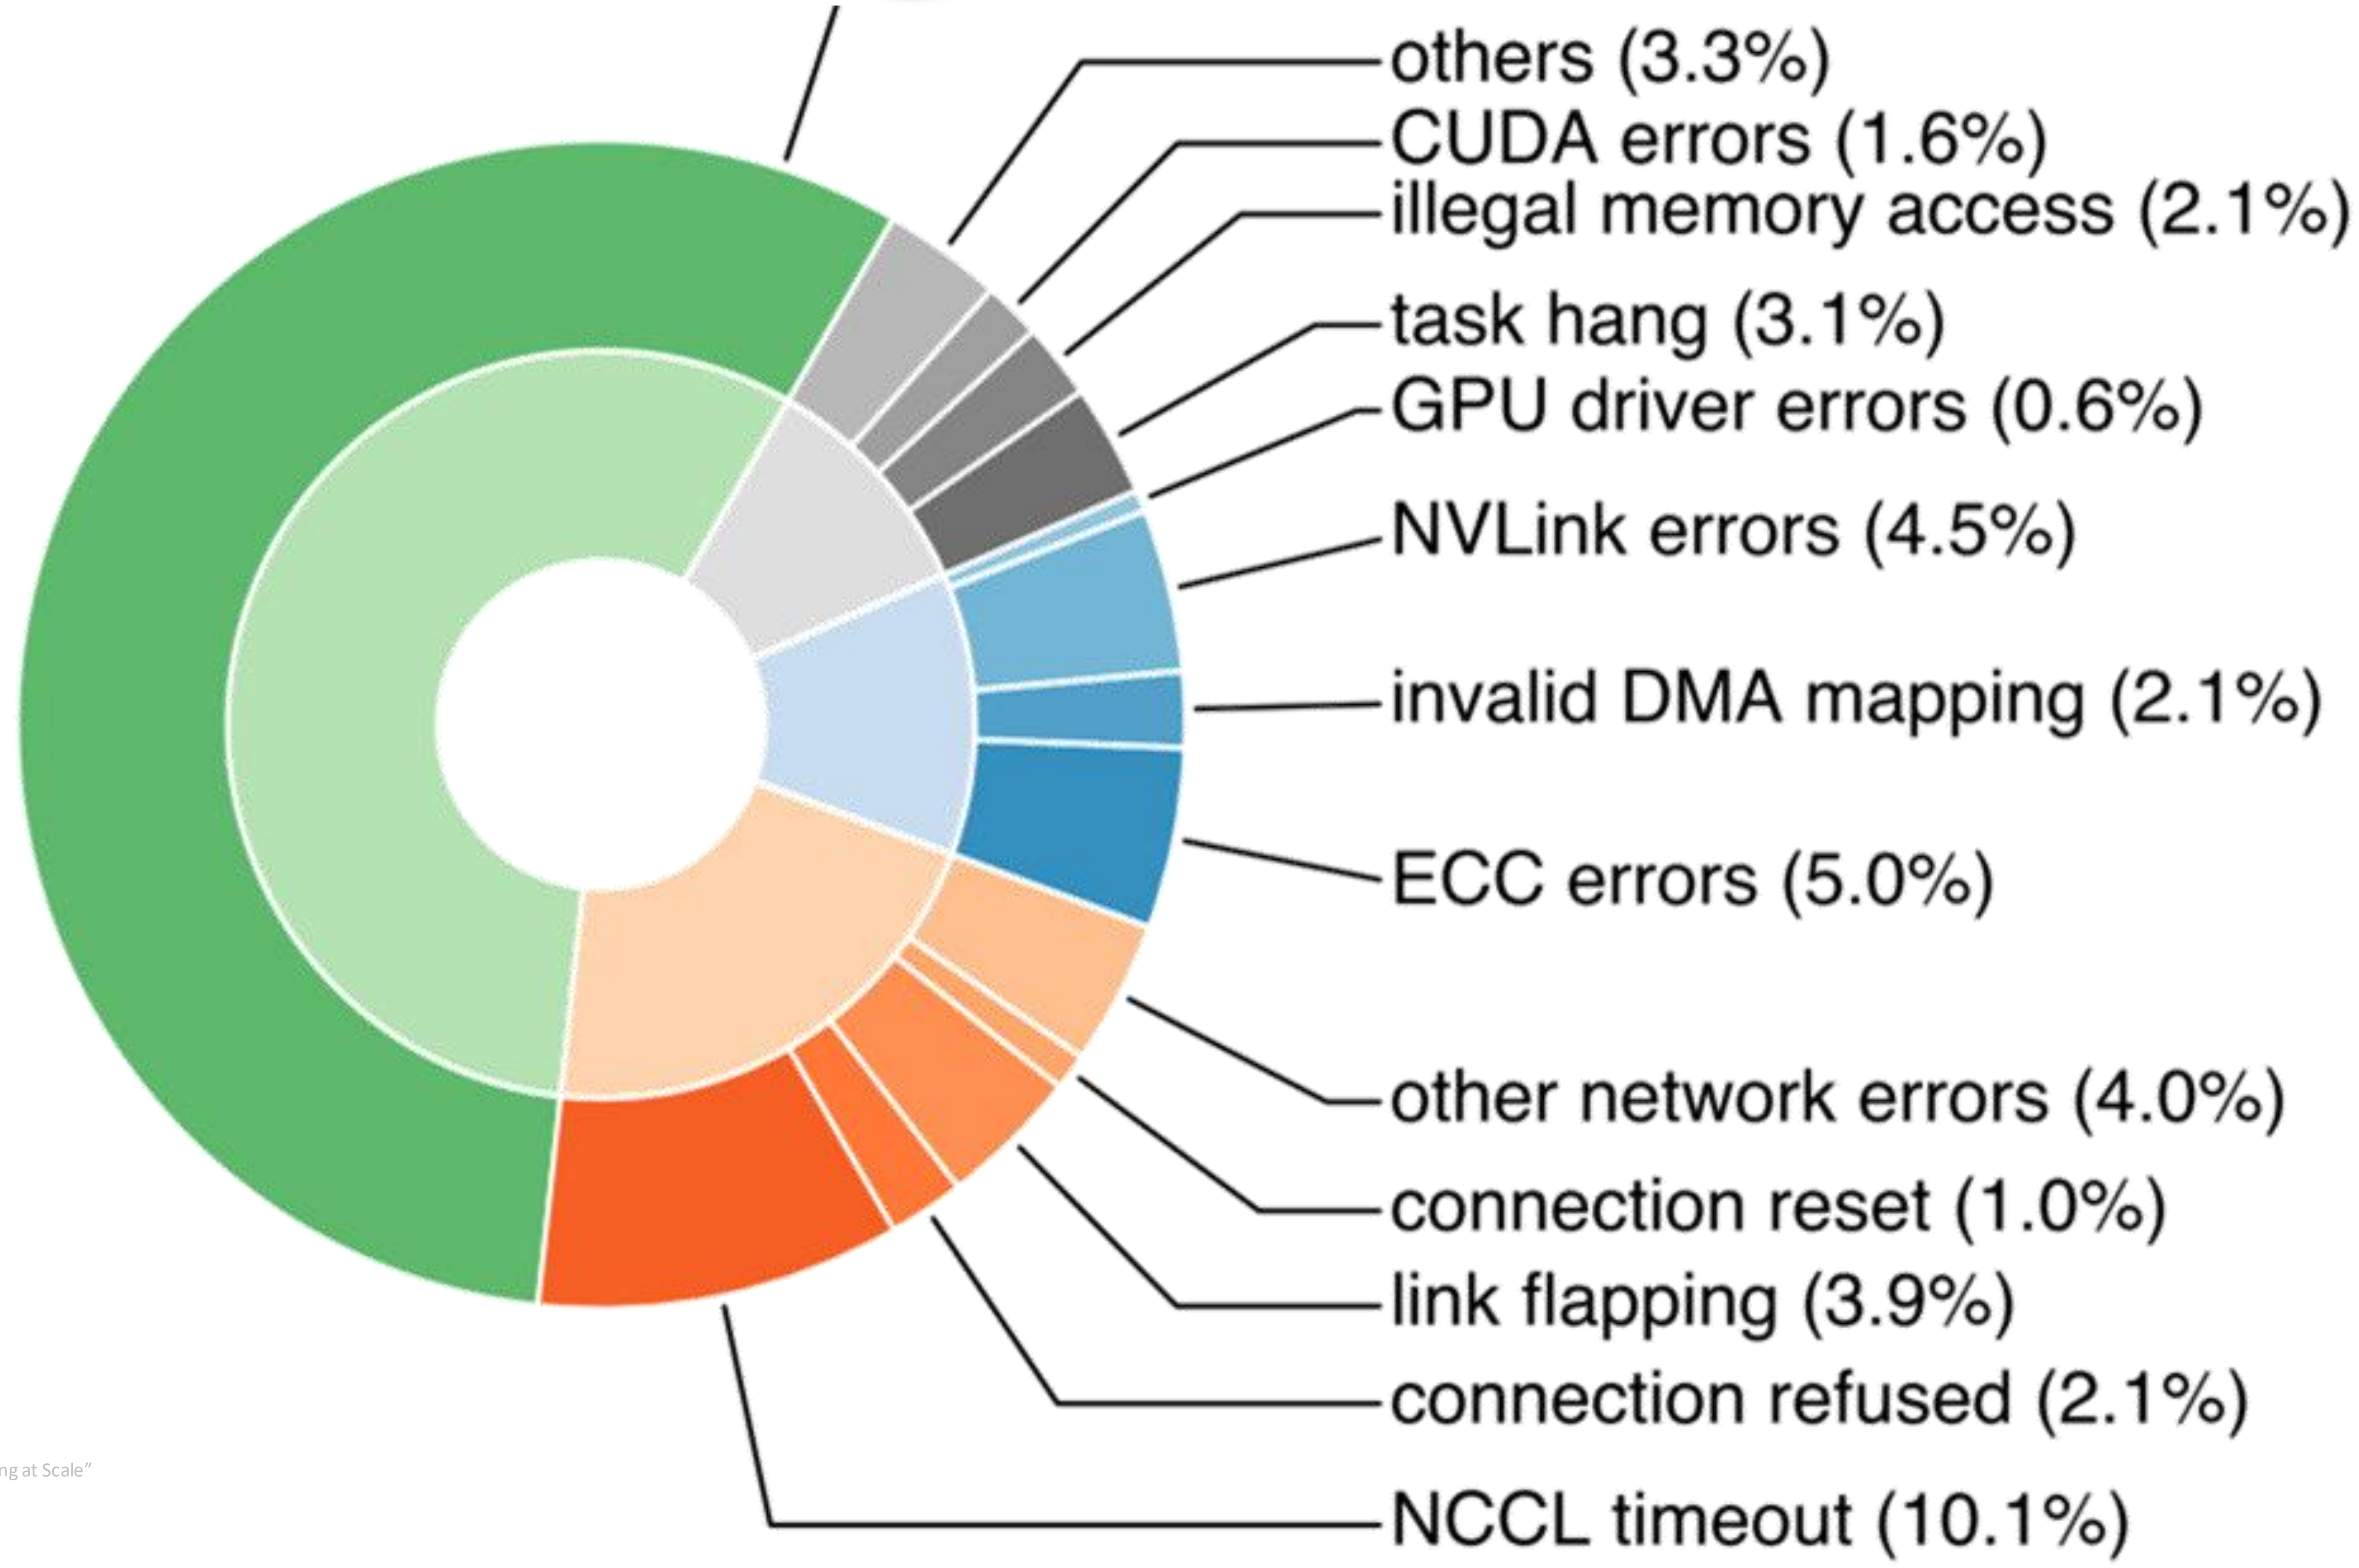

* "Unicron: Economizing Self-Healing LLM Training at Scale" https://doi.org/10.48550/arXiv.2401.00134

😡Up to 43.4% AI training job starts end in FAILURE!

Source: Meta

...resulting in idle clusters

Synchronization may be playing a large and sometimes-silent role in failures and training performance.

turba

Synchronization Error Based Data Inconsistency and Corruption

Time disagreement can lead to inconsistent views of state in a distributed AI system and contributes to reproducibility errors.

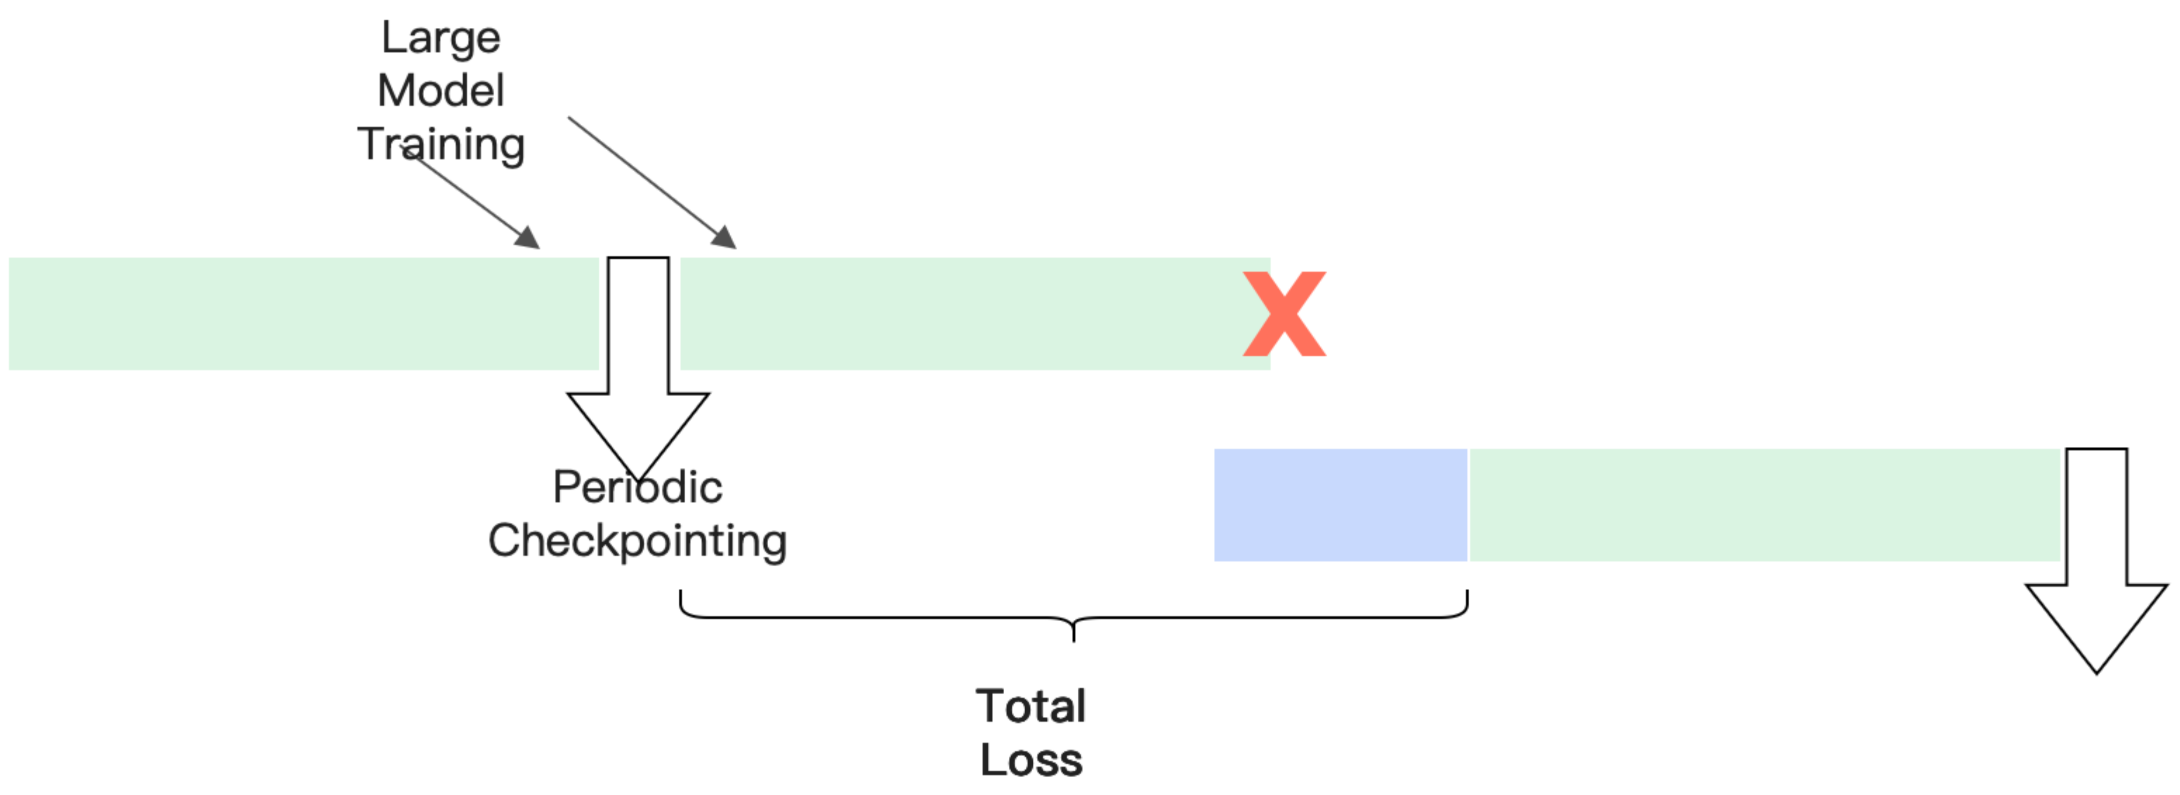

- **In Checkpointing–** (at Left) The result may be a checkpoint that can't be correctly resumed or consumes excess resources.

- **Causality Violation–** Similarly, consider failure synching distributed database used in an AI pipeline driven by clock skew
- **Reproducibility Failures–** Unreliable clocks introduce nondeterminism

https://www.alibabacloud.com/help/en/pai/user-guide/easyckpt

Synchronization Security Example

Kill Chain: GNSS spoofing ➜ Rogue Grand-master ➜ Boundary-clock replay ➜ 100 μs skew ➜ Training stalls / incorrect ordering / financial audit failure

- **Delay attack:** 140 μs skew achieved in 90 s by replaying Sync/F-Up packets on a production PTP net
- **GNSS spoofing:** Open-sourced L1 repeater can shift a datacenter by seconds in minutes
- **Regulatory parallel:** MiFID II already mandates ≤100 μs traceable to UTC; medical & automotive drafts are heading to 10 μs

At these precisions, **one malicious packet can move the entire cluster back in time…"**

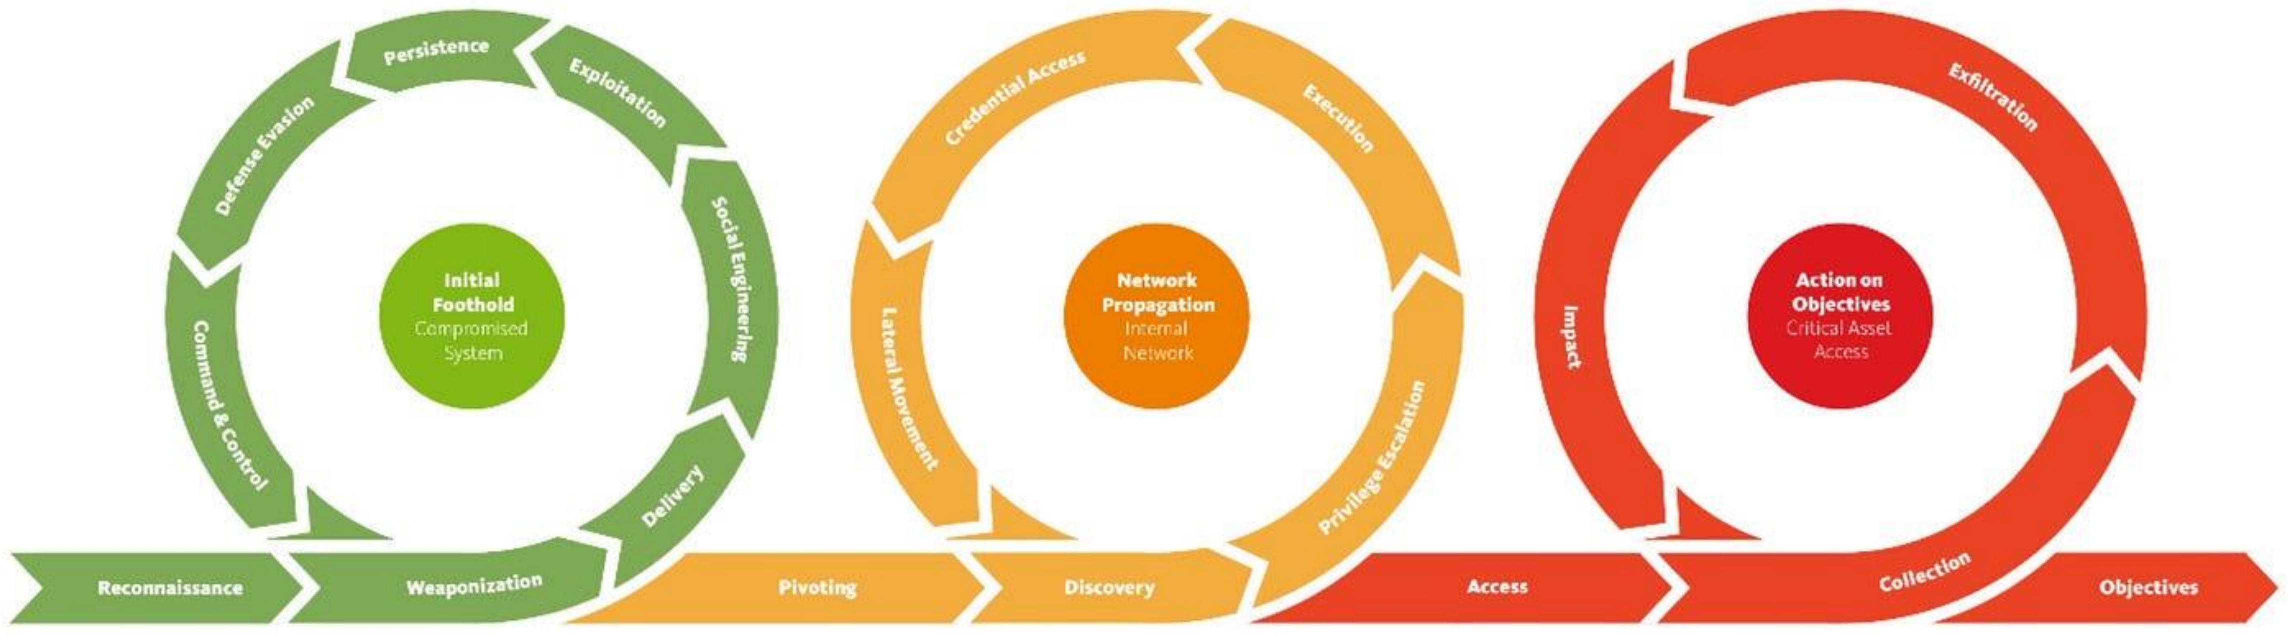

https://warnerchad.medium.com/unified-kill-chain-in-cyber-threat-intelligence-b577bf340ceb

turba

Next Steps

Solve Future Needs and Advanced Solutions

- **Hybrid Electronic/Photonic Computing-** Increasingly, electrical/photonic transitions in system, interconnect requires microsecond level scheduling. Optical Circuit Switching requires sub 1ns level of recovery and alignment resolution.
- **Memory-Centric and Disaggregated Architectures-** Memory coherence requires highly reliable synchronization.
- **Heterogeneous Compute Orchestration-** Make time awareness a feature of cluster schedulers and AI orchestration frameworks
- **Resilience and Security-** Common availability of cost-efficient Optical atomic clocks and frequency combs (chip scale?)

Close Research & Engineering Gaps

Gap	Why today's solutions fall short	Near-term direction
Dynamic path asymmetry	PTP assumes symmetric delay; ECMP plus congestion breaks it, injecting 10–50 μs error bursts under load .	In-band telemetry + per-flow skew estimation; Kalman-filter PTP variants are under study .
Trust & provenance	PTP security extensions cover only a subset of management TLVs; attackers can still rewrite clock trees .	Signed Sync/F-Up packets, remote-attestation of grandmaster firmware, and on-NIC anomaly ML.
Hold-over & multi-layer sync	Loss of GNSS forces quartz to drift; CSACs extend hold-over to <1 μs/day but cost/power hinder rack-scale deployment .	Hybrid: rack CSAC disciplining, plus optical frequency-comb distribution for campus backbone.
Cross-domain orchestration	Schedulers (K8s/SLURM) are clock-agnostic; they pad latency buffers.	APIs exposing *TimeSync Quality* so orchestrators can co-locate latency-sensitive tasks only on nodes within ±X ns.

We're heading toward a **"clock everywhere"** paradigm: every device, every process in an AI workflow will be aware of a global time and coordinate accordingly.

https://docs.graphcore.ai/projects/tf-model-parallelism/en/latest/pipelining.html

Develop Requirements for Next Applications

- **Photonic and analog AI fabrics** – Optical circuit switches re-route wavelengths in sub 10 ns windows; end hosts must react with timely switchover to avoid data loss.
- **Disaggregated memory (CXL 3.x pools)** – Cache-coherent loads across racks need read-modify-write round-trips < 1 µs. Any clock drift bigger than the bus retry window causes protocol back-off.
- **Federated / edge inference** – 5G TSN slices guarantee ≤1 µs end-to-end time error, yet the radio access network itself still relies on vulnerable PTP packets; leap-second smear bugs or spoofing can crash an entire edge cluster in seconds.
- **Deterministic Applications** – Finance already demands ≤100 µs traceable to UTC; emerging medical-device and automotive standards are heading to **10 µs** or tighter, forcing cloud-hosted AI to reduce its uncertainty budget, not just its mean offset.

It's the Stability, scalability, and security that are **lagging**.

turba

Improve Security: Zero Trust Synchronization

Five Pillars of a Zero-Trust Time Synchronization System

Pillar	Control example	Std / Tech leverage
Strong Identity	X.509 certs on grand-master & boundary clocks	IEEE 1588-2019 Security TLVs
Authenticated Channels	MAC-sec / IPsec + NTS-for-PTP	NTS (RFC 8915)
Least-Privilege Topology	Clock micro-segmentation; no lateral BMC discovery	NIST SP 800-207 Zero-Trust ZTA principles
Continuous Attestation	On-NIC drift sensors, signed telemetry every 1 s	"PTP in zero-trust world" best practice
Resilient Hold-over	CSAC/OCXO disciplining + multipath fiber + GNSS	Cisco/Juniper timing guides

Assume every time source and path is hostile until proven safe & continuously verified.

Call to Action

Call to Action

Let’s increase precision, RAS and security!

- **Engineers:** Participate in IEEE/OCP to further new ideas and develop reference implementations (ex: photonic combs, signed Sync/F-Up, Zero Trust, hardware root-of-trust, etc…).

- **Vendors:** ship comb-disciplined NICs with drift attestation hooks.

- **Regulators:** adopt time-uncertainty budgets in AI safety/compliance checks.

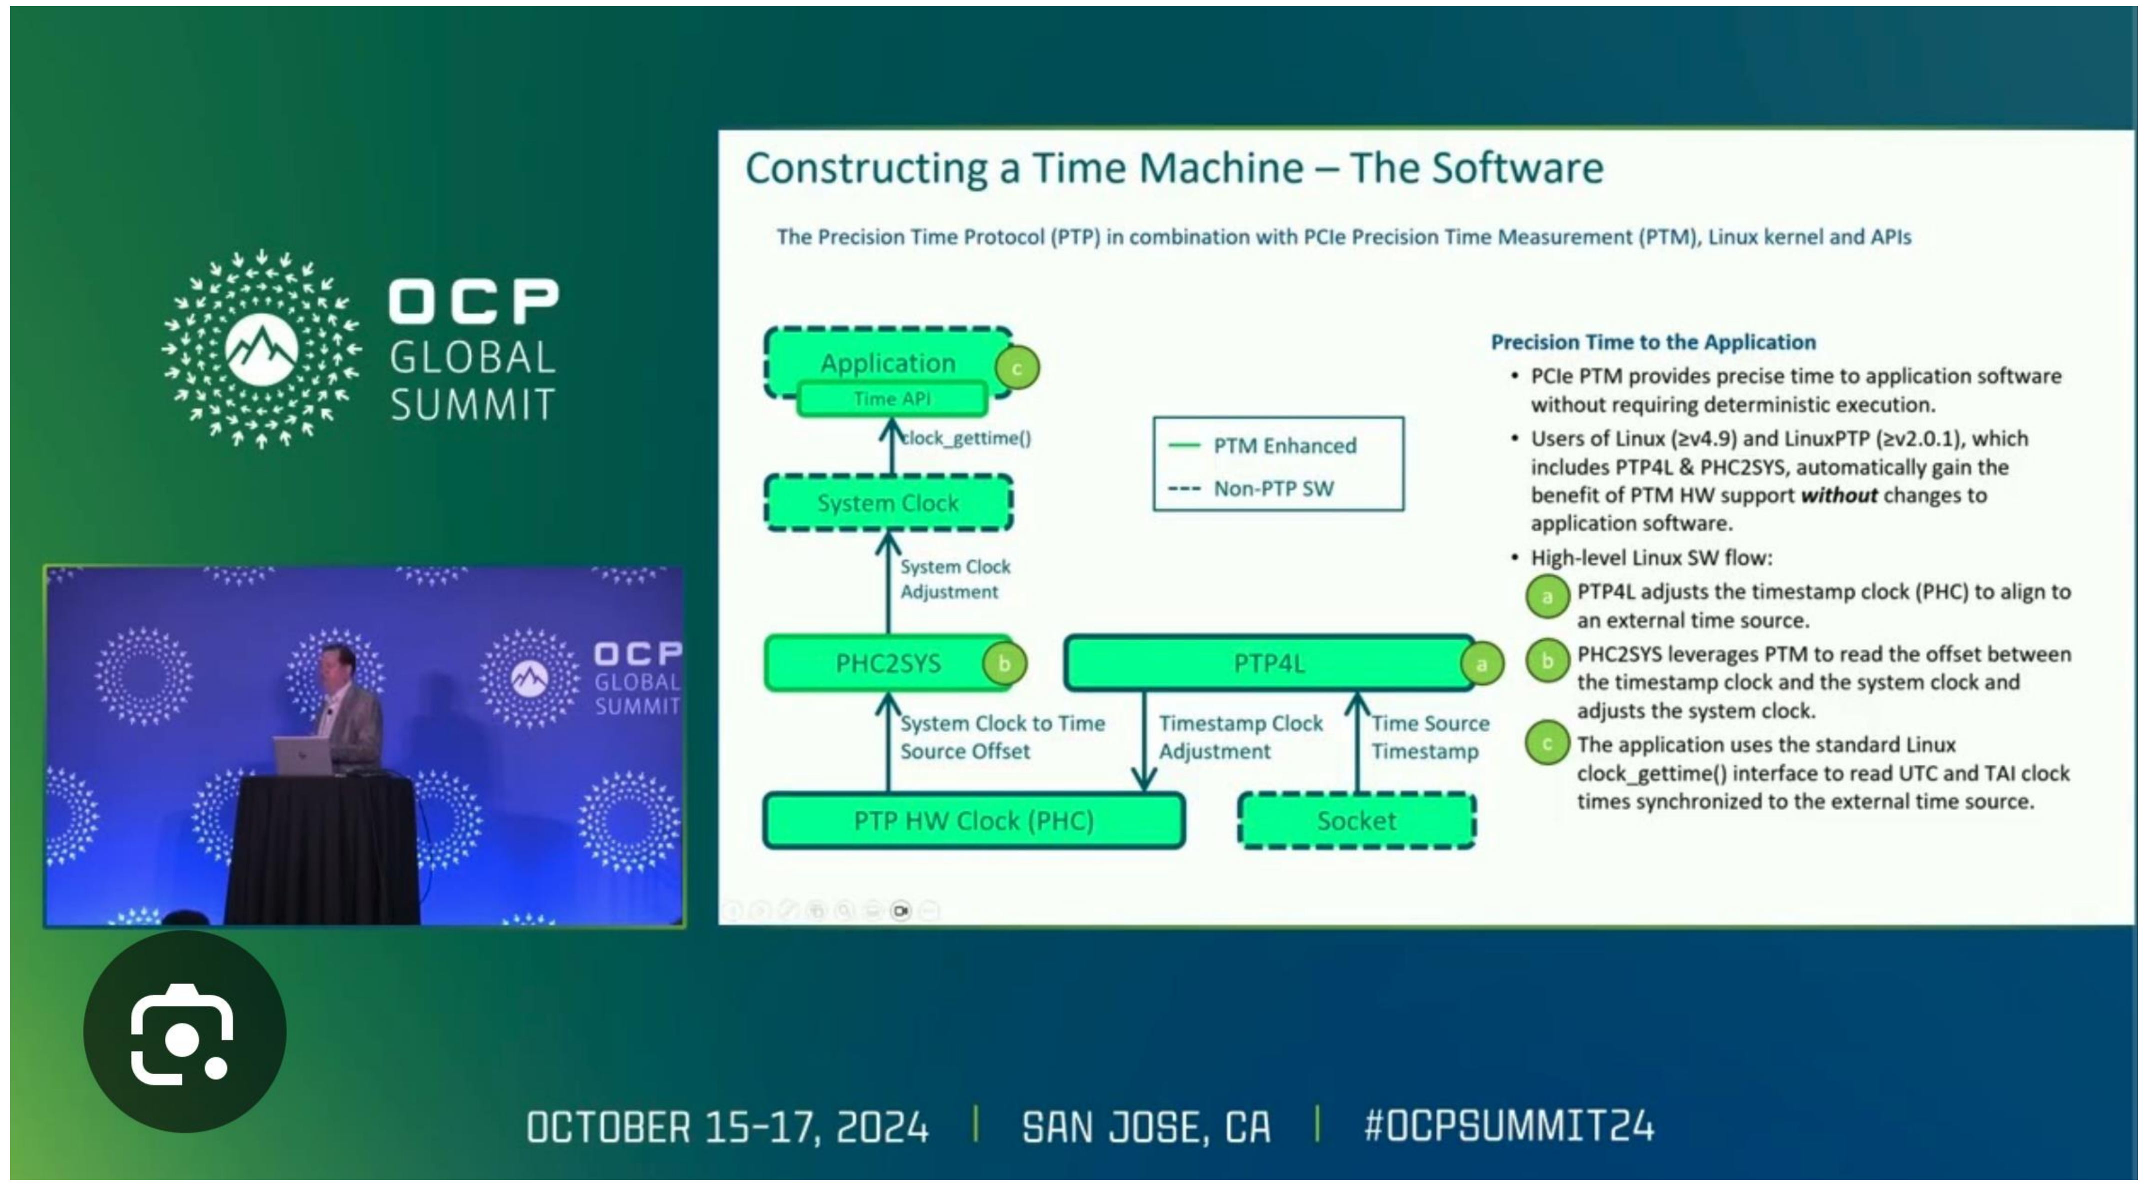

Precision was yesterday’s hurdle; provably secure precision is tomorrows.
Let’s write that playbook together.

Conclusion

PTP solved microseconds; AI is already asking for nanoseconds—with proofs.

Time is not just metadata, it's infrastructure.

Thank you

turba

Back up

May 2025

Past, Present, Future

Past: Evolution of timing in data centers; from ToD to loose synchronization using NTP to widespread adoption

Present: PTP for determinism and network coordination. Precision timing for WAN and AI workloads; distributed training, inference at the edge, and model/data consistency across GPU clusters.

Future: Time-aware infrastructure in:

- AI-native data centers
- Memory-centric computing
- High-performance data fabrics
- Tightly coupled hybrid (electronic and photonic) interconnects
- Multi-tenant, multi-GPU scheduling and orchestration

AI Workload

The existence of **PTP (µs-level)** and **White Rabbit (sub-ns-level)** does not mean synchronization is a solved problem; in practice, every time we ratchet precision up, new bottlenecks appear.
These broad realities keep the problem open:

Layer	What improves with today's tech	What still breaks	Why it matters to AI/HPC
Physical / Network	Hardware timestamping in NICs plus transparent clocks gives sub-µs accuracy at light-to-moderate load.	• Accuracy degrades by two–three orders of magnitude under micro-bursts or buffer contention. One study measured **360 µs–1 ms jitter** through commodity switches when queues build up • Fiber delay changes ±tens ns/km / °C; long White Rabbit links drift unless continuously temperature-trimmed	AI fabrics routinely push Tb/s across leaf–spine networks. Congestion episodes of a millisecond shove GPUs into idle-wait states, cutting utilization.
Scale / Topology	Open Time Server + PTP4U has shown **>1 M clients** on a single grand-master tree .	• Boundary-clock hierarchies add noise each hop; large brown-field networks can't enable BC mode safely • Today's designs assume homogeneous Ethernet; coming CXL/PCIe-over-fabric, optical circuit switches, and wireless edge tiers break that assumption.	Next-gen clusters mix rack-scale optical fabrics with electrical backplanes and remote memory pools. They need picosecond-class alignment across heterogeneous media—well beyond current PTP profiles.
Security / Resilience	MACsec, ACLs, and PTP security TLVs help authenticate masters.	• **Delay-and-spoof attacks** can offset a 5G base-station or O-RAN cell by >2 • Insider spoofing can silently skew slaves by >100 µs without tripping alarms • If GNSS is jammed, quartz clocks walk off by >1 ms inside an hour; even large hyperscalers now bolt CSACs onto grandmasters for hold-over	AI governance and forensics depend on "ground truth" ordering. A malicious 100 µs skew can reorder trades, invalidate model audit trails, or mis-attribute sensor fusion events in autonomy scenarios.

IOT ANALYTICS

April 2025

Enterprise generative AI application landscape

Mapping 530 real world projects across 68 common business activities

XX% = Share of GenAI projects that support the respective activity

Supporting activities

IT/technology	Software develop.* 15%	IT support 8%	Cybersecurity 5%	Network mgmt. 2%	
R&D	Product design 8%	Prototyping 8%	Feasibility studies 8%	Collaboration 2%	IP management -
Legal/compliance	Regulatory compl.* 5%	Contract drafting 3%	Risk assessment 2%	Litigation mgmt. <1%	IP protection <1%
Human resources	Recruitment 3%	Onboarding 3%	Engagement 3%	Training 2%	Performance mgmt. 2%
Corporate strategy	Growth planning 3%	Competitive analysis 2%	Market expansion 1%	Partnership develop. <1%	M&A -
Finance/account.*	Financial reporting 3%	Accounts mgmt. <1%	Tax compliance <1%	Payroll <1%	Budgeting -
Procurement	Sourcing <1%	Supplier relationship <1%	Contract negotiation <1%	Supplier selection -	Purchase processing -
Facility mgmt.*	Safety management <1%	Energy management <1%	Building maintenance -	Security systems -	Space planning -

Primary activities

Supply chain	Operations/mfg.*	Marketing	Sales	Customer support	Maintenance/FS*
Demand forecasting 2%	Process optimization 11%	Content creation 17%	Lead generation 3%	Issue resolution 35%	Inspections <1%
Order fulfillment 2%	Resource planning <1%	Digital marketing 7%	Account management 2%	Inquiry handling 34%	Preventive maintenance <1%
Inventory management <1%	Quality control <1%	Market research 6%	Negotiation <1%	Post-sale support 19%	Word order management <1%
Transportation management <1%	Production scheduling <1%	Campaign planning 4%	Pipeline management <1%	Feedback collection 7%	Reporting <1%
Warehouse management -	-	Marketing analytics 3%	Forecasting <1%	SLA management <1%	Asset management -

* Account. = Accounting, Mgmt. = Management, develop. = Development, Comp. = Compliance, Mfg. = Manufacturing, FS = Field service. **Source:** IoT Analytics Research - Generative AI Market Report 2025-2030, List of Generative AI Projects 2025. We welcome republishing of images but ask for source citation with a link to the original post and company website.

Low — Adoption level — High