TENSOR NETWORKS

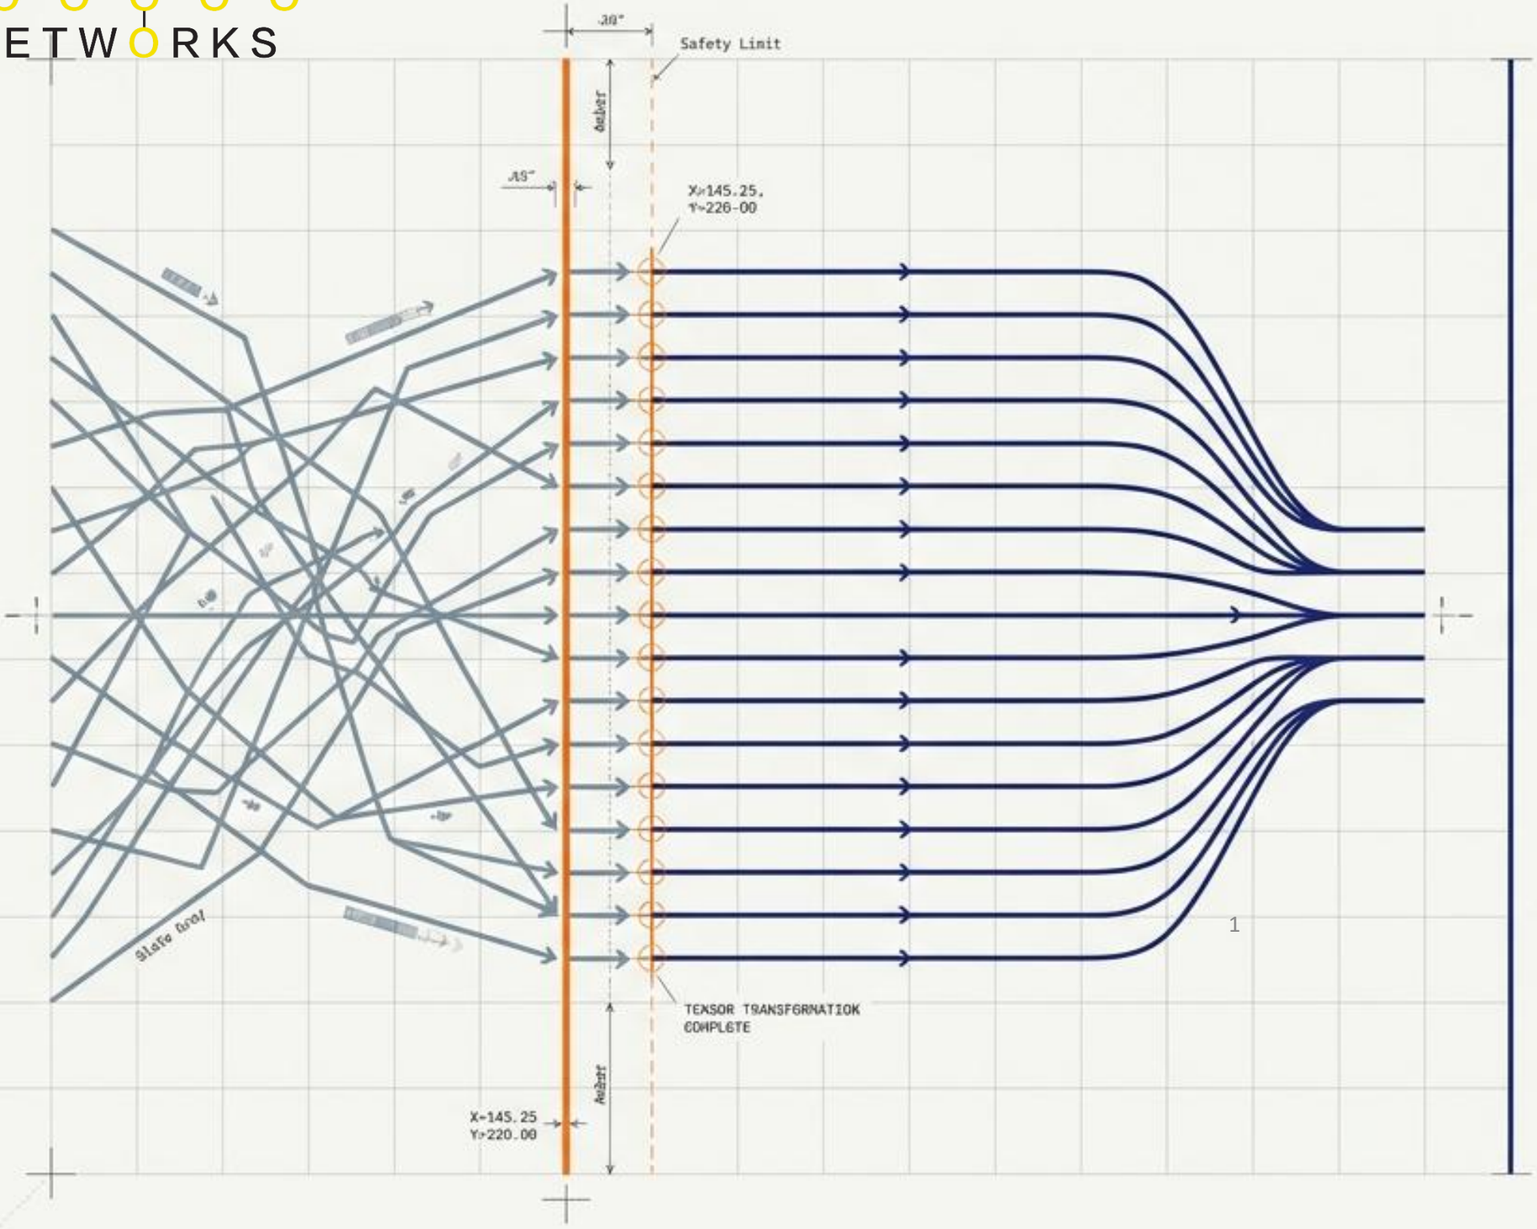

Professional Services & Scope of Work for PTCP Integration

A Strategic Deployment Framework by Tensor Networks, Inc.

801c: 28.8 XE:-172:3.57.99 Data: 16.52.2023 Tags: IBM Plex Mono
Data: 2222.155 Tag: IBM Plex Mono

SCOPE OF SERVICES:

- Phase 1: Architectural Assessment & PTCP Feasibility Study
- Phase 2: Tensor Network Design & Mathematical Optimization
- Phase 3: Controlled Deployment & Safety Validation
- Phase 4: Ongoing Support & Performance Monitoring

TECHNICAL DELIVERABLES:

- Bounded System Architecture Blueprints (DWG/PDF)
- Mathematical Safety Limit Definitions (JSON/XML)
- Optimized Routing Algorithms (Python/C++)
- Comprehensive Performance Report (PDF)

Disclaimer: Forward-Looking Statements This presentation may contain forward-looking statements regarding future product plans, expected features, or anticipated release dates.
These statements outline our general product direction and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for our products remain at our sole discretion and are subject to change without notice. These statements do not represent a commitment, promise, or legal obligation to deliver any specific material, code, or functionality

REACTIVE DATA ROUTING

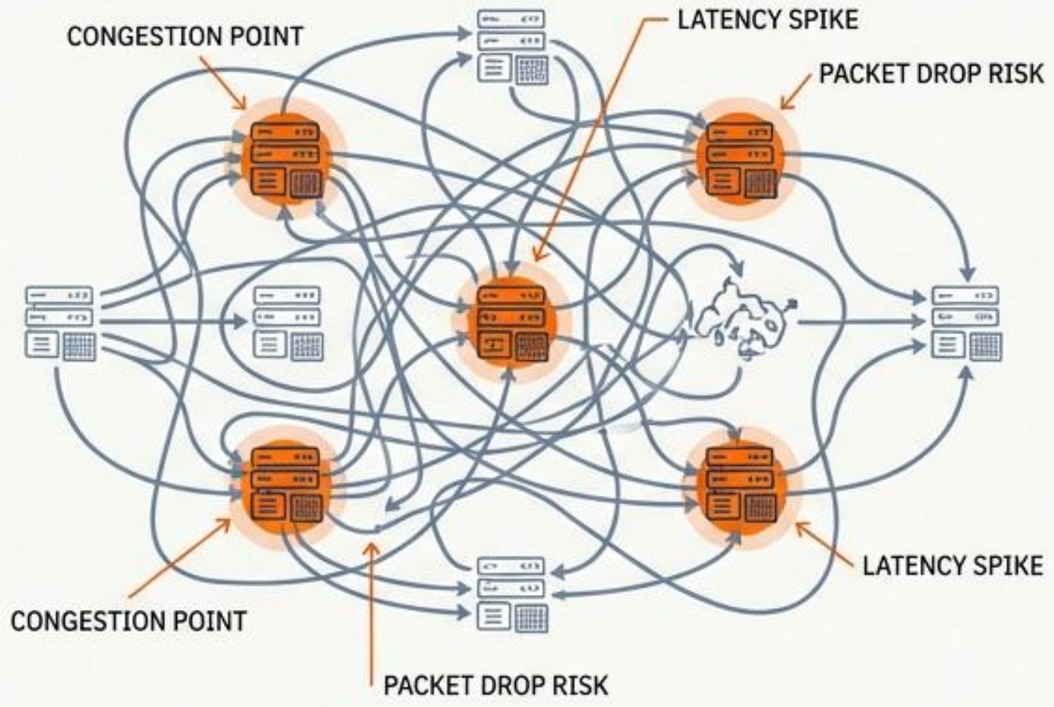

PREDICTIVE ORCHESTRATION

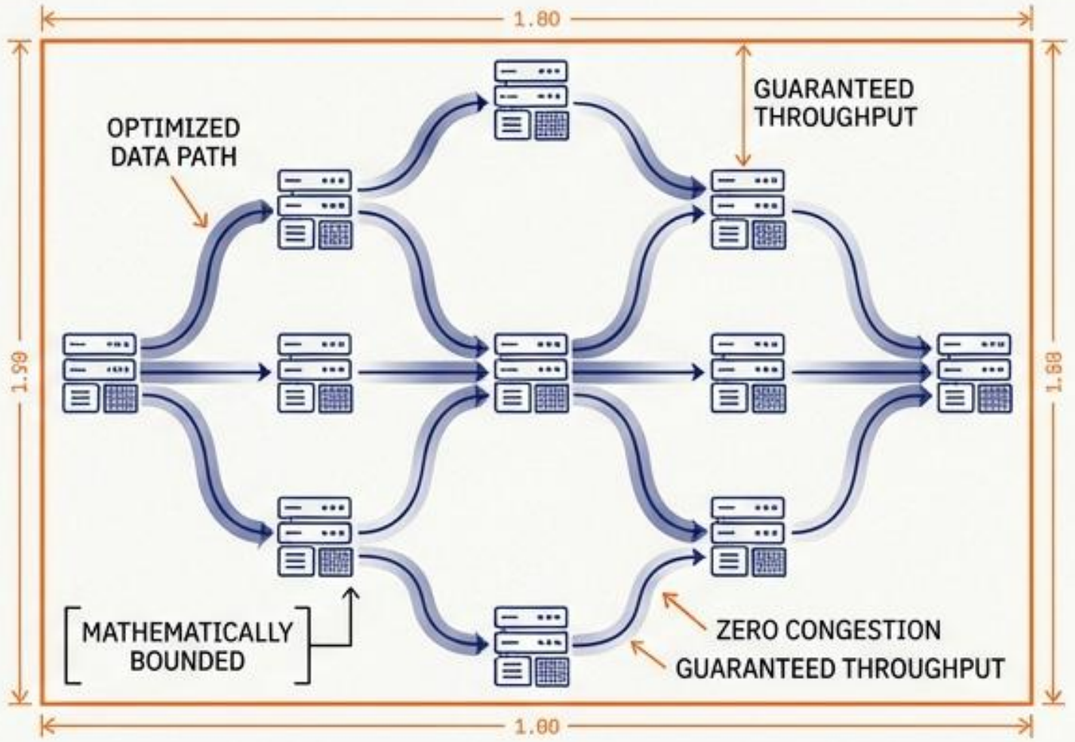

3

The Transition to **Mathematically Bounded Orchestration**

Core Thesis

Deploying a predictive infrastructure fabric requires a structural shift from reactive routing to pre-emptive, mathematically bounded control. This ensures operational integrity under all conditions.

Key Guarantee

This framework is engineered to integrate the Predictive Tensor Control Plane (PTCP) with absolute safety, seamlessness, and continuous operational excellence. All parameters are predefined and bounded.

The 20-Week Integration Methodology

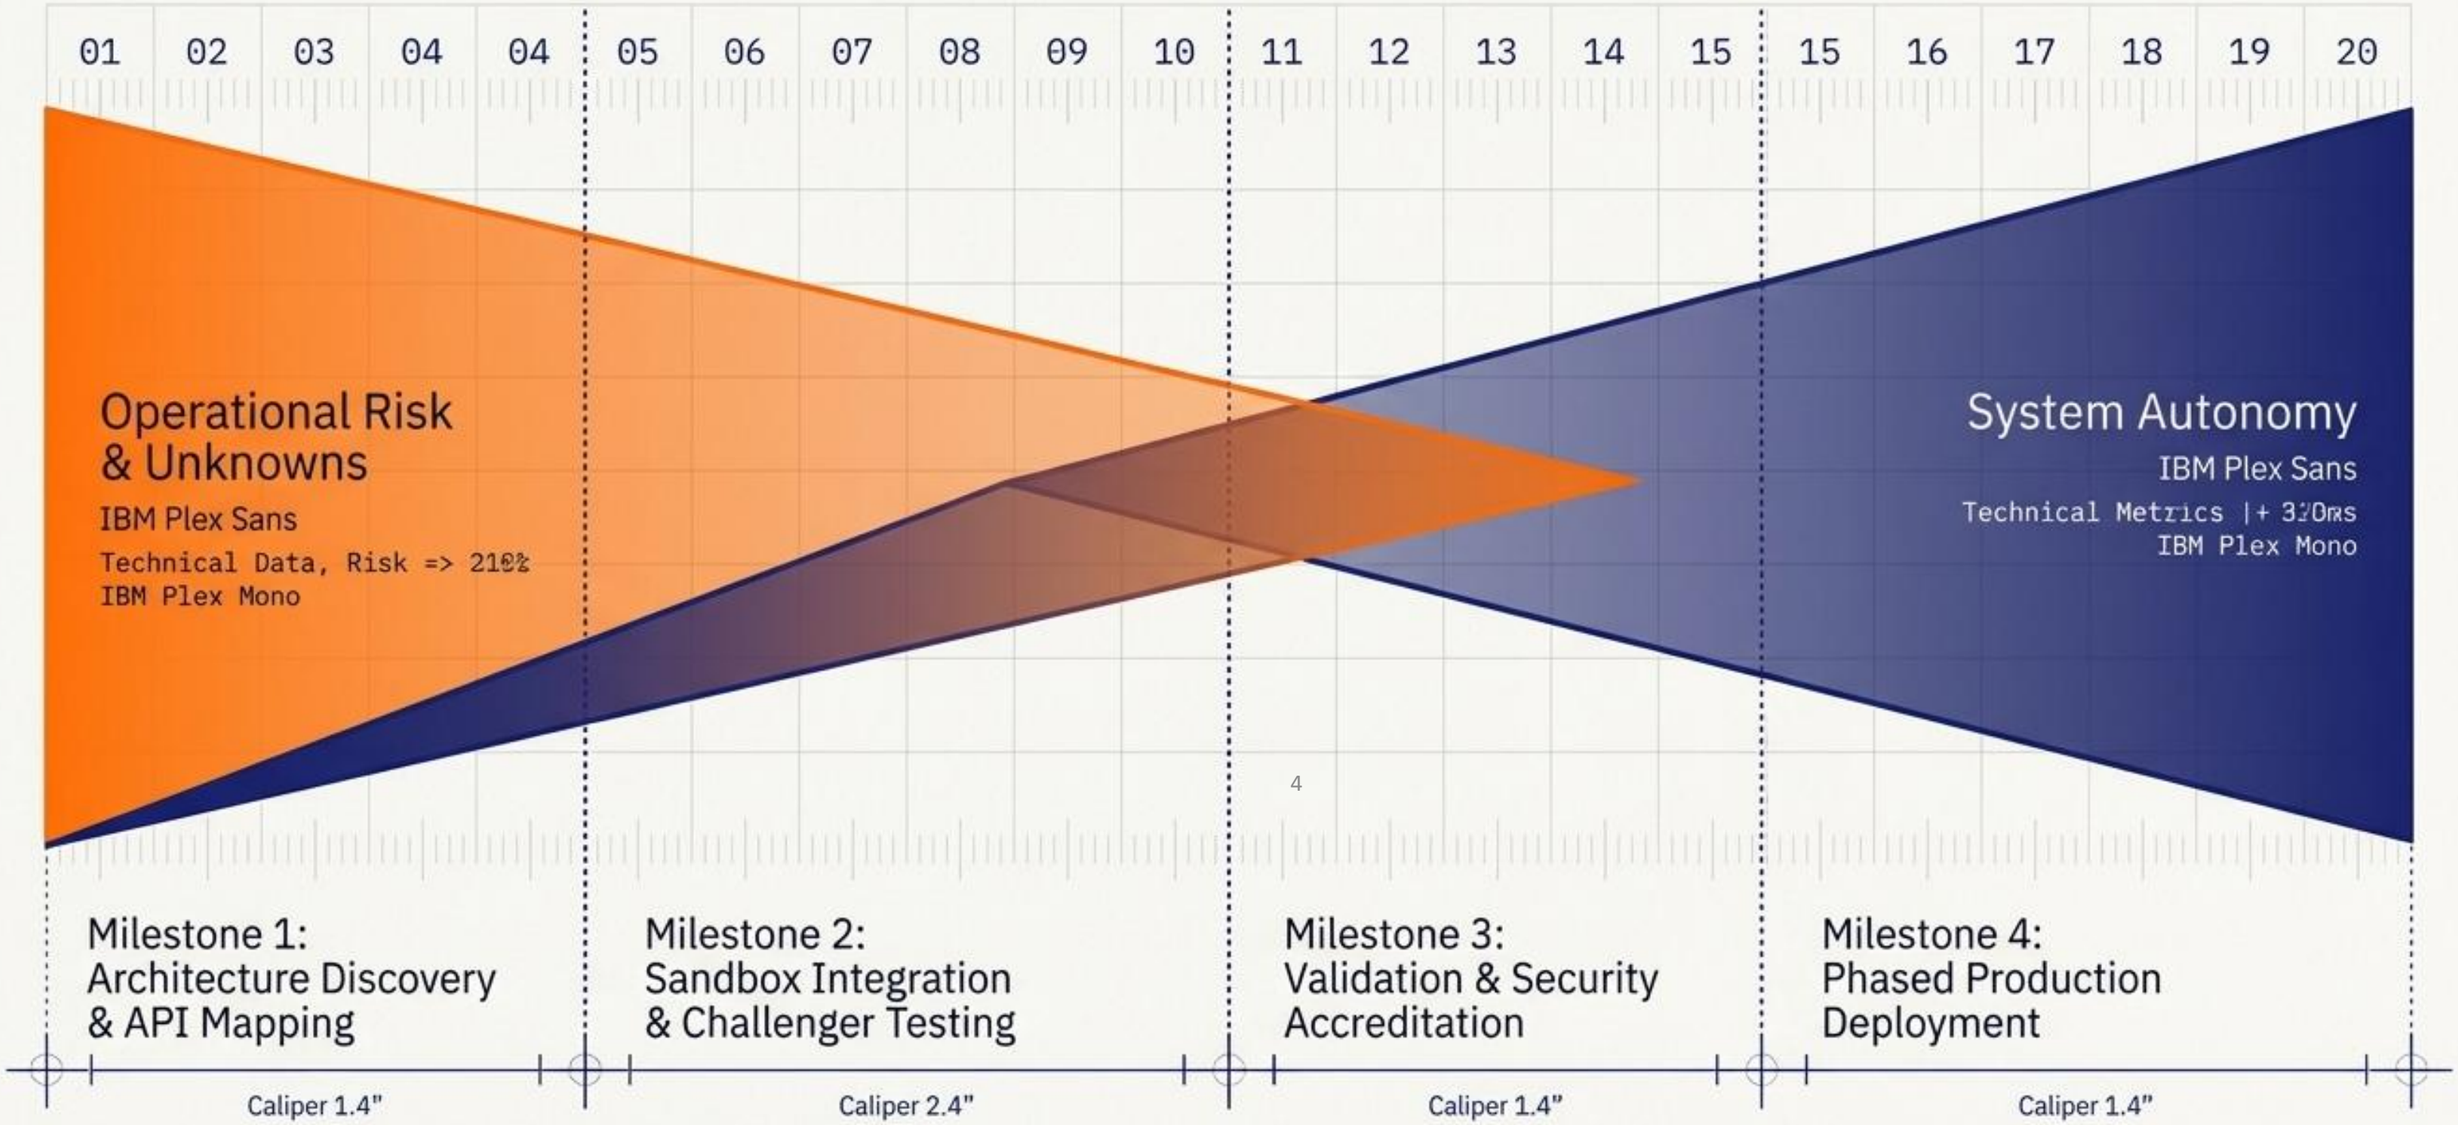

Phase 1: Architecture Discovery & Boundary Definition

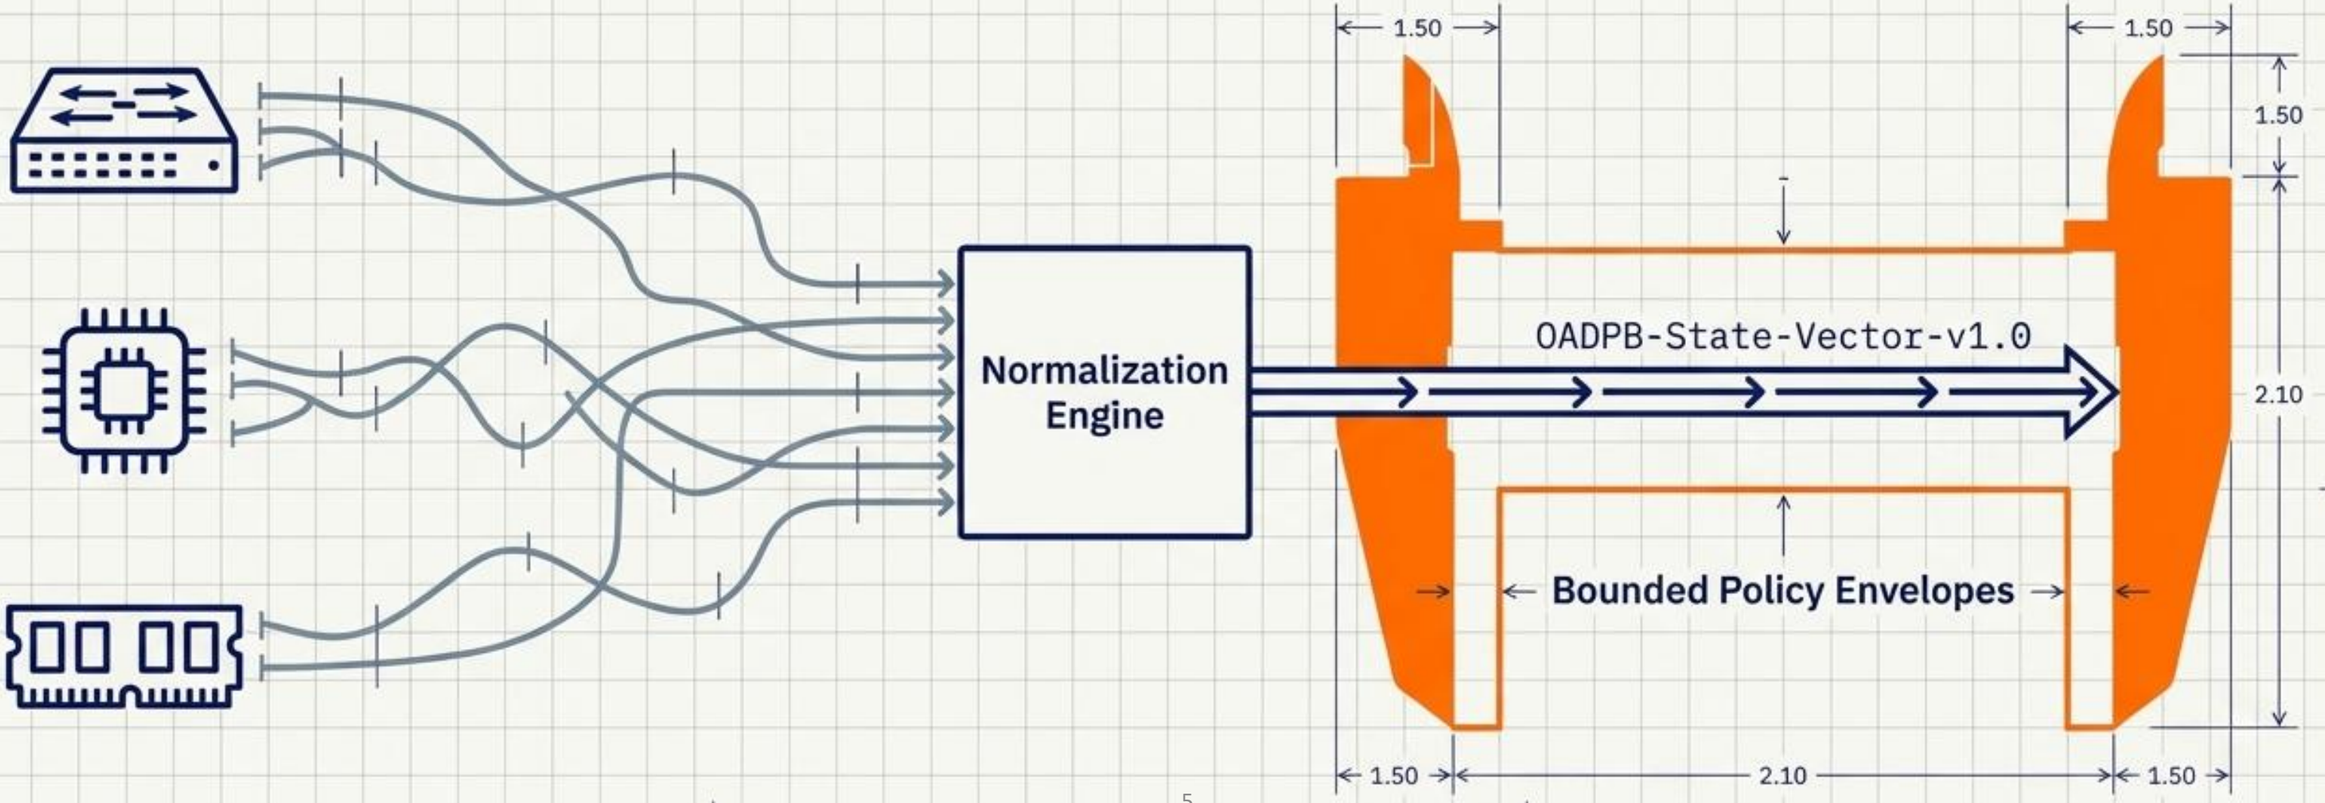

Hardware Analysis

Identifying target COTS appliances, DPUs (e.g., Broadcom Thor), and CXL controllers (e.g., Astera Labs Leo).

Telemetry Normalization

Translating raw queue depths and workload phases into the standardized OADPB schema.

Policy Envelopes

Defining the absolute mathematical limits (maximum traffic shifting, minimum SLA latencies) before deployment.

Phase 2: Sandbox Integration & Challenger Testing

Agent Deployment: Lightweight edge agents installed on target controllers.

Quarantined Learning: The PoL-TT model ingests live telemetry to build its predictive density matrix with zero live actuation.

Advisory Verification: Proving predictive accuracy and ROI by comparing model conditionals against actual historical network congestion.

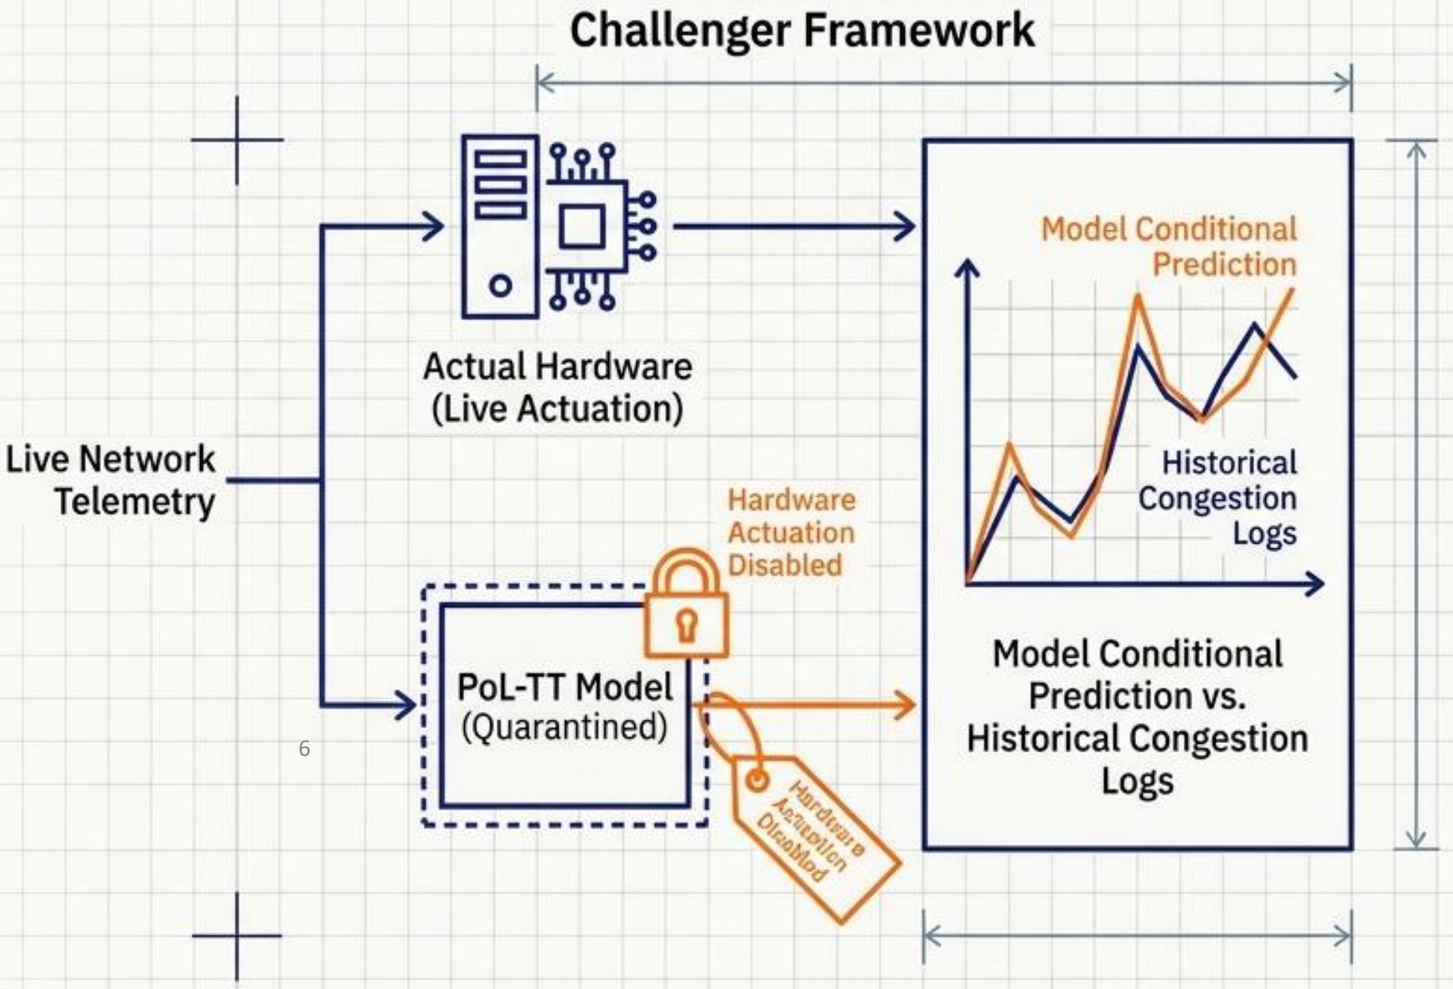

Phase 3: Validation, Security & Accreditation

1. Cryptographic Provenance

Validating telemetry and policy envelopes with signatures from Trusted Execution Environments (TEEs).

2. Tenancy Scrubbing

Stripping proprietary application signatures and tenant IDs into compact mathematical summaries.

3. Anomaly Quarantine

Simulating massive cyberattacks and Electronic Warfare (EW) jamming to ensure pathological data is quarantined, not absorbed.

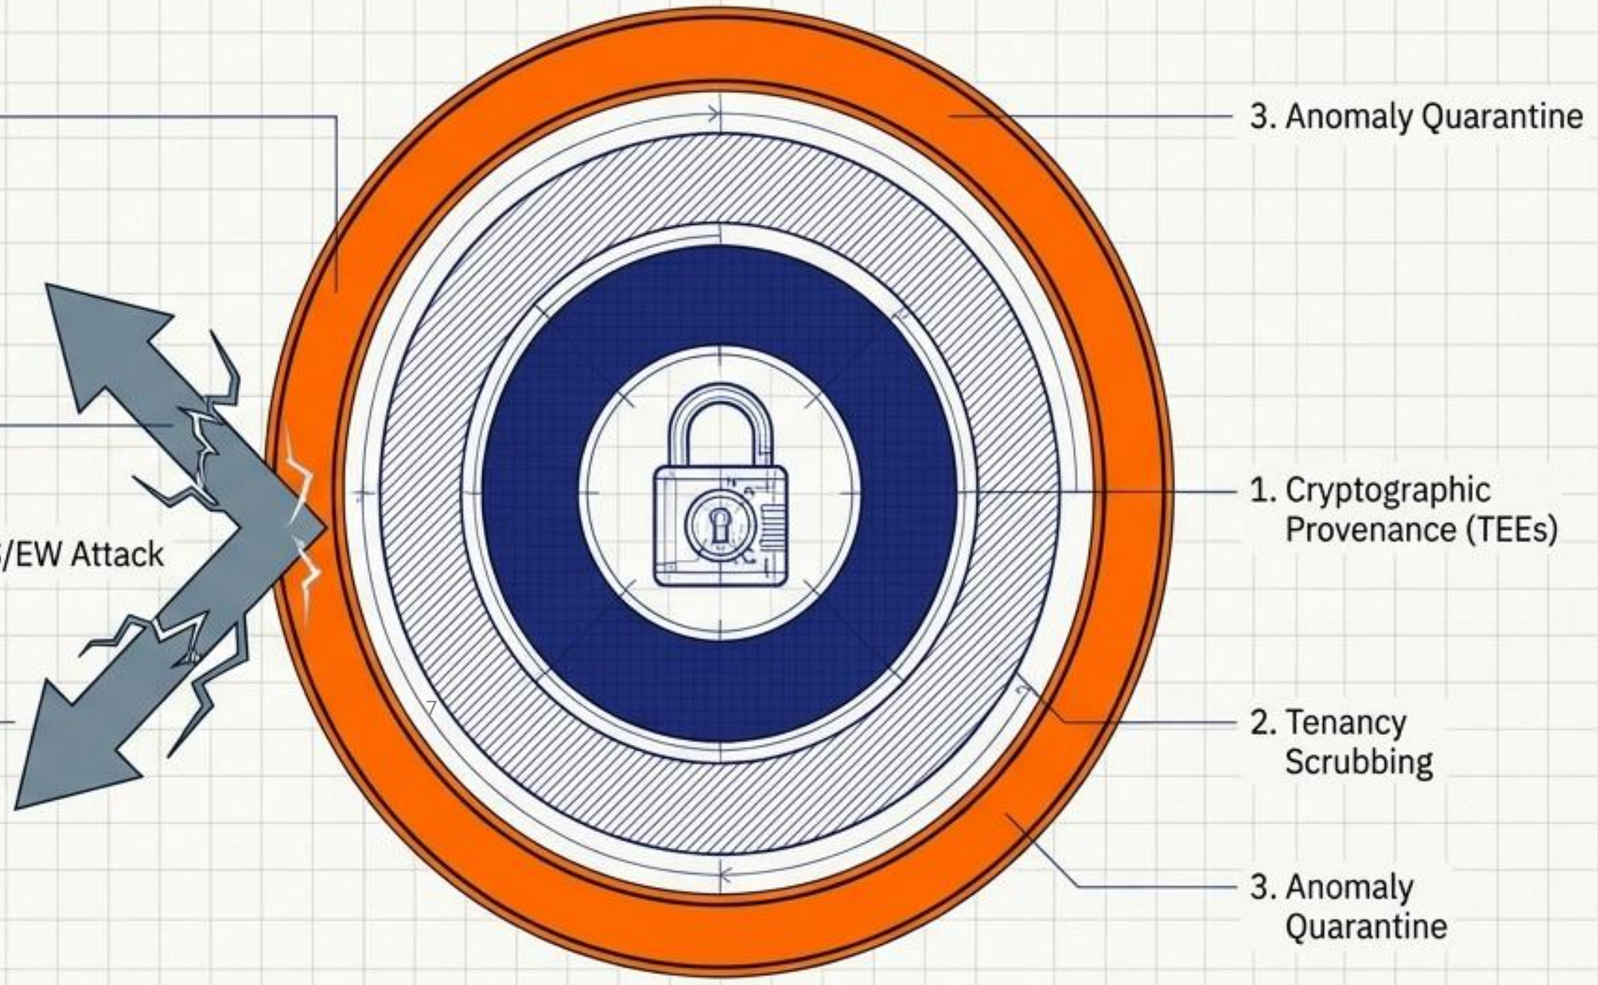

Phase 4: Phased Production Handover

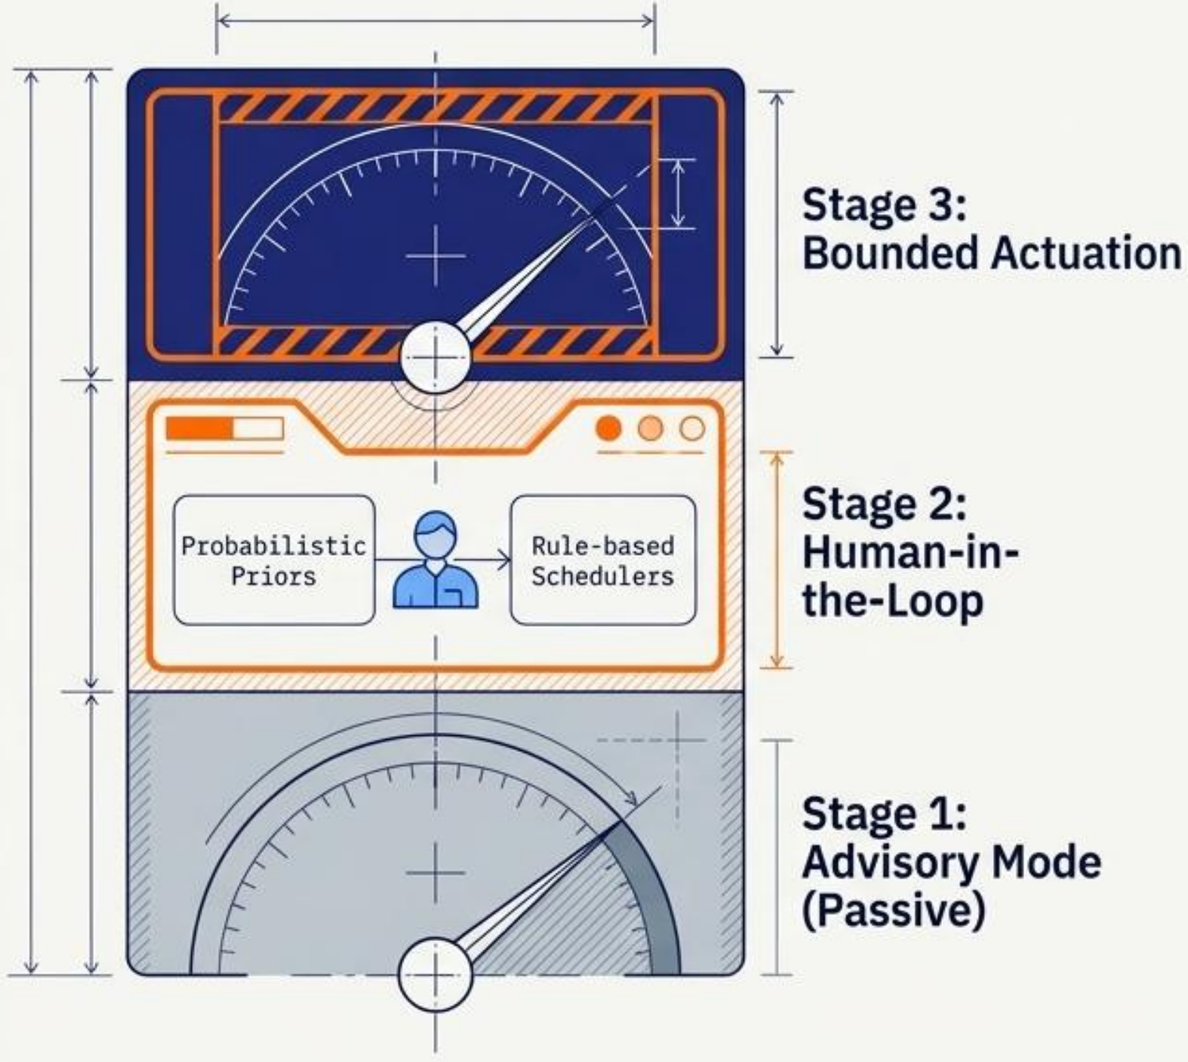

Human-in-the-Loop (Stage 2):

PTCP feeds predictive priors into existing schedulers. Manual approval required for critical infrastructure actuation.

Bounded Actuation (Stage 3):

Autonomous control granted for pre-emptive traffic pacing, predictive CXL tiering, and cross-layer checkpointing—strictly within Phase 1 hardcoded limits.

Industry-Specific Customization Matrix

Domain	Primary Focus	Integration Target	Key Deliverables
OEM Partners	Deep firmware integration	ARM cores & matrix accelerators	Mapping TT-cross mathematical updates for hyperscaler GTM.
Defense & DoD	Air-gapped, SWaP-constrained	CJADC2 edge	IL6/FedRAMP, EW anomaly quarantine, functioning over jammed networks.
Telecommunications	5G/6G & O-RAN	RAN Intelligent Controller (RIC)	Near-RT xApp deployment, enforcing URLLC slicing against eMBB noise.
SLED & Utilities	Legacy hardware maximization	SCADA grids	Pacing non-critical telemetry during weather events to ensure real-time visibility.
Enterprise/Financial	Microsecond latency	HFT & Life Sciences	Pre-warming CXL/HBM tiers with historical tick or genomic sequence data.

Dedicated Managed Services: Resident Engineering

The Guarantee

Predictive AI fabrics are living systems; we do not deploy and abandon. Highly specialized Resident Engineers (REs) embed directly on-site or at C2 facilities.

Continuous Model Tuning

Utilizing Singular Value Decomposition (SVD) for background recompression of Tensor Train cores, preventing mathematical rank inflation.

Policy Envelope Adjustments

Manually adjusting bounded action spaces as hardware scales or workloads shift (e.g., from LLM Training to Inference).

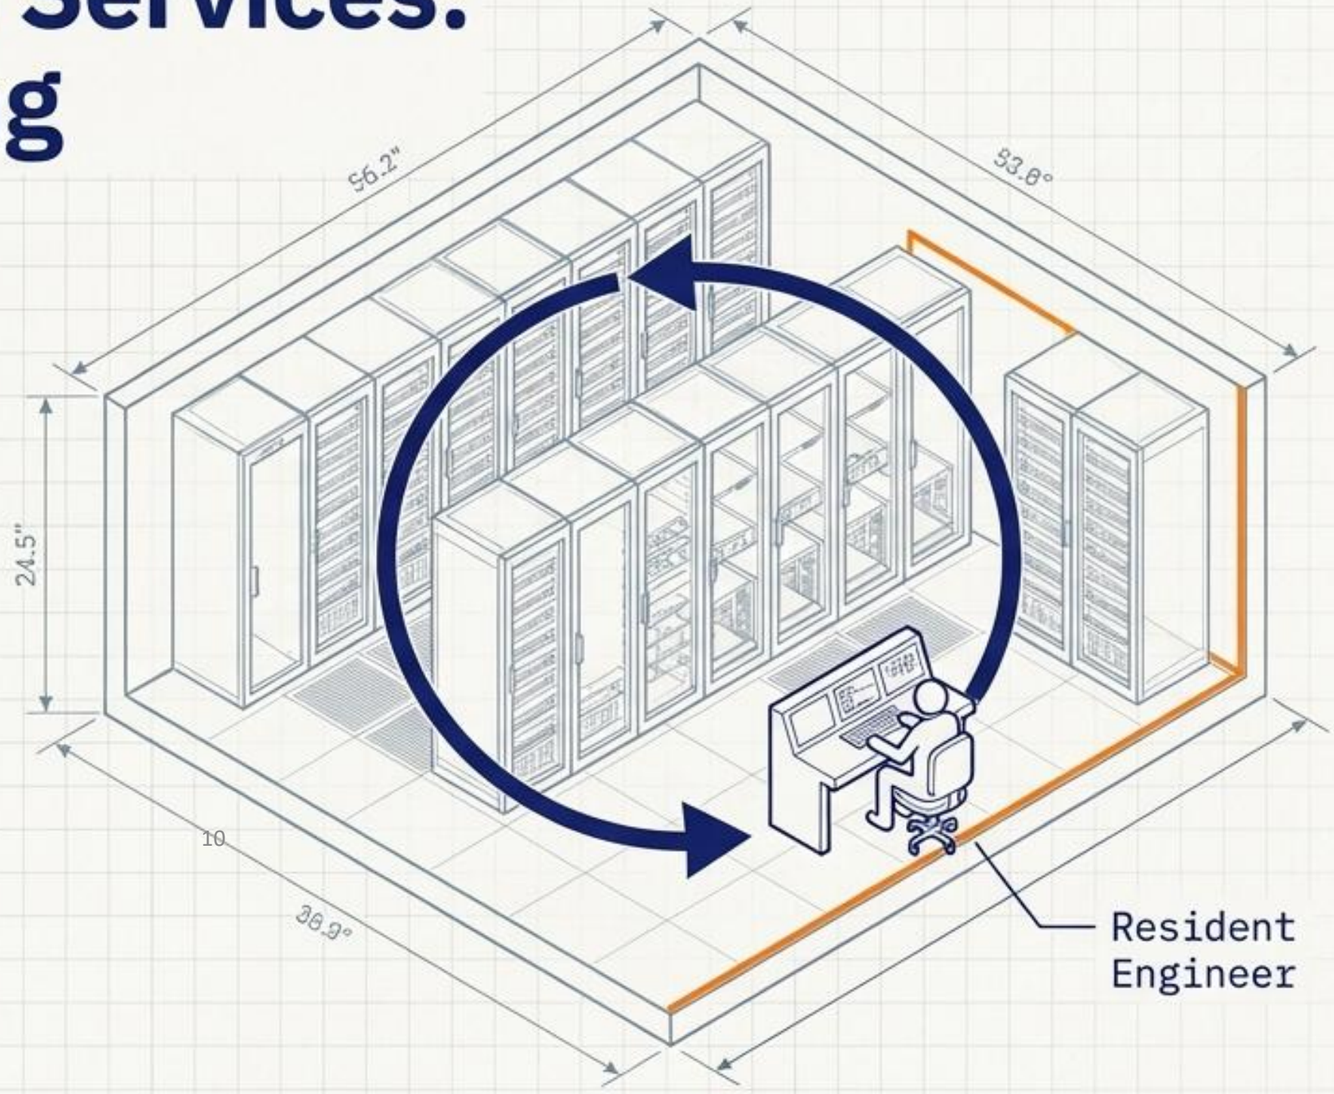

Global Threat Intelligence & Tier-3 Support

Macro-Level Drift Analysis

Analyzing compact mathematical summaries centrally to detect slow-moving baseline drift across multi-site infrastructures.

Federated Security Updates

Pushing signed, encrypted model updates to immunize local deployments against newly discovered, AI-driven network attacks.

Tier-3 SLA Guarantee

24/7/365 escalation to the core architectural engineers (U.S. Patent No. 11,308,384 B1 authors), complete with bi-weekly CapEx and AI yield (Tokens/s) performance reviews.

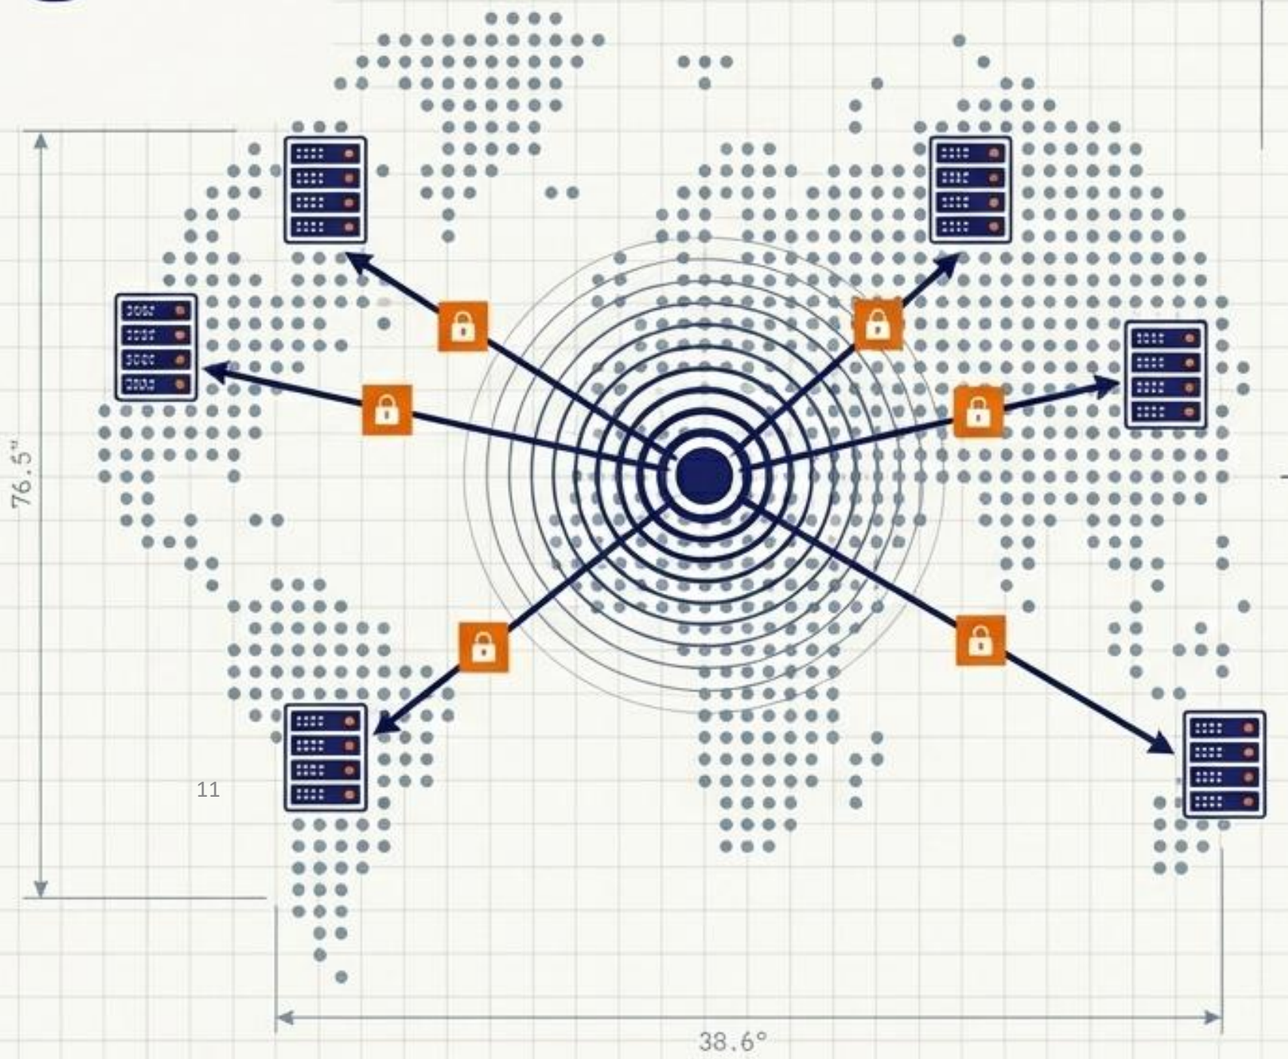

PTCP Software Topology & Licensing Map

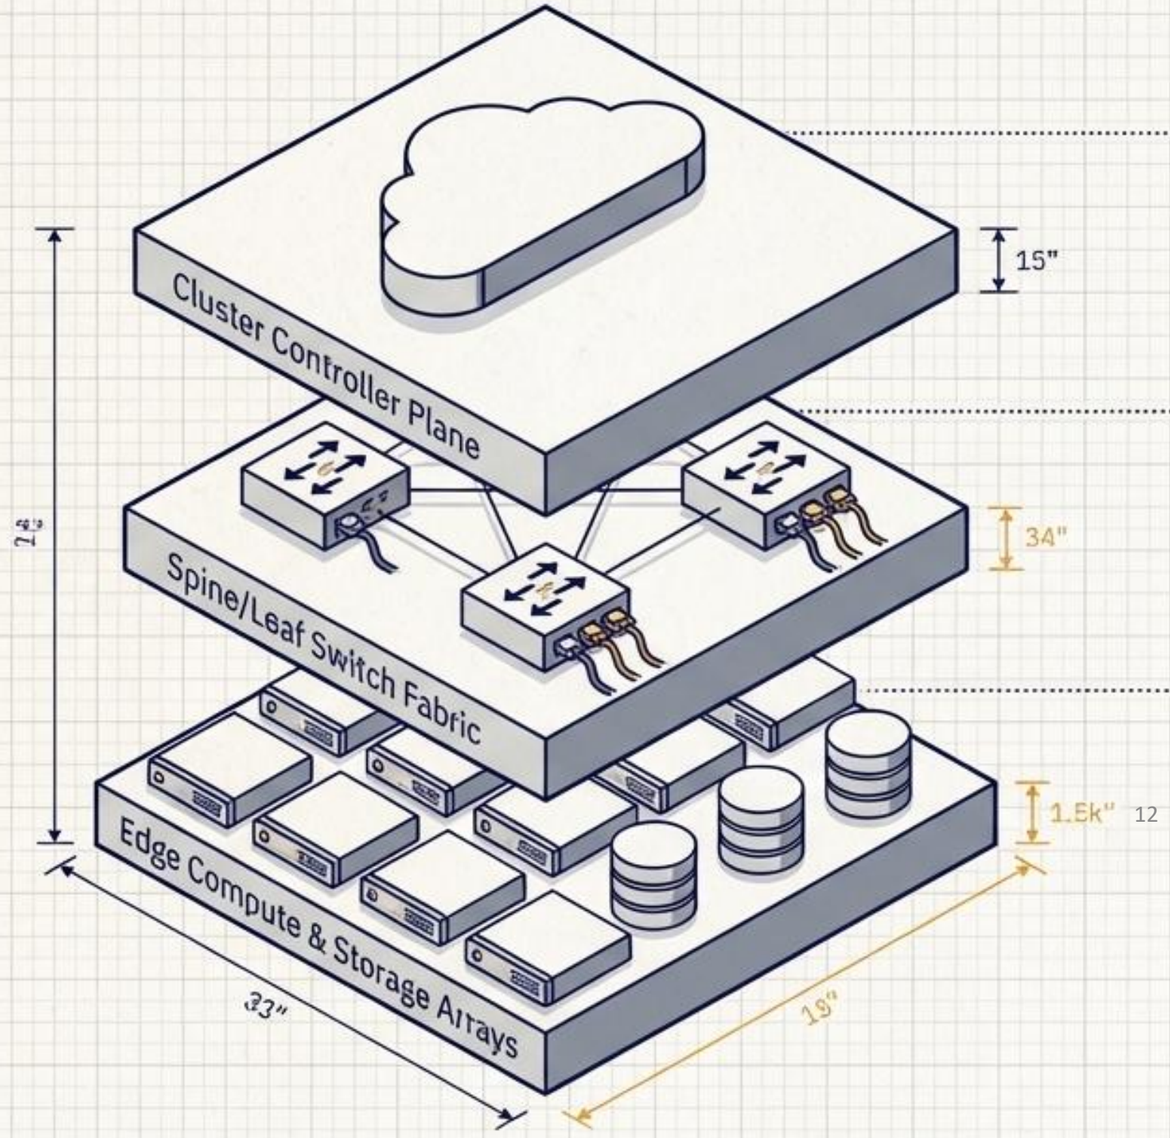

116"

2"

Cluster Controller ($150k/yr)

Overarching semantic engine managing global policy envelopes and threat quarantine (VCF/Central RIC).

Network Fabric Agent ($4k/yr per switch)

PoL-TT fast-path execution pacing network congestion (e.g., Tomahawk).

Edge Compute / DPU Agent ($1.5k/yr per node)

Localized traffic shaping on SmartNICs/DPUs (e.g., Thor, BlueField).

CXL / Storage Agent ($800/yr per controller)

Predictive cache tiering and checkpoint buffering (e.g., Astera Leo)

Note: 15% discount applied >1,000 edge units, 25% >5,000 units.

NRE & Deployment Services Modules

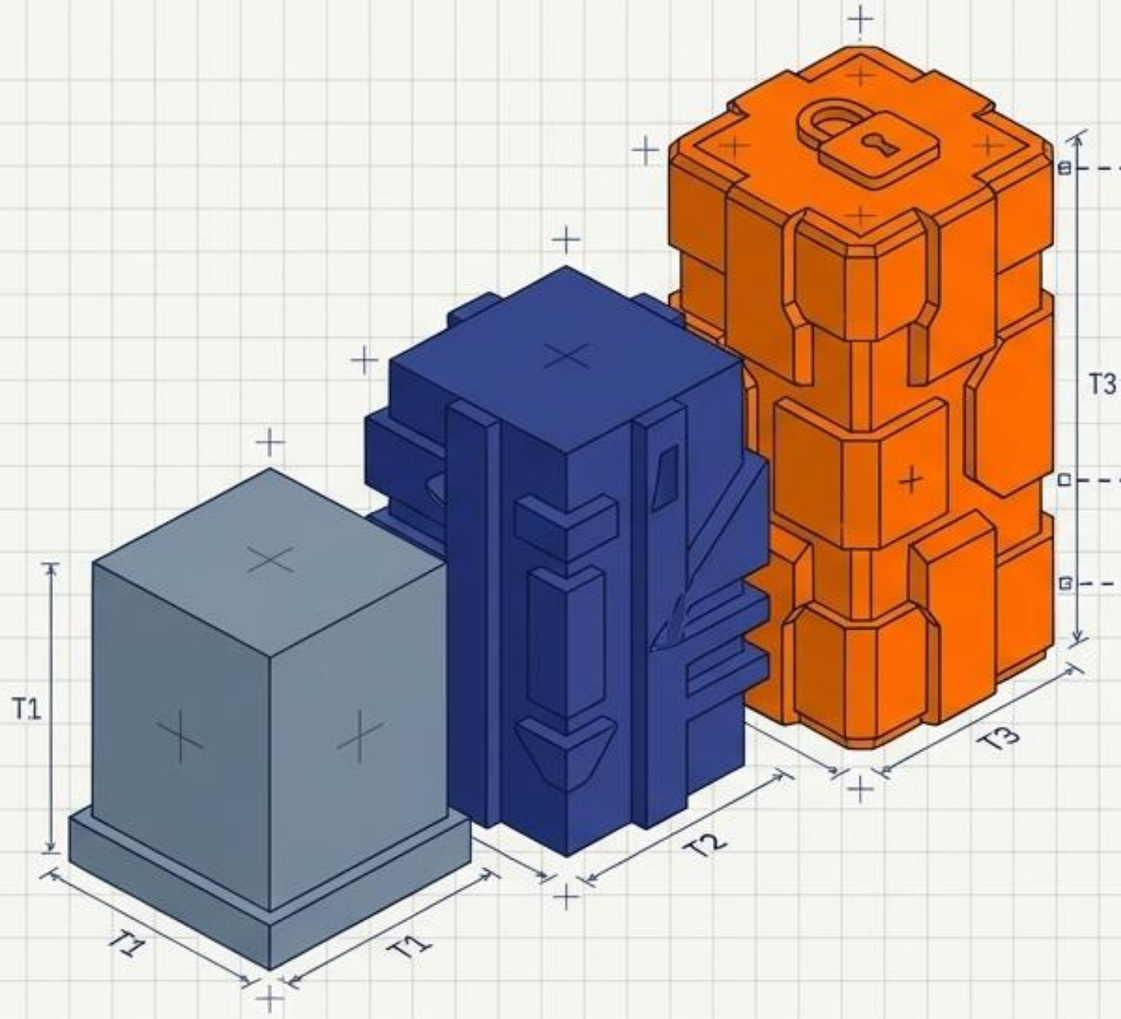

Tier 1: Enterprise/SLED ($250k - $400k).

Standard COTS Data Centers, Shared HPC clusters.
Basic telemetry normalization and standard envelopes.

Tier 2: Telecom/Utility ($500k - $750k).

5G/6G MEC, RIC integration, SCADA grids. Custom
O-RAN xApp development, high-availability sandbox.

Tier 3: Defense/OEM ($850k - $1.2M+).

Silicon-level integration, Air-gapped DoD, CJADC2.
Custom ARM mapping, IL6/FedRAMP hardening, EW testing.

*Note: Billed on a milestone basis (25% increments across execution, post-sandbox, post-validation, and production).

Total Cost of Ownership (TCO) Baseline

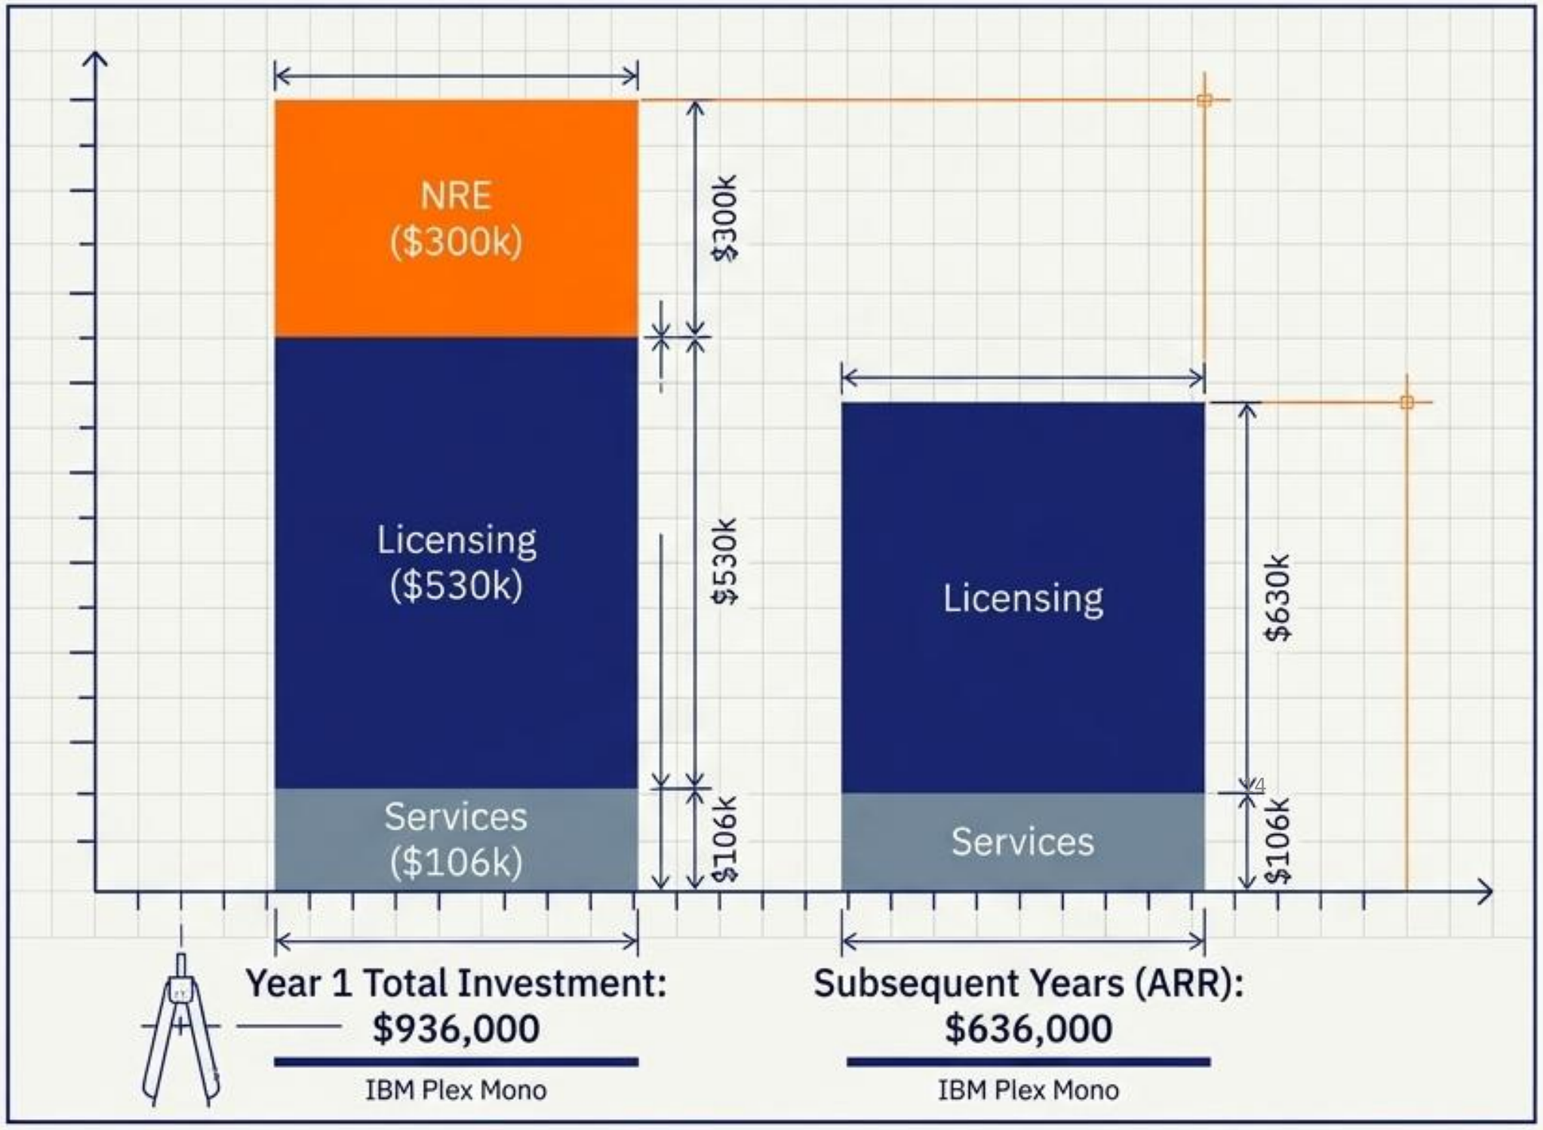

Spec-sheet

- **Deployment Profile:** Mid-sized Enterprise AI Factory.

IBM Plex Mono

- **Infrastructure Scale:** 1 Central Controller, 20 Network Switches, 200 Compute Nodes.
- **Year 1 Breakdown:** Implementation NRE ($300k) + Software Licensing ($530k) + Standard Remote Managed Services ($106k).

IBM Plex Sans

IBM Plex Mono

Thank You for Your Partnership We look forward to building the future of networking together. Please don't hesitate to reach out with any questions.

Tensor Networks, Inc. Contact Us:

information@tensornetworks.net