carahsoft

TENSOR NETWORKS

Red-Team-Resilient Post-Quantum Mission Fabric

Thank you for your interest in exploring this content.

Carahsoft is the **Trusted Government IT Solutions Provider**® supporting a broad portfolio of industry-leading technologies and a wide range of contract vehicles.

As the **Master Government Aggregator**®, Carahsoft connects government agencies, industry partners, and technology providers to deliver innovative, mission-focused solutions.

In partnership with Tensor Networks, we provide technology solutions that drive modernization, strengthen operations, and ensure compliance with evolving government standards.

To learn more about how Carahsoft can support your technology needs, please visit carahsoft.com

Explore More Resources:
carah.io/TensorNetworksResources

Join Events & Webinars:
carah.io/Events

Discover Technology Solutions:
carah.io/TensorNetworks

Learn About Procurement:
carah.io/TensorNetworksContracts

Connect With Our Team:
TensorNetworks@Carahsoft.com
888-662-2724

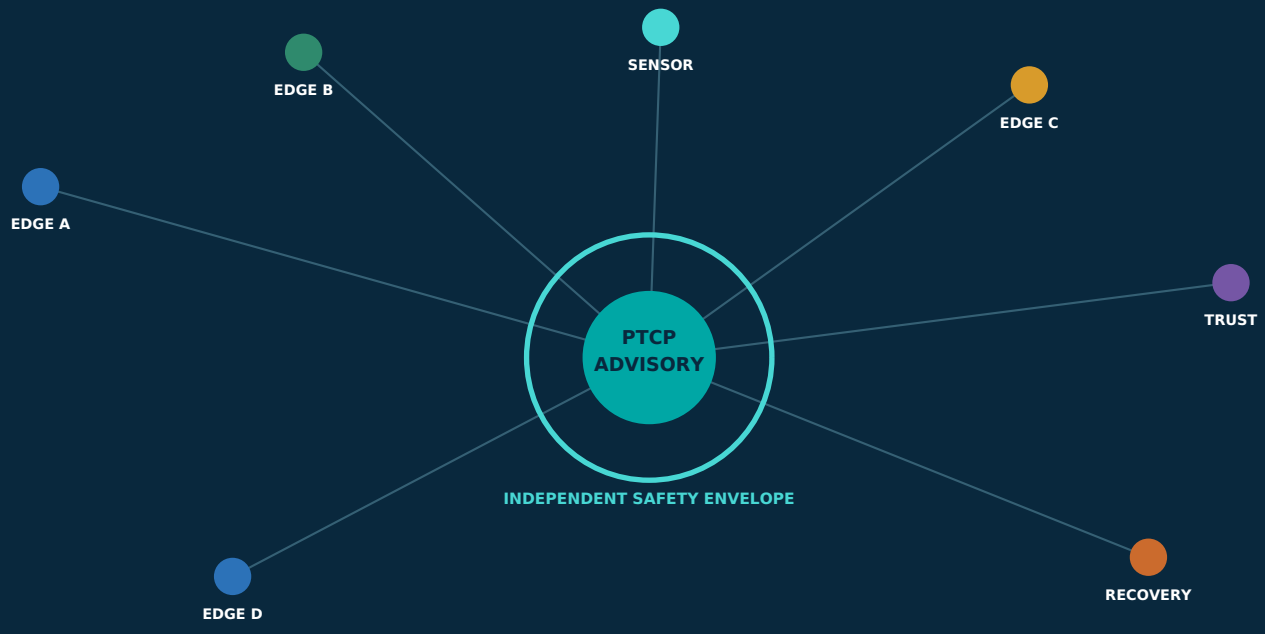

PUBLIC-SOURCE TECHNICAL WHITE PAPER | JULY 2026

Red-Team-Resilient Post-Quantum Mission Fabric

Predictive Tensor Control Plane architecture fast/slow paths with fail-safe local autonomy

POL-TT density modeling, Dtopo topology security, bounded quarantine, controller-compromise containment, and classical PQC cryptographic agility

SCIENTIFIC AND OPERATIONAL STATUS

PTCP.ai publishes a fast-path / predictive-path implementation concept and edge demonstration materials. Independent operational benchmarks were not identified; all authority remains outside PTCP analytics.

Prepared as	Postdoctoral-level, source-grounded systems analysis
Research basis	Julia Ochoa PTCP and TNQG supplied manuscripts; primary standards
Distribution	Public-source; defensive architecture; not an operational order

Version 2.0 | Evidence checked through July 17, 2026

Executive abstract

A predictive network controller can improve routing and anomaly awareness, but it also concentrates influence over the topology it observes. If the controller, its telemetry, or its learning process is compromised, a system designed to defend the mission can become a mechanism for traffic redirection, selective isolation, trust suppression, or denial of service. The design problem is therefore not simply to make PTCP more accurate. It is to make the mission fabric remain safe, useful, and cryptographically governed when PTCP is wrong, isolated, deceived, or hostile.

This white paper defines a **red-team-resilient, post-quantum mission fabric** in which the Predictive Tensor Control Plane (PTCP) is a bounded analytics and recommendation service, never the sole authority for route acceptance, quarantine, identity, key management, or recovery. The Pattern-of-Life Tensor Train (POL-TT) estimates a compressed joint density over normalized telemetry and produces a negative-log-likelihood anomaly score. The source-defined Dtopo score then combines regional density anomaly, graph-curvature gradients, and normalized cut-capacity shifts. These signals are fused with deterministic route-security evidence, provenance quality, persistence, and independent observers before any high-impact action. [1]

The revision adds a deployable two-timescale implementation. A computationally heavier **slow path** maintains the global density model, forecasts, Dtopo evidence, policy candidates, and champion/challenger promotion. A deterministic **fast path** at each assurance cell evaluates local events against precompiled, signed, expiring action templates and executes only after local safety checks. PTCP.ai publicly describes this design as a brain/reflex architecture, reports a microsecond-class edge-reflex concept, and states that PTCP can operate on existing SONiC leaf/spine hardware. These are vendor-published implementation claims; the public materials reviewed do not provide independent reproduction, raw benchmark data, or operational certification. [39-42]

Problem statement

Mission networks must distinguish a true topology attack from equipment failure, legitimate maintenance, demand shifts, disconnected operations, and incomplete observation. A centralized predictive controller may itself be poisoned, captured, replayed, or cut off. Route leaks and trust degradation can redirect traffic or manufacture false bottlenecks, while indiscriminate quarantine can remove the last viable mission path. At the same time, migration to post-quantum cryptography changes message sizes, compute cost, protocol negotiation, certificate handling, and key-management dependencies. A safe architecture must detect manipulation without allowing either a model or a single controller credential to become a mission-wide kill switch.

Solution statement

Separate observation, inference, authorization, enforcement, cryptography, and recovery into independently attestable planes. Use POL-TT and Dtopo as explainable anomaly evidence; bind every recommendation to signed source evidence, uncertainty, expiry, and a mission authority contract; require deterministic validation or independent corroboration for destructive containment; project every action into a local safety envelope; and retain a known-good, controller-independent operating mode. Provide classical PQC through profile-pinned crypto brokers, validated modules, cryptographic bills of materials, and explicit downgrade policy. During controller loss or compromise, local cells continue only pre-authorized functions within time, scope, route-change, trust, and data-release budgets, then reconcile through a separately keyed recovery plane.

Principal conclusions

- POL-TT is useful because it scores joint behavior rather than isolated counters, but its output is evidence, not proof of attack. Density models are vulnerable to poisoning, evasion, concept drift, and observation gaps; online learning must freeze or roll back when provenance or residual tests fail. [1,29]

- Dtopo can expose topology deformation that scalar alarms miss: a route leak may be ordinary at one interface yet abnormal in prefix fan-out, AS-path relationships, graph curvature, and cut redundancy across a region. The original Dtopo definition should be preserved, calibrated, and accompanied by a visible evidence vector rather than hidden inside an opaque composite. [1,6,7,21]
- Deterministic routing controls remain authoritative. RPKI origin validation, BGP Roles and the Only-to-Customer attribute, prefix/AS-path policy, maximum-prefix limits, and path-validation mechanisms provide protocol evidence that PTCP should consume, not replace. [22-28]
- Quarantine is a reversible state machine, not a binary classifier. Escalation requires confidence, persistence, corroboration, and a survivability check; de-escalation uses lower thresholds, dwell time, attestation, and controlled replay to prevent oscillation and premature trust restoration.
- A controller signature proves which credential authorized a message, not that the message is safe. Local enforcers must validate sequence, freshness, scope, mission authority, cryptographic profile, and safety invariants; high-impact actions require a separately administered co-authorization or deterministic hard-fail condition.
- Bounded local autonomy is defined by contracts: explicit action sets, time-to-live, route-change budgets, blast-radius limits, minimum service, data-age limits, trust ceilings, and fail-safe outcomes. Local autonomy cannot create a new trust root, weaken cryptography, expand cross-domain release, or silently acquire broader mission authority.
- Post-quantum cryptography is classical cryptography executed on conventional systems. FIPS 203, 204, and 205 standardize ML-KEM, ML-DSA, and SLH-DSA; safe migration requires inventory, profile management, validated implementations, interoperability testing, and authenticated downgrade resistance. [12-19]
- PTCP.ai publishes fast-path, autonomic-control, SONiC/existing-hardware, and benchmark-validation materials. The same materials distinguish theoretical complexity from validation protocols still required. Treat them as vendor evidence of an implementation direction, not independent proof of operational performance. [39-42]
- Initial deployment should be offline replay and shadow mode; controlled actuation follows only after false-containment cost, controller-compromise behavior, fast-path timing, slow-path convergence, and independent fallback performance are demonstrated. [1,42]

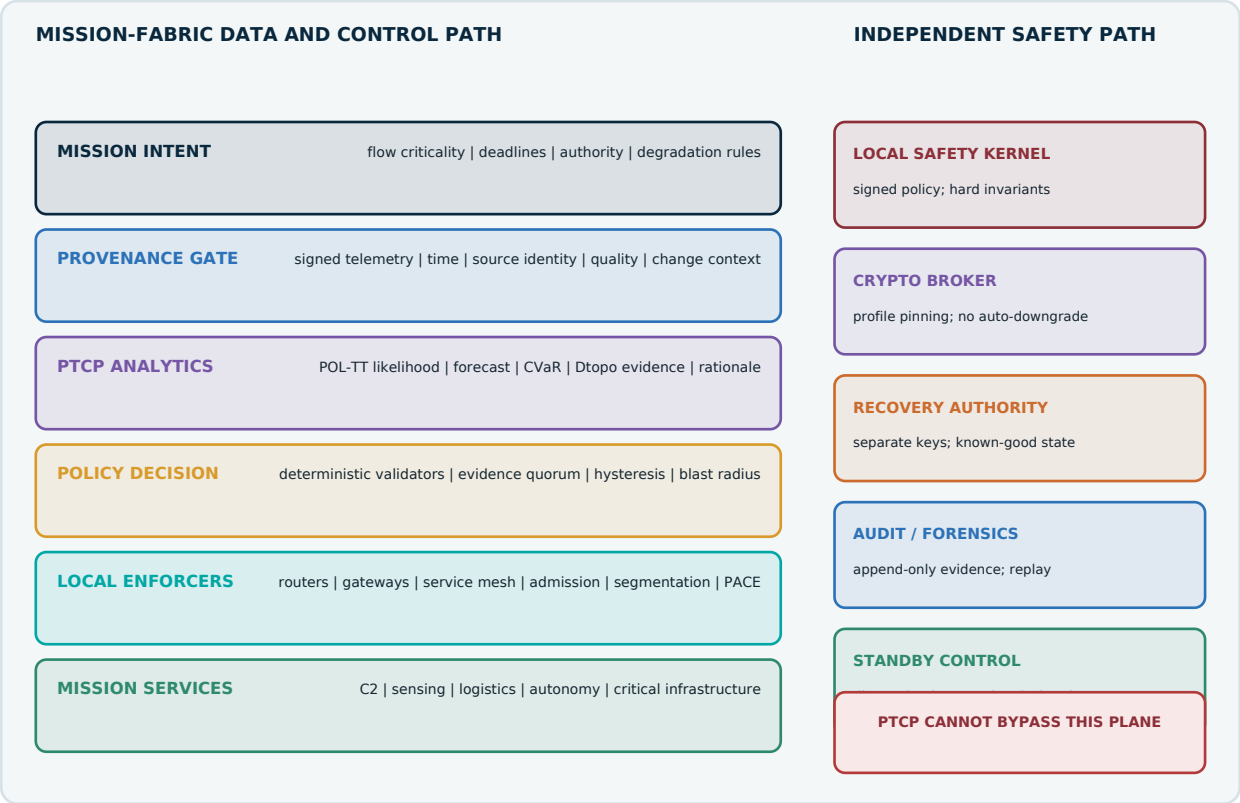

Figure 1. Reference separation of authority. PTCP supplies evidence and bounded recommendations; independent local safety, cryptographic, and recovery functions retain final enforcement authority.

Design axioms

Axiom	Architectural consequence	Failure prevented
No sole predictive authority	PTCP cannot directly create trust roots, change crypto profiles, or issue unrestricted quarantine.	Controller capture becomes advisory compromise rather than universal control.
Evidence before action	Every action carries source lineage, confidence, freshness, model version, rationale, and expiry.	Untraceable or replayed recommendations are rejected.
Deterministic controls first	Protocol-invalid or policy-invalid routes are handled by standards-based controls; PTCP adds cross-signal detection.	Statistical uncertainty cannot override explicit invalidity.
Unknown is not normal	Low coverage or conflicting observers trigger constrained mode, sampling, and review.	Telemetry suppression cannot manufacture confidence.
Local safety is authoritative	Each enforcement cell projects recommendations into a signed, versioned safety envelope.	A remote optimizer cannot violate minimum service or blast-radius limits.
Recovery uses separate authority	Known-good images, revocation, and standby control are administered with independent keys and paths.	A compromised controller cannot approve its own recovery.
No silent crypto downgrade	Algorithm changes are policy-authorized, signed, observable, and reversible.	Isolation cannot force quantum-vulnerable or unauthenticated fallback.

Document control, terminology, and method

Field	Value
Title	Red-Team-Resilient Post-Quantum Mission Fabric: PTCP, POL-TT, Dtopo, and Fail-Safe Local Autonomy
Purpose	Define a testable defensive architecture, fallback design, development scope, and verification program.
Research basis	Julia Ochoa PTCP and TNQG manuscripts supplied by the user; PTCP.ai public implementation materials; current primary public standards and guidance.
Evidence date	External standards and policy checked through July 17, 2026.
Classification	Public-source technical analysis; no classified parameters, target data, exploit instructions, or operational orders.
Scientific status	PTCP has vendor-published implementation and demonstration claims, but no independently reproduced operational benchmark was identified. TNQG is conjectural and is not an operational dependency.
Intended readers	Network architects, cyber-resilience engineers, cryptographic migration leads, test authorities, and research program managers.

Terminology choices

- **POL-TT** denotes the manuscript's Pattern-of-Life Tensor Train probability-density model. The manuscript alternates capitalization as PoL-TT; this paper uses POL-TT consistently.
- **Dtopo** denotes the source-defined topology defect score composed of regional anomaly, graph-curvature-gradient, and normalized cut-shift terms. Extensions are shown as a separate evidence vector so they are not misrepresented as part of the source equation.
- **Quarantine** means a reversible, scoped restriction such as route de-preference, admission reduction, traffic drain, credential restriction, or segment isolation. It does not imply physical destruction or permanent exclusion.
- **Classical PQC** means quantum-resistant algorithms implemented on classical computing platforms. It is distinct from quantum key distribution, quantum networking, and TNQG.
- **Bounded local autonomy** means pre-authorized operation within explicit limits during degraded connectivity. It does not confer new command, release, legal, or cross-domain authority.

Method and claim discipline

The paper first reconstructs the PTCP proposal from its equations and algorithms, then subjects the controller, evidence pipeline, and recovery plane to a defensive adversary model. Architectural additions are derived engineering recommendations, not claims made by the Ochoa manuscript. Standards-based routing and cryptographic controls are treated as the baseline; predictive analytics augment them. Quantitative values are framed as program-calibrated parameters unless an authoritative source defines a requirement. The design intentionally avoids operational network identifiers, exploitable thresholds, weapons-employment logic, and classified mission data.

Claim boundary

The supplied PTCP paper defines algorithms and a validation protocol but reports no experimental results. PTCP.ai separately publishes an autonomic brain/reflex concept, a microsecond-reflex presentation, and statements about operation on existing SONiC leaf/spine hardware. Those public materials establish vendor intent and claimed implementation direction; they do not, by themselves, establish independent operational effectiveness. This paper therefore attributes website claims to the vendor and reserves terms such as validated or proven for reproduced test evidence. The supplied TNQG paper is conjectural; no production security or communications function depends on it. [39-42]

Contents

1. Problem statement and design objectives

1.1 The control-plane paradox

PTCP is intended to observe a high-dimensional network, forecast likely future conditions, rank candidate paths by expected cost and tail risk, and recommend topology-security actions. Those functions are valuable precisely because they aggregate a view broader than any one router. The same aggregation creates a high-value target. A compromised observer can falsify the state; a poisoned model can redefine normality; a captured controller can sign unsafe recommendations; and an isolated controller can leave dependent enforcers without current intent. The architecture must therefore assume compromise of any single analytics, management, identity, time, or transport dependency.

1.2 Mission consequences

Adversarial effect	Network symptom	Mission-level risk	Required property
Route leak or unauthorized redistribution	Unexpected prefix fan-out, path relationship violation, centrality shift	Traffic redirection, overload, observation, or black hole	Deterministic validation plus cross-observer anomaly detection
Trust degradation	Attestation failure, credential revocation, posture decline, logging loss	Compromised nodes remain preferred or legitimate nodes are falsely excluded	Trust-source diversity, freshness, and scoped response
Telemetry poisoning	Biased quantiles, forged counters, coordinated low-and-slow drift	POL-TT learns adversarial behavior as normal	Signed provenance, guarded learning, holdout replay, rollback
Controller compromise	Validly signed but unsafe actions, selective suppression, rapid policy churn	Mission-wide self-denial or covert redirection	Local safety kernel, dual authorization, independent revocation
Controller isolation	Expired intent, stale model, lost global coordination	Essential services stop or operate without bounds	Known-good local policy, PACE, store-forward, timed autonomy
Crypto downgrade or key-service loss	Handshake failures, incompatible profiles, forged negotiation	Loss of confidentiality, integrity, or availability	Profile pinning, CBOM, approved fallback, local key continuity

1.3 Design objectives

- Detect topology manipulation from a combination of statistical, graph-structural, cryptographic, and policy evidence while making uncertainty explicit.
- Contain suspected regions without removing every path required for essential service, recovery, or operator reach-back.
- Keep local cells useful when wide-area coordination is unavailable, but prevent isolated cells from expanding trust, disclosure, or actuation authority.
- Prevent a compromised PTCP instance, model artifact, signing key, or orchestration service from bypassing independent safety and cryptographic policy.
- Maintain classical PQC agility across constrained devices, routers, service meshes, control messages, software/model signing, and recovery artifacts.
- Produce auditable evidence for operators, incident responders, accreditation authorities, and post-event model improvement.
- Demonstrate safe failure through red-team testing before any closed-loop deployment.

1.4 Non-objectives

The design does not replace BGP, IGPs, segment routing, software-defined networking, RPKI, BGP Roles, cross-domain guards, zero-trust policy, key-management systems, or incident response. It does not assert that a statistical anomaly is malicious, promise zero false positives, or guarantee uninterrupted service. It does not use TNQG to provide security, faster-than-light communication, or operational spacetime control.

2. PTCP scientific basis and claim boundaries

2.1 Source-defined telemetry and density model

The PTCP manuscript represents each edge by a multimodal feature vector containing available bandwidth, queue occupancy, jitter, loss, energy or operating cost, trust, and additional variables. Robust online quantiles map each coordinate into [0,1]. The global state is discretized into bins and approximated by a Tensor Train with bounded internal rank. For at most n bins per dimension and rank r, the manuscript states O(d n r^2) storage rather than the dense O(n^d) tensor. A point-density query produces the negative log-likelihood anomaly score. [1,3,4]

$$x\_tilde(e,j,t) = clip\_[0,1] ((x(e,j,t) - q(j,0.05)) / (q(j,0.95) - q(j,0.05) + epsilon\_q))$$

$$p\_Theta(s\_t) = max(P\_hat[q\_1(s\_t,1), ..., q\_d(s\_t,d)], p\_min) / Z\_Theta$$

$$S\_anom(s\_t) = -log\ p\_Theta(s\_t)$$

Compression makes joint monitoring computationally plausible, but it also introduces approximation error, rank-selection risk, bin-boundary effects, and model-state dependency. A low estimated density means the current state is unusual under the retained model and context. It does not identify intent or prove attack.

2.2 Predictive routing and bounded actuation

PTCP converts normalized link attributes into a dimensionless effective capacity and inverse-capacity link length. It forecasts scenarios, evaluates candidate paths over a horizon, and selects by expected cost plus Conditional Value-at-Risk (CVaR), with a penalty for route churn. The selected recommendation is projected into a policy envelope. The manuscript explicitly describes PTCP as an overlay that can publish weights, labels, and recommendations to existing control systems. [1,5,36-38]

$$u\_t^* = Pi\_Omega\ arg\ min\ E[J] + rho\ CVaR\_alpha(J)$$

Security interpretation

The projection operator is the essential opening for safe engineering. In this paper, Pi_Omega is not implemented inside the same trust boundary as the optimizer. A separately implemented local safety kernel verifies scope, sequence, expiry, authority, minimum service, route diversity, trust policy, and cryptographic profile before actuation.

2.3 Source-defined Dtopo score

For monitored region R, the manuscript defines Dtopo as the sum of the median regional anomaly score, a weighted graph-curvature-gradient magnitude, and a weighted normalized cut-capacity shift. The curvature term is intended to reveal local changes in graph geometry; the cut term compares current cut capacity or conductance with a robust historical window. Forman-Ricci or Ollivier-Ricci curvature are offered as candidate estimators. [1,6,7]

$$D\_topo(R,t) = median\_R(S\_anom) + nu\ ||grad\_G\ kappa(R,t)||\_2 + xi\ Delta\ C\_cut(R,t)$$

$$a\_t^* = Pi\_Omega\_safe\ (a\_raw(D\_topo))$$

The manuscript proposes two score thresholds: suspicious regions receive increased sampling, stronger authentication, and reduced route preference; higher scores may produce quarantine, segmentation, or traffic drain. It also requires rate limits, maximum blast-radius constraints, minimum service availability, human approval thresholds, rollback, and audit. The present architecture retains those concepts and adds explicit evidence quality, deterministic validators, controller-compromise containment, recovery authority, and a multi-stage state machine.

2.4 Scientific separation from TNQG

The separate TNQG manuscript uses tensor-network cuts, capacities, and reconstructed geometry as a conjectural quantum-gravity research program. Its own limitations state that the geometry dictionary is model-dependent, generic tensor networks need not have a smooth dual, and Lorentzian causal structure requires additional assumptions. [2] PTCP's use of geometric terminology is operational and information-theoretic; the manuscript explicitly says its geodesics are minimum-risk network paths and that no physical AdS/CFT equivalence is assumed. This white paper therefore has no TNQG production dependency.

2.5 PTCP.ai implementation claims and evidence status

The public PTCP.ai site adds an implementation narrative that is not present in the supplied six-page manuscript. It describes a biological 'brain and reflex' pattern: a global density view and predictive control function prepare decisions, while a distributed edge reflex evaluates local probability and performs a bounded hardware action. A public presentation labels this a PTCP fast path and depicts a local hardware processor or ARM service core. The site also states that the overlay works on existing SONiC leaf/spine switches and is hardware-agnostic. [39-41]

Public claim	Evidence class	What it supports	What remains to be demonstrated
Brain/reflex and distributed edge fast path	Vendor architecture presentation	A concrete two-timescale design intent and candidate placement	Executable artifacts, reproducible latency distribution, overload behavior, safety isolation
Microsecond reflex arc / edge evaluation	Vendor-reported demonstration concept	Feasibility hypothesis for compact local scoring and action	Clocked end-to-end measurement, hardware/software bill, p99/p99.9 latency, failure modes
Works on SONiC leaf/spine / existing hardware	Vendor product statement	Candidate integration target and lower transition burden	Named versions, adapters, conformance, switch coverage, rollback and scale results
POL-TT O(d n r^2) storage and O(d r^2) query	Analytic complexity claim	Asymptotic resource rationale under bounded rank	Measured constants, rank growth, accuracy, update cost, adversarial robustness
Wire-speed or guaranteed quarantine	Vendor performance/security claim	A testable performance objective	Independent precision/recall, time-to-detect/contain, false-containment cost, safety proof
Validation protocols are required	Vendor technical overview	Recognition that controlled benchmarking is still necessary	Public datasets, code/configuration, third-party reproduction, operational test authority

Evidence conclusion

PTCP.ai can reasonably be cited as evidence that Tensor Networks, Inc. has articulated an implementable fast/slow architecture and a target deployment on existing network hardware. It cannot be cited, from the public materials reviewed, as independent proof that POL-TT/Dtopo achieve operational detection accuracy, that quarantine is mathematically infallible, or that every action completes at wire speed. This program converts those claims into testable acceptance criteria. [39-42]

3. Adversary and failure model

3.1 Assumed adversary capabilities

- Control one or more routers, gateways, telemetry collectors, service identities, or management accounts, but not every independently administered component at once.
- Advertise or redistribute unauthorized routes, alter path attributes, create plausible trust degradation, replay old controller commands, or suppress selected telemetry.
- Poison online model updates, shape behavior to evade thresholds, induce coordinated low-rate drift, or exploit maintenance and partition events as cover.
- Compromise a PTCP process or its ordinary signing credential and issue syntactically valid but unsafe recommendations.

- Deny, delay, reorder, or partition connectivity between central analytics, local enforcement cells, cryptographic services, and recovery services.
- Exploit software, firmware, model, dependency, or build-pipeline weaknesses; attempt algorithm downgrade or substitution during PQC migration. [29-31]

3.2 Assumptions and trust anchors

No distributed architecture can remain secure if every root of trust, every observer, and every recovery channel is simultaneously controlled. The design assumes at least one independently managed local safety root, one recoverable known-good software/policy image, and a means to authenticate recovery authority. Those assumptions must be made visible in the system security plan and tested. Availability of external time, positioning, enterprise PKI, or wide-area management is not assumed during isolation.

Trust anchor	Administration	What it may authorize	What it may not authorize
PTCP analytics identity	Network analytics operations	Signed evidence and bounded recommendations	Safety-policy change, trust-root creation, crypto downgrade, recovery approval
Local safety root	Platform or site authority	Enforcement within installed mission envelope	Enterprise-wide policy, new cross-domain release, unlimited autonomy
Cryptographic policy authority	Independent key/security management	Approved algorithm profiles, key lifetimes, revocation	Routing optimization or mission prioritization
Recovery authority	Separated incident/recovery administration	Controller revocation, known-good restore, staged rejoin	Self-approval by the suspected controller
Mission authority	Authorized operational chain	Flow priority, minimum service, degradation and human-control rules	Cryptographic implementation exceptions without security approval

3.3 Attack paths and defensive invariants

Attack path	Detection evidence	Invariant
Route leak / hijack	RPKI state, BGP Role/OTC, AS-path and prefix-fan-out change, observer disagreement, Dtopo	Protocol-invalid evidence cannot be normalized away by the model.
Trust sabotage	Attestation, certificate status, posture, logging continuity, clock quality, independent reachability	One trust source cannot quarantine every peer or remove all recovery paths.
POL-TT poisoning	Update-source quality, holdout loss, rank growth, drift residual, replay divergence	Suspicious data may be scored but cannot update the trusted baseline.
Controller capture	Command-rate anomaly, invariant violation, co-signer absence, diverse-controller disagreement	Local safety rejects unsafe commands even when the controller signature is valid.
Recovery impersonation	Separate chain, hardware-backed key, ceremony record, staged attestation	A quarantined controller cannot authorize its own credential restoration.
PQC downgrade	Signed offer transcript, pinned minimum profile, CBOM and module identity	Failure to negotiate does not silently enable a weaker asymmetric algorithm.

4. Solution statement and reference architecture

4.1 Federated assurance cells

The mission fabric is partitioned into independently enforceable cells aligned to platforms, sites, enclaves, or service regions. Each cell contains a local policy decision and enforcement function, a cryptographic broker, a minimal evidence cache, a last-known-good policy set, and an append-only audit buffer. Enterprise PTCP services correlate broad telemetry and publish signed recommendations. Cells accept only recommendations that match local authority, freshness, evidence, and safety constraints. A cell remains capable of deterministic routing, admission, segmentation, and mission-flow prioritization when enterprise analytics are unavailable.

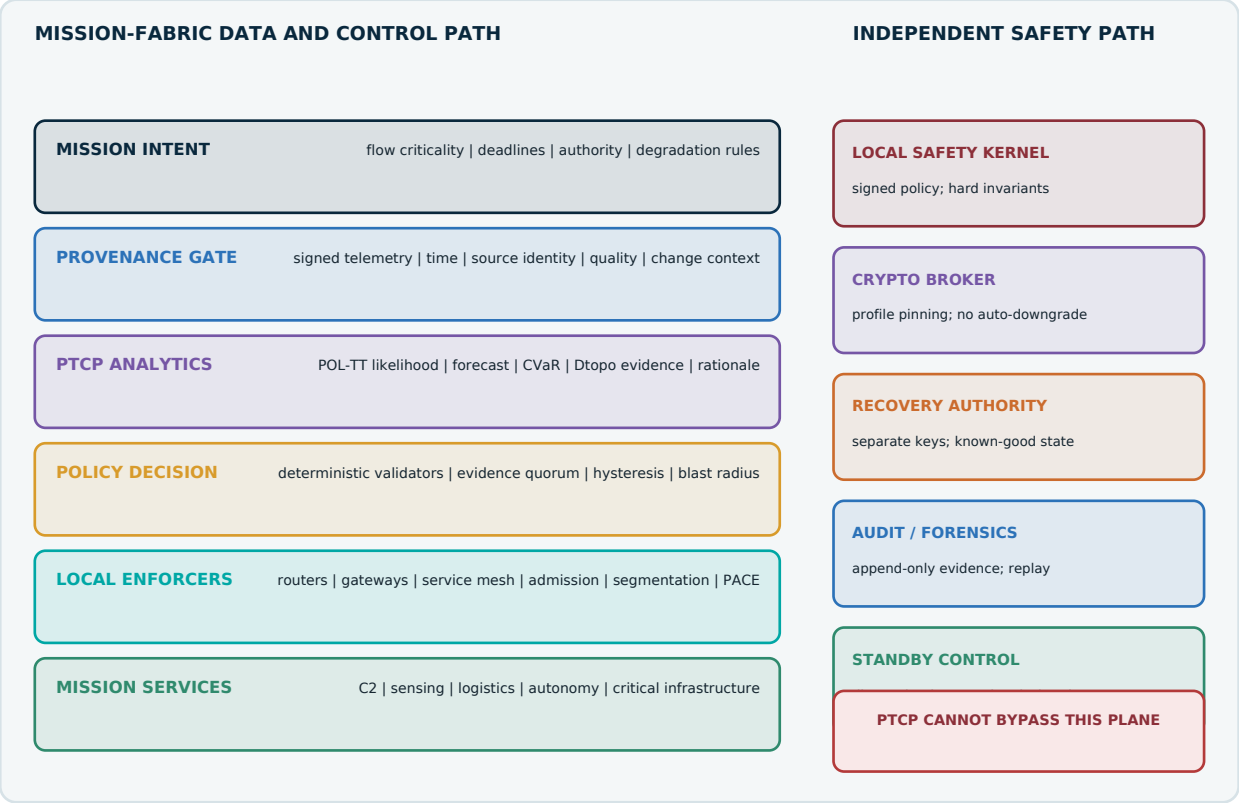

Figure 2. Federated mission-fabric reference architecture. Safety, cryptography, and recovery are independently enforceable and remain available when PTCP is isolated or rejected.

4.2 Planes and responsibilities

Plane	Primary functions	Isolation boundary	Required evidence
Observation	Collect route, link, identity, posture, service, time, and mission-flow telemetry	Read-only from enforcement; no policy write	Source identity, signature/MAC, sequence, timestamp, quality, schema
Inference	POL-TT density, forecasts, CVaR routing, curvature/cut analysis, Dtopo	No direct device actuation	Model/version hash, training window, residuals, uncertainty, rationale
Policy	Combine deterministic validators, anomaly evidence, mission contracts, and operator rules	Separate implementation and signing authority	Decision record, corroboration, blast-radius and availability proof
Enforcement	Install route preference, admission, segmentation, credential, and service controls	Local safety kernel has final authority	Signed scoped action, expiry, monotonic sequence, rollback token
Cryptographic	Profiles, modules, keys, attestation, revocation, downgrade prevention	PTCP cannot change profiles or trust roots	CBOM, policy signature, module identity, key status, transcript
Recovery	Known-good restore, controller revocation, evidence replay, staged rejoin	Separate keys, staff roles, and reachability	Incident authorization, attestation, image/policy hash, reconciliation proof

4.3 Mission assurance contract

Every controlled flow is associated with a machine-readable assurance contract. The contract is the durable authority boundary across connected, degraded, and isolated modes. It defines producer and authorized consumers; classification and releasability; priority; deadline and maximum data age; minimum acceptable service; confidentiality and integrity profile; permitted paths or domains; degraded behavior; local-autonomy permissions; retention; audit; and rejoin rules. PTCP optimizes only inside the feasible set created by these contracts.

Central architectural invariant

Optimization selects among authorized feasible alternatives. It does not redefine what is authorized, create an otherwise forbidden path, weaken a cryptographic profile, or expand the set of consumers. If no feasible alternative remains, the system reports infeasibility and enters the contract's pre-defined degraded or safe state.

4.4 Trust and command flow

- PTCP publishes an **advice envelope**: action proposal, target scope, expected benefit, tail-risk estimate, Dtopo evidence, confidence/coverage, model version, expiry, rollback, and signature.
- The policy plane attaches deterministic route-security results, mission constraints, and any required independent co-authorization.
- The local safety kernel recomputes safety predicates using local state. It may narrow, delay, substitute, or reject the proposal but cannot broaden it.
- The enforcement adapter performs an atomic or staged change, verifies resulting state, and rolls back if the post-condition or minimum-service proof fails.
- The audit plane records proposal, evidence, decision, actuation, observed result, and reversal using monotonic sequence and integrity protection for later reconciliation.

4.5 Fast-path / slow-path implementation architecture

The production design converts PTCP.ai's brain/reflex concept into a strict two-timescale contract. The slow path may be centralized at enterprise scale or federated by region, but it is never synchronous with packet forwarding. It consumes signed telemetry, updates candidate density models, evaluates trusted POL-TT likelihood, constructs forecast scenarios, calculates CVaR path rankings and Dtopo evidence, checks mission policy, and publishes a signed bundle. The fast path runs at a local assurance node or isolated service core adjacent to the router, gateway, DPU, or host. It performs a bounded local density query or threshold lookup, deterministic validation, template matching, safety-envelope projection, transactional enforcement, and post-condition verification. [39-42]

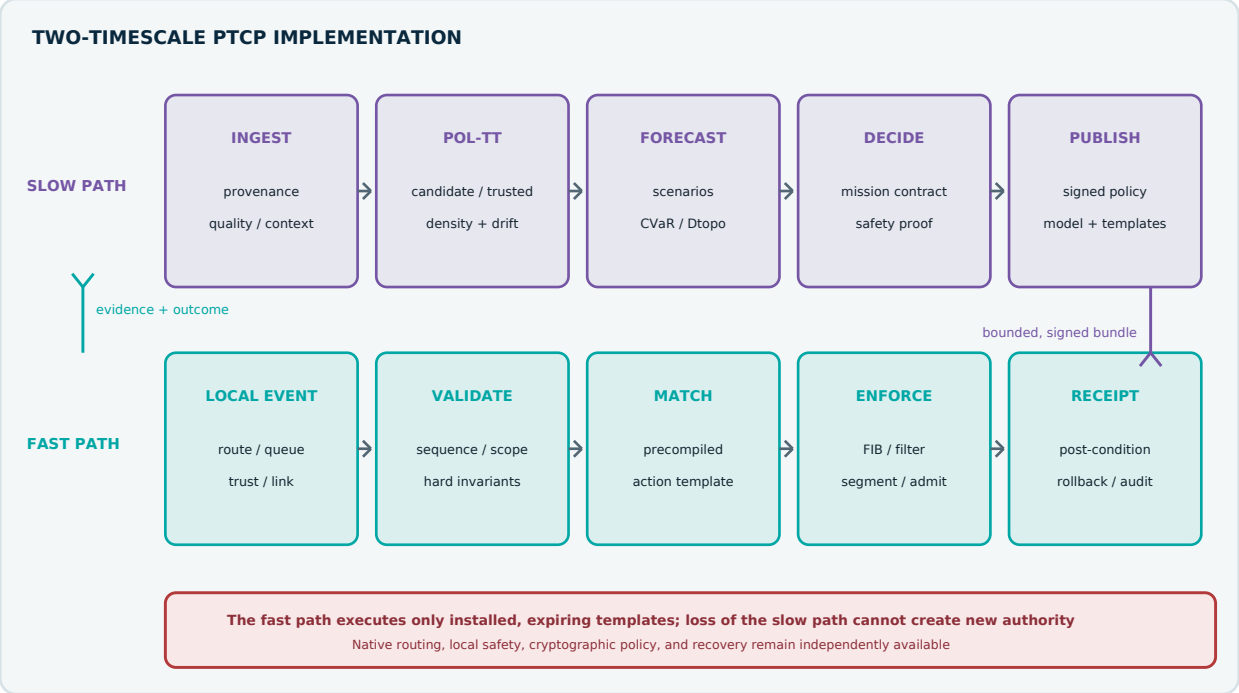

Figure 3. Implementable two-timescale control loop. The slow path creates evidence, forecasts, and bounded policy bundles; the local fast path executes only pre-authorized templates and remains subordinate to independent safety, cryptographic, and recovery authority.

Property	Slow path - predictive / cognitive	Fast path - deterministic / reflex
Placement	Enterprise or regional analytics cluster; diverse standby possible	Assurance node, switch service core, DPU/SmartNIC, gateway, or host agent beside enforcement
Input	Global telemetry, route views, trust/posture, mission context, model history, change records	Local route/link/queue/trust events, installed bundle, local counters, cached validation state
Computation	POL-TT candidate/trusted maintenance, forecasts, CVaR, Dtopo, calibration, champion/challenger tests	Signature/sequence/scope validation, compact density query, threshold state machine, hard safety predicates
Output	Signed PolicyBundle, ModelBundle, action templates, thresholds, expiry, rollback, explanation	Atomic local action or rejection, post-condition, DecisionReceipt, audit and evidence summary
Timing	Application-specific seconds-to-minutes or event-driven cadence; measured, not assumed	Bounded per-device target derived from control deadline; never placed in the packet-forwarding critical path unless hardware proves it
Failure result	Last-known-good bundle remains until explicit expiry; no new authority is created	Native routing and installed local policy remain; unsafe or unverifiable changes are rejected
Update authority	May propose new model/policy candidate; promotion requires governance and evidence	Cannot train trusted baseline, change crypto profile, create trust root, or approve recovery

4.6 Signed control contracts between paths

The interface is artifact-based rather than remote-procedure trust. A **PolicyBundle** binds bundle identifier, issuer and co-authorizer, target cells, valid-from/expiry, mission-contract version, model/feature hashes, deterministic preconditions, evidence thresholds, action templates, cumulative action budgets, minimum-service invariants, approved cryptographic profile, rollback plan, and signatures. A separate **ModelBundle** contains only the compact read-only cores, quantiles/bins, contexts, confidence limits, and health metadata needed locally. The fast path returns an immutable **DecisionReceipt** with observed event, local evidence references, accepted/narrowed/rejected disposition, enforcement transaction, post-condition, rollback, resource timing, and monotonic sequence.

Interface	Allowed direction	Acceptance checks	Security failure behavior
TelemetryEnvelope	Fast/local to slow	Identity, schema, sequence, event/receipt time, clock quality, lineage, classification, quality	Reject or mark unknown; preserve raw record; freeze affected learning

Interface	Allowed direction	Acceptance checks	Security failure behavior
PolicyBundle	Slow/policy authority to fast	Issuer, co-authorization, scope, version, expiry, invariants, crypto profile, action budget, signature	Reject bundle; continue valid last-known-good or native mode
ModelBundle	Model registry to fast	Feature/transform hash, bounded ranks/resources, validation status, target runtime, signature	Do not load; retain prior validated model or deterministic-only operation
DecisionReceipt	Fast to audit/slow	Local sequence, evidence/action hashes, actual state, post-condition, timing, rollback outcome	Buffer locally; never infer success from missing receipt
RecoveryPackage	Independent recovery to cell	Separate chain, hardware-backed key, incident authority, known-good hashes, staged scope	Ordinary PTCP credentials have no permission to invoke it

4.7 Implementation stack and integration sequence

- **Observe:** export RIB/FIB and BGP updates through BMP where applicable; ingest streaming telemetry over authenticated gNMI or equivalent; collect RPKI, BGP Role/OTC, identity, posture, service, time, and change-context evidence. No production writes.
- **Normalize and model:** retain raw evidence, calculate versioned robust quantiles, build context-specific candidate POL-TT models, score trusted models, and run density/curvature/cut ablations against conventional baselines.
- **Compile:** transform an approved high-level recommendation into a finite catalog of device-specific templates such as route reject/de-preference, maximum-prefix guard, next-hop removal, segment admission restriction, traffic drain, or credential scope reduction.
- **Distribute:** sign and publish bundles through a mutually authenticated, rate-limited channel. SONiC leaf/spine is a vendor-stated initial target; adapters can use gNMI/gNOI, NETCONF/RESTCONF, routing-policy APIs, P4Runtime, eBPF, service-mesh APIs, or platform-native interfaces as independently tested. [27,39-41]
- **Enforce:** the local safety kernel validates authority and simulates or stages the proposed delta, proves minimum-service and rollback conditions, commits transactionally, verifies actual RIB/FIB/policy state, and rolls back or substitutes a safer action on failure.
- **Learn:** the slow path consumes receipts and outcomes. It may train a challenger offline or in candidate state, but only an evidence-driven promotion gate can replace the trusted champion. A compromised or isolated controller cannot self-promote.

4.8 Failure ownership and timing budgets

Each deployment defines separate timing budgets for observation freshness, local detection, policy validation, device commit, convergence, post-condition verification, rollback, slow-path recomputation, and operator decision. A single 'wire-speed' claim is insufficient because detection, authorization, and network convergence are distinct intervals. Tail percentiles and deadline-miss probability are reported under load, telemetry loss, route churn, PQC negotiation, controller partition, and rollback. Safety is assessed by the latest time at which a bounded action can still preserve the mission contract, not by average API latency.

5. Evidence collection, provenance, and trust

5.1 Evidence before inference

A density model cannot distinguish a real network from a convincingly forged one unless the evidence pipeline exposes source identity, collection path, freshness, missingness, and conflict. Every telemetry record is therefore wrapped in a provenance envelope before normalization. The envelope carries source and collector identities, hardware or workload attestation state where available, schema version, monotonic sequence, event and receipt time, time-quality estimate, measurement unit and sampling policy, transformation lineage, classification/releasability label, and an integrity value. Batch signatures or message authentication may be used according to the approved cryptographic profile; raw high-rate counters need not each carry a separate public-key signature.

Evidence family	Representative features	Independent corroboration	Primary manipulation risk
Routing state	Adj-RIB-In/Out, selected route, AS path, communities, BGP Role/OTC, withdrawals, prefix fan-out	Multiple BMP collectors, route reflectors, external/coalition view where authorized	Forged or suppressed route view; collector impersonation
Route authorization	RPKI origin state, cache serial, trust anchor, BGPsec/path status when deployed	Diverse validators and cache paths	Stale or poisoned cache; invalid-to-valid rewriting
Link performance	Utilization, queue, delay, jitter, loss, errors, interface state	Peer-side counters, active probes, packet/flow summaries	Counter falsification; selective sampling; probe bias
Trust and identity	Attestation, certificate status, device/workload identity, posture, privilege, policy version	Independent identity and endpoint services; local hardware root	Mass trust suppression; stale revocation; credential theft
Topology and service	Adjacency, reachability, dependency graph, cut capacity, service health	Control-plane databases, service mesh, synthetic transactions	False dependency or artificial bottleneck
Context	Approved change, mission phase, maintenance window, mobility, weather/environment, partition state	Signed change authority and local operator confirmation	Abuse of legitimate-change labels to mask attack

5.2 Routing evidence is deterministic where possible

RFC 7908 defines a route leak as propagation beyond intended scope. Protocol and operational controls already provide evidence that should not be replaced by a learned model: RPKI origin validation verifies whether an origin AS is authorized for a prefix; RPKI-to-router mechanisms distribute validated payloads; BGP Roles and the Only-to-Customer attribute support route-leak prevention and detection; BGPsec can provide signed path authorization where deployed; and BGP operational guidance covers filtering, maximum-prefix limits, session protection, and related controls. BMP provides a standard monitoring feed for route views. [21-28]

> **Precedence rule**
>
> A standards-based invalid or explicit local-policy violation is a distinct evidence class. It may invoke the configured deterministic disposition even if POL-TT considers the event statistically common. Conversely, a route that is origin-valid is not automatically safe: origin validation does not prove intended propagation scope, benign path choice, trustworthy telemetry, or mission suitability.

5.3 Evidence quality and the unknown state

For each monitored region and decision epoch, the policy plane computes a visible evidence-quality vector rather than one hidden confidence number. Components include source provenance, freshness, coverage, time quality, collector diversity, observer agreement, transformation integrity, and context authorization. A summary C_evidence may be calculated for prioritization, but the individual terms remain auditable. Low evidence quality never converts a high anomaly into normality. It limits destructive automation and selects a conservative response: preserve minimum service, freeze learning, increase local observation, require stronger authentication, and request independent confirmation.

$$C_{evidence} = f(C_{provenance}, C_{freshness}, C_{coverage}, C_{time}, C_{diversity}, C_{consensus})$$

Condition	Interpretation	Permitted automatic response
High anomaly, high-quality corroborated evidence	Probable topology defect or policy violation	Scoped constrain or quarantine inside safety envelope; preserve recovery/minimum service
High anomaly, low coverage	Possible attack or observation failure	Freeze learning, raise sampling, constrain new trust/route changes, no broad quarantine
Low anomaly, deterministic invalidity	Known invalid event may be common or model-poisoned	Apply deterministic policy; mark model disagreement for investigation
Low anomaly, severe observer disagreement	Potential suppression, collector compromise, or split view	Enter degraded-evidence mode; favor local known-good policy
Normal evidence after authorized change	Expected transition, still subject to safety invariants	Allow bounded adaptation and continue heightened observation

5.4 Time, ordering, and replay resistance

Topology decisions are time-sensitive, but wide-area time may be degraded or manipulated. Records therefore carry both event time and receipt time, source clock quality, monotonic sequence, and a bounded-age policy. Local enforcers reject expired or repeated commands using monotonic counters and retained sequence windows even when wall-clock confidence is poor. Where a decision depends on cross-source ordering, the system records uncertainty rather than fabricating a precise total order. Reconciliation preserves conflicts for later analysis.

6. POL-TT density modeling and adversarial robustness

6.1 Model objective

POL-TT approximates the joint distribution of selected network and mission-context features. The benefit over independent thresholds is relational sensitivity. A moderate prefix increase, moderate queue growth, and small trust decline may each be individually ordinary yet jointly improbable for the current mission phase. Tensor Train factorization can make this joint density tractable when ranks remain bounded. The model should be trained and evaluated by region, service class, and relevant context so that legitimate mobility or maintenance is not forced into one global stationary distribution. [1,3,4]

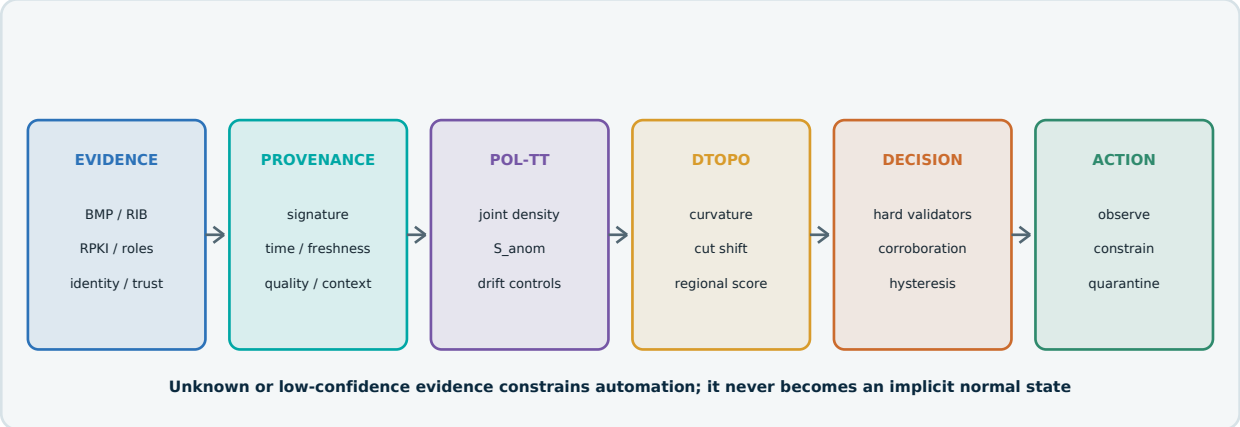

Figure 4. Evidence-to-action pipeline. POL-TT and Dtopo are two stages in a broader assurance chain that preserves deterministic validation, provenance, uncertainty, and bounded enforcement.

6.2 Feature model

Feature group	POL-TT coordinates	Why joint density helps	Guardrail
Route propagation	Prefix count/fan-out, update/withdraw rate, AS-path length/change, Role/OTC result, origin state	Detects relationship and propagation patterns that a local valid route may not reveal	Retain raw validator result; do not encode invalid as merely high cost
Topology geometry	Degree, centrality proxy, curvature, cut conductance/capacity, path diversity	Detects concentration, bridge formation, and redundancy loss	Use stable region definitions and estimator versioning
Performance	Bandwidth, queue, jitter, loss, reliability, energy/cost	Separates expected congestion from unusual topology-performance combinations	Robust scaling; units and sampling policy preserved
Trust	Attestation, identity health, privilege, posture, certificate/revocation, log continuity	Finds correlated trust decline around a path or service	Independent trust sources; missingness is explicit
Context	Mission/service class, maintenance/change, mobility, partition mode, time-quality state	Reduces false anomaly from legitimate regime changes	Context labels are signed and cannot bypass invariants
Model health	Rank, truncation error, held-out loss, update volume, residual drift	Makes model degradation itself observable	Model-health features do not train the same model without separation

6.3 Two-speed learning and baseline promotion

A continuously adapting anomaly model can be trained by the attacker it is meant to detect. The architecture therefore maintains at least two logical baselines. The **trusted slow model** changes only after provenance, holdout, stability, and review criteria are met. A **candidate fast model** may track nonstationarity for comparison but cannot immediately redefine trusted normality or authorize quarantine. Promotion is signed, versioned, reversible, and evaluated against immutable replay sets containing normal operations, authorized changes, partitions, and prior red-team cases.

Control	Mechanism	Failure addressed
Ingestion quarantine	Untrusted or low-quality records may be scored but are excluded from trusted updates	Direct data poisoning
Rank and update budget	Cap TT rank growth, update magnitude, and accepted sample volume per source/region	Resource exhaustion and rapid baseline capture
Holdout gate	Require non-regression on protected normal, failure, and adversarial replay sets	Evasion hidden by aggregate training loss
Residual monitor	Compare candidate likelihood, forecast error, calibration, and downstream decision disagreement	Silent drift and misspecification
Diverse model comparison	Compare TT result with simple robust baselines and deterministic features	Common-mode model error and complexity camouflage
Signed snapshot/rollback	Retain model, quantiles, bins, code, dependencies, and training manifest	Irrecoverable poisoning and unverifiable model state
Promotion separation	Model author cannot unilaterally sign production promotion	Insider or build-pipeline compromise

6.4 Missingness, conflict, and nonstationarity

Missing values are security-relevant observations, not values to be silently imputed. The feature schema distinguishes not observed, stale, not applicable, policy-withheld, collector-failed, and cryptographically invalid. Multiple contextual models or explicit context coordinates handle mission phases and maintenance regimes. A model-change detector may recommend retraining, but an abrupt distribution shift simultaneously raises the possibility of attack; it cannot self-authorize wholesale baseline replacement.

6.5 Calibration and explanation

Raw negative log-likelihood is not comparable across every model, rank, context, or feature set. Each production candidate therefore requires calibrated percentiles or conformal/empirical coverage on held-out data, documented p_min and normalization behavior, and sensitivity analysis. Explanations identify the feature groups and conditional relationships most responsible for low density, the nearest known regimes, missing evidence, and model disagreement. Explanation is an operator aid and test artifact, not causal proof.

7. Dtopo topology security and dynamic detection

7.1 Why topology deformation matters

A route manipulation can be designed to look locally plausible. Its broader graph effect may be less plausible: a new bridge appears, shortest paths concentrate through one region, a previously redundant cut loses capacity, or a set of services acquires correlated trust and reachability changes. Dtopo is intended to summarize this deformation. The density term captures behavioral surprise; the curvature gradient captures local structural discontinuity; and the cut term captures change in regional separation or capacity. [1,6,7]

Dtopo component	Security meaning	Route-leak example	Principal ambiguity
Median regional S_anom	Many features in the region jointly depart from normal behavior	Unusual prefix fan-out plus path, queue, and trust combination	Legitimate surge, mobility, maintenance, model drift
||grad_G kappa||	Local graph geometry changes sharply across adjacent edges or regions	An unexpected transit bridge changes neighborhood structure	Planned reroute, topology restoration, estimator instability
Delta C_cut	Normalized cut capacity/conductance shifts against a robust historical window	Traffic or reachability collapses onto fewer crossing paths	Link failure, scheduled isolation, measurement loss

7.2 Preserve Dtopo; expose the evidence vector

The source equation should remain identifiable for reproducibility. Deterministic route-security results, trust-source conflicts, and evidence confidence should not be silently folded into recalibrated weights and still called Dtopo. The decision service instead evaluates a structured vector E_topo. This preserves the research claim while allowing operationally necessary evidence.

$$E_{topo}(R,t) = \{D_{topo}, V_{origin}, V_{role\_OTC}, V_{path}, V_{local\_policy}, T_{trust}, C_{evidence}, P_{persistence}, I_{mission}\}$$

$$decision_t = g(E_{topo}(R,t), safety\_envelope, mission\_contract, current\_mode)$$

V terms are categorical or scored deterministic-validation results; T_trust summarizes independent identity/posture evidence without replacing its components; C_evidence exposes quality; P_persistence tracks duration and repeated confirmation; and I_mission represents consequence, minimum-service, and recovery constraints. The policy function g is versioned, testable, and substantially simpler than the predictive model.

7.3 Dynamic route-leak detection sequence

Step	Processing	Output / safety rule
1. Acquire	Collect pre- and post-policy route views, validation state, path/role evidence, link/trust telemetry, and change context	Reject unverifiable records from authoritative decisions; retain for forensics
2. Validate	Evaluate origin authorization, BGP Role/OTC, prefix/AS-path filters, max-prefix, and available path-security evidence	Deterministic result remains separate and cannot be overwritten by model score
3. Score density	Query the context-appropriate trusted POL-TT model; compute S_anom and contribution analysis	Freeze model update if evidence or residual health is degraded
4. Define region	Map affected peers, prefixes, services, and dependencies into stable monitored region R	Region expansion is bounded to prevent attacker-driven blast-radius inflation
5. Compute Dtopo	Estimate curvature gradient and normalized cut shift; combine with median regional anomaly	Record estimator version, window, coverage, and uncertainty
6. Corroborate	Compare diverse collectors, local device state, trust services, active probes, and authorized-change records	Conflict invokes unknown/degraded mode, not automatic normality
7. Decide	Apply state-machine thresholds, persistence, mission consequence, minimum service, and co-authorization	High-impact action must pass local safety projection
8. Verify	Observe post-action reachability, route state, mission-flow delivery, and collateral effects	Atomic/staged rollback on failed post-condition

7.4 Trust-degradation detection

Trust is modeled as a set of independently sourced assertions rather than one mutable scalar. Device attestation, workload identity, certificate and revocation state, software/model provenance, privilege, logging continuity, time quality, and network reachability may degrade for different reasons. A coordinated decline along a path can lengthen PTCP's effective geometry and increase Dtopo-related evidence, but the enforcement policy distinguishes cryptographic invalidity, stale evidence, administrative revocation, posture failure, and observation loss. A failure of the enterprise trust service cannot by itself revoke every locally attested identity or strand the recovery plane.

7.5 Thresholding, hysteresis, and rate control

Thresholds are calibrated per context and consequence class, not published as universal constants. Escalation uses an upper threshold, minimum persistence, evidence-quality condition, and action-rate budget. De-escalation uses a lower threshold, successful attestation, stable post-condition, and dwell time. This hysteresis reduces route and quarantine oscillation. The system rate-limits both recommendations and enforcement changes, and it aggregates correlated alerts so an attacker cannot exhaust operator attention by fragmenting one event.

8. Quarantine decision logic and bounded response

8.1 Reversible state machine

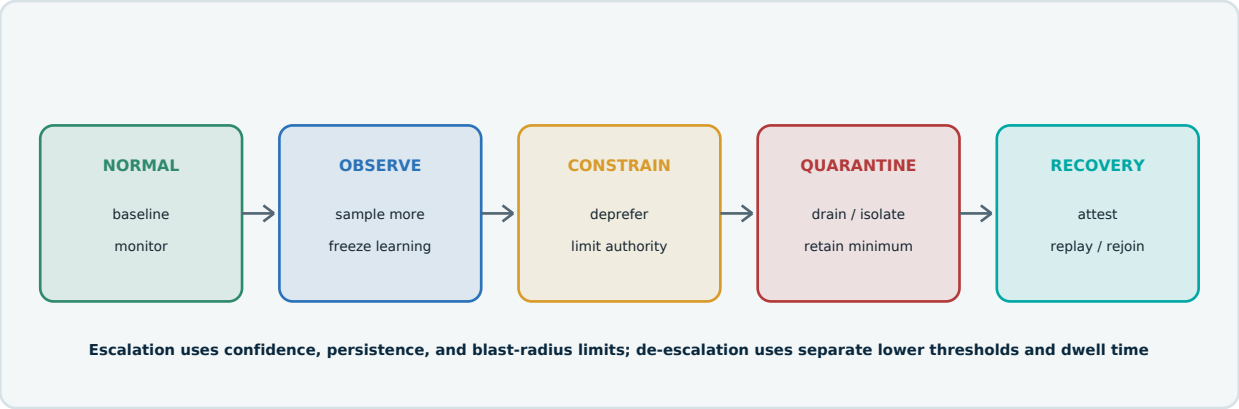

Figure 5. Quarantine is an evidence-gated state machine. Progression is reversible, rate-limited, and subject to minimum-service and recovery-path constraints.

State	Entry condition	Allowed action	Exit evidence
Normal	Expected state; deterministic controls satisfied	Ordinary routing and observation	Anomaly, invalidity, conflict, or evidence degradation
Observe	Unusual behavior or degraded evidence without sufficient containment confidence	Increase sampling, freeze learning, require stronger authentication, preserve current paths	Resolution to normal or corroborated persistence
Constrain	Persistent/corroborated suspicion or policy concern	De-preference, admission limits, restrict new sessions/privilege, pin known-good paths	Clear evidence and dwell time, or escalation proof
Quarantine	Deterministic hard failure or high-confidence corroborated defect with survivability proof	Scoped traffic drain, segmentation, route rejection, credential restriction	Attestation, remediation, replay review, stable minimum service
Recovery	Remediation authorized through separate recovery plane	Known-good restore, staged connectivity, shadow observation, limited privilege	Rejoin criteria and post-condition hold for dwell period

8.2 Safety-envelope predicates

Before actuation, each local cell evaluates predicates using locally available state. The recommendation is rejected or narrowed if any hard predicate fails. Soft predicates may require operator or independent co-authority approval. The envelope is signed by the appropriate safety authority, versioned, cached locally, and cannot be modified by ordinary PTCP administration.

- The target is within the controller's authorized region, resource class, and action type.
- The recommendation and all required evidence are fresh, correctly sequenced, and bound to the current policy/model versions.
- At least the configured minimum independent paths, recovery channels, and essential services remain available after the proposed action.
- The action does not create a new cross-domain release, consumer, trust root, algorithm profile, credential scope, or mission authority.
- The cumulative number, frequency, and blast radius of active changes remain within the local budget.
- A rollback or safe substitute exists and its trigger can be observed locally.

- Actions with designated human or dual-control thresholds carry valid, independently administered authorization.

8.3 Graduated response catalog

Response	Effect	Appropriate evidence	Fail-safe constraint
Raise observation	Higher sampling, additional collectors/probes, preserve raw evidence	Anomaly or evidence conflict	Must not overload control plane or expose sensitive telemetry
Increase authentication	Shorter sessions, re-attestation, reduced privilege	Trust staleness or moderate decline	Preserve recovery and essential local identities
Route de-preference	Keep path available but select alternatives first	Moderate Dtopo or uncorroborated leak concern	No path churn beyond budget; minimum diversity retained
Admission/throttle	Limit new low-priority flows or affected source	Congestion plus suspicious propagation	Priority/minimum-service contracts remain satisfied
Traffic drain	Move existing flows before isolation	High-confidence defect with viable alternate	Verify target path and data-age consequences before cutover
Segment quarantine	Block or isolate scoped region/service/identity	Deterministic invalidity or corroborated severe defect	Recovery path, operator access, and forensics remain reachable
Controller rejection	Ignore PTCP advice identity or model version	Compromise evidence or repeated invariant violation	Local known-good mode and separate recovery authority activate

8.4 Avoiding false-containment cascades

The policy engine evaluates counterfactual reachability before broad changes: what essential flows, key services, identity dependencies, monitoring paths, and recovery channels would remain if the action succeeds exactly as proposed? It also considers correlated recommendations from the same evidence source as one blast-radius event. Quarantine cannot recursively consume its own evidence channel without a designated out-of-band alternative. If predicted collateral effects exceed policy, the system selects a narrower constrain state or requests human/independent review.

9. Compromised-controller containment and recovery

9.1 Why command signatures are insufficient

A valid signature authenticates the credential used to produce a command. If the controller or its key is compromised, unsafe commands can still be correctly signed. The receiver must therefore verify semantic authorization and safety independently. Each advice envelope contains action scope, evidence references, expected post-condition, risk, expiry, monotonic sequence, model/code identity, policy version, and rollback. High-impact recommendations additionally require a separate policy or mission co-authorization whose key is not available to the PTCP runtime.

Compromise indicator	Local reaction	Enterprise reaction
Invariant-violating signed command	Reject, retain evidence, enter M1 degraded or M3 compromise mode	Alert independent recovery/security authority; compare other cells
Excessive action rate or scope	Apply rate/blast-radius budget; reject excess	Suspend controller action profile pending review
Model/version not approved	Reject advice; continue last-known-good policy	Revoke artifact or distribution channel; verify supply chain
Diverse controller disagreement	Use deterministic/local policy; no tie broken by privilege alone	Investigate evidence/model commonality; promote no controller automatically
Controller credential revoked	Reject new messages; preserve audit and bounded local operation	Activate separately keyed standby/recovery process
Controller unreachable	Expire advice by TTL; transition M1 to M2 by contract	Restore reachability without overriding local safety state

9.2 Independent controller health and diversity

Controller health is measured externally through signed heartbeats, attested workload identity, output-invariant monitoring, model and dependency hashes, command-rate and target-distribution statistics, and comparison with a simpler reference policy. A standby controller should differ in failure domain and, where feasible, implementation or model family. Diversity is not assumed to provide Byzantine consensus; it is an additional source of disagreement evidence. Automatic leadership transfer occurs only under a signed recovery policy and cannot override local safety rejection.

9.3 Revocation and recovery sequence

Stage	Required action	Safety condition
Contain	Local cells reject suspected controller identity/model version; freeze nonessential automation	Known-good local policy and audit remain available
Preserve	Snapshot commands, evidence, models, logs, sequences, and enforcement state	Forensics cannot block minimum service
Authorize	Separate recovery authority validates incident status and signs revocation/restore plan	Suspected controller cannot co-sign its own reinstatement
Restore	Boot attested known-good controller/policy/model and refresh cryptographic material as required	Artifact, configuration, and dependency hashes match approved manifest
Shadow	Reconstructed controller observes and recommends without actuation	Outputs are compared with local and reference policies
Reconcile	Resolve divergent local logs/state with explicit conflicts and signed replay	No blind last-writer-wins for trust, release, or safety state
Re-enable	Restore authority gradually by action class and region	Dwell period, post-condition metrics, and independent approval satisfied

9.4 Recovery-plane independence

Recovery must not share every dependency with the primary controller. At minimum, recovery credentials, policy store, known-good images, and a management reachability option are separately administered. The recovery plane is narrow: it can revoke controller identities, restore approved artifacts, initiate staged rejoin, and obtain audit evidence. It cannot become an unrestricted alternate optimizer or persistent bypass around normal policy. Its use is logged and subject to role separation.

10. Bounded local autonomy during isolation

10.1 Operating modes

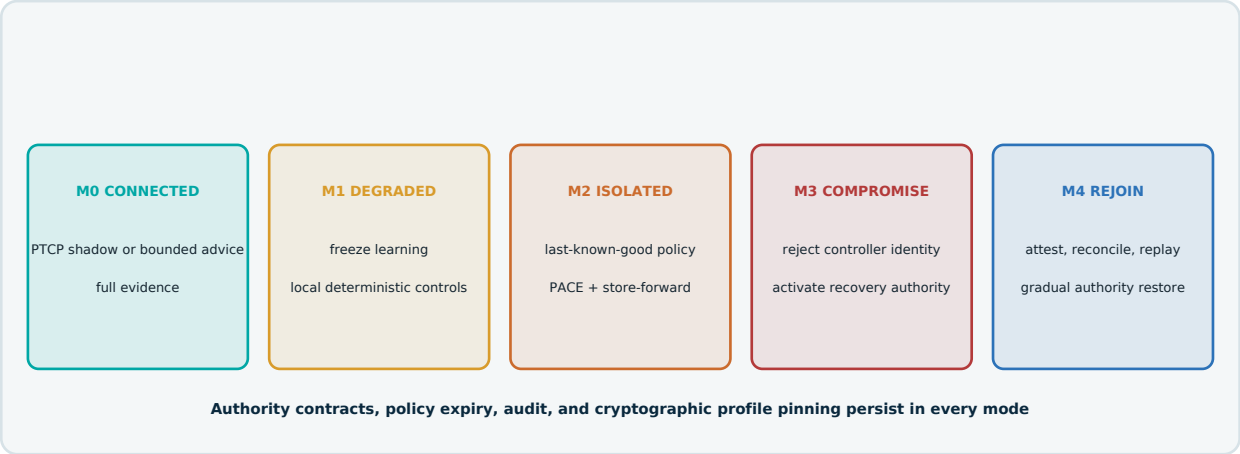

Figure 6. Controller connectivity and trust modes. Isolation and suspected compromise are distinct: compromise revokes controller authority, while isolation expires it and continues within local contracts.

Mode	Trigger	PTCP authority	Local behavior	Exit
M0 connected	Healthy evidence and controller	Shadow or bounded action classes	Normal deterministic protocols plus safety projection	Evidence/controller degradation or incident
M1 degraded	Low coverage, conflict, stale trust, partial loss	Advice may be accepted only for non-destructive classes or disabled	Freeze learning, tighten change budget, favor known-good paths	Evidence recovery or isolation threshold
M2 isolated	Controller/enterprise unreachable beyond contract timer	None after advice expiry	Local routing, PACE, priority, caching/store-forward, approved sessions	Authenticated reachability and staged reconciliation
M3 compromise	Controller identity/model/output deemed unsafe	Explicitly rejected	Known-good policy, controller revocation, recovery channel only	Separate recovery authority and shadow validation
M4 rejoin	Connectivity/authority restored but state diverged	Shadow, then staged action classes	Attest, reconcile, replay, resolve conflicts, restore budgets gradually	Stable post-condition and approval

10.2 Local authority contract

A local autonomy contract is signed and installed before disconnection. It is not improvised after the controller disappears. The contract specifies a finite action vocabulary and quantitative budgets appropriate to the platform or site. Budgets are policy parameters, not universal values, and are tested against mission threads and failure cases.

Bound	Contract content	Fail-safe behavior when exhausted
Time	Maximum disconnected duration by function, credential, policy, and data class	Expire affected privilege or enter designated safe/degraded state
Scope	Permitted devices, services, prefixes, paths, tenants, and action types	Reject out-of-scope action; do not infer expansion
Change rate	Route/policy/admission changes per interval; churn and retry limit	Hold last safe state and request operator intervention
Blast radius	Maximum concurrent affected flows, nodes, services, or regions	Select narrower constraint or no action
Minimum service	Essential flows, recovery, operator access, identity/key continuity, and required path diversity	Preserve minimum set even when anomaly remains
Trust	Which local attestations/credentials remain acceptable and for how long	Deny new privilege; preserve explicitly approved essential sessions
Data	Maximum data age, release policy, retention, store-forward and conflict rules	Mark stale, withhold, or degrade according to mission contract
Crypto	Pinned algorithm/profile/module set and authorized emergency exceptions	Fail closed or use pre-approved limited fallback; never silent downgrade

10.3 Local functions that continue

- Standards-based local routing and forwarding, including installed filters and deterministic invalid-route disposition.
- Priority and admission enforcement for mission-flow contracts using local measurements and conservative capacity assumptions.
- PACE path selection among pre-authorized primary, alternate, contingency, and emergency transports; no creation of a new unauthorized transport relationship.
- Local identity and service authorization using cached, short-lived, explicitly scoped credentials or locally anchored attestations as approved by policy.
- Caching, replication, and store-and-forward with data-age, integrity, classification, and conflict metadata.
- Local anomaly observation and evidence preservation; POL-TT scoring may continue on a trusted snapshot, but online promotion and broad quarantine remain disabled when evidence is inadequate.

- Operator visibility, manual override within role, rollback, and an out-of-band or alternate recovery channel where the platform supports one.

10.4 Functions that do not expand locally

Isolation does not authorize a cell to create new trust roots, accept an unapproved controller, lower classification/releasability barriers, extend credentials indefinitely, weaken public-key policy, install unsigned software or models, broaden quarantine beyond its region, or assume a higher command/release authority. Where an essential function cannot continue inside these constraints, the contract defines a safe state and makes the loss visible rather than hiding it behind unauthorized automation.

11. Classical PQC cryptographic agility

11.1 Standards baseline

Post-quantum cryptography uses classical hardware and software. NIST finalized ML-KEM in FIPS 203 for key establishment, ML-DSA in FIPS 204, and SLH-DSA in FIPS 205 for signatures. NIST SP 800-227 provides final KEM guidance. The June 2026 update to NIST CSWP 39 defines crypto agility as the capability to replace and adapt algorithms across protocols, applications, software, hardware, firmware, and infrastructure while preserving security and operations. Executive Order 14412 directs accelerated federal PQC migration and cryptographic inventory activities; National Security Systems follow applicable NSA/CNSS policy. [12-20]

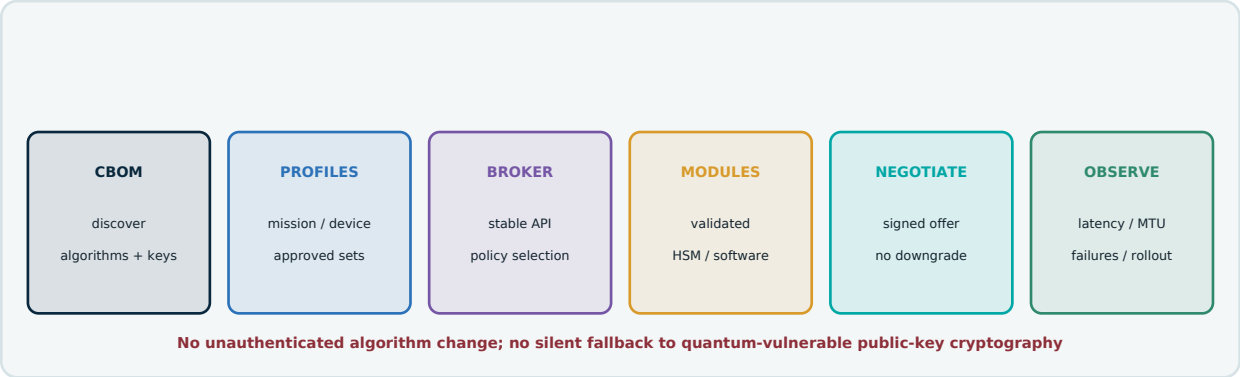

Figure 7. Cryptographic-agility chain. PTCP may observe cryptographic health and cost, but algorithm profiles, trust roots, and downgrade policy remain under independent cryptographic authority.

11.2 Cryptographic architecture

Component	Function	Red-team-resilient property
Cryptographic bill of materials (CBOM)	Inventory algorithms, protocols, certificates, libraries, modules, keys, firmware, data sensitivity/lifetime, and owners	Makes hidden quantum-vulnerable and non-agile dependencies discoverable
Profile registry	Maps mission/device/protocol contexts to approved algorithm sets, parameters, lifetimes, and transitions	PTCP cannot invent or lower a profile; versions are signed and pinned
Crypto broker/API	Stable application interface to key establishment, signatures, encryption, MAC, hashing, and attestation	Algorithms can change without application rewrites; policy remains centralized but locally cached
Validated modules	Approved software/hardware implementations, entropy, key protection, self-tests	Module identity and validation status are measured; failure is explicit
Key management	Generation, distribution, storage, activation, rotation, revocation, destruction, and recovery	Separate administrative roles and local continuity prevent PTCP key capture from becoming universal
Negotiation guard	Authenticated, transcript-bound selection inside pinned minimum policy	Rejects stripping, substitution, and silent downgrade
Telemetry and audit	Records profile, module, handshake outcome, latency, bytes, failures, and downgrade attempts	Supports migration evidence without exposing secret material

11.3 Profile strategy

The program defines profiles by mission consequence and platform capacity rather than enabling every algorithm everywhere. ML-KEM is the NIST-standardized general KEM baseline; ML-DSA and SLH-DSA provide different signature properties and performance tradeoffs. Exact parameter sets, certificate formats, protocols, and CNSA applicability are selected by the responsible authority. Hybrid classical/PQC modes may be used only where an approved profile and protocol specify combination, transcript binding, failure handling, and retirement. More algorithms do not automatically create more security.

Artifact / channel	PQC use	Fallback principle
PTCP advice and policy packages	PQC-capable signatures; model/code/policy hash bound to authorization	Reject unverifiable package; continue last-known-good local policy
Telemetry batches	PQC-capable source/collector authentication where public-key signatures are required; efficient symmetric integrity after approved establishment	Buffer, mark unverifiable, or use locally established approved integrity; do not treat unauthenticated data as trusted
Controller-enforcer sessions	Approved KEM/session protocol and authenticated transcript	Preserve existing approved session only within lifetime; otherwise enter local mode
Software, firmware, and models	PQC-capable release signatures and provenance chain	Boot/install known-good signed artifact; no unsigned emergency update
Root and recovery artifacts	Long-lived, conservatively protected signatures; offline or hardware-backed administration	Use pre-positioned approved recovery material and ceremony; never self-signed by suspect controller
Stored mission data	Quantum-resistant key protection based on sensitivity and confidentiality life	Rewrap/re-encrypt by migration policy; preserve integrity and version history

11.4 Isolation and cryptographic fallback

Failure	Permitted continuity	Prohibited behavior
Enterprise key service unreachable	Use locally cached active keys/credentials only within signed scope and lifetime; local hardware-backed operations continue	Indefinite credential extension or creation of an unapproved local CA
PQC module unavailable	Use a separately validated approved module/profile if installed; preserve last approved sessions when policy permits	Automatic downgrade to quantum-vulnerable public-key establishment
PQC negotiation failure	Fail closed for new protected association or enter a pre-approved limited mission mode	Retry into weaker unauthenticated algorithm without explicit policy
Revocation service unreachable	Apply cached status, short lifetime, local deny lists, and risk-specific limits	Assume all stale credentials are valid or revoke every identity blindly
Controller key suspected	Reject controller advice key; local and recovery keys remain separate	Allow controller to sign its own replacement or recovery authorization
Clock quality degraded	Use monotonic sequence, bounded receipt age, local secure counters, and conservative expiry	Disable expiry/replay checks because wall-clock time is uncertain

11.5 Performance and interoperability evidence

PQC migration changes key, ciphertext, and signature sizes as well as compute, memory, power, handshake, certificate, and fragmentation behavior. Each device/protocol profile is tested under realistic loss, latency, MTU, concurrency, boot, rekey, recovery, and degraded-time conditions. Measurements include p50/p95/p99 establishment latency, CPU/accelerator utilization, memory, energy, message expansion, fragmentation/reassembly, failure rates, key rotation time, and impact on mission-flow deadlines. Cryptographic cost may be exposed to PTCP as a read-only telemetry feature, but PTCP cannot lower the profile to improve routing cost.

12. Hardware architecture and deployment scope

12.1 Hardware reference scope

Hardware element	Minimum role	Security and resilience properties	Sizing evidence
Enterprise analytics cluster	POL-TT training/query, forecasts, curvature/cut analysis, replay	Isolated control and data interfaces; measured boot; workload identity; resource quotas; no direct unmediated device write	Feature rate, dimensions/bins/rank, region count, scenario horizon/samples, replay window
Local assurance node	Evidence cache, deterministic validators, safety kernel, local policy, audit buffer	Rugged/appropriate environment; secure boot; hardware root; protected monotonic state; fail-safe I/O	Local flow/device count, isolation duration, change rate, log retention, power budget
Routers/switches/gateways	Existing routing/forwarding, filters, segment/admission enforcement, telemetry	Control-plane policing; signed configuration channel; atomic/staged commit and rollback; independent local policy	RIB/FIB scale, update rate, policy objects, telemetry export, convergence targets
Security module / HSM	Key generation/storage/use, module identity, attestation, recovery key protection	Applicable validated implementation; tamper response as required; role separation; audit	Concurrent handshakes/signatures, key count, latency, endurance, environmental constraints
DPU/SmartNIC or accelerator (optional)	Telemetry preprocessing, flow observation, cryptographic or policy offload	Signed firmware; isolated DMA/IOMMU; fail-open/closed defined per service; software fallback	Packet/flow rate, message size, algorithm workload, power/thermal envelope
Independent recovery controller	Known-good restore, revocation, staged rejoin, evidence retrieval	Separate credentials, administration, and reachability; normally quiescent; narrowly scoped commands	Number of cells, recovery RTO/RPO, image/policy volume, out-of-band bandwidth
Evidence store	Immutable or append-only raw/derived evidence, model and decision lineage	Integrity protection, access separation, retention tiers, write-once or object-lock where appropriate	Event rate, retention, compression, classification partitions, replay throughput
Time and sequence source	Clock quality, secure counters, event ordering support	Diverse time inputs where available; spoof/jam detection; local monotonic fallback	Holdover duration, drift tolerance, timestamp precision, sequence capacity
Lab and emulation range	Digital twin, route/telemetry injection, PQC and failure testing	Separated from production; reproducible traffic; evidence capture; safe synthetic adversary scenarios	Topology size, link emulation, telemetry rate, hardware-in-loop, concurrent experiments

12.2 Enforcement-cell pattern

A local assurance node may be integrated with a router, gateway, platform computer, or site security appliance, but the logical safety boundary remains explicit. The node stores the current mission contract, local autonomy contract, cryptographic profile, deterministic routing policy, last-known-good configuration, authorized controller identities, action budget, and rollback state. If the analytics process shares hardware, mandatory isolation prevents it from modifying these artifacts or enforcement interfaces outside the safety API.

12.3 Hardware roots and boot chain

- Measured and secure boot cover firmware, bootloader, operating system/hypervisor, enforcement agent, policy bundle, crypto broker, and model runtime as applicable.
- Remote or local attestation reports component identity and version without treating attestation alone as authorization.
- Monotonic counters or protected sequence state support replay resistance when external time is uncertain.
- Recovery and ordinary controller keys reside in separate logical or physical security domains; compromise of one does not enable the other.

- Hardware replacement, depot maintenance, and constrained-platform provisioning have documented re-key, re-attestation, and evidence-preservation procedures.

12.4 Availability and environmental design

Mission-fabric hardware is sized for tail demand and degraded operation, not only average telemetry. Essential local functions retain reserved compute, memory, storage, and power so a telemetry flood or model workload cannot starve routing, safety, key use, or recovery. Thermal, power, vibration, radiation, electromagnetic, and maintenance constraints are platform-specific verification inputs. Loss of an optional accelerator invokes a measured software or reduced-function mode; it cannot silently bypass policy or cryptographic checks.

13. Software, firmware, data, and interface scope

13.1 Software components

Component	Responsibilities	Required defensive behavior
Telemetry adapters	BMP, routing, streaming telemetry, identity, posture, service, and mission-context ingestion	Mutual authentication where supported; schema validation; backpressure; source quotas; raw evidence preservation
Provenance and quality service	Envelope verification, time/sequence, lineage, missingness, conflict, context authorization	Reject or downgrade unverifiable evidence; never silently repair security-significant fields
Feature/normalization service	Robust quantiles, dimensionless mapping, binning, region/context assignment	Version every transform; prevent current suspicious stream from instantly redefining quantiles
POL-TT runtime	Density query, candidate training, rank control, model health, explanations	Trusted/candidate separation; signed snapshots; update freeze; bounded resources; deterministic replay mode
Forecast and risk engine	Transition scenarios, candidate paths, expected cost and CVaR	Authorized candidate set only; route-churn penalty; infeasibility visible; no direct enforcement
Dtopo engine	Curvature, cut baseline, regional score, estimator uncertainty	Stable region/version metadata; coverage checks; robust windows; component-level explanation
Deterministic validation service	RPKI, Role/OTC, prefix/path/local policy, identity and artifact validity	Separate result channel and precedence; independent data sources where required
Policy decision service	Evidence vector, state machine, mission contract, co-authorization, action proposal	Simple/versioned rules; hysteresis; action budgets; survivability and rollback proof
Local safety kernel	Scope, authority, sequence, expiry, minimum service, blast radius, crypto profile	Small trusted code base; fail-safe defaults; cannot be updated by ordinary PTCP credential
Enforcement adapters	Translate authorized action to router, SDN, service mesh, gateway, admission, and identity controls	Transactional/staged change, post-condition verification, rollback, idempotence
Crypto broker and key client	Profile selection, module dispatch, negotiation guard, key lifecycle, telemetry	Stable API; signed profiles; no silent downgrade; secrets excluded from analytics/audit
Recovery manager	Revocation, known-good restore, shadow re-entry, reconciliation	Separate authority and path; narrow interface; conflict-preserving replay
Audit/evidence ledger	Advice, evidence, decisions, actions, outcomes, model/code/config lineage	Append-only integrity, access separation, retention, redaction, cross-domain handling
Simulation and test harness	Replay, topology emulation, fault injection, synthetic adversary scenarios, metric collection	Deterministic seeds/manifests; production isolation; reproducible expected outcomes

13.2 Canonical data objects

Object	Required fields	Security property
Telemetry envelope	Source/collector identity, schema, sequence, event/receipt time, clock quality, value/unit, sampling, missingness, lineage, labels, integrity	Evidence cannot be detached from provenance or context
Model manifest	Model type, TT ranks/cores hash, quantiles/bins, features, contexts, training window, code/dependencies, metrics, approvers	Reproducible trusted baseline and rollback
Dtopo record	Region, density statistics, curvature estimator/result, cut window/result, weights, coverage, uncertainty, source references	Score is decomposable and comparable across versions
Advice envelope	Proposal, scope, expected outcome, tail risk, evidence, confidence, model/policy version, expiry, rollback, signature	Recommendation is bounded, attributable, and replay-resistant
Decision record	Deterministic results, state transition, mission/safety predicates, co-authorizations, accepted/narrowed/rejected action	Local authority and rationale are auditable
Crypto profile	Algorithms/parameters, protocols, modules, key/cert lifetimes, negotiation, fallback, retirement, authority signature	No algorithm ambiguity or controller-defined downgrade
Reconciliation package	Local state/log range, hashes, conflicts, proposed replay, trust and policy deltas, recovery authorization	Rejoin preserves conflict and prevents unaudited overwrite

13.3 Interface principles

- Advice and enforcement APIs are capability-scoped, versioned, mutually authenticated, rate-limited, and idempotent.
- Read telemetry, recommend action, authorize action, execute action, change crypto profile, and recover controller are distinct permissions and interfaces.
- Schemas distinguish absent, stale, invalid, withheld, and not-applicable values; units and normalization versions are mandatory.
- All actions include pre-conditions, post-conditions, expiry, rollback, and a correlation identifier that links evidence to outcome.
- Adapters expose actual resulting state rather than merely reporting that an API request succeeded.
- Cross-domain and coalition interfaces minimize and label evidence; sensitive topology and trust details are released only by applicable policy.

13.4 Secure development and supply chain

The software and model lifecycle follows a secure development framework, threat modeling, dependency and artifact provenance, protected build/signing, reproducible or independently verified releases, vulnerability response, and supplier risk management. Models, quantile state, region definitions, policy, and cryptographic profiles are treated as security-relevant artifacts, not ordinary configuration. NIST adversarial-ML guidance informs poisoning/evasion tests; NIST SSDF and supply-chain guidance inform development and acquisition controls. [29-31]

14. Mission and critical-infrastructure applications

14.1 Contested joint and distributed command-and-control

A distributed C2 fabric spans heterogeneous terrestrial, tactical, commercial, coalition, and space-dependent transports. POL-TT can model joint behavior across route propagation, link quality, service dependencies, trust, and mission-flow age. Dtopo can highlight structural concentration or cut degradation that precedes outright outage. Local assurance cells preserve priority, deterministic filtering, PACE, data-age limits, and signed store-forward during partition. The expected benefit is improved detection and bounded continuity, not elimination of propagation delay, jamming, physical loss, or command judgment.

14.2 Autonomous and remotely operated platforms

Autonomous platforms need communications behavior that is explicit before loss of link. The mission contract defines which sensing, navigation-support, coordination, return, hold, or safe-state functions continue locally; PTCP optimizes network resources only inside those bounds. A platform rejects controller advice that would remove required command/recovery reachability or weaken identity/crypto. Rejoin requires state and authority reconciliation. The networking controller is not a weapons-release or target-selection authority.

14.3 Space-terrestrial and intermittently connected systems

Space and remote systems exhibit planned contact windows, high delay variance, constrained power, and long software/crypto lifecycles. Context-partitioned POL-TT avoids treating every planned disconnect as an attack while still detecting unexpected reachability, path, trust, or cut changes. Pre-positioned policy, crypto profiles, cached validation data, local sequence protection, store-forward, and bounded autonomy are essential. PQC overhead must be measured against link, MTU, compute, power, and certificate constraints rather than assumed acceptable.

14.4 Logistics, ports, bases, and industrial control

Mission networks depend on energy, fuel, transport, maintenance, communications, building systems, and operational technology. A PTCP-informed fabric can correlate enterprise routes with service and OT dependencies, but safety controllers and protective functions remain outside predictive actuation. For OT, observation is often safer than control; quarantine may mean broker isolation, unidirectional data flow, maintenance access restriction, or IT-side traffic drain rather than changing a protective relay or process controller. NIST cyber-resilience, zero-trust, and OT guidance support segmentation, least privilege, recovery, and safe engineering. [8-11,33-35]

14.5 Coalition and cross-domain exchange

Coalition and cross-domain fabrics combine different trust anchors, release rules, cryptographic timelines, and visibility. Advice is therefore computed over the evidence legitimately available to the decision domain; absence caused by policy is labeled rather than treated as sensor failure. Quarantine cannot expand data release or expose sensitive topology to obtain better model accuracy. Federated cells exchange signed, minimized security labels and route health rather than unrestricted raw telemetry.

14.6 Civil critical infrastructure and emergency continuity

Energy, communications, water, transportation, healthcare, and emergency services can use the same architecture pattern without adopting military command structures: independent local safety, topology anomaly evidence, standards-based routing controls, PQC inventory/agility, bounded site autonomy, and separately authorized recovery. Sector-specific safety, regulation, privacy, and availability requirements govern the local contract. The architecture must not centralize sensitive industrial telemetry beyond need.

14.7 Military deployment blueprint

A military deployment is a federation of assurance cells, not one global controller. Enterprise instances provide model governance, software/model registry, cryptographic-policy distribution, and broad replay. Geographic or functional theater instances provide the regional slow path and a diverse standby. Mission-partner environments, bases, ships, aircraft, ground formations, space gateways, and deployed data centers host local fast paths next to their existing routers or gateways. Each cell receives only the mission contracts, route views, model subset, action templates, and cryptographic profile authorized for that domain. Cross-domain guards and coalition release controls remain external decision points.

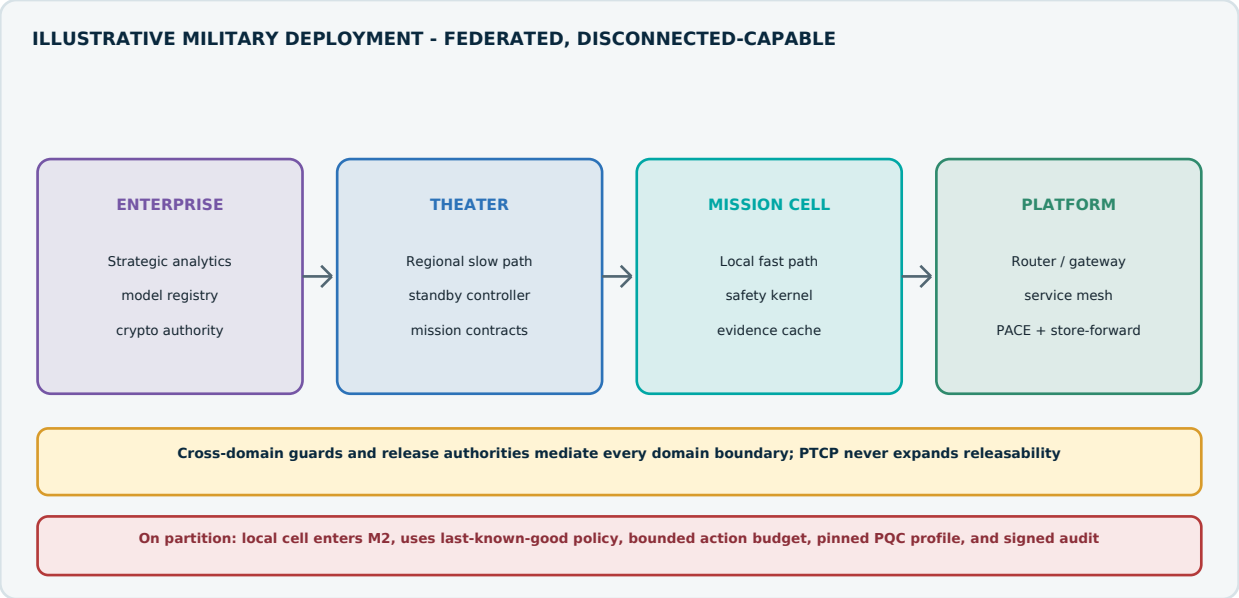

Figure 8. Military deployment pattern. Global awareness is useful but not required for local safety: every mission cell can reject the controller, continue native routing and essential services, and reconcile later through separate recovery authority.

Deployment tier	Installed components	Normal operation	Partition or compromise behavior
Enterprise / strategic	Slow-path analytics cluster; immutable evidence lake; trusted/candidate registries; policy and crypto authorities; recovery coordination	Cross-theater modeling, replay, software/model promotion, profile lifecycle, strategic trend evidence	No direct packet dependency; compromise causes identity/model revocation and standby activation
Theater / regional	Regional slow path; route collectors; regional POL-TT/Dtopo contexts; policy compiler; diverse standby	Correlates route, transport, trust, service, and mission-flow evidence; publishes scoped bundles	Freezes learning, uses last trusted model, or loses authority at expiry; subordinate cells continue locally
Mission enclave / command node	Local assurance node; fast-path runtime; safety kernel; crypto broker; evidence/audit cache; enforcement adapters	Validates bundles, executes reversible templates, protects minimum C2/recovery paths, sends receipts	M2 isolation or M3 compromise; local contract, PACE, pinned PQC, action budget, no authority expansion
Platform / tactical edge	Compact agent/service core; native router policy; read-only model subset; local keys; store-forward	Local anomaly checks and deterministic response within compute/power/link constraints	Falls back to native known-good policy and platform safe state; preserves command/recovery reachability
Coalition / cross-domain gateway	Domain-specific fast path; release-minimized evidence adapter; guard/authorization integration	Shares signed health summaries and authorized route evidence; applies each domain's policy	No inference may bypass guard or broaden release; domain remains independently operable
Recovery enclave	Quiescent known-good controller, images/bundles, separate keys, out-of-band path and procedures	Exercises restore and evidence retrieval; does not optimize traffic	Revokes suspect controller/model, restores known-good state, conducts staged shadow rejoin

Military fielding sequence

Increment	Deployment activity	Decision evidence
M-0 range integration	Connect vendor PTCP prototype and government-owned replay harness to emulated and hardware-in-loop routers; verify fast/slow interfaces and negative permissions	Reproducible build, interface conformance, no unmediated writes, timing decomposition, baseline comparisons
M-1 observe-only enclave	Install route/telemetry collectors and local assurance nodes at a non-mission-critical enclave; no actuation	Evidence quality, coverage, model stability, operator workload, storage/resource tails
M-2 shadow mission thread	Issue advice envelopes and compare with operator/native routing decisions across planned outages and adversarial route/trust scenarios	Detection calibration, explanation, false-containment cost, mission-flow age/deadline metrics
M-3 reversible local actions	Authorize route de-preference, sampling, re-authentication, or bounded admission at selected cells	Safety-kernel rejection, minimum-service proof, commit/rollback tails, independent override
M-4 disconnected exercise	Isolate slow path and selected support services; exercise M2/M3 modes, local PQC continuity, store-forward, recovery and rejoin	Useful isolated duration, authority conformance, key/profile continuity, conflict-preserving reconciliation
M-5 federated operational evaluation	Federate enterprise, theater, mission and coalition cells with scoped models and diverse controllers	Mission-thread benefit, red-team closure, accreditation evidence, sustainment and multi-vendor portability

Military safety boundary

PTCP controls networking and communications resources only. It does not select targets, authorize force, generate rules of engagement, override platform flight/navigation safety, or replace command judgment. Network isolation must preserve the command, recovery, safety, and legal-control paths specified by the mission contract.

14.8 Critical-infrastructure deployment blueprint

Critical infrastructure uses the same two-timescale logic but a different authority model. The enterprise security operations center or network operations center hosts the broad slow path. A regional operations center may host sector or geography contexts. Each plant, substation, water facility, pipeline station, port, hospital, data center, or communications site hosts a local assurance node at the enterprise/OT boundary. Passive sensors and brokers supply permitted evidence. PTCP actions terminate at IT routing, remote-access, DMZ, application-gateway, or identity controls unless an asset-specific safety case explicitly authorizes more. Process controllers, protective relays, safety instrumented systems, emergency shutdown, and manual safety functions remain independent. [34,35]

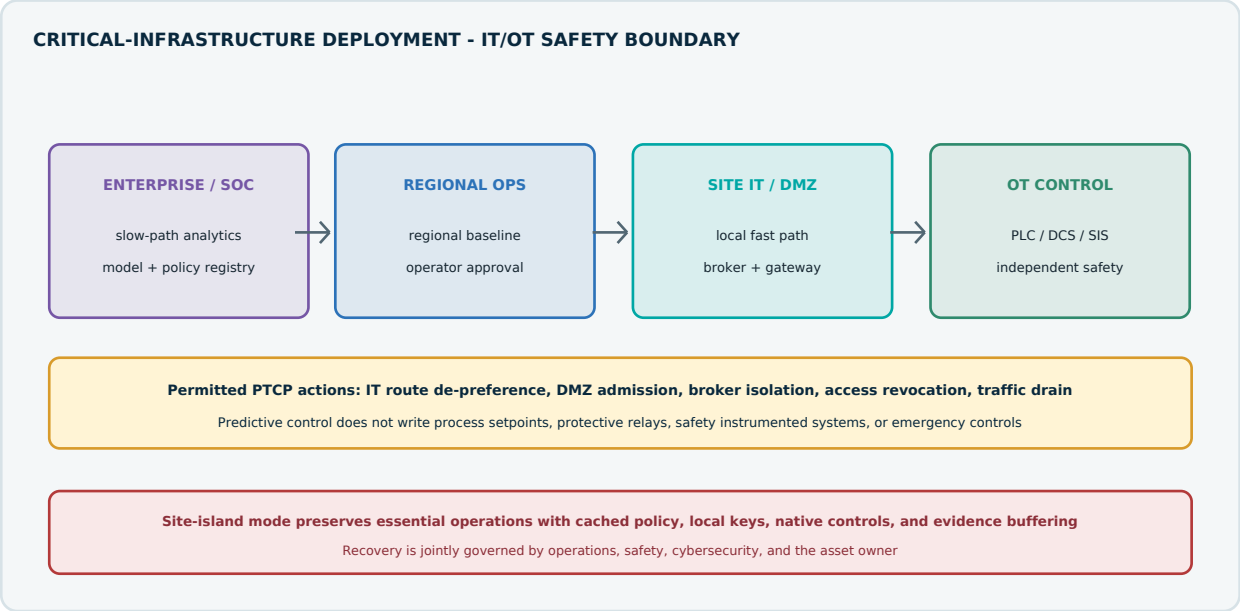

Figure 9. Critical-infrastructure deployment pattern. PTCP may protect connectivity surrounding an industrial process, but the process-safety hierarchy remains authoritative and controller-independent.

Zone	Fast/slow placement	Permitted action examples	Non-negotiable boundary
Enterprise IT / cloud	Primary slow path, evidence lake, model/policy registry, SOC console	Recommend WAN path, isolate IT workload, constrain identity, prioritize emergency communications	Enterprise model cannot directly write OT process or safety state
Regional operations	Regional slow path or read replica; sector-specific contexts; human approval	Coordinate regional traffic drain, carrier alternative, field-access restriction, evidence sharing	Owner/operator and regulatory authority govern service and disclosure
Site IT / industrial DMZ	Local fast path, safety kernel, crypto broker, gateway/broker adapters	Route de-preference, DMZ admission, jump-host revocation, broker isolation, maintenance-session termination	No automatic bypass of change management, maintenance authority, or emergency access policy
OT supervisory network	Primarily passive observation; tightly scoped enforcement only after safety case	Increase monitoring, restrict untrusted remote path, preserve historian/control availability	No optimization of process setpoints; latency/jitter changes require engineering validation
Basic/process control	No PTCP runtime dependency	Native segmentation and approved access controls only	DCS/PLC logic remains governed by process engineering and validated change control
Safety and protection	Physically/logically independent	PTCP may observe health through a one-way or isolated interface if authorized	SIS, relays, emergency shutdown, manual control and protective functions cannot depend on PTCP
Site recovery	Local known-good bundles, config backups, recovery keys, append-only evidence	Restore gateway/IT network, revoke suspect controller, stage reconnect	Recovery requires asset owner, operations/safety, and cybersecurity decision rights

Critical-infrastructure fielding sequence

- Begin with asset inventory, data-flow and dependency mapping, network/identity/crypto inventory, and a joint operations-safety-cyber hazard analysis. Define which evidence may leave a site and at what granularity.
- Deploy passive telemetry and deterministic route/access validation at the enterprise and industrial DMZ. Baseline planned maintenance, seasonal load, storms, startup/shutdown, backup communications, and emergency exercises as explicit contexts.
- Run POL-TT/Dtopo offline and in shadow mode through representative process and carrier events. Compare with existing SOC/NOC and OT monitoring; measure false positives by safety and service consequence, not only alert count.
- Authorize low-risk reversible actions first: increased sampling, operator notification, IT route de-preference, maintenance-access re-authentication, DMZ traffic drain, or broker isolation with a verified alternate path.
- Exercise island mode with enterprise controller, WAN, time, PKI status, and selected key services unavailable. The site must retain native process control, approved local identity/key continuity, audit buffering, and manual recovery communications.
- Scale regionally only after independent test shows that correlated quarantine cannot cascade across sites, emergency communications remain available, and sector/asset safety authority can override or disable predictive actuation locally.

14.9 Common deployment package and ownership

Package element	Military owner pattern	Critical-infrastructure owner pattern
Mission/service contracts	Operational communications authority with cyber, platform, coalition and safety stakeholders	Asset owner and operations with safety/process engineering, network, cyber, privacy and regulatory stakeholders
Slow-path service	Enterprise/theater network operations with separate model-risk and recovery authority	Enterprise/regional NOC-SOC with sector and site governance
Fast-path assurance node	Platform/site communications authority; locally enforceable safety contract	Site IT/OT boundary owner; local operations/safety veto
Crypto profile and keys	Independent cryptographic/key-management authority; platform-specific cache and recovery	Enterprise/site PKI and key-management owners; device/protocol migration plan
Deterministic policy	Routing, cross-domain and coalition authorities	Network engineering, identity, remote-access and industrial-zone policy owners
Recovery	Separately credentialed incident/recovery chain with known-good artifacts	Joint asset-owner, operations/safety and cybersecurity recovery authority
Acceptance evidence	Mission-thread, disconnected-operation, red-team, cyber, safety and interoperability test	Service continuity, safety case, sector compliance, incident response, cyber and vendor interoperability test

Use case	PTCP-derived benefit to evaluate	Required safeguard	Primary metric
Multi-path C2	Earlier detection of structural path concentration and tail-risk-aware alternative ranking	Minimum command/recovery paths; controller-independent routing	Deadline miss and age-of-information tails
Autonomous platform group	Joint bandwidth/trust/context anomaly and bounded network coordination	Installed autonomy contract; no authority expansion	Mission-flow completion during partition
Space/remote site	Contextual distinction between planned and abnormal disconnect	Local crypto/key continuity and store-forward	Useful isolated duration; rejoin integrity
Base/port logistics	Dependency-aware detection of route/trust degradation	OT protective systems excluded from predictive actuation	Essential service retention and recovery time
Coalition gateway	Correlated topology evidence across permitted views	Release-minimized evidence and independent domain policy	Policy-conformant delivery and false containment
Critical infrastructure	Regional anomaly and cut-redundancy awareness	Sector safety case, privacy minimization, local authority	Service availability and bounded blast radius

15. Development and transition roadmap

15.1 Program structure

The program is organized as independently testable work packages with government or enterprise-owned interface specifications, evidence schemas, safety contracts, and conformance tests. PTCP does not begin as a production controller. The initial product is an evidence pipeline, offline replay corpus, and shadow analytics service. Closed-loop privileges are earned by action class and region through measured gates.

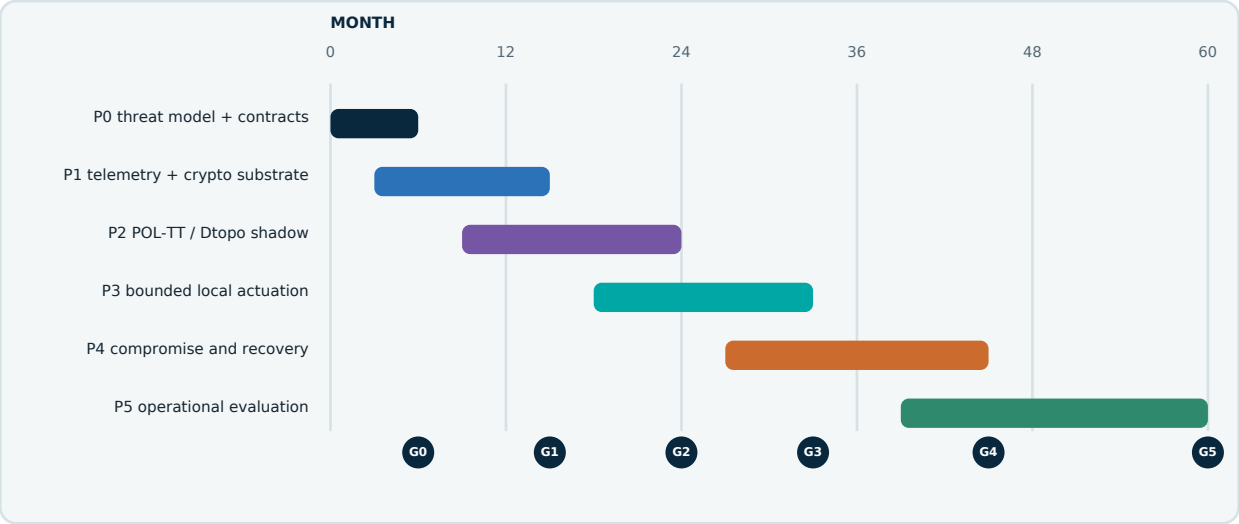

Figure 10. Illustrative 60-month evidence-driven roadmap. Gates authorize the next risk class; calendar progress alone does not.

15.2 Work packages

Work package	Scope	Primary deliverable
WP1 assurance architecture	Threat model, authority separation, mission/local contracts, safety predicates, recovery roles	Approved reference architecture and formal interface/authority specification
WP2 evidence substrate	Adapters, provenance envelope, time/sequence, deterministic validators, evidence quality	Conformance-tested telemetry and route-security evidence pipeline
WP3 POL-TT research	Feature/context selection, TT implementation, calibration, two-speed learning, model health	Reproducible offline/shadow evidence with baseline comparisons
WP4 Dtopo research	Region definitions, curvature/cut estimators, calibration, explanation, leak/trust cases	Detection evidence package with uncertainty and ablation results
WP5 policy and safety	State machine, evidence quorum, hysteresis, blast radius, minimum service, rollback	Independent local safety kernel and machine-checkable policy suite
WP6 PQC agility	CBOM, profiles, broker, modules, key lifecycle, negotiation, constrained-platform tests	Approved profiles and measured interoperability/performance evidence
WP7 local autonomy	PACE, local identity/key continuity, store-forward, action budgets, rejoin	Disconnected-operation safety case and hardware-in-loop evidence
WP8 compromise/recovery	Controller health, revocation, standby diversity, known-good restore, reconciliation	Controller-compromise containment and recovery demonstration
WP9 test and accreditation	Digital twin, red team, safety, cyber, crypto, performance, operator studies	Release criteria, test data, residual-risk record, authorization evidence

15.3 Phase gates

Gate	Evidence required	Authority unlocked
G0 architecture	Threat model, claims ledger, safety/authority contracts, data governance, test plan	Build and collect in isolated range
G1 substrate	Provenance, deterministic validators, crypto profiles, local known-good mode, interface conformance	Shadow collection on representative systems
G2 analytics	POL-TT/Dtopo calibration, baseline comparison, poisoning/evasion evaluation, explanation and resource bounds	Shadow recommendations to operators
G3 bounded actuation	False-containment cost, rollback, minimum-service proof, safe rate/blast-radius tests	Limited reversible actions in controlled environment
G4 compromise/recovery	Controller/key/model compromise containment, isolation duration, recovery/rejoin evidence, red-team closure	Operational evaluation by selected cells/action classes
G5 transition	Mission-thread benefit, operator workload, accreditation, supply-chain sustainment, residual-risk acceptance	Scaled fielding within approved profiles

15.4 Acquisition and sustainment principles

- Acquire interfaces, evidence, test corpora, policy semantics, and reproducibility - not only a proprietary model service.
- Require multi-vendor enforcement adapters and an independently implementable local safety kernel.
- Treat TT models, region definitions, cryptographic profiles, and safety policy as controlled configuration items with sustainment owners.
- Preserve a non-PTCP known-good mode and an exportable audit/replay format through every upgrade.
- Budget for cryptographic migration, module validation, device refresh, certificate/key lifecycle, and long-lived constrained platforms.
- Use red-team findings and operational false-containment events as gated change inputs, not informal model retraining data.

16. Verification, red-team evaluation, and evidence

16.1 Verification ladder

Level	Environment	Questions answered	Exit artifact
L0 specification	Architecture and policy models	Are authority, state transitions, expiry, minimum service, and rollback unambiguous?	Reviewed requirements, invariants, model/property checks
L1 component	Unit and property tests	Do parsers, validators, TT queries, estimators, crypto broker, and safety predicates handle bounds and malformed input?	Automated test evidence and coverage
L2 replay	Curated normal/failure/adversarial corpus	Are scores calibrated, reproducible, and better than simple baselines without unacceptable false containment?	Versioned dataset, metrics, ablation, error analysis
L3 emulation	Digital twin / network range	How does the closed architecture behave under route, trust, telemetry, timing, controller, and link faults?	End-to-end traces and mission-flow metrics
L4 hardware-in-loop	Representative routers, gateways, crypto modules, edge compute	Do timing, MTU, compute, power, transactional change, and fallback meet constraints?	Platform-specific performance and safety case
L5 shadow	Representative operational network, no actuation	What would PTCP recommend; how often would it be useful, wrong, stale, or unsafe?	Operator-reviewed recommendation corpus
L6 bounded pilot	Limited reversible action classes	Does independent safety prevent harm and does measured mission value justify added complexity?	Operational evaluation and residual-risk decision

16.2 Defensive red-team scenario families

Scenario family	Controlled injection	Expected safe behavior	Evidence
Route leak / relationship violation	Synthetic unauthorized propagation and policy mismatch in isolated range	Deterministic validator flags; POL-TT/Dtopo add regional context; scoped response preserves minimum path	Detection/containment time, precision, collateral route loss
Origin/path evidence conflict	Diverse validators or collectors disagree	Unknown/degraded mode; no false normal and no broad destructive quarantine	Conflict visibility, decision consistency, operator workload
Low-and-slow topology drift	Gradual path concentration, cut loss, and trust decline	Drift/model-health evidence rises; trusted baseline resists capture	Time-to-detect, baseline contamination, rank/update behavior
Telemetry poisoning/evasion	Forged, replayed, suppressed, or context-mislabeled observations	Provenance/sequence failure, learning freeze, holdout protection, conservative action	Poison acceptance rate, replay rejection, false containment
Controller compromise	Validly signed invariant-violating or excessive-scope recommendations	Local safety rejects; PTCP identity/model is contained; known-good mode continues	Unsafe-action acceptance rate (target zero), containment time, service retention
Controller/key isolation	Partition analytics, policy, key, and time services in combinations	Advice expires; local contract and pinned crypto continue within bounds	Useful autonomy duration, expired privilege, rejoin conflicts
Crypto downgrade/substitution	Manipulate offered profiles, module identity, or negotiation transcript	Authenticated negotiation rejects; no silent weaker public-key fallback	Downgrade rejection, interoperability, outage behavior
Recovery-plane abuse	Unauthorized restore, self-reinstatement, stale known-good package	Separate authorization rejects; audit preserved; staged rejoin required	Unauthorized recovery acceptance (target zero), RTO/RPO
False-containment cascade	Correlated alerts target multiple critical paths and evidence channels	Blast-radius aggregation and minimum-service proof narrow or block action	Maximum affected mission flows, rollback time, surviving recovery path

16.3 Core measurement framework

Dimension	Metrics	Interpretation
Detection	Precision/recall by scenario, time-to-detect, calibration, false alarm rate, coverage/conflict rate	Separate model discrimination from evidence availability
Containment	Time-to-constrain/quarantine, false-quarantine probability, collateral flows/nodes, minimum-service violations	A fast detector is not valuable if it self-denies the mission
Routing	p50/p95/p99 latency and age, loss/jitter, route churn, tail outage, path diversity, deadline misses	Evaluate expected and tail mission performance against baselines
Model	Likelihood/calibration, TT rank/storage/query cost, truncation error, drift, poisoning acceptance, rollback fidelity	Test whether compression and learning remain stable and useful
Controller resilience	Unsafe-command acceptance, isolation transition time, local useful duration, state conflict, recovery RTO/RPO	Prove PTCP is non-essential to minimum safe function
PQC	Handshake/signature latency, CPU/memory/power, message expansion/MTU, rekey/revocation, module failure, downgrade rejection	Measure platform and protocol consequences of migration
Human/system	Alert load, explanation usefulness, review time, override/rollback success, procedure error	Automation must reduce rather than hide operational risk

16.4 Baselines, ablation, and statistics

Every claimed benefit is compared with standards-based deterministic controls and simpler statistical baselines, not only with an unprotected network. Ablation removes POL-TT, curvature, cut, trust, or corroboration terms to determine marginal value. Experiments predefine datasets, splits, scenarios, seeds, tuning boundaries, confidence intervals, consequence-weighted false-containment cost, and stopping rules. Test data containing sensitive topology is governed and minimized; publication can use synthetic or abstracted results.

17. Governance, safety, and residual risk

17.1 Decision rights

Decision	Accountable authority	Required separation
Mission-flow priority and minimum service	Mission/service owner	PTCP may estimate benefit but cannot redefine authority
Deterministic route policy	Network/security authority	Model team cannot suppress invalidity or local filters
Model promotion	Analytics owner plus independent test/security approval	Training operator cannot self-approve production baseline
Quarantine action class	Security/mission authority according to consequence	High-impact action requires independent authorization or hard deterministic rule
Cryptographic profile and exception	Cryptographic/key-management authority	PTCP and ordinary network operators cannot weaken profile
Controller revocation and recovery	Independent incident/recovery authority	Suspected controller cannot approve reinstatement
Residual-risk acceptance	Authorizing executive/system owner	Evidence producer does not accept its own unresolved safety risk

17.2 Safety case structure

- Claim: the mission fabric preserves specified minimum safe functions when PTCP is unavailable, compromised, or rejected.
- Argument: authority separation, local contracts, safety predicates, deterministic controls, cryptographic profile pinning, and independent recovery break single-controller failure paths.
- Evidence: requirements traceability, specification/property tests, component results, replay/emulation, hardware-in-loop, shadow data, red-team results, operator exercises, and known limitations.
- Conditions: stated trust-anchor assumptions, platform/environment limits, supported profiles, maximum isolation duration, evidence coverage, maintenance procedures, and authorized action classes.
- Residual risk: simultaneous compromise of independent roots, unobserved common-mode supply-chain faults, novel model failures, incomplete topology, and physical/transport loss remain possible.

17.3 Privacy, classification, and intelligence value

A mission-fabric evidence graph reveals topology, dependencies, trust, mission priority, and failure modes. It is therefore a high-value intelligence object. Collection is purpose-limited; raw data and derived features are labeled, access-controlled, compartmented, retained by policy, and minimized across organizational or coalition boundaries. Explanations and audit exports are filtered without destroying accountability. Model artifacts and synthetic data are evaluated for leakage of sensitive operational patterns.

17.4 Principal residual risks

Residual risk	Why it remains	Mitigation / acceptance evidence
Common-mode compromise	Diverse logical planes may share supplier, build, hypervisor, identity, or staff	Dependency map, supplier diversity where justified, independent builds/tests, recovery exercises
False quarantine	Attack, failure, maintenance, and observation gaps can look similar	Graduated response, corroboration, minimum-service proof, operator authority, measured consequence
Model capture or evasion	Adaptive adversary and nonstationary network exceed training coverage	Guarded learning, baselines, drift/residual tests, red team, shadow deployment
Stale local autonomy	Policy, keys, routes, and mission context age during long isolation	Explicit TTL and safe state, local observation, prepositioned alternatives, exercises
PQC implementation defect	New code and protocol integrations may introduce ordinary vulnerabilities or interoperability failures	Validated modules, KEM guidance, test vectors, side-channel and platform testing, rollback
Recovery abuse	Emergency mechanisms are privileged and less frequently used	Narrow scope, separate keys/roles, ceremony, periodic test, immutable evidence
Resource exhaustion	Telemetry, graph, forecast, crypto, or audit workload can starve essential functions	Quotas, reserved resources, graceful reduction, admission, tail-load testing

17.5 Ethical and operational constraint

This architecture is defensive network assurance. It intentionally excludes target selection, weapons-release logic, operational exploitation, and instructions for manipulating real networks. Local network autonomy remains subordinate to applicable law, policy, safety, and human authority. Predictive confidence never substitutes for command responsibility or evidence required by another operational process.

18. Conclusions and recommended program decision

PTCP offers a testable research proposition: compress joint telemetry into a bounded-rank density model, forecast network state, rank paths by expected and tail risk, and detect regional topology defects through density, curvature, and cut changes. The proposition becomes suitable for mission-fabric evaluation only when embedded in an architecture that assumes the predictive controller can fail or become hostile. The decisive design move is authority separation.

POL-TT and Dtopo should supply decomposable evidence. Deterministic routing validation retains precedence. Evidence quality, persistence, and independent corroboration govern progression through observe, constrain, quarantine, and recovery. Local safety kernels enforce minimum service, scope, expiry, change budgets, and rollback. Controller compromise activates a separately keyed recovery plane; isolation activates pre-installed bounded autonomy. Cryptographic profiles, keys, and downgrade rules remain independent of PTCP and support a measured migration to classical PQC.

Recommended decision

Authorize a gated research and prototyping program beginning with evidence architecture, deterministic route-security integration, classical PQC inventory/agility, and controller-independent local safety. Implement POL-TT and Dtopo in offline replay and shadow mode. Do not authorize broad closed-loop quarantine until the program demonstrates calibrated detection, acceptable false-containment cost, unsafe-command rejection, useful isolated operation, no-silent-downgrade behavior, and independently controlled recovery on representative hardware and mission threads.

Gate 0 entry criteria

Criterion	Minimum evidence before build
Authority architecture	Signed allocation of observation, analytics, policy, enforcement, crypto, recovery, mission, and residual-risk decision rights
Threat model	Controller, telemetry, model, routing, trust, key, supply-chain, time, partition, and recovery threats with stated assumptions
Mission contracts	Representative flows with priority, deadline/data age, minimum service, release, crypto, local-autonomy, and rejoin rules
Safety invariants	Machine-testable scope, sequence, expiry, minimum-service, blast-radius, trust, crypto, and rollback predicates
Evidence plan	Canonical provenance schema, deterministic validator integration, collection authority, privacy/classification, and replay corpus strategy
PQC plan	CBOM method, approved profile authority, module/protocol candidates, constrained-platform test matrix, and downgrade policy
Exit metrics	Predefined detection, containment, mission, model, local-autonomy, recovery, PQC, and operator thresholds for each gate

Appendix A. Derived system requirements

The following technology-agnostic requirements are derived from the architecture. Values such as timeouts, action budgets, minimum paths, and evidence thresholds are supplied by the applicable mission, platform, security, and cryptographic authorities and verified on representative systems.

ID	Requirement	Verification
ARCH-001	The system shall implement logically distinct observation, inference, policy, enforcement, cryptographic, and recovery authorities with separately assignable credentials and permissions.	Architecture review; privilege and compromise tests
ARCH-002	PTCP analytics shall have no unmediated production write path to routers, gateways, identity stores, key services, crypto profiles, recovery services, or safety policy.	Interface inventory; negative authorization tests
IMP-001	The implementation shall separate a global/regional slow path from a local deterministic fast path, with packet forwarding and native routing able to continue independently of both.	Process/interface mapping; component isolation and failure tests
IMP-002	The slow path shall distribute signed, scoped, expiring PolicyBundle and ModelBundle artifacts rather than confer an unrestricted remote command capability.	Schema/signature tests; capability and expiry negative tests
IMP-003	The fast path shall execute only installed action templates after local authority, precondition, minimum-service, blast-radius, crypto-profile, and rollback validation.	Template conformance; validly signed unsafe-action campaign
IMP-004	The fast path shall emit a DecisionReceipt binding evidence, disposition, actual enforcement state, post-condition, rollback outcome, timing, and local monotonic sequence.	Receipt completeness, loss/replay, and reconciliation tests
IMP-005	Fast-path detection, validation, device commit, convergence, verification, and rollback timing shall be measured separately at tail percentiles under representative load and fault conditions.	Instrumented benchmark; p50/p95/p99/p99.9 and deadline-miss evidence
IMP-006	Loss, compromise, or expiry of the slow-path service shall not create new local authority or disable native known-good routing, cryptographic policy, safety, or recovery functions.	Partition, controller capture, bundle expiry, and native-mode tests
AUTH-001	Every action shall identify the authorizing role, scope, action class, target, policy version, sequence, expiry, and rollback reference.	Schema conformance; replay/expiry tests
AUTH-002	A receiver shall validate semantic authority and local safety independently of the controller's message signature.	Validly signed unsafe-command test
EVD-001	Every authoritative telemetry record or batch shall bind source/collector identity, schema, sequence, event/receipt time, quality, lineage, labels, and integrity.	Schema, signature/MAC, and tamper tests
EVD-002	The schema shall distinguish absent, stale, invalid, withheld, not-applicable, collector-failed, and cryptographically unverifiable evidence.	Missingness and conflict test corpus
EVD-003	High-impact topology actions shall require the configured independent observer diversity or a deterministic hard-fail condition.	Single-source compromise and quorum tests
EVD-004	Low evidence coverage or observer conflict shall enter an explicit degraded/unknown state and shall not be represented as normal.	Suppression and split-view tests
RTG-001	The system shall ingest and preserve deterministic route-security results separately from statistical scores.	RPKI/Role/OTC/local-policy integration test
RTG-002	A statistical model shall not override a deterministic invalidity or explicit route-policy disposition.	Model-poisoning disagreement test
POL-001	Each POL-TT model shall be bound to versioned features, transforms, quantiles, bins, ranks/cores, contexts, code, dependencies, training manifest, and approvers.	Manifest and deterministic replay
POL-002	Trusted and candidate online-learning states shall be separated; candidate updates shall not directly redefine production normality or authorize high-impact action.	Privilege and promotion workflow test
POL-003	The trusted POL-TT update path shall freeze when provenance, coverage, drift, residual, resource, or compromise criteria fail.	Fault and poisoning injection
POL-004	POL-TT rank, memory, query time, update magnitude, and accepted sample rate shall be bounded per deployment profile.	Resource-stress and rank-growth test
POL-005	Anomaly scores shall be calibrated and reported with context, model health, coverage, and component/contribution explanation.	Held-out calibration and operator review

ID	Requirement	Verification
TOP-001	Dtopo records shall expose regional density, curvature-gradient, cut-shift, weights, estimator versions, windows, coverage, and uncertainty as separate fields.	Schema and recomputation test
TOP-002	Monitored-region definitions and expansion rules shall be versioned and bounded independently of current anomaly input.	Region-inflation and blast-radius test
TOP-003	Deterministic route evidence, trust evidence, confidence, persistence, and mission consequence shall accompany rather than silently redefine the source Dtopo equation.	Claims and data-model review
QTN-001	Quarantine shall implement normal, observe, constrain, quarantine, and recovery states with distinct entry/exit conditions, hysteresis, and dwell time.	State-machine model and transition tests
QTN-002	Every constrain/quarantine action shall preserve configured minimum service, required path diversity, operator/recovery reachability, and applicable safety functions.	Counterfactual and live post-condition tests
QTN-003	Containment actions shall be reversible, scoped, expiring, rate-limited, and bound to a verified rollback or safe substitute.	Rollback, expiry, and idempotence tests
QTN-004	Correlated actions derived from common evidence shall count against a cumulative blast-radius budget.	False-containment cascade test
CTL-001	The system shall reject advice from unapproved controller identities, models, code, policy versions, sequences, or expired envelopes.	Identity/model/sequence negative tests
CTL-002	High-impact actions shall require independent co-authorization unless an approved deterministic hard-fail policy explicitly applies.	Dual-control bypass test
CTL-003	External controller-health monitoring shall detect invariant violations, anomalous action rate/scope, attestation failure, and diverse-controller disagreement.	Controller compromise campaign
REC-001	Controller revocation, known-good restore, and staged rejoin shall use authority and keys unavailable to the ordinary PTCP runtime.	Credential compromise and recovery test
REC-002	Known-good recovery packages shall bind software, firmware, model, policy, configuration, dependencies, crypto profile, and provenance.	Artifact substitution and stale-package test
REC-003	Rejoin shall preserve divergent evidence and conflicts; trust, release, safety, or key state shall not use blind last-writer-wins reconciliation.	Partition/rejoin conflict test
AUT-001	Each local cell shall hold a signed autonomy contract defining permitted scope, time, change, blast radius, minimum service, trust, data, crypto, and safe-state rules.	Contract conformance under isolation
AUT-002	Isolation shall not authorize creation of trust roots, policy expansion, crypto weakening, unsigned updates, or broader cross-domain/mission authority.	Negative authority tests
AUT-003	Local privileges, credentials, policies, and stored data shall expire or degrade according to explicit rules when their evidence or time limits are exceeded.	Long-duration isolation test
DEP-001	A military deployment shall federate enterprise, theater/region, mission-enclave, and platform cells such that loss of any superior PTCP tier does not remove subordinate native routing, safety, cryptographic, or recovery authority.	Multi-tier partition and compromise exercise
DEP-002	PTCP advice at coalition or cross-domain boundaries shall remain subordinate to the applicable guard, release, mission-partner, and domain policy and shall not expand information releasability.	Cross-domain negative authorization and release tests
DEP-003	A critical-infrastructure deployment shall terminate predictive write authority at the approved IT/OT boundary unless an asset-specific hazard analysis and safety case authorize a narrower action.	Interface inventory; hazard and negative-write tests
DEP-004	Process controllers, protective relays, safety instrumented systems, emergency shutdown, and manual safety functions shall not depend on PTCP availability or prediction correctness.	Loss/compromise test with independent safety verification
PQC-001	The program shall maintain a CBOM covering algorithms, protocols, modules, libraries, firmware, keys/certificates, data life, owners, and migration status.	Inventory sampling and automated evidence
PQC-002	Cryptographic profiles shall be signed, versioned, context-specific, locally cached, and controlled independently of PTCP.	Profile substitution and isolation test
PQC-003	Algorithm negotiation and fallback shall be authenticated/transcript-bound and shall reject silent downgrade to disallowed public-key cryptography.	Downgrade/stripping campaign
PQC-004	PQC implementations shall use approved modules and parameter sets applicable to the system and shall expose measured platform overhead and failure behavior.	Module identity, KAT, load, MTU, power tests
PQC-005	Compromise of a controller signing key shall not compromise local safety, cryptographic policy, or recovery signing keys.	Key-separation and compromise test
SW-001	Software, firmware, models, policies, schemas, and dependencies shall follow secure development, provenance, signing, vulnerability response, and controlled-update practices.	Supply-chain and release audit
AUD-001	The system shall record evidence, advice, decision, authorization, action, post-condition, rollback, model/code/config versions, and operator intervention in an integrity-protected ledger.	Audit completeness and tamper test
TST-001	PTCP shall remain shadow-only until calibrated detection, false-containment cost, resource bounds, explanation, and baseline superiority are demonstrated.	Gate evidence review
TST-002	Closed-loop privileges shall be granted by action class and region only after unsafe-command rejection, minimum-service, rollback, isolation, PQC, and recovery tests pass.	Operational test readiness review

Appendix B. Detection feature and evidence matrix

Manipulation	POL-TT joint features	Dtopo effect	Deterministic / independent evidence	Preferred initial state
Provider-to-provider or peer leak	Role/OTC result, AS-path relationship, prefix fan-out, update rate, next-hop, utilization	Bridge/curvature change; cut concentration	RFC 9234 role/OTC; prefix/path policy; diverse route views	Reject if det erministic; otherwise o bserve/const rain
Unauthorized origin	Origin state, prefix length, new origin, fan-out, service reachability	Regional path and cut change	RPKI origin validation; local authorization	Deterministi c disposition plus scoped impact analysis
More-specific/internal-prefix leak	Prefix-length distribution, source role, volume, duration, reachability	New local bridges and path concentration	Prefix filters, max-prefix, route policy	Reject/const rain with mi nimum-servi ce check
Path manipulation	AS/path sequence, path-length delta, communities, path-validation state, latency	Curvature discontinuity and centralization	BGPsec/path evidence where deployed; policy and observer comparison	Observe/con strain; quarantine on corrobora tion
Trust degradation	Attestation, certificate/revocation, posture, privilege, log continuity, route/performance	Low-trust edges lengthen; cut capacity falls	Independent identity, hardware root, local deny list, operator evidence	Observe, re-authenticate , constrain privilege
Telemetry suppression	Missingness type, collector diversity, receipt gap, sequence, traffic/route residual	Apparent or uncertain topology shift	Peer-side counters, active probes, alternate collector	Degraded/u nknown; freeze learning
Telemetry forgery/replay	Signature/MAC, sequence, time quality, source behavior, cross-counter mismatch	Potential false deformation	Provenance verification and monotonic state	Reject evidence; contain source if cor roborated
Low-and-slow model poisoning	Update source mix, quantile drift, likelihood residual, rank, holdout loss	Dtopo baseline slowly shifts	Immutable replay, candidate/trusted divergence, approver workflow	Freeze promotion; investigate
Legitimate maintenance	Signed change, affected scope/time, expected route/link/trust changes	Potential high Dtopo but predicted region	Change authority, operator confirmation, work order	Observe with change -specific bounds
Physical/link failure	Interface errors/state, peer counters, environmental/power, route convergence	Cut loss and curvature shift	Device alarms, active probes, independent path state	Constrain/re route; do not infer malice
Controller compromise	Action rate/scope, target distribution, invariant failures, model identity, output divergence	May manufacture or suppress topology response	Local safety, attestation, diverse controller, recovery authority	M3 compro mise; reject controller

Interpretation rule

No row is a fixed signature. The matrix defines evidence expectations and safe initial posture. Exact region definitions, thresholds, persistence, and actions are calibrated in a controlled range and approved by the responsible network, security, mission, and safety authorities.

Appendix C. Fallback-mode decision matrix

Function	M0 connected	M1 degraded evidence	M2 isolated	M3 controller compromise	M4 rejoin
PTCP inference	Trusted model; shadow/bounded advice	Score on snapshot; learning frozen or candidate-only	Optional local snapshot; no enterprise authority	Controller/model rejected and preserved for forensics	Shadow compare before staged restoration
Deterministic routing	Full policy and validators	Continue; prefer local corroborated evidence	Continue from local policy/cache within validity	Continue independently of PTCP	Revalidate affected state before replay
Quarantine	All approved states/actions	Observe/constrain; broad action needs independent evidence	Local scope and budget only	Block PTCP-origin actions; local deterministic rules remain	No replay of expired containment; staged review
Mission prioritization	Current enterprise/local contract	Conservative capacity and tighter change budget	Installed contract; PACE/store-forward	Installed known-good contract	Resolve contract/version conflict explicitly
Identity/trust	Enterprise plus local evidence	Shorter scope/lifetime; more local attestation	Cached/local trust within contract	Controller identity revoked; other roots separate	Refresh and reconcile revocation carefully
PQC profile	Current signed profile	Pinned; no network-driven downgrade	Locally cached profile and approved material	Controller cannot alter profile or recovery keys	Authenticate profile/version before new sessions
Keys/credentials	Normal lifecycle and revocation	Conservative issuance; preserve essential sessions per policy	No new enterprise privilege; local continuity within TTL	Rotate suspected controller material through recovery authority	Re-establish and revoke stale material
Software/model update	Signed controlled update	Defer nonessential updates	Known-good only; no unsigned emergency package	Freeze suspect channel/artifacts	Attest, verify provenance, staged deployment
Audit	Stream and retain	Prioritize security/action evidence	Local append-only buffer	Preserve incident snapshot and rejected commands	Signed upload; conflict-preserving merge
Operator control	Normal role and override	Raised visibility and review	Local authorized roles only	Independent recovery/security role activates	Dual review until stable

Conflict precedence during rejoin

State type	Default reconciliation
Safety or release policy	Most restrictive valid state remains until authorized conflict resolution; no timestamp-only overwrite.
Controller identity/revocation	Independent recovery/security authority decides; suspect controller messages do not restore trust.
Cryptographic profile	Use the newest mutually authenticated approved profile that satisfies local minimum; reject unknown downgrade.
Route and service state	Revalidate current state; replay only idempotent, unexpired, still-authorized changes.
Model/quantile state	Do not merge online updates blindly; compare candidate histories and retrain/promote through normal gate.
Audit evidence	Preserve both histories, ordering uncertainty, gaps, and conflict markers; deduplicate by immutable identifiers.
Mission data	Apply data-age, provenance, ownership, and application conflict policy; do not assume last writer is authoritative.

Appendix D. Glossary

Term	Definition
Advice envelope	Signed PTCP recommendation containing scope, evidence, expected outcome, risk, model/policy version, expiry, and rollback; it is not direct authority.
Blast radius	Cumulative set of flows, nodes, services, regions, or dependencies affected by active or correlated changes.
Bounded local autonomy	Pre-authorized local operation inside explicit time, scope, change, trust, data, cryptographic, and safety limits.
CBOM	Cryptographic bill of materials: an inventory of algorithms, protocols, modules, libraries, keys/certificates, firmware, data life, ownership, and migration state.
CVaR	Conditional Value-at-Risk, a tail-risk measure used in the PTCP candidate-path objective.
Dtopo	PTCP topology defect score combining regional anomaly, graph-curvature-gradient, and normalized cut-shift components.
Deterministic validator	Standards- or policy-based check such as RPKI origin validation, BGP Role/OTC, prefix/path policy, credential validity, or artifact signature.
Evidence quality	Visible properties of provenance, freshness, coverage, time, diversity, agreement, transformation integrity, and authorized context.
Mission assurance contract	Machine-readable rules for producer/consumer, authority, priority, deadline/data age, minimum service, release, cryptography, degradation, audit, and recovery.
POL-TT	Pattern-of-Life Tensor Train, the source-defined bounded-rank joint probability-density model used to estimate normal network behavior.
PQC	Post-quantum cryptography: classical cryptographic algorithms designed to resist attacks by classical and quantum computers.
PTCP	Predictive Tensor Control Plane, Ochoa's proposed overlay for tensor-compressed telemetry, risk-aware routing, and topology-security evidence.
Quarantine	Reversible scoped restriction such as de-preference, admission reduction, traffic drain, credential restriction, route rejection, or segment isolation.
RPKI	Resource Public Key Infrastructure used to support verifiable Internet-number-resource and route-origin authorizations.
Safety envelope	Independently administered local constraints that project a recommendation into an authorized, survivable, reversible action or reject it.
TNQG	Tensor-Network Quantum Gravity, Ochoa's conjectural operational reconstruction research program; not a dependency of this architecture.
Unknown state	Explicit state for inadequate, stale, conflicting, or unverifiable evidence; it limits destructive automation and is not equivalent to normal.
Zero trust	Architecture that grants no implicit trust based solely on network location and evaluates identity, device/workload, resource, and policy for access. [8]

Appendix E. PTCP.ai claim and verification ledger

This ledger prevents a public vendor claim, an analytic complexity result, a laboratory demonstration, and an independently reproduced operational result from being treated as equivalent. The status column reflects only the sources reviewed for this paper as of July 17, 2026; it is not a conclusion that no non-public evidence exists.

Claim / proposition	Source and current status	Acceptance evidence required for fielding
PTCP implements an autonomic brain/reflex architecture	PTCP.ai public architecture presentation; vendor-documented design [40]	Versioned component diagram, code/artifact provenance, interface conformance, independent architecture review
A local fast path can evaluate compact probability and act at microsecond class	PTCP.ai presentation; vendor-reported concept/demonstration, no raw timing corpus identified [40]	Clock definition and instrumentation; hardware/software bill; p50/p95/p99/p99.9; overload, cache miss, rollback and failure results
PTCP can operate on existing SONiC leaf/spine hardware	PTCP.ai product statement; target platform claim [39]	Named SONiC releases and ASICs; adapter source/binary provenance; topology/scale matrix; commit, rollback, failover and coexistence tests
POL-TT reduces dense O(n^d) storage to O(d n r^2) at bounded rank	Supplied PTCP manuscript and PTCP.ai technical overview; analytic/asymptotic claim [1,42]	Measured memory/update/query constants; rank/accuracy curves; concept drift, poisoning, missingness and resource exhaustion
POL-TT point-density query is O(d r^2)	PTCP.ai technical overview; analytic claim [42]	Implementation benchmark over d, r, contexts and hardware; deterministic output and numerical-stability evidence
Dtopo can identify route leaks and trust degradation	Source algorithm and proposed validation protocol; no independent operational accuracy found [1,42]	Labeled multi-topology corpus; precision/recall, calibration, TTD, attack/maintenance separation, ablation and independent reproduction
Geometric quarantine can occur at wire speed without unsafe blast radius	Vendor objective/claim; not established by reviewed public evidence [39-42]	End-to-end detect/authorize/commit/converge/verify timing; false-containment cost; minimum-service and correlated-cascade tests
Predictive routing improves tail latency and accelerator utilization	Vendor value proposition and manuscript algorithm [1,39,41]	Controlled comparison with ECMP/traffic-engineering baselines; p99 flow completion, churn, deadline misses, utilization, workload completion and energy
Controller compromise or isolation is survivable	Derived architecture requirement in this paper; not a vendor proof claim	Validly signed unsafe commands, key/model capture, telemetry split view, long partition, bundle expiry, native fallback, recovery and rejoin campaigns
Classical PQC agility can coexist with fast/slow PTCP	Derived architecture using NIST standards/guidance [12-20]	Per-platform/protocol interoperability, module validation, MTU/fragmentation, handshake/load/power tails, downgrade and isolated-key-continuity tests

Minimum reproducibility package

- Versioned source or independently inspectable executable, build manifest, software bill of materials, compiler/runtime, container or image hash, and signing provenance.
- Named switch/router/DPU/CPU models, firmware and SONiC/network-OS versions, topology, link rates, protocol configuration, telemetry sampling, traffic/workload generator, clock and measurement method.
- POL-TT feature schema, contexts, quantiles/bins, ranks/cores, training/test split, update policy, random seeds, model-health thresholds, and challenger-promotion record.
- Dtopo region definitions, curvature/cut estimators, windows, weights, coverage/uncertainty rules, attack and non-attack labels, deterministic validator results, and action state-machine configuration.
- Raw and summarized results with confidence intervals, p50/p95/p99/p99.9 tails, resource utilization, false-containment and mission/service cost, negative results, failure injection, and baseline/ablation comparisons.
- Signed PolicyBundle/ModelBundle/DecisionReceipt examples, local safety rules, audit/reconciliation artifacts, controller-compromise campaign, and independent replication report.

References

[1] J. Ochoa, Predictive Tensor Control Plane (PTCP): Tensor-Train Telemetry, Risk-Aware Geodesic Routing, and Topology-Native Security, Tensor Networks, Inc. (2026), supplied six-page manuscript.

[2] J. Ochoa, Tensor-Network Quantum Gravity as an Operational Reconstruction Program: Definitions, Algorithms, and Continuum-Limit Conjectures, Tensor Networks, Inc. (2026), supplied six-page manuscript.

[3] I. V. Oseledets, "Tensor-Train Decomposition," SIAM Journal on Scientific Computing 33(5), 2295-2317 (2011). doi:10.1137/090752286.

[4] T. G. Kolda and B. W. Bader, "Tensor Decompositions and Applications," SIAM Review 51(3), 455-500 (2009). doi:10.1137/07070111X.

[5] R. T. Rockafellar and S. Uryasev, "Optimization of Conditional Value-at-Risk," Journal of Risk 2, 21-41 (2000). doi:10.21314/JOR.2000.038.

[6] R. Forman, "Bochner's Method for Cell Complexes and Combinatorial Ricci Curvature," Discrete and Computational Geometry 29, 323-374 (2003). doi.

[7] Y. Ollivier, "Ricci Curvature of Markov Chains on Metric Spaces," Journal of Functional Analysis 256(3), 810-864 (2009). doi.

[8] S. Rose et al., Zero Trust Architecture, NIST SP 800-207 (August 2020). doi:10.6028/NIST.SP.800-207.

[9] R. Ross et al., Developing Cyber-Resilient Systems: A Systems Security Engineering Approach, NIST SP 800-160 Vol. 2 Rev. 1 (December 2021). csrc.nist.gov.

[10] Joint Task Force, Security and Privacy Controls for Information Systems and Organizations, NIST SP 800-53 Rev. 5, Update 1 (2020/2021). csrc.nist.gov.

[11] A. Nelson et al., Incident Response Recommendations and Considerations for Cybersecurity Risk Management, NIST SP 800-61 Rev. 3 (April 2025). csrc.nist.gov.

[12] National Institute of Standards and Technology, Module-Lattice-Based Key-Encapsulation Mechanism Standard, FIPS 203 (August 2024). csrc.nist.gov.

[13] National Institute of Standards and Technology, Module-Lattice-Based Digital Signature Standard, FIPS 204 (August 2024). csrc.nist.gov.

[14] National Institute of Standards and Technology, Stateless Hash-Based Digital Signature Standard, FIPS 205 (August 2024). csrc.nist.gov.

[15] G. Alagic et al., Recommendations for Key-Encapsulation Mechanisms, NIST SP 800-227 (September 2025). csrc.nist.gov.

[16] E. Barker et al., Considerations for Achieving Crypto Agility: Strategies and Practices, NIST CSWP 39upd1 (June 29, 2026). csrc.nist.gov.

References - continued

[17] The White House, Executive Order 14412, Securing the Nation Against Advanced Cryptographic Attacks (June 22, 2026). whitehouse.gov.

[18] National Security Agency, Post-Quantum Cybersecurity Resources, including CNSS Policy 15/CNSA 2.0 pointers and QKD limitations (accessed July 2026). nsa.gov.

[19] National Institute of Standards and Technology, Security Requirements for Cryptographic Modules, FIPS 140-3 (2019). csrc.nist.gov.

[20] E. Barker, Recommendation for Key Management: Part 1 - General, NIST SP 800-57 Part 1 Rev. 5 (May 2020). csrc.nist.gov.

[21] K. Sriram et al., Problem Definition and Classification of BGP Route Leaks, RFC 7908 (June 2016). rfc-editor.org.

[22] A. Azimov et al., Route Leak Prevention and Detection Using Roles in UPDATE and OPEN Messages, RFC 9234 (May 2022). rfc-editor.org.

[23] P. Mohapatra et al., BGP Prefix Origin Validation, RFC 6811 (January 2013). rfc-editor.org.

[24] R. Bush and R. Austein, The Resource Public Key Infrastructure (RPKI) to Router Protocol, Version 1, RFC 8210 (September 2017). rfc-editor.org.

[25] M. Lepinski and K. Sriram, BGPsec Protocol Specification, RFC 8205 (September 2017). rfc-editor.org.

[26] J. Durand, I. Pepelnjak, and G. Doering, BGP Operations and Security, RFC 7454 / BCP 194 (February 2015). rfc-editor.org.

[27] J. Scudder, R. Fernando, and S. Stuart, BGP Monitoring Protocol, RFC 7854 (June 2016). rfc-editor.org.

[28] M. Lepinski and S. Kent, An Infrastructure to Support Secure Internet Routing, RFC 6480 (February 2012). rfc-editor.org.

[29] A. Vassilev et al., Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations, NIST AI 100-2 E2025 (March 2025). csrc.nist.gov.

[30] M. Souppaya, K. Scarfone, and D. Dodson, Secure Software Development Framework (SSDF) Version 1.1, NIST SP 800-218 (February 2022). csrc.nist.gov.

[31] J. Boyens et al., Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations, NIST SP 800-161 Rev. 1, Update 1 (May 2022/November 2024). csrc.nist.gov.

[32] E. Tabassi, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1 (January 2023). nist.gov.

References - continued

[33] National Institute of Standards and Technology, Implementing a Zero Trust Architecture, NIST SP 1800-35 (June 2025). csrc.nist.gov.

[34] K. Stouffer et al., Guide to Operational Technology Security, NIST SP 800-82 Rev. 3 (September 2023). doi.

[35] National Institute of Standards and Technology, The NIST Cybersecurity Framework (CSF) 2.0 (February 2024). doi:10.6028/NIST.CSWP.29.

[36] E. W. Dijkstra, "A Note on Two Problems in Connexion with Graphs," Numerische Mathematik 1, 269-271 (1959).

[37] J. Y. Yen, "Finding the K Shortest Loopless Paths in a Network," Management Science 17(11), 712-716 (1971). doi.

[38] D. P. Bertsekas, Dynamic Programming and Optimal Control, 4th ed., Athena Scientific (2017).

[39] Tensor Networks, Inc., Predictive Tensor Control Plane (PTCP) product site, including brain/reflex, global density, existing-hardware, and SONiC leaf/spine statements (accessed July 17, 2026). ptcp.ai.

[40] Tensor Networks, Inc., The Autonomic Blueprint: Building the Nervous System for AGI, 14-page PTCP.ai presentation describing distributed reflexes, local hardware processing, microsecond reflex arcs, anomaly filtering, and champion/challenger learning (accessed July 17, 2026). ptcp.ai PDF.

[41] Tensor Networks, Inc., PTCP Solutions, vendor description of predictive routing, topology-security, payload-blind analysis, and intended AI/critical-infrastructure applications (accessed July 17, 2026). ptcp.ai.

[42] Tensor Networks, Inc., The Empirical Tensor: Technical Overview - An Overview of PTCP & TNQG Benchmarks, Geometric Security, and Technical Implementation (May 21, 2026), including analytic complexity statements and validation protocols required for production readiness. ptcp.ai PDF.

Source and evidence note

External standards and policy were checked against primary NIST, NSA, White House, and RFC Editor sources through July 17, 2026. The supplied Ochoa manuscripts contain no DOI, journal citation, arXiv identifier, or peer-review status in the provided documents and are cited as supplied 2026 manuscripts. The PTCP paper reports definitions, algorithms, complexity, limitations, and a proposed validation protocol but no experimental results. PTCP.ai publishes a brain/reflex architecture, fast-path and existing-hardware claims, intended applications, and a technical overview; these are cited as vendor sources. No public raw benchmark corpus, executable reproduction package, third-party replication, or operational certification was identified in the reviewed materials. Program thresholds, mission parameters, platform profiles, microsecond timing, accuracy, and wire-speed containment in this white paper are therefore requirements to measure and verify, not asserted operational performance.