PTCP Zero-Day

Payload-Blind, Topology-Native Defense Against Living-off-the-Land and Non-Semantic Zero-Day Threats

Postdoctoral Technical Whitepaper
Commercial market rationale, security architecture, differentiators, and implementation guidance for enterprise, government, critical-infrastructure, AI-fabric, and regulated environments.

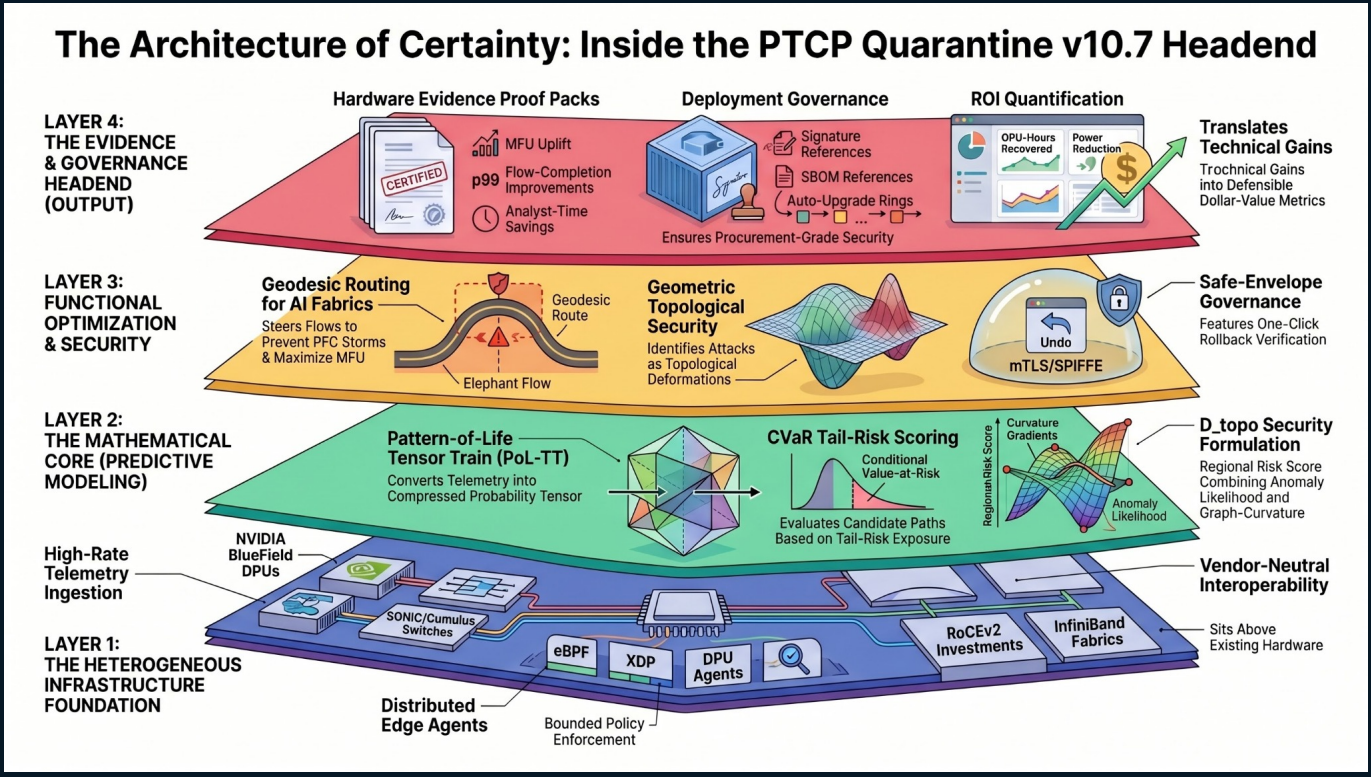

Figure 1. PTCP Zero-Day inherits the PTCP Quarantine headend architecture: heterogeneous telemetry foundation, predictive mathematical core, functional security optimization, and evidence/procurement governance layers.

Payload-blind
No packet content, commands, prompts, secrets, or raw logs are required for the core decision.

Topology-native
Models attack activity as graph deformation, not only as malware identity.

Safe-envelope
Stages containment by default, gates active enforcement behind approval and rollback.

Prepared as marketing and technical collateral for PTCP Zero-Day v11.x. Claim boundary: this whitepaper describes a defensive architecture and commercial evaluation rationale; no security product can guarantee prevention of every zero-day in every deployment.

Executive Abstract

PTCP Zero-Day is a payload-blind, topology-native security product designed for a class of attacks that conventional semantic security controls struggle to adjudicate: zero-day exploitation, Living-off-the-Land (LotL) activity, valid-credential abuse, encrypted exfiltration, control-plane compromise, and AI-agent boundary crossing. These attacks do not always present a malicious file, a recognizable packet payload, or a signature. They often look like legitimate administration until viewed as a deformation of how identities, endpoints, services, and flows relate to one another. The product's core thesis is simple: when an adversary moves laterally, escalates privilege, stages encrypted exfiltration, or rewires control-plane state, the content of a packet may remain opaque, but the network's local geometry changes. PTCP Zero-Day turns that geometry into a measurable security signal. It fuses Pattern-of-Life normalcy, normalized telemetry, graph curvature, cut-capacity shifts, relative information drift, and tail-risk scoring into a bounded D_topo decision. When risk crosses the customer-calibrated threshold, PTCP renders a safe-envelope Geometric Kill Switch plan that can be enforced by existing EDR/XDR, NDR, firewall, microsegmentation, SOAR, Kubernetes, DPU, eBPF/XDP/TC, Windows WFP, SONiC, or cloud-flow control points. The market value is that PTCP adds a missing layer rather than replacing the security stack. EDR tells the organization what a process did. NDR tells it what metadata changed. SIEM/SOAR tells it what alerts and workflows exist. Microsegmentation enforces known boundaries. PTCP tells the organization when the infrastructure itself is geometrically inconsistent with the learned pattern of life and what minimal containment action would reduce blast radius while preserving rollback evidence.

Problems solved
Semantic blind spots, LotL detection gaps, encrypted-flow ambiguity, overbroad quarantine, SOC evidence debt.

Unique mechanism
Pattern-of-Life + graph curvature + cut capacity + QRE-style drift + CVaR tail risk.

Commercial value
Lower dwell-time risk, safer containment planning, privacy-preserving evidence, reusable integrations.

Contents

1. The Market Problem: Non-Semantic Attacks Defeat Semantic Inspection

Modern cyber defense has become highly effective at many semantic tasks: recognizing known malware, parsing command lines, inspecting packet payloads, correlating indicators, and mapping endpoint behavior to known tactics. The strategic problem is that high-end intrusions increasingly avoid giving defenders a clean semantic object to inspect. LotL intrusions use legitimate binaries, built-in administrative tools, approved remote-management channels, and valid credentials. Zero-day intrusions may begin before any signature exists. Exfiltration may occur over encrypted or allowed channels. AI-agent and cloud-control-plane attacks can change policy, identity, or workflow state without dropping conventional malware.

CISA and partner agencies have highlighted Living-off-the-Land activity as a serious threat pattern in critical infrastructure operations, and MITRE ATT&CK; treats valid accounts as a technique that can support initial access, persistence, privilege escalation, and defense evasion. NIST's zero trust architecture guidance also rejects implicit trust based solely on network location, noting that perimeter models are insufficient once attackers have breached the perimeter. Those positions converge on one conclusion: the market needs controls that can assess behavior even when the actor appears authenticated and the payload is unavailable or unhelpful.

> Core market gap: An attacker can be semantically clean while being geometrically wrong. PTCP Zero-Day is designed for that gap: it asks whether the relationships among identities, nodes, services, routes, control-plane objects, and flows have become improbable under the learned pattern of life.

Why semantic controls are not enough

Legacy assumption	What advanced attacks do	Consequence
Malware can be recognized by file, hash, signature, or behavior tree.	Use native tools, trusted remote admin, cloud APIs, and signed binaries.	Endpoint semantics may look administrative instead of malicious.
Packet payloads can be inspected and classified.	Use encryption, allowed tunnels, cloud services, and payload minimization.	DPI can be blind, unavailable, or privacy-inappropriate.
Perimeter location implies relative trust.	Compromise valid accounts and move laterally inside the environment.	The attacker inherits legitimate permissions and path access.
Static segmentation boundaries match attack boundaries.	Pivot through least-monitored relationships, shared credentials, and service dependencies.	Containment is either too late, too broad, or operationally risky.

This creates a product category opportunity: a decision layer that is neither an endpoint signature engine nor a packet-inspection engine, but a mathematical topology engine that consumes non-semantic telemetry and produces bounded, auditable, rollback-aware containment plans.

2. The Specific Problems PTCP Zero-Day Solves

Problem	Why it matters	PTCP Zero-Day solution
Living-off-the-Land activity	Adversaries can use native tools and legitimate administration channels, reducing the usefulness of malware signatures.	Scores non-semantic deviations: valid-credential anomalies, remote execution geometry, fan-out, unusual administrative reach, and lateral-edge changes.
Valid-credential abuse	Authentication success is not equivalent to mission-authorized behavior.	Compares identity, device, service, geography, time, trust, privilege, and route behavior against Pattern-of-Life baselines.
Encrypted exfiltration	Organizations often cannot or should not decrypt every stream.	Detects flow-shape, egress asymmetry, cut perturbation, elephant-flow concentration, and tail-risk exposure without reading payloads.
Static segmentation limits	Predefined zones cannot anticipate every compromise boundary.	Computes topology-aware minimum-cut containment and renders least-blast-radius action plans.
Unsafe response automation	A wrong quarantine can disrupt critical operations.	Projects actions through a safe envelope: dry run by default, approval gates, change ticket, dual control, rollback verification, and action hashes.
SOC evidence debt	Teams need defensible, repeatable rationale, not just alerts.	Produces proof packs, deterministic decision hashes, component scores, denial reasons, and control mappings.
AI-agent and control-plane drift	Non-human actors can change policies, routes, identities, or runtime boundaries.	Scores provenance, policy signature, image/SBOM references, mTLS/SPIFFE posture, and boundary crossings.
Privacy and data-retention risk	Payload collection creates legal, operational, and privacy burdens.	Rejects raw semantic fields by design and uses normalized metadata, hashes, and calibrated features.

The important distinction is that PTCP is not merely another detector. It is a decision system: the detection product is not only a score, but a containment geometry, an evidence bundle, and a governance state. The product answers the buyer's hardest operational question: What can we safely do now, and why is that the minimal action?

3. The Technical Thesis: Attacks Warp Topology Before They Reveal Content

PTCP Zero-Day operationalizes two research ideas. The TNQG paper supplies a disciplined geometry vocabulary: edge capacities, cut areas, inverse-capacity distances, information metrics, curvature diagnostics, and falsifiable reconstruction tests. The PTCP paper translates that language into network control: high-dimensional telemetry is normalized, compressed into a Pattern-of-Life Tensor Train, converted into risk-aware geometric link lengths, and used to compute D_topo security decisions.

The product does not claim that enterprise networks are physical spacetime, nor that generic tensor networks prove a quantum-gravity theory. The geometry is operational: a mathematical way to map dynamic infrastructure relationships into distances, cuts, curvature, and information drift. That claim discipline matters commercially because it makes the product testable, falsifiable, and procurement-ready rather than merely metaphorical.

Operational translation of the papers into the product

Research construct	Operational meaning inside PTCP Zero-Day	Security value
Dimensionless capacity	Normalize telemetry such as bandwidth, queue, jitter, loss, energy cost, identity trust, provenance, and workload posture.	Avoids mixing units and reduces brittle thresholding.
Inverse-capacity distance	High-trust/high-capacity relationships become shorter; low-trust or congested relationships become longer.	Creates a graph where attacks alter path geometry.
Cut area / min cut	Boundary capacity around a suspicious region is measured and minimally severed when necessary.	Targets containment to the topological blast radius.
Curvature diagnostics	Sharp local changes in graph structure indicate deformation around a region.	Detects lateral movement and policy drift as shape changes.
Information metric / relative entropy	Feature distributions are compared to Pattern-of-Life baselines through QRE-style drift.	Detects payload-opaque anomalies.
CVaR tail risk	Paths and actions are evaluated by expected cost plus worst-tail exposure.	Optimizes for mission risk, not only average risk.

Mathematical summary

$$S_{anom}(s_t) = -\log p_\theta(s_t)$$

$$w_e(t) = l0 / (\epsilon_w + \log \chi_e(t)) + \lambda_{deny} * 1[T_e(t) < T_{min}]$$

$$D_{topo}(R,t) = \text{median}(S_{anom}) + \nu * ||\nabla_G \kappa(R,t)||_2 + \xi * \Delta C_{cut}(R,t)$$

$$a^*_t = \text{Projection}_{safe}(a_{raw}(D_{topo}, QRE_{drift}, CVaR_{tail\_risk}))$$

The equations are deliberately practical. They are not a replacement for the organization's existing controls; they generate risk-weighted labels, routing weights, quarantine recommendations, and safe action plans that feed those controls.

4. Product Architecture and Operating Model

PTCP Zero-Day is best understood as a four-layer overlay above heterogeneous infrastructure. It begins with telemetry normalization, transforms that telemetry into a predictive mathematical core, optimizes security decisions against functional constraints, and emits evidence for governance, audit, procurement, and rollback.

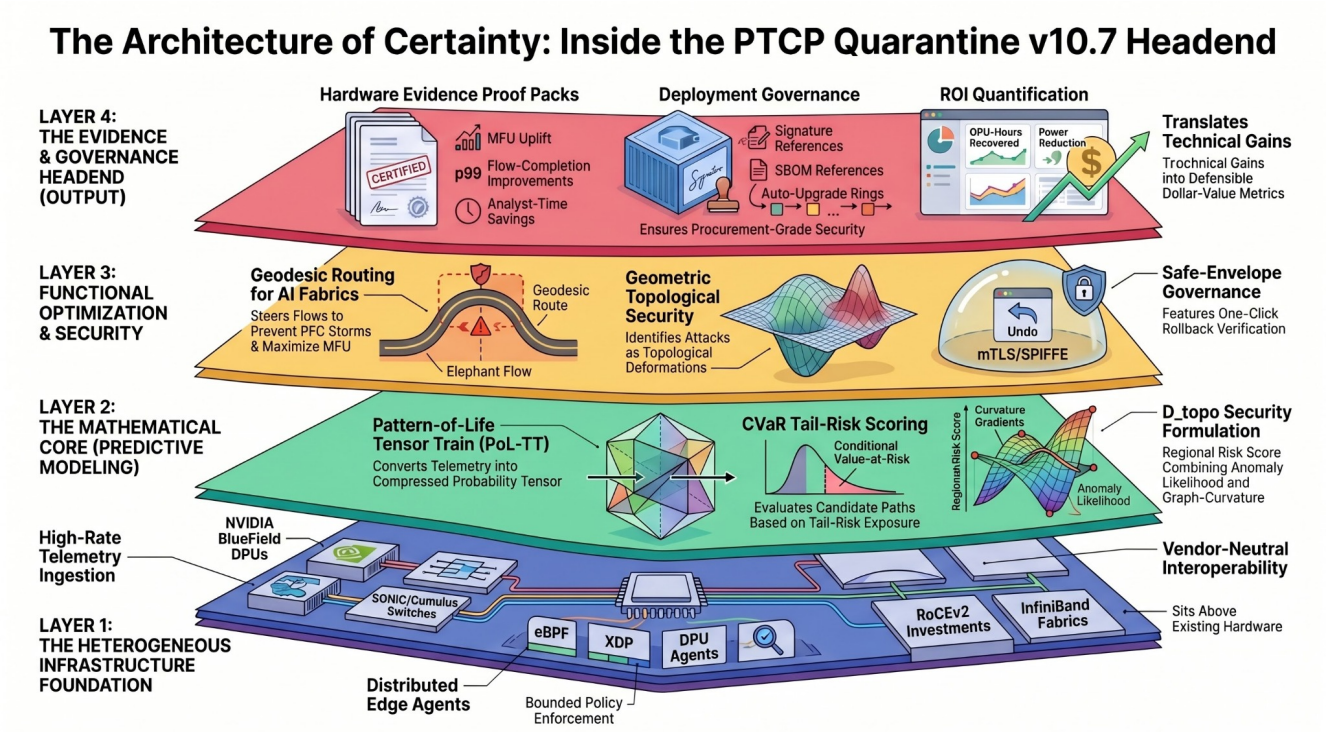

Figure 2. The PTCP Quarantine v10.7 headend architecture that PTCP Zero-Day extends for LotL and non-semantic zero-day defense.

Layer	Function	Commercial implication
Telemetry foundation	Accepts non-semantic signals from endpoints, flow logs, EDR/NDR/SIEM/SOAR, cloud-flow, Kubernetes audit, OT gateways, DPUs, eBPF/XDP/TC, Windows WFP, and fabric devices.	Can be deployed as an overlay without replacing existing tools.
Predictive core	PoL-TT baselines, QRE-style drift, D_topo scoring, graph-window correlation, CVaR tail risk.	Provides a differentiated decision science layer.
Security optimization	Geodesic routing, dynamic minimum-cut quarantine planning, rollback-safe policy rendering.	Turns alerts into bounded actions.
Evidence and governance	Decision hashes, proof packs, connector certification, SBOM/provenance, signed policy promotion, approval records.	Supports SOC, GRC, procurement, and regulated operations.

Design rule: observe everywhere, act narrowly, verify continuously, and never retain secrets or packet payloads by default.

5. Detection Model: From Pattern of Life to Zero-Day Geometry

PTCP Zero-Day does not require the adversary's tool to be known. It asks whether the current network state is improbable under the learned normal state and whether the improbable change creates a high-risk topology. That logic is particularly well-suited to zero-day and LotL operations because those operations often alter relationships before they produce recognizable signatures.

Primary scoring components

Component	Non-semantic indicators	Attack pattern addressed
Credential legitimacy abuse	MFA gap, privilege jump, impossible or unusual time/geography, device-trust mismatch, new admin reach.	Valid accounts, insider-like abuse, identity takeover.
Lateral movement geometry	New east-west edges, fan-out, remote admin protocols, service-to-service path expansion.	LotL lateral movement and reconnaissance.
Native tool / LOLBin context	Approved tools appearing in unusual topology context without retaining command bodies.	PowerShell, WMI, SSH, RDP, kubectl, cloud CLI, and similar administrative utilities.
Exfiltration geometry	Egress asymmetry, encrypted tunnel concentration, elephant-flow deformation, route/region shift.	Encrypted exfiltration and payload-opaque staging.
Trust collapse	Identity, endpoint, workload, service, or control-plane trust drops below calibrated baselines.	Credential compromise, device compromise, privilege escalation.
Control-plane provenance	Unsigned policy, route or ACL change, image/SBOM drift, mTLS/SPIFFE gap, firmware/kernel drift.	Cloud, Kubernetes, network, and fabric control-plane compromise.
AI-agent boundary crossing	Toolchain privilege drift, model-agent action reach, policy writes, data movement against workload normalcy.	Agentic automation risk and AI operations abuse.

Why payload-blindness is a feature, not a compromise

Payload-blind operation creates three advantages. First, it works when payloads are encrypted or unavailable. Second, it reduces privacy and data-retention exposure because sensitive content is not required for the scoring path. Third, it gives the product a cleaner integration boundary: connectors can map non-semantic evidence without sharing raw commands, prompts, secrets, or packet bodies.

Practical example: A legitimate administrator account uses SSH or WinRM to touch a new set of hosts, opens a new service dependency, and produces an egress imbalance. A signature engine may see authorized access. PTCP sees a sudden decrease in local topological plausibility, a fan-out deformation, and a cut-capacity change around the region.

6. Response Model: The Geometric Kill Switch and Safe-Envelope Governance

The product's response model is intentionally conservative. It does not assume that every high-risk score should immediately sever production traffic. Instead, it renders a Geometric Kill Switch plan: a ranked, auditable, rollback-aware set of candidate actions that existing enforcement points can apply after governance gates pass.

Containment workflow

Step	Description	Outcome
1. Localize deformation	Identify the region whose anomaly, curvature gradient, cut shift, and QRE-style drift exceed thresholds.	A region-level topological defect label.
2. Compute boundary	Find the boundary edges and minimum-cut containment alternatives that reduce blast radius.	A small action set instead of broad isolation.
3. Project into safe envelope	Apply policy: dry-run default, active gate, change ticket, recorded approval, dual control, blast-radius limit, rollback verification.	No unsafe automation by default.
4. Render connector action	Generate policy-as-code, firewall, microsegmentation, DPU/eBPF, Windows WFP, Kubernetes/Cilium, SOAR, or ticketing plan.	Existing controls become enforcement surfaces.
5. Produce proof pack	Record component scores, decision hash, plan hash, denial reasons, rollback plan, and evidence summary.	Audit and procurement evidence.

Governance gates

- Dry-run by default: active enforcement is disabled until explicitly enabled by customer policy.
- Change-ticket and approval gates: high-risk actions require recorded approval, dual control, and change context.
- Blast-radius guard: maximum nodes, actions, and boundary edges can be capped per tenant or fabric.
- Rollback verification: every boundary drop, node freeze, or controller action must have a reversible plan or remain advisory.
- Action hashing: deterministic hashes support after-action review, evidence retention, and change reconciliation.

Why this matters to buyers: PTCP aims to reduce both security risk and response risk. The value is not only faster containment; it is safer, more explainable containment that can survive change-review, audit, and regulator scrutiny.

7. Market Differentiation and High-Value Buyer Outcomes

PTCP Zero-Day is positioned as a complement to major security platforms, not as a replacement. The most compelling commercial posture is to make existing controls smarter: feed them topology-native evidence and bounded action plans when semantic tools cannot prove what happened.

The current market already includes strong EDR/XDR, NDR, AI-SOC, and microsegmentation platforms. Microsoft Defender XDR describes incident-level attack disruption across many signals; CrowdStrike Falcon Insight XDR emphasizes endpoint and beyond detection and response; Palo Alto Cortex XDR unifies endpoint, network, cloud, identity, and email sources; SentinelOne positions Singularity as unified endpoint, cloud, identity, and data protection; Illumio emphasizes Zero Trust Segmentation for breach containment; and Darktrace positions NETWORK around AI-assisted network detection and autonomous response. PTCP's point of differentiation is narrower and additive: it supplies payload-blind topology evidence and rollback-safe kill-switch geometry when these platforms need a non-semantic containment trigger.

Category	What the category usually does well	Where PTCP is different
EDR / EPP / XDR	Endpoint telemetry, process trees, file reputation, behavior rules, containment at host.	Detects topology warp when binaries and credentials are legitimate; operates without raw command-body retention.
NDR	Network metadata, flow analytics, encrypted-traffic behavior, east-west visibility.	Adds cut capacity, curvature, QRE-style drift, CVaR tail risk, and minimum-cut containment planning.
SIEM / SOAR	Correlation, search, incident workflow, case management, playbooks.	Supplies deterministic geometric evidence, proof packs, denial reasons, action hashes, and minimal blast-radius plans.
Microsegmentation	Workload isolation and enforcement policy.	Computes dynamic containment boundaries from current topology deformation rather than only predefined zones.
NGFW / SASE / ZTNA	Access enforcement, traffic control, policy boundaries, threat-prevention services.	Acts as a payload-blind decision layer that can recommend when and where enforcement should change.
Cloud/Kubernetes security	Posture, configuration, runtime, API, and workload protections.	Models control-plane provenance, policy drift, service graph changes, and agentic boundary crossing as geometry.

High-value buyer outcomes

Risk reduction
Earlier detection of valid-credential, encrypted, control-plane, and lateral movement anomalies.

Operational safety
Containment planning with blast-radius limits, rollback, and approval evidence.

Stack leverage
Uses existing enforcement points instead of requiring rip-and-replace.

- For CISOs: address a class of threats that evade signatures and payload inspection while improving governance of response automation.
- For SOC leaders: give analysts a concise, explainable topology story instead of another unprioritized alert stream.
- For network and platform teams: generate minimum-cut or route-risk recommendations that respect service availability and rollback.

- For GRC and procurement: create reusable proof packs, SBOM/provenance artifacts, control mappings, and deterministic evidence hashes.
- For privacy and legal teams: minimize sensitive content collection by rejecting raw payloads, commands, prompts, tokens, and secrets.

8. Deployment Blueprint and Proof-Pack Evidence

The recommended deployment pattern is shadow-first, evidence-rich, and integration-oriented. PTCP should initially observe and score without active enforcement, then compare decisions against incident history, purple-team scenarios, and customer-calibrated baselines. After calibration, the organization can promote specific action types into approval or active modes with tight blast-radius budgets.

Phase	Tasks	Exit criteria
Discovery	Map telemetry sources, enforcement points, identity systems, control planes, and sensitive-data boundaries.	Connector inventory, telemetry contract, payload-field rejection policy.
Shadow mode	Ingest non-semantic telemetry, learn Pattern-of-Life baselines, run D_topo/QRE/CVaR scoring.	False-positive review, threshold calibration, proof-pack template.
Purple-team validation	Inject labeled LotL, valid-credential, encrypted exfiltration, control-plane, cloud/K8s, and OT scenarios.	Precision/recall, time-to-detect, false-positive quarantine impact, rollback evidence.
Approval mode	Render action plans into SOAR/ticketing/firewall/microsegmentation/Kubernetes/DPU/eBPF endpoints without active execution.	Change records, action hashes, safe-envelope denial reasons.
Controlled activation	Enable narrow action classes for approved regions and fabrics only.	Dual approval, rollback tests, monitored p95/p99 latency, periodic drift review.

Evidence artifacts that matter commercially

- Proof pack: region ID, D_topo components, QRE drift, tail-risk, dominant driver, decision hash, action hash, and rollback plan.
- Readiness pack: OpenAPI, SBOM, deployment manifest, connector certification, security documentation, and release provenance.
- Benchmark pack: precision, recall, F1, ROC/PR analysis where available, time-to-detect, false-positive quarantine impact, and analyst-time savings.
- Policy pack: Sigma, OCSF/STIX mappings, OPA/Rego, Kubernetes/Cilium, firewall, SOAR, and microsegmentation action exports.

Commercial adoption principle: PTCP Zero-Day should be sold and deployed as a decision/evidence overlay that increases the value of the buyer's current security stack. The fastest path to adoption is usually integration with existing SIEM/SOAR, EDR/XDR, NDR, firewall, ZTNA, and segmentation investments.

9. Validation, Limitations, and Responsible Claim Discipline

A world-class security product is not distinguished only by strong claims; it is distinguished by testable claims. PTCP Zero-Day should be evaluated with customer-local, non-semantic telemetry and labeled scenarios. The central validation question is not whether PTCP can name malware. The question is whether PTCP can detect and localize topological deformation earlier and recommend safer containment than baseline workflows.

Metric	Why it matters
Precision	Measures whether recommended defective regions are truly defective in labeled replay.
Recall	Measures coverage of injected or historically labeled compromise regions.
Time to detect / plan	Quantifies how quickly PTCP moves from telemetry to staged containment plan.
False-positive quarantine impact	Captures the operational cost of over-containment.
Tail-risk reduction	Assesses whether PTCP reduces worst-case mission exposure versus baseline response.
Analyst-time savings	Turns explainability and action planning into SOC labor value.
Rollback verification	Confirms every active-capable action has a recovery path.
Payload-blind contract compliance	Confirms no raw payloads, commands, prompts, tokens, passwords, or secrets enter the scoring path.

Limitations and boundaries

- No product can guarantee prevention of every zero-day or every malicious insider/valid-credential operation.
- Customer-specific calibration is required because normal topology differs by sector, architecture, identity model, and workload.
- Payload-blind operation is a deliberate design boundary; PTCP should complement, not replace, semantic tools when payloads or endpoint semantics are available and authorized.
- Active enforcement should remain disabled until the organization validates thresholds, rollback, change control, and operational impact in its own environment.
- The TNQG and PTCP research language should remain operational and testable: geometry is a modeling discipline, not an unsupported physics claim for cybersecurity.

10. Why the Market Should Strongly Consider PTCP Zero-Day

The strongest reason to consider PTCP Zero-Day is that the attack surface has moved beyond signatures and payloads. Organizations cannot assume that the next critical incident will arrive with a recognizable malicious binary, a readable packet body, or an obviously unauthorized login. In high-end intrusions, the decisive signal may be relational: who touched what, from where, through which trust boundary, with what fan-out, against which learned service dependency, and with which effect on the local graph of the infrastructure.

PTCP Zero-Day converts that relational signal into a product category. It gives the market a payload-blind, topology-native decision layer that can identify non-semantic deformations, quantify tail risk, and render bounded containment actions through existing controls. This is valuable because it aligns with how enterprises actually operate: they already have EDR, NDR, SIEM/SOAR, firewalls, identity, ZTNA, microsegmentation, Kubernetes, cloud, and OT tooling. PTCP makes those controls more decisive when content-centric evidence is weak or unavailable.

The product's strategic differentiation

Not a signature product
It does not depend on known malware, known CVEs, or readable packet content.

Not alert-only
It produces containment geometry, action plans, proof packs, and rollback evidence.

Not rip-and-replace
It complements current platforms and turns them into enforcement surfaces.

Decision statement for buyers

A security leader should evaluate PTCP Zero-Day when the organization has one or more of the following conditions: encrypted east-west traffic; heavy use of remote administration; high-value OT, AI, or regulated environments; prior concerns about LotL and valid-credential attacks; a need for safer automated containment; or a requirement to produce evidence-grade incident and procurement artifacts. In those settings, PTCP's topology-native view may fill a security gap that neither semantic endpoint tools nor static segmentation can fully cover.

The recommendation is therefore not to replace the current stack, but to run PTCP Zero-Day in shadow mode, validate it against customer-local labeled scenarios, integrate it with existing enforcement points, and then promote narrow action classes through the safe envelope. That path gives the buyer a low-risk way to test whether topology-native security materially reduces dwell time, blast radius, analyst burden, and response uncertainty.

References and Source Grounding

The whitepaper uses public threat and architecture guidance, the two attached technical papers, and the PTCP product artifacts supplied in the working set. These references are included for traceability; they are not legal, certification, or procurement advice.

[1] CISA, NSA, FBI, and international partners. PRC State-Sponsored Actors Compromise and Maintain Persistent Access to U.S. Critical Infrastructure, AA24-038A, Feb. 7, 2024. https://www.cisa.gov/news-events/cybersecurity-advisories/aa24-038a

[2] CISA and partner agencies. Identifying and Mitigating Living Off the Land Techniques, Feb. 2024. https://www.cisa.gov/resources-tools/resources/identifying-and-mitigating-living-land-techniques

[3] MITRE ATT&CK.; Valid Accounts, Technique T1078, Enterprise. https://attack.mitre.org/techniques/T1078/

[4] NIST Special Publication 800-207. Zero Trust Architecture, Aug. 2020. https://doi.org/10.6028/NIST.SP.800-207

[5] Julia Ochoa, Tensor Networks, Inc. Tensor-Network Quantum Gravity as an Operational Reconstruction Program: Definitions, Algorithms, and Continuum-Limit Conjectures. Attached Paper A.

[6] Julia Ochoa, Tensor Networks, Inc. Predictive Tensor Control Plane (PTCP): Tensor-Train Telemetry, Risk-Aware Geodesic Routing, and Topology-Native Security. Attached Paper B.

[7] Tensor Networks, Inc. PTCP Quarantine v10.7 product codebase and headend architecture overview. Uploaded working artifact.

[8] Tensor Networks, Inc. PTCP Zero-Day v11.x implementation and commercial readiness artifacts generated in this engagement. Uploaded working artifacts.

[9] Microsoft Learn. Automatic attack disruption in Microsoft Defender XDR. https://learn.microsoft.com/en-us/defender-xdr/automatic-attack-disruption

[10] CrowdStrike. Falcon Insight XDR product material. https://www.crowdstrike.com/en-us/platform/endpoint-security/falcon-insight-xdr/

[11] Palo Alto Networks. Cortex XDR product material. https://www.paloaltonetworks.com/cortex/cortex-xdr

[12] SentinelOne. Singularity Platform product material. https://www.sentinelone.com/

[13] Illumio. Zero Trust Segmentation product/resource material. https://www.illumio.com/

[14] Darktrace. NETWORK product material. https://www.darktrace.com/products/network

Abbreviations

Term	Meaning
CVaR	Conditional Value-at-Risk; a tail-risk metric used to account for rare but severe outcomes.
D_topo	PTCP topology defect score combining anomaly likelihood, graph-curvature gradient, and cut-capacity shift.
LotL	Living off the Land; adversary use of legitimate tools, credentials, or administrative channels.
PoL-TT	Pattern-of-Life Tensor Train; compressed density model for normal multi-dimensional telemetry.
QRE	Quantum Relative Entropy in the research analogy; in PTCP Zero-Day, an operational information-divergence proxy over payload-blind features.
Safe envelope	The governance boundary that projects raw recommendations into allowed, auditable, rollback-safe actions.