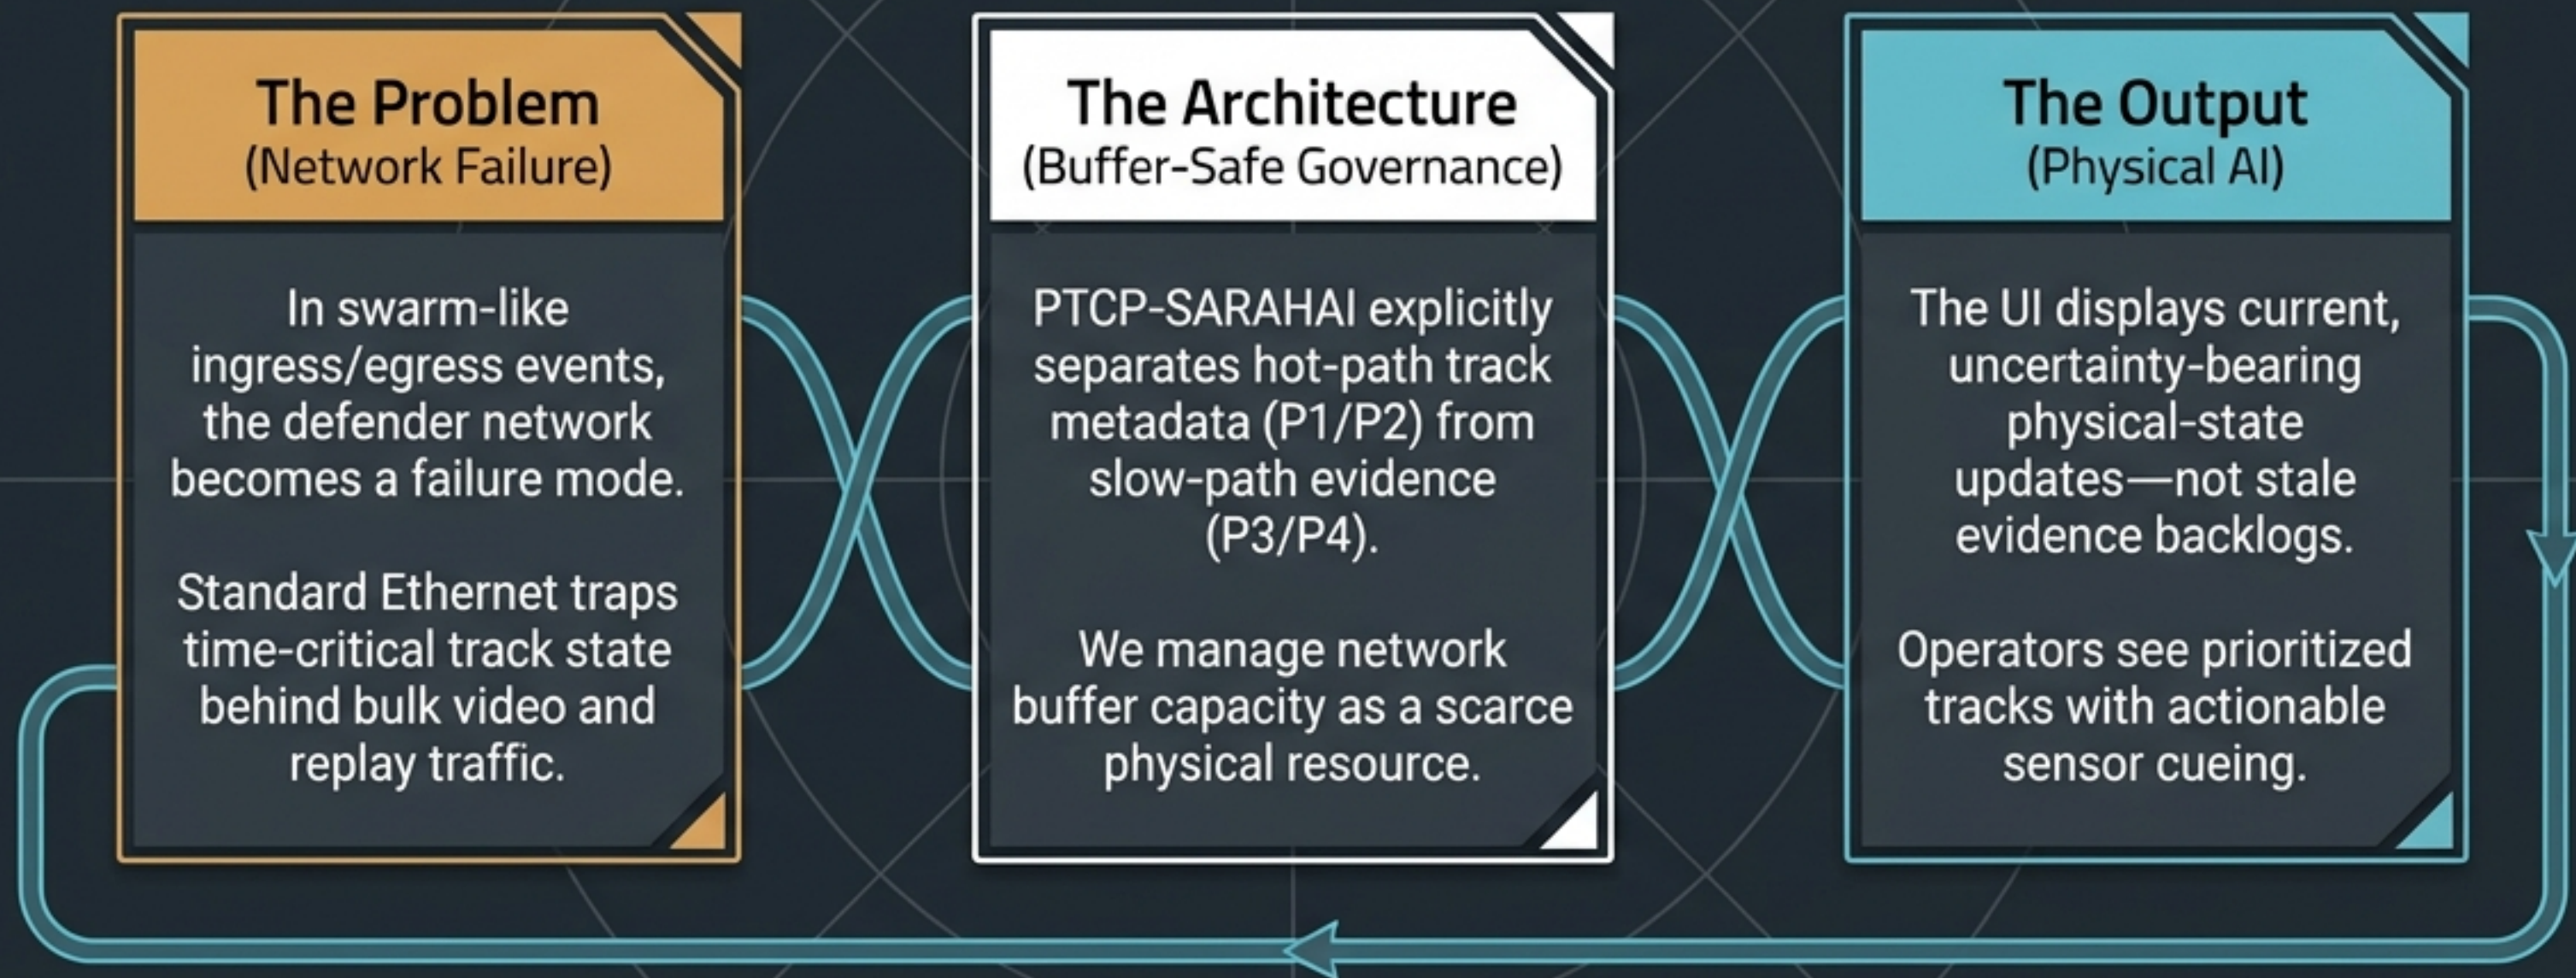




# Buffer-Safe Physical AI for Air-Domain Awareness

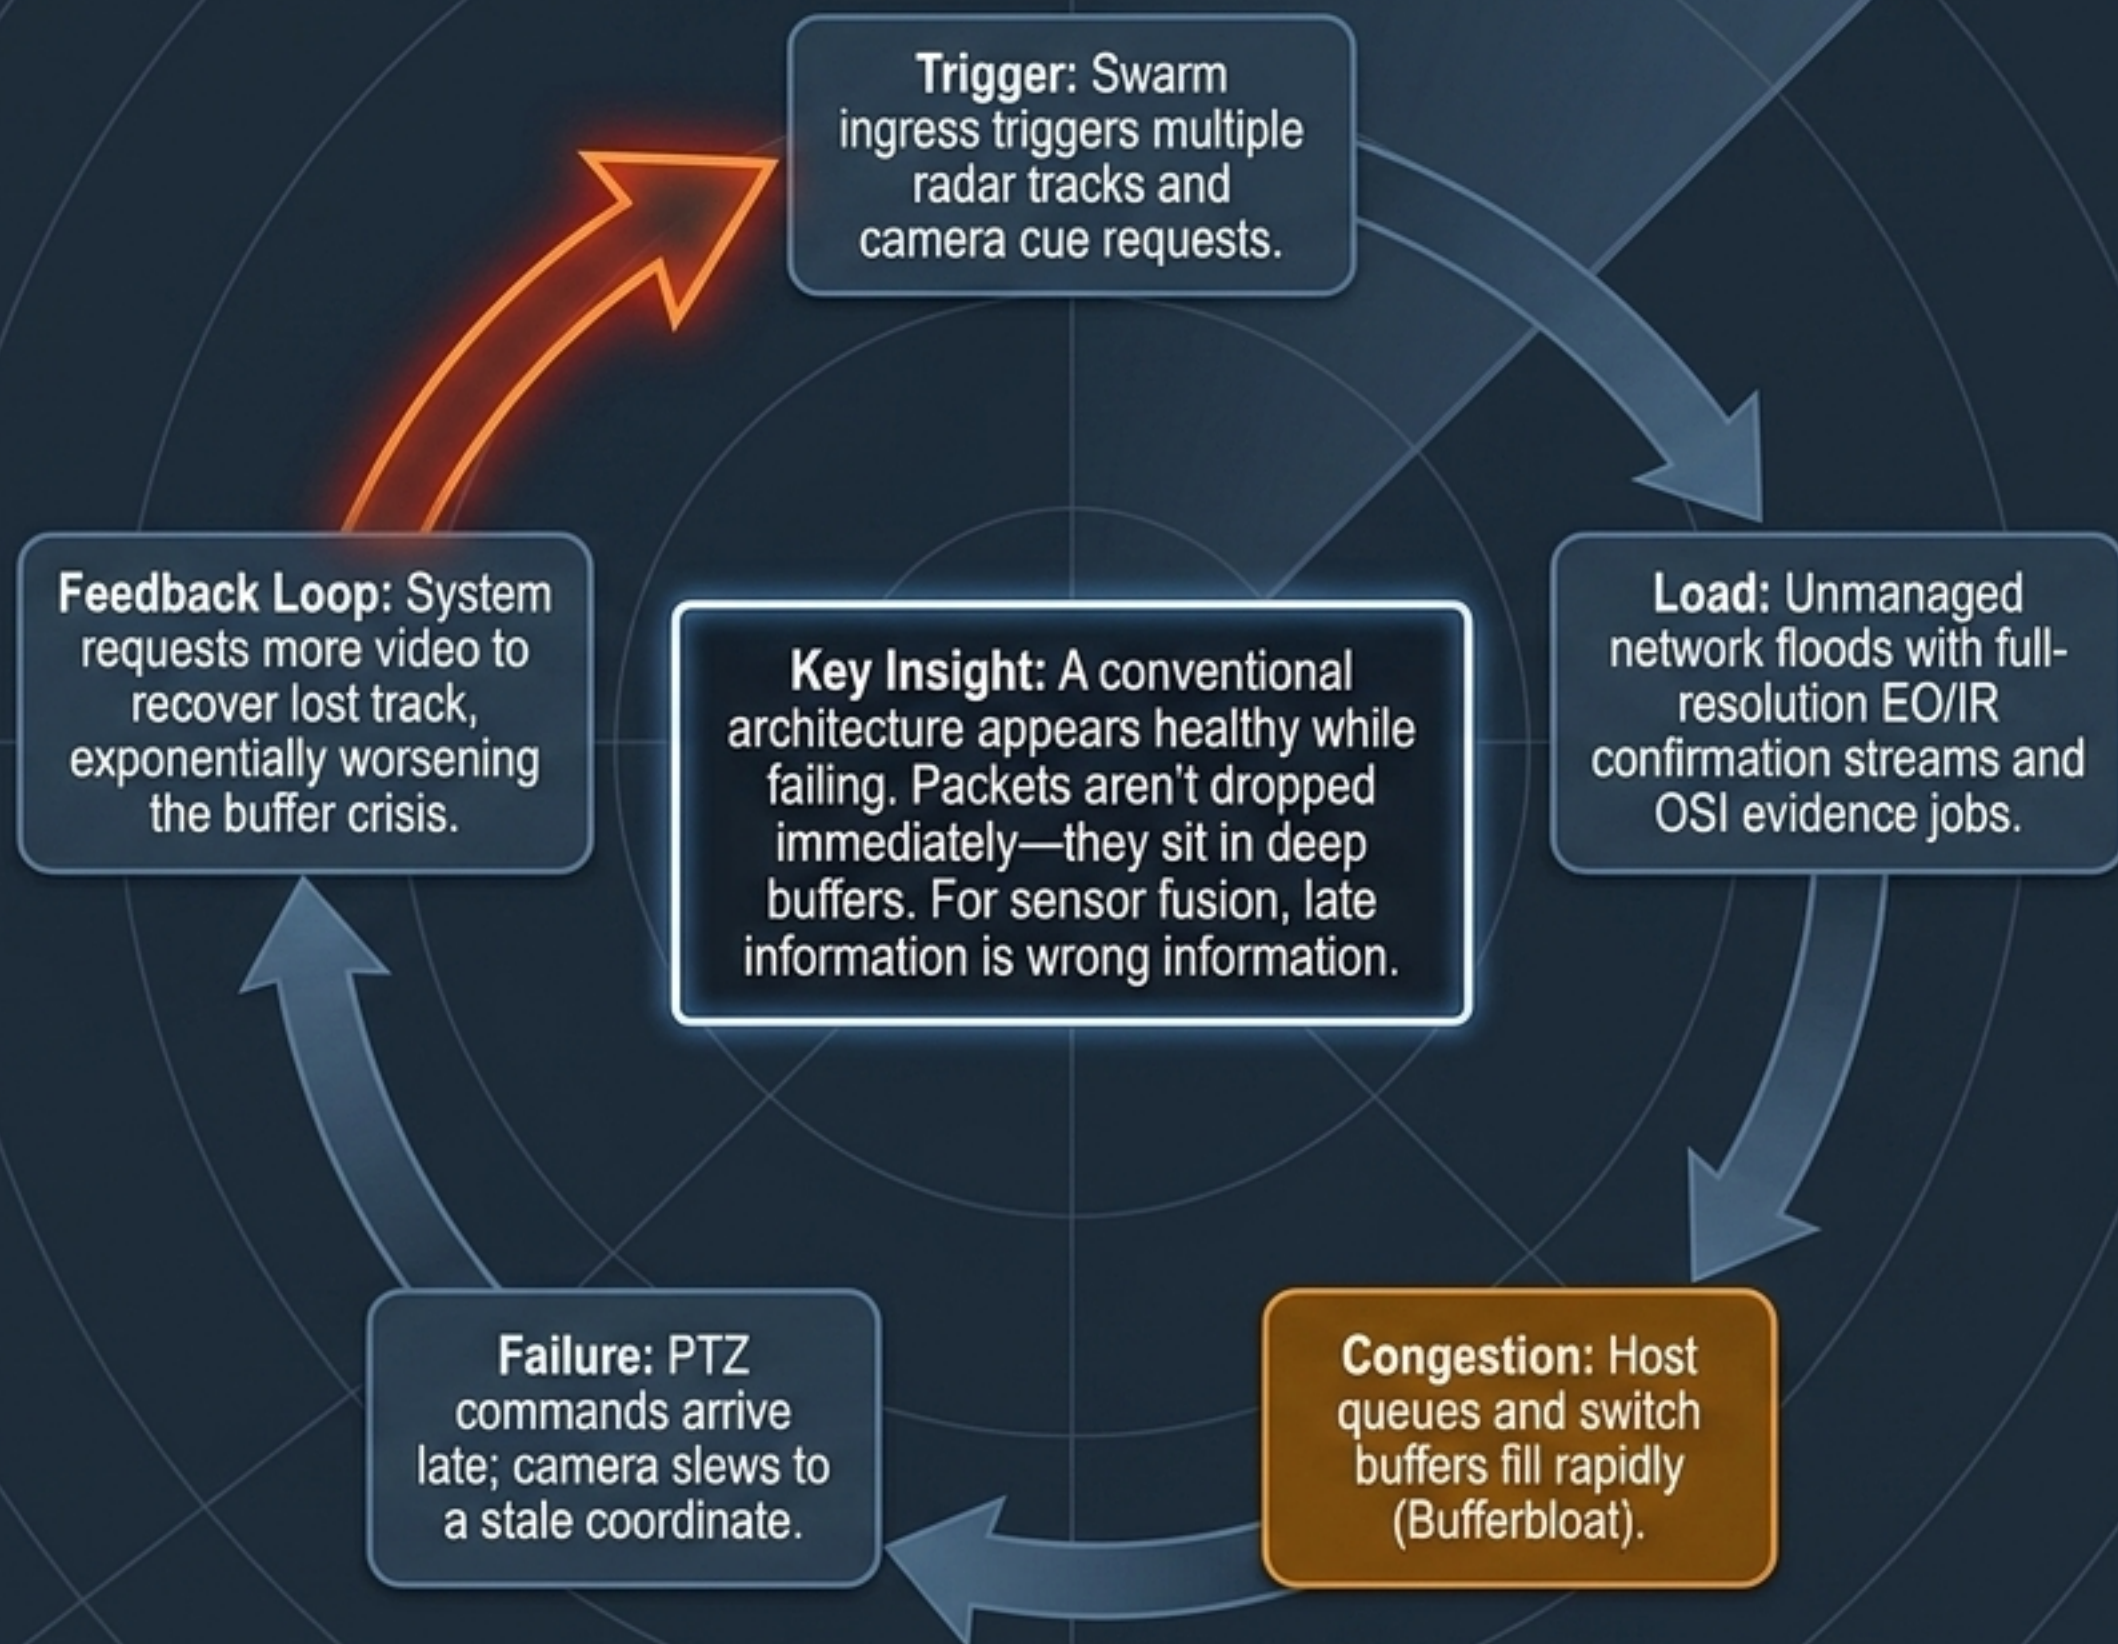
PTCP-SARAHAI: Fusing radar, EO/IR, and predictive network control to preserve cue-to-slew tracking under swarm ingress conditions.

# The Awareness Imperative: Network as a Perception Layer

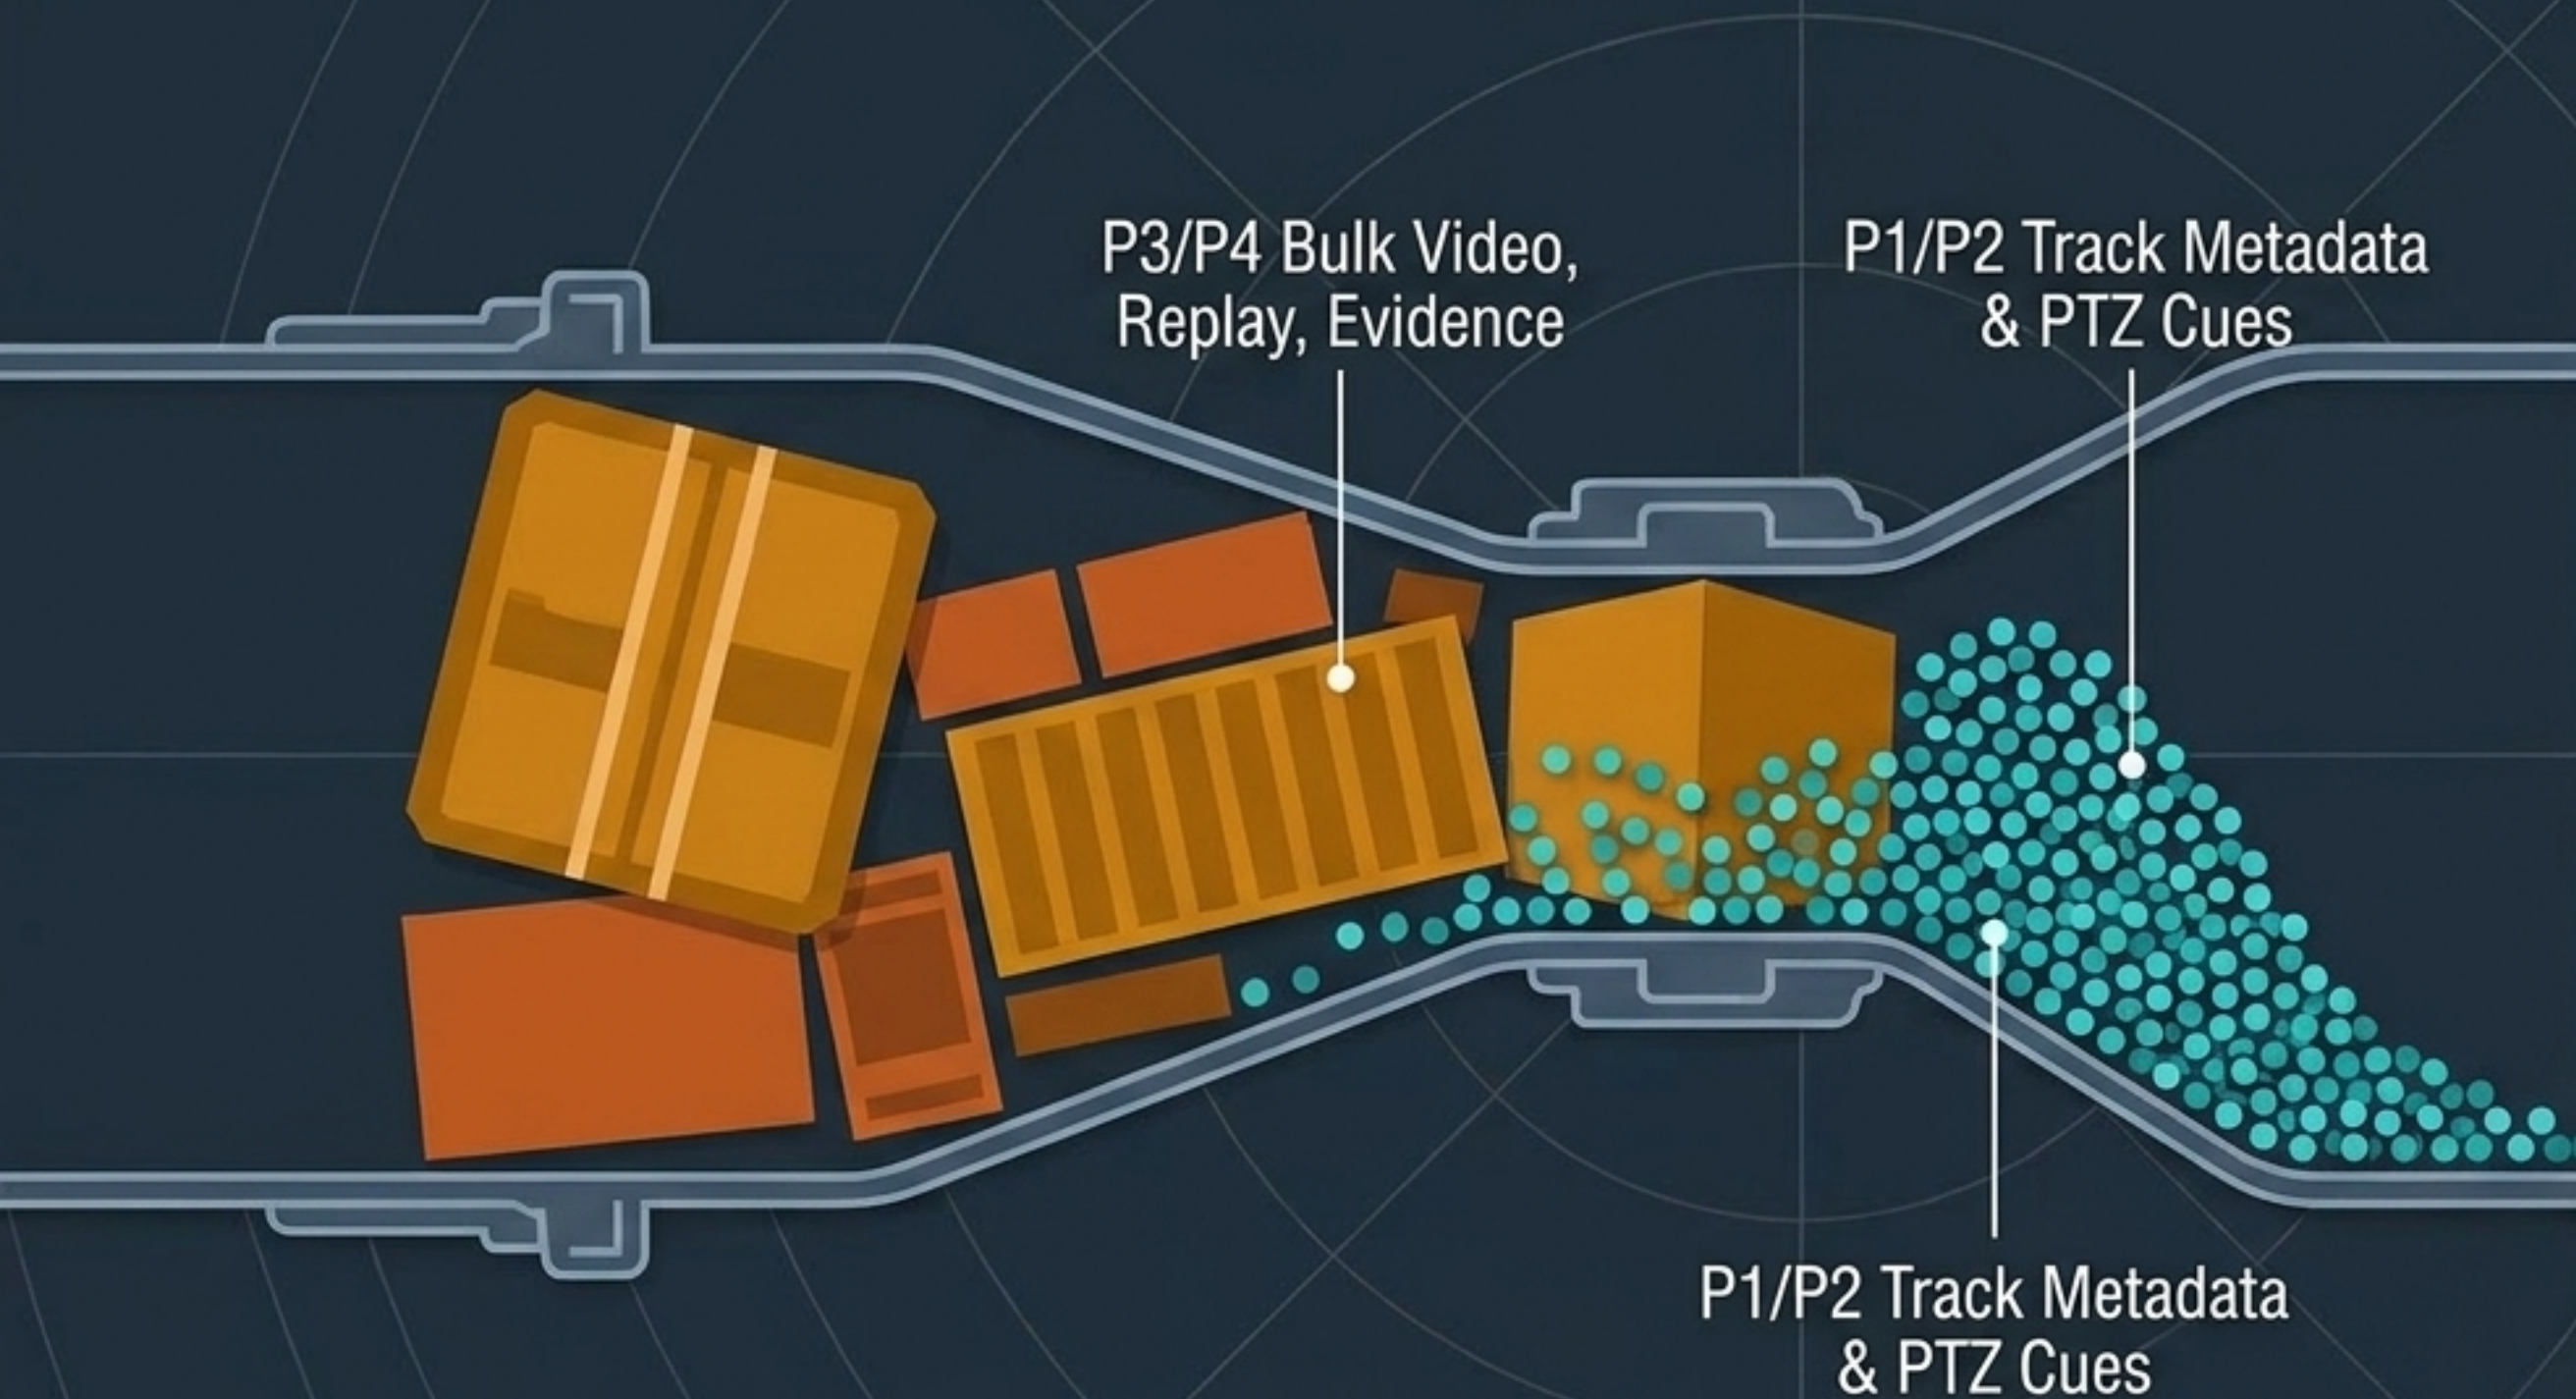


**PTCP-SARAHAI converts high-rate sensor streams into prioritized, covariance-bearing physical-state updates while preserving evidentiary integrity.**

# The Swarm Paradox: More Sensors Create More Blind Spots



# Anatomy of a Failure: Bufferbloat



## Impact Metrics


### Stale Cue-to-Slew

PTZ cameras chase old coordinates.




### Widened Covariance

Track uncertainty explodes due to delayed updates.




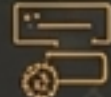

### Operator Overload

Wall of video freezes or lags at the critical moment.






# The Paradigm Shift: From Video Management to Physical AI

## Conventional VMS

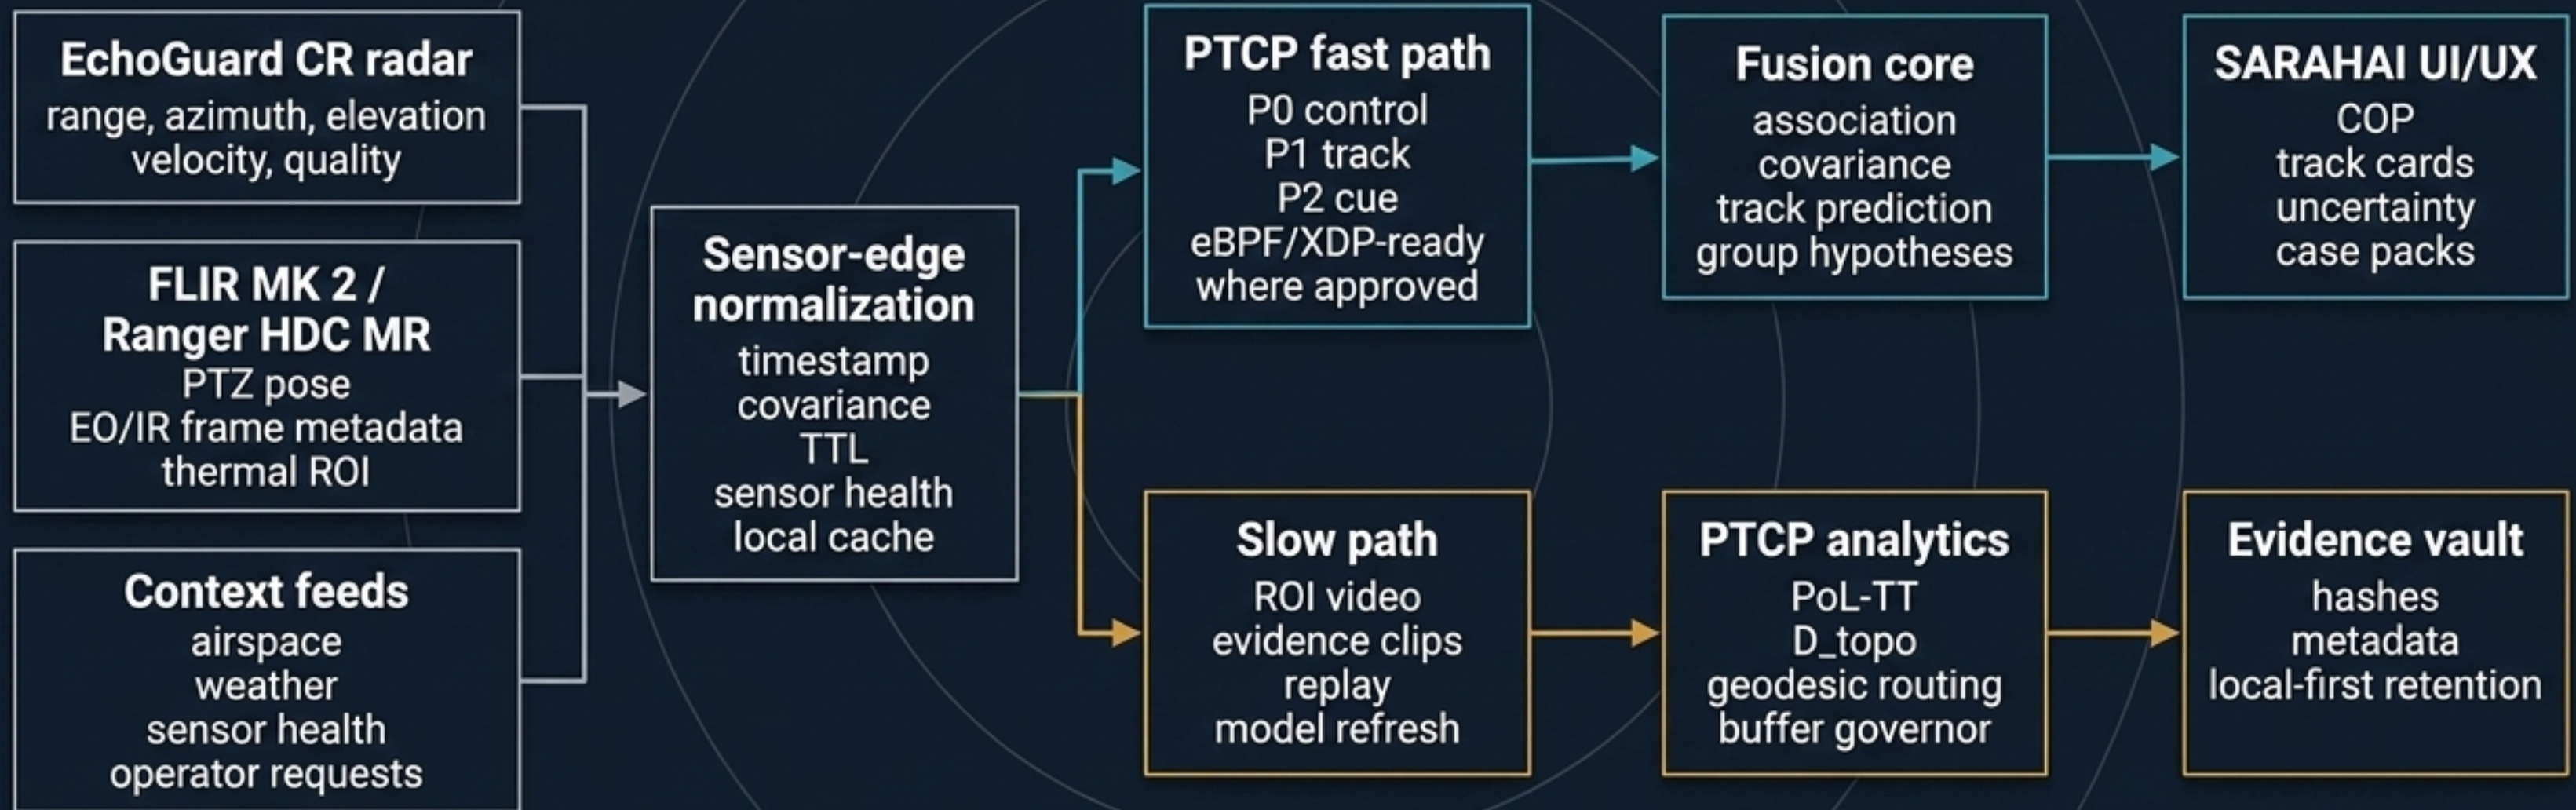
1	<b>Focus</b> Optimizes for showing every possible frame to every possible operator.	
2	<b>Method</b> Best-effort transport; camera browsing.	
3	<b>Result in Crisis</b> Human cognitive overload; network saturation; stale data.	

## Physical AI

	<b>Focus</b> Optimizes for maintaining a coherent physical-state estimate under constrained compute, network, and human attention.
	<b>Method</b> Senses, models, predicts, and governs physical-world state.
	<b>Result in Crisis</b> Current, explainable, uncertainty-bearing track cards with actionable sensor cueing.

Physical AI closes the loop between physical sensors, physical tracks, physical network resources, and operator decision workflows.

# PTCP-SARAHAI Architecture: Governing the Dataflow



## Buffer-safety invariant

Predicted P1/P2 track and cue traffic stays below reserved service capacity; P3/P4 bulk traffic is compressed, stored locally, or delayed before queues saturate.

# Architectural Isolation: The Hot Path vs. The Slow Path

## Payload contents:

```
`track_id`, `timestamp`, `position`,  
`velocity`, `covariance`, `classification hint`,  
`TTL`, `reason codes`
```

Time-critical metadata and cueing commands.  
Compact, prioritized, and uncertainty-bearing.

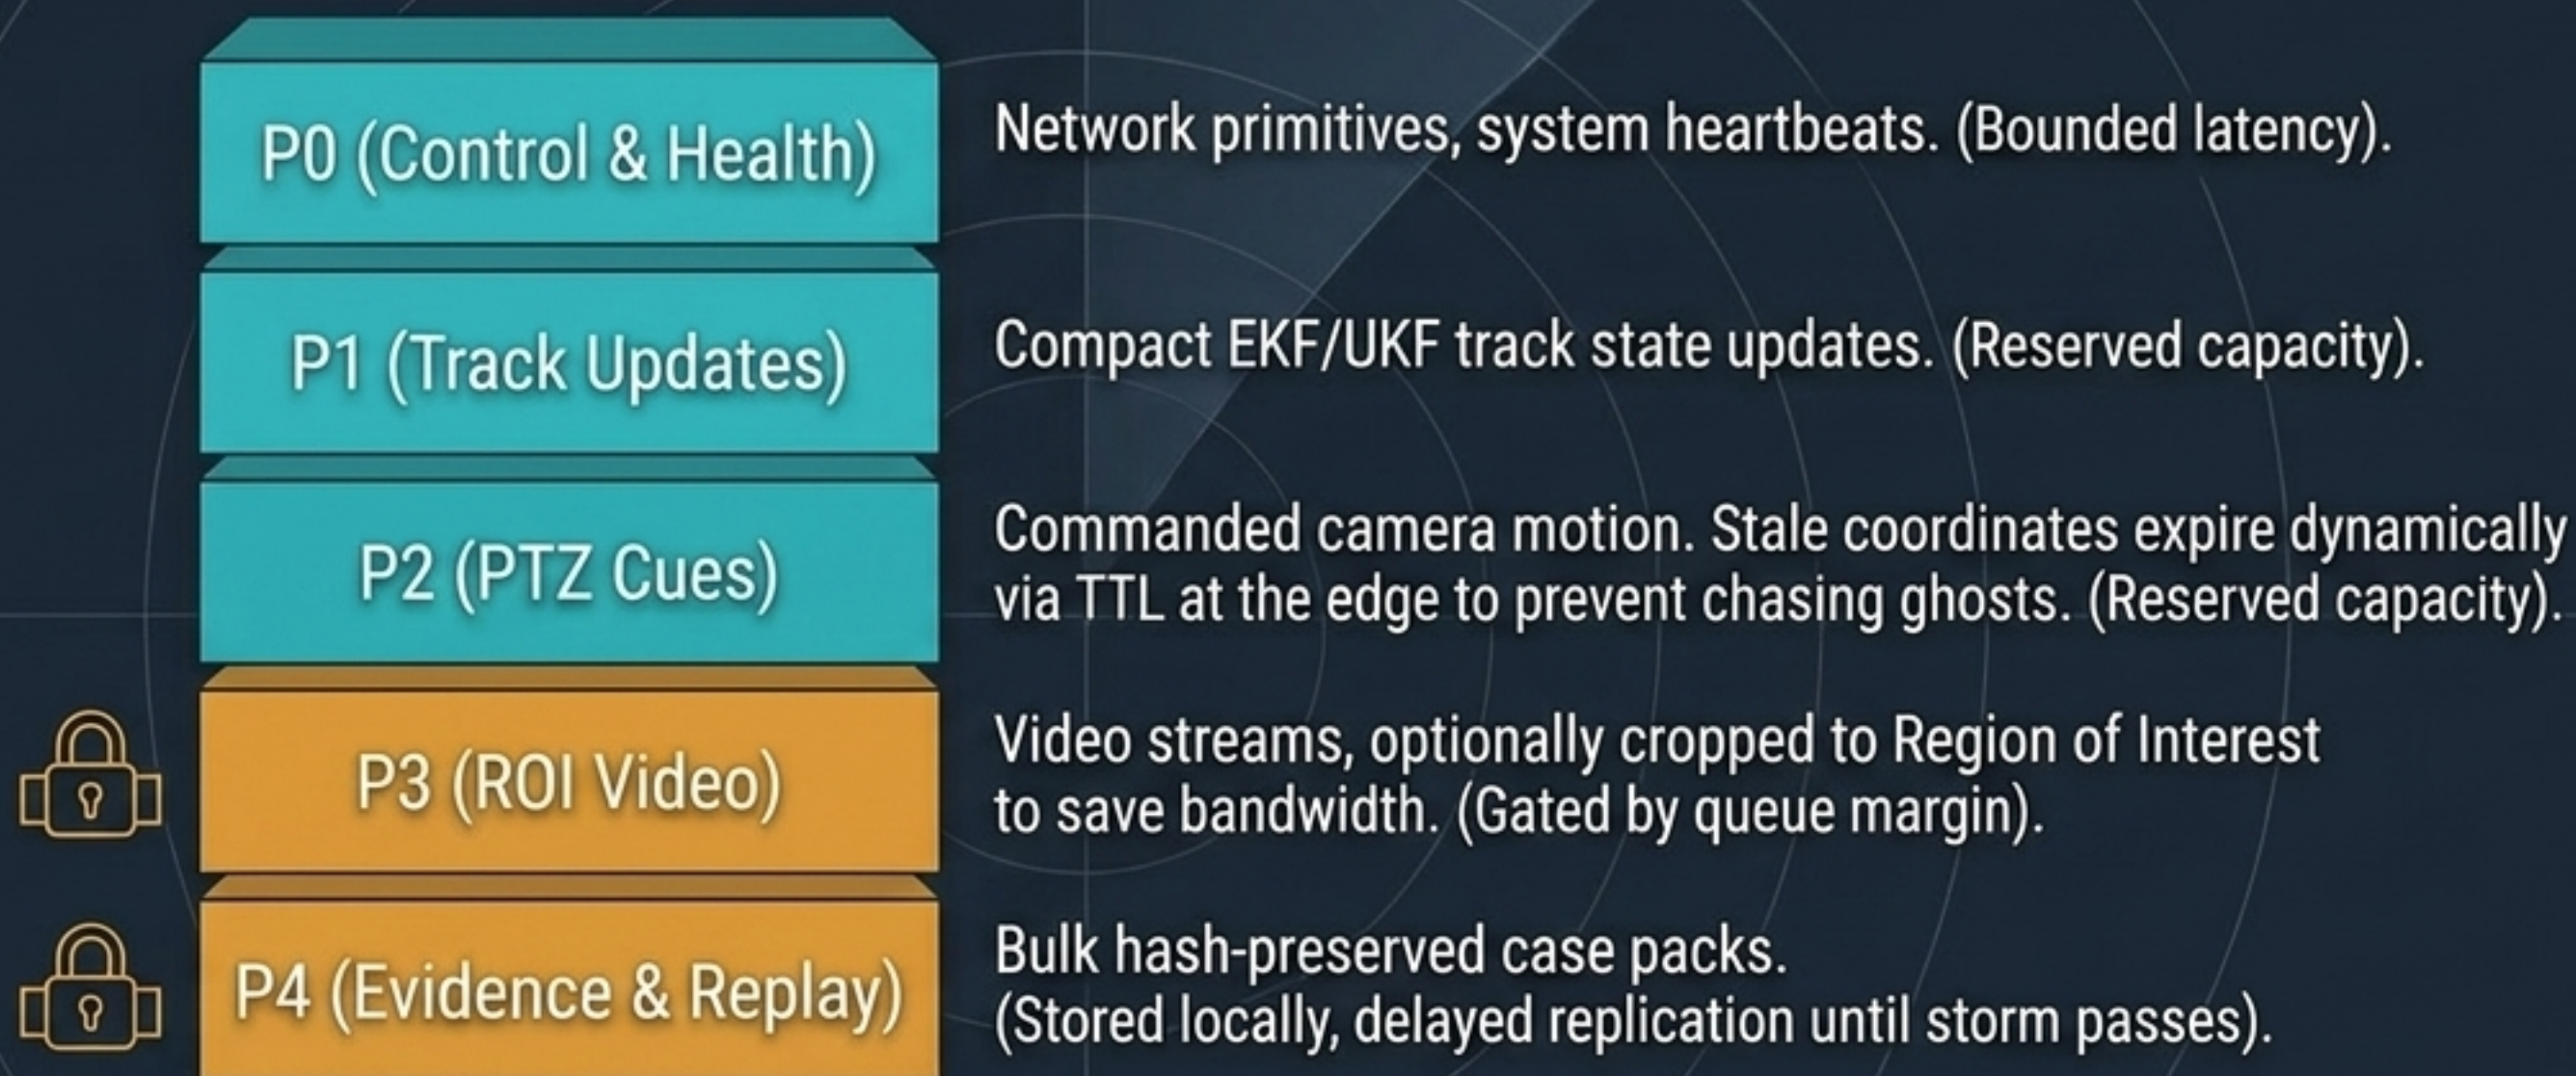
The Hot Path

The Slow Path

Payload contents: High-rate evidence, full-resolution video frames, replay workloads, model refreshes

High-volume payloads. Gated based on real-time network margin to protect the hot path.

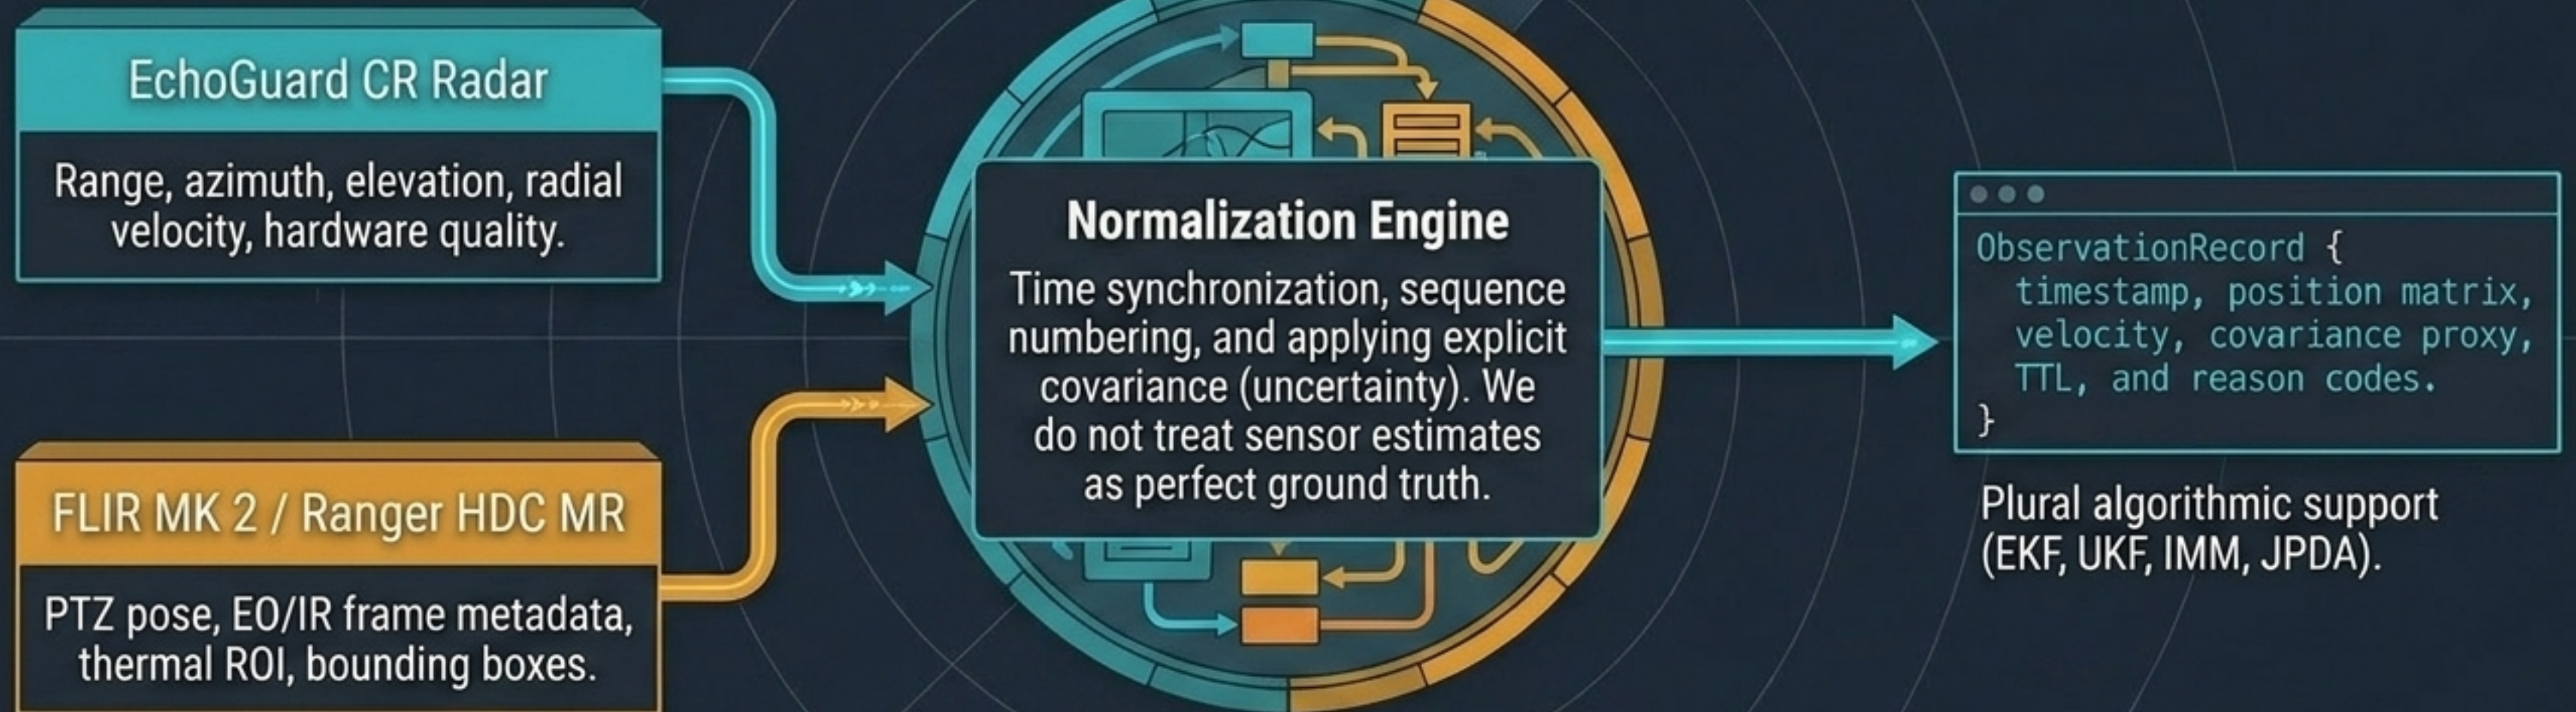
# Buffer Governance & Priority Lanes (P0 – P4)



## Takeaway

By enforcing class-specific invariants, the system guarantees the physics of tracking remain isolated from the physics of data replication.

# Physical Sensing & Kinematic Normalization



# Optical-Flow PTZ Tracking: Isolating the Signal



**Total Optical Flow:**  
Target + background moving together as the camera pans.

**Commanded PTZ Motion:**  
Known pan/tilt mechanical commands sent to the camera.

**Residual Object Motion:**  
Only the true moving target remains highlighted.

By subtracting expected background flow based on known PTZ commands, the system isolates residual object motion. Centroids feed directly into Kalman prediction for continuous, ultra-fast camera cueing.

# Predictive Network Control: PoL-TT & Topology Scoring

## PoL-TT (Pattern-of-Life Tensor Train)

- **Concept:** Compresses high-dimensional telemetry (queue depth, sensor slew rates, GPU pressure) into a bounded-rank tensor representation  $O(d n r^2)$ .
- **Action:** Forecasts network demand and computes an anomaly score (an operational surprise measure, not an automatic threat conclusion).

## $D_{\text{topo}}$ & Geodesic Routing

- **Concept:** Queue growth is treated as physical geometry. A link with rising queue depth literally “lengthens” mathematically.
- **Action:**  $D_{\text{topo}}$  scores topology-native anomalies (curvature gradients, cut-capacity shifts).
- **Result:** Routes metadata dynamically based on expected cost + Conditional Value-at-Risk (CVaR), avoiding links that are good on average but fail during burst loads.

# SYNTHESIS: The Closed-Loop Awareness Engine

## 1. Physical Sensing

Radars & Optics measure the world and assign uncertainty.



## 2. Algorithmic Fusion

Kinematics build tracks, predict motion, and generate cue commands.

## 4. Predictive Network Governance

PTCP shapes network capacity, reserves P1/P2 lanes, and routes around queue depth to ensure Phase 1 data reaches Phase 2 instantly.

## 3. Operator Action

UI renders prioritized tracks and evidence; operator requests insights.

**Core Insight:** These are not separate features. High-rate sensors will crush the network (breaking Phase 4) unless the AI governs the buffer. Algorithmic Fusion (Phase 2) is useless if the network delivers stale data. It is a single, interlocking capability.

# Operator UX: The Common Operating Picture (COP)

Top Telemetry Bar: D\_topo (network geometry anomaly score) and **Buffer State (Safe/Degraded)** sit next to **Active Tracks**, treating network health as a primary operational metric.



The Radar Display:  
Shows covariance  
rings and velocity  
vectors, not just dots.  
Explicitly visualizes  
uncertainty.

Priority Lanes:  
Makes invisible  
network pressure  
visible. Operators see  
exactly how much  
P3/P4 bulk data is  
being gated to  
protect P1/P2 tracks.

# Cognitive Offloading: Track Cards Over Video Walls

## Legacy Video Walls



During a multi-target swarm, human attention is the scarcest resource. Forcing an operator to scan raw video across sectors causes task saturation.

## Governed Track Cards

- 1 Classification & Confidence:**  
T-041 UAS-like object | confidence 0.92  
**Source Fusion:** sources: radar + EO/IR
- 2 Reason Codes & State:**  
Stable covariance; PTZ cue active.
- 3 Group Hypotheses:**  
Ingress and egress split under review.  
Operators evaluate intent, not just blips.

**Result:** The operator acts as an executive decision-maker reviewing a governed physical-state estimate, rather than a raw data processor.

# The Local-First Evidence Vault

**The Conflict:** High-fidelity recording and cloud replication consume the exact network capacity needed for live tracking during an event.

## The PTCP-SARAHAI Solution



### Local-First Retention

Raw clips, track histories, and model states are immediately stored on local NVMe drives.



### Immutable Case-Packs

Evidence is cryptographically hashed to ensure forensic integrity and chain-of-custody (crucial for OSI/investigative analysts).

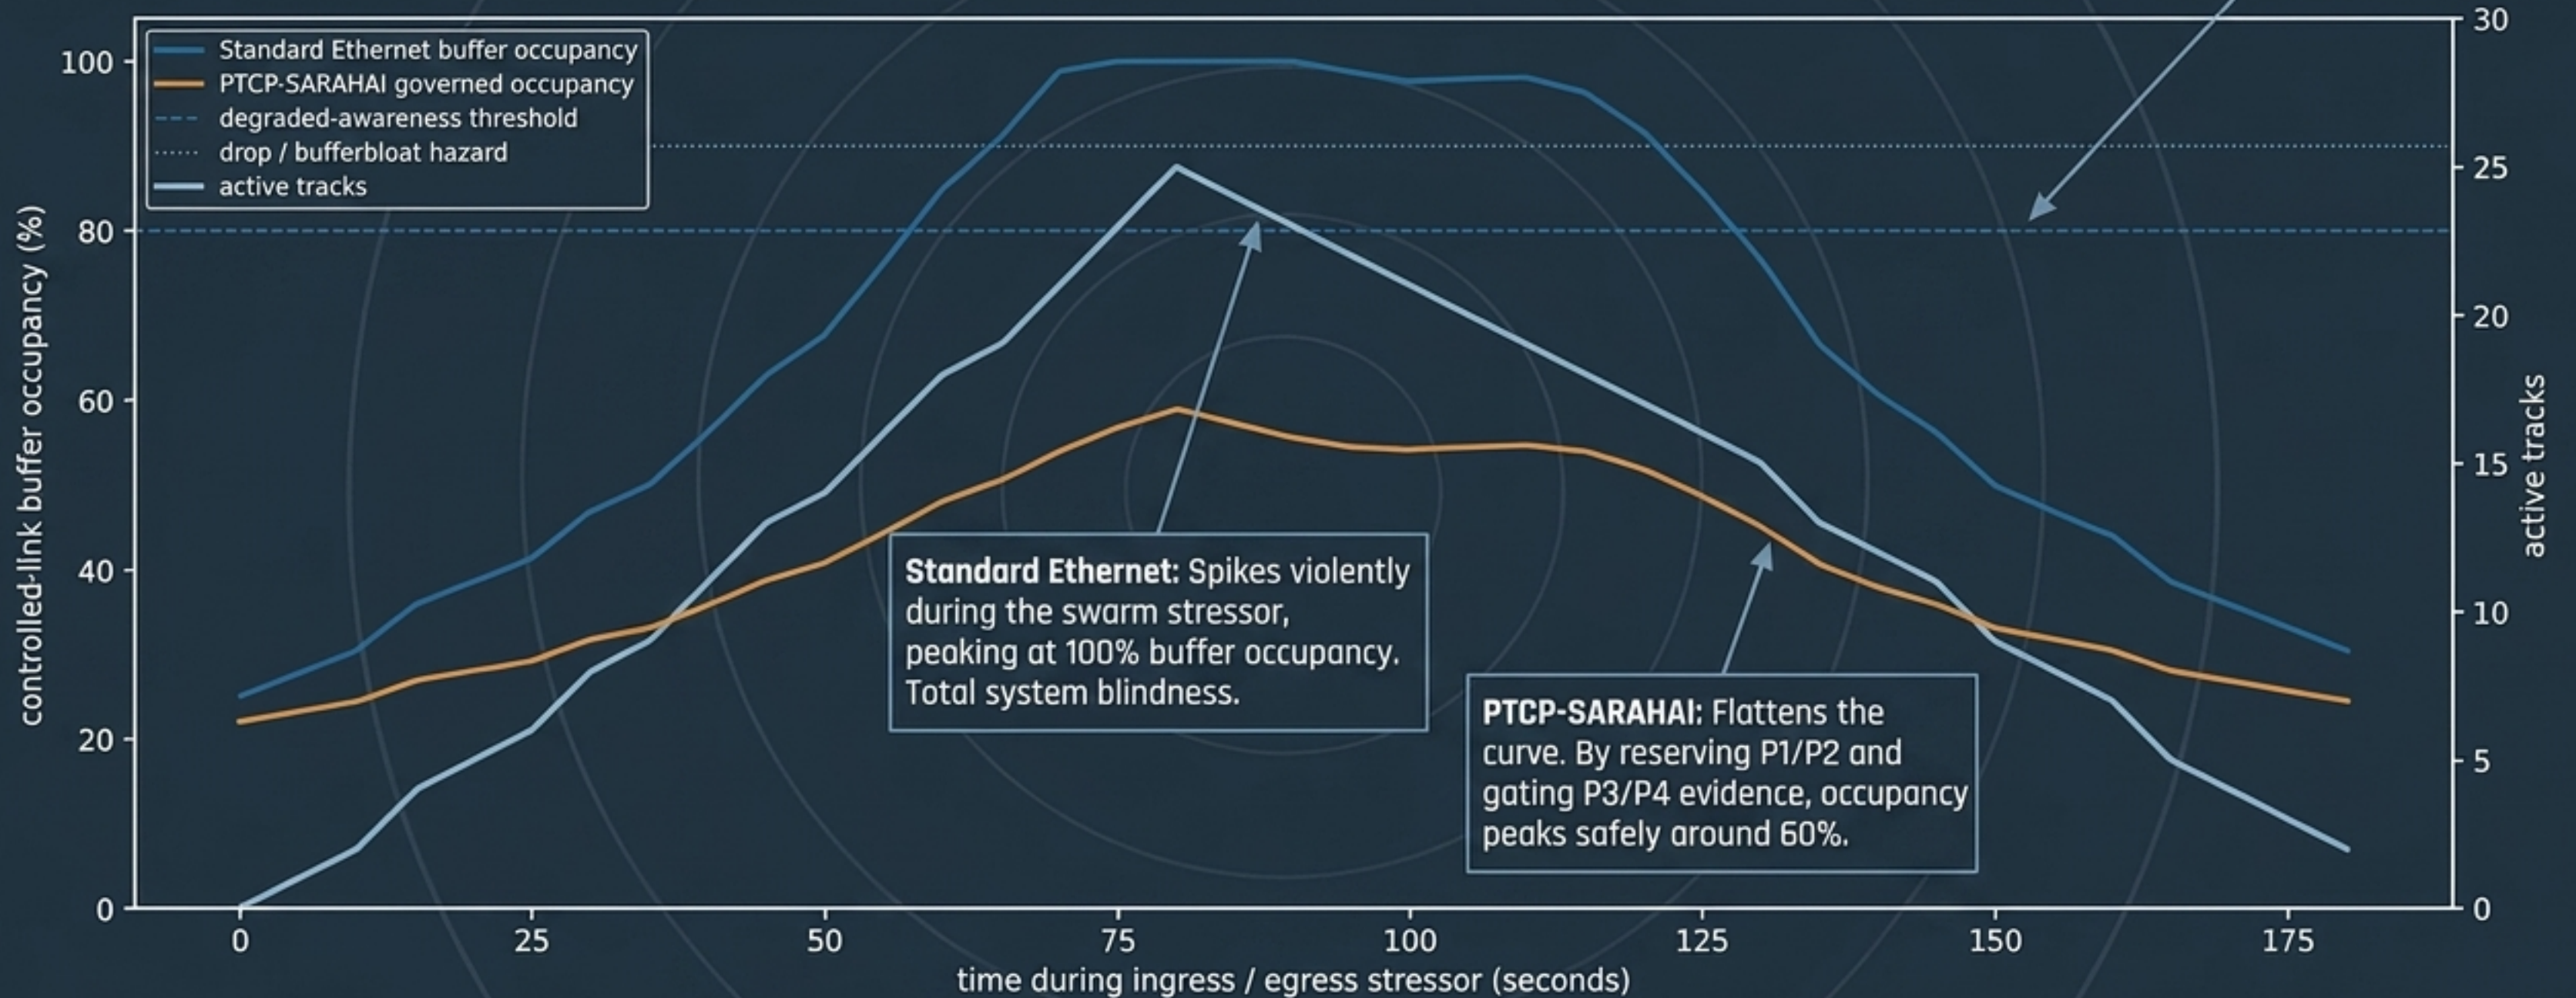


### Queue-Margin Replication

The system delays network replication. It waits until the swarm event ends and the network queue margin returns to normal before uploading the heavy case-packs.

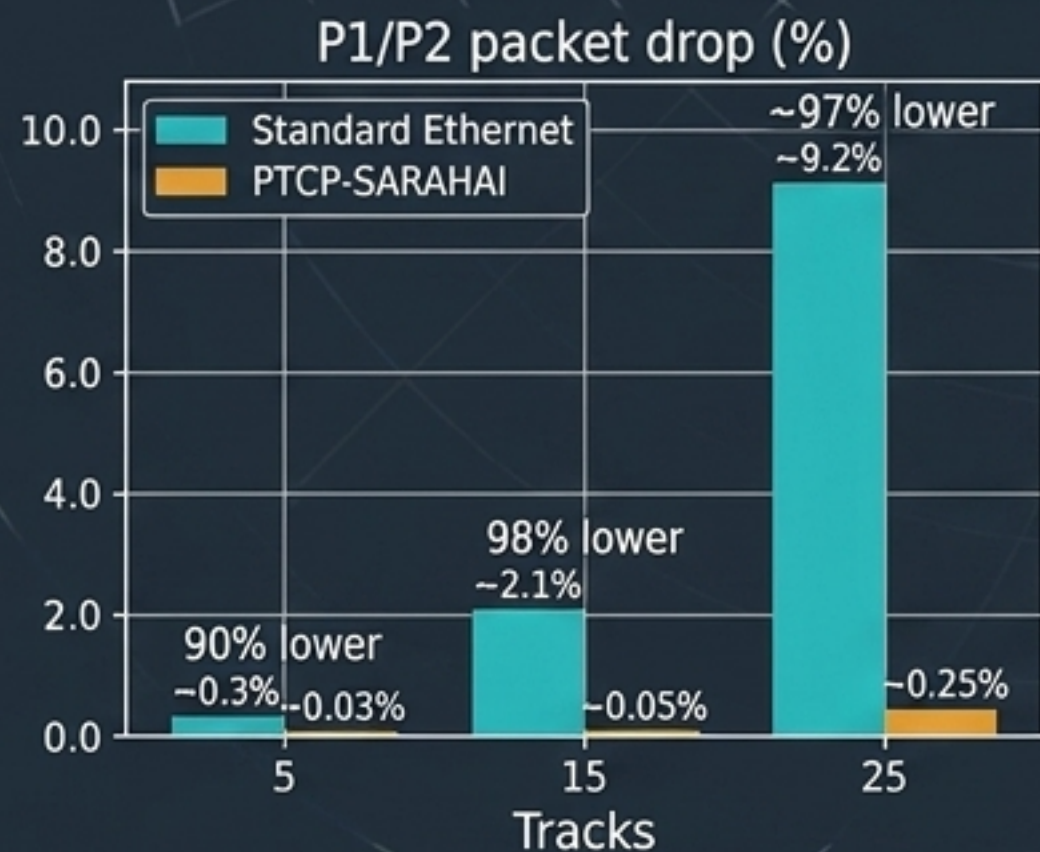
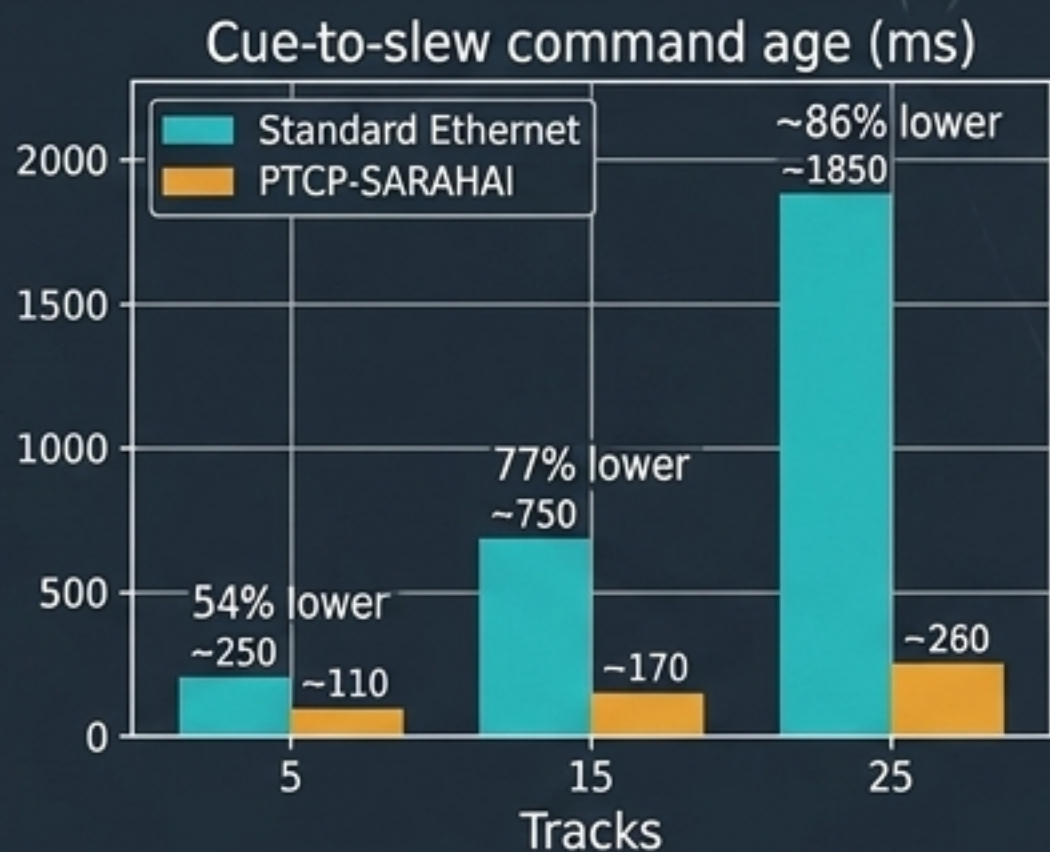
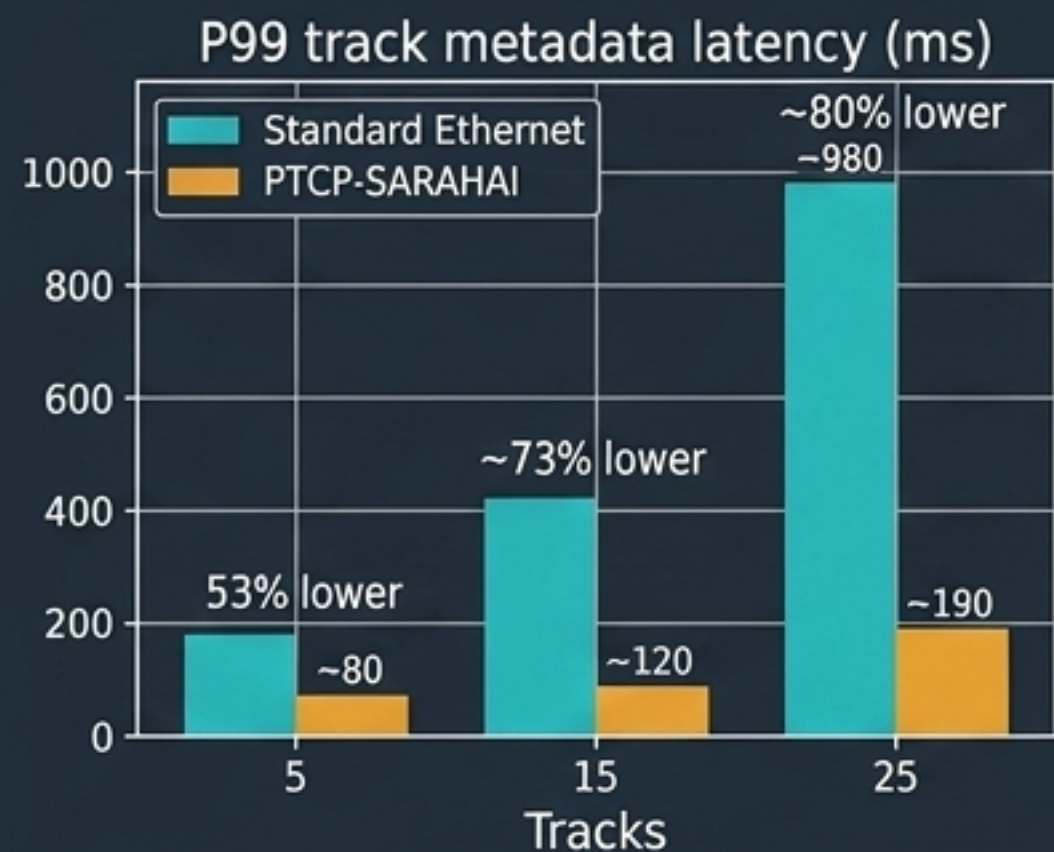
Live awareness is never sacrificed for evidence completion.

# The Stress Test: Buffer Occupancy Under Swarm Load



**Conclusion:** PTCP acts as a physical pressure valve, guaranteeing the tracking network survives the very event it was built to monitor.

# Scaling the Envelope: P99 Latency & Packet Drop



## Key Benchmark Targets (at 25 modeled tracks):

### P99 Track Metadata Latency:

Standard Ethernet spikes to ~1000ms.

PTCP-SARAHAI remains flat at < 200ms (80% lower). Protects track continuity.

### P1/P2 Packet Drop: Standard Ethernet hits

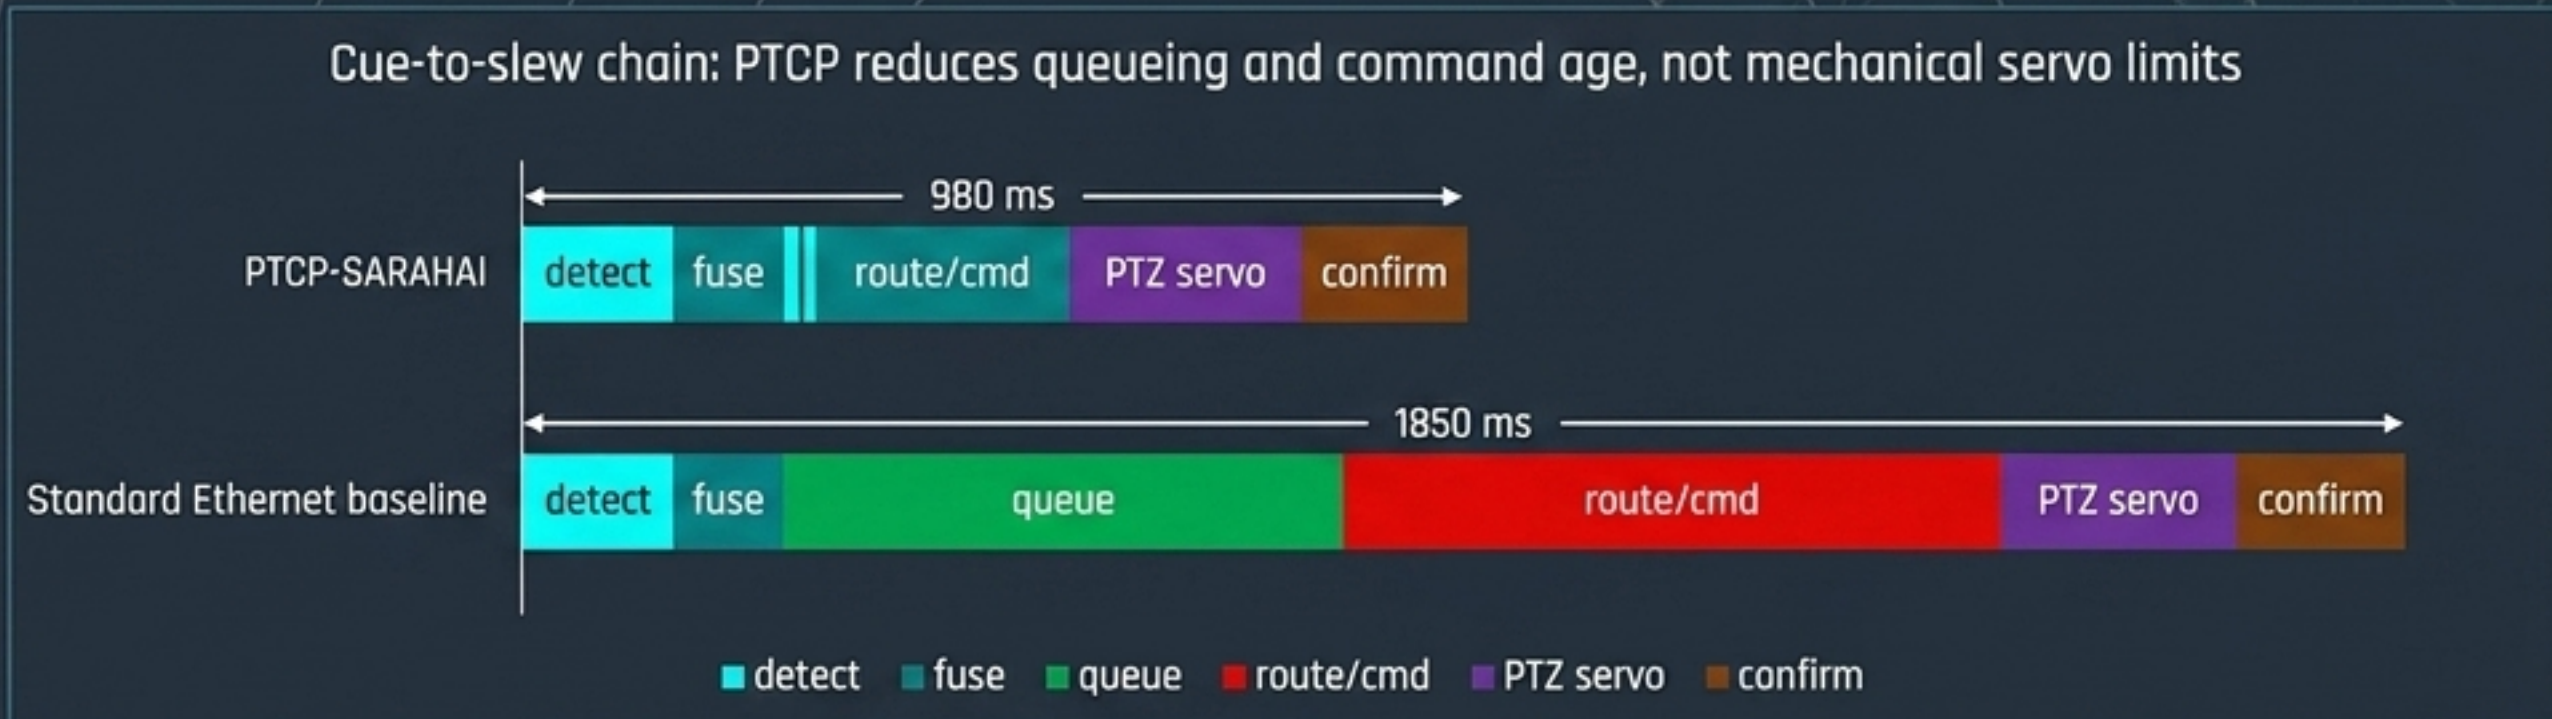
~9% drop rate (fatal for state estimation).

PTCP-SARAHAI maintains <= 0.25% drop rate.

# Protecting the Cue-to-Slew Timeline

PTCP-SARAHAI Total: 980 ms.  
The green queue delay is virtually eliminated.

Standard Baseline Total: 1850 ms.  
Massive chunks of time wasted in green "Queue" and red "Route/Cmd" delays.



**Honest Engineering**  
PTZ servo time and sensor refresh times remain identical. We do not claim to eliminate mechanical limits. We eliminate the artificial software/network drag that makes those mechanical limits worse.

# The Architectural Divide: VMS vs. Physical AI

Dimension	Standard Baseline	PTCP-SARAHAI
Transport Priority	Best-effort bulk video.	Hot-path (metadata) vs Slow-path (evidence).
Cue-to-Slew Mechanics	Prone to stale coordinates and chasing ghosts.	TTL-gated; guaranteed sub-300ms command age at 25 tracks.
Swarm Behavior	Network saturation; cascading failure.	Deterministic degradation; non-critical data shed automatically.
Evidence Retention	Streams everything to cloud immediately.	Local-first hash vaults; replicated only when safe.
Operator UX	Exhausting video walls; high cognitive load.	Prioritized track cards; explicit uncertainty (covariance).

The same passive Physical AI stack applies wherever high-rate sensors, scarce attention, and overloaded networks converge.

### Airports & heliports

runway perimeter, wildlife, drone alerts, camera cueing

### Ports & maritime terminals

quay security, small-craft correlation, container-yard monitoring

### Energy & utilities

substations, pipelines, solar/wind farms, water assets

### Data centers & campuses

perimeter, rooftops, parking, yard robotics, evidence workflows

### Industrial yards & mining

vehicle/UAS safety, high-wall monitoring, autonomous equipment zones

### Disaster response

wildfire perimeter, flood watch, temporary airspace awareness

### Rail & logistics

right-of-way monitoring, yard ingress/egress, thermal/visual confirmation

### Stadiums & public venues

non-kinetic perimeter awareness, operator triage, incident evidence


Commercial value: faster triage, lower backhaul cost, fewer stale alerts, more complete evidence, and a benchmarkable upgrade path.

# Commercial Value Translation



## Lower Backhaul Costs

By prioritizing metadata and keeping heavy evidence local, remote sites (like pipelines/substations) can operate over LTE/SATCOM without continuous high-bandwidth video streaming costs.



## Faster Triage

Operators review 3 track cards instead of 30 camera feeds, radically reducing response times during critical perimeter breaches.



## Fewer False Alerts

Radars and optical flow combined with PoL-TT anomalies filter out environmental noise, eliminating "alert fatigue."



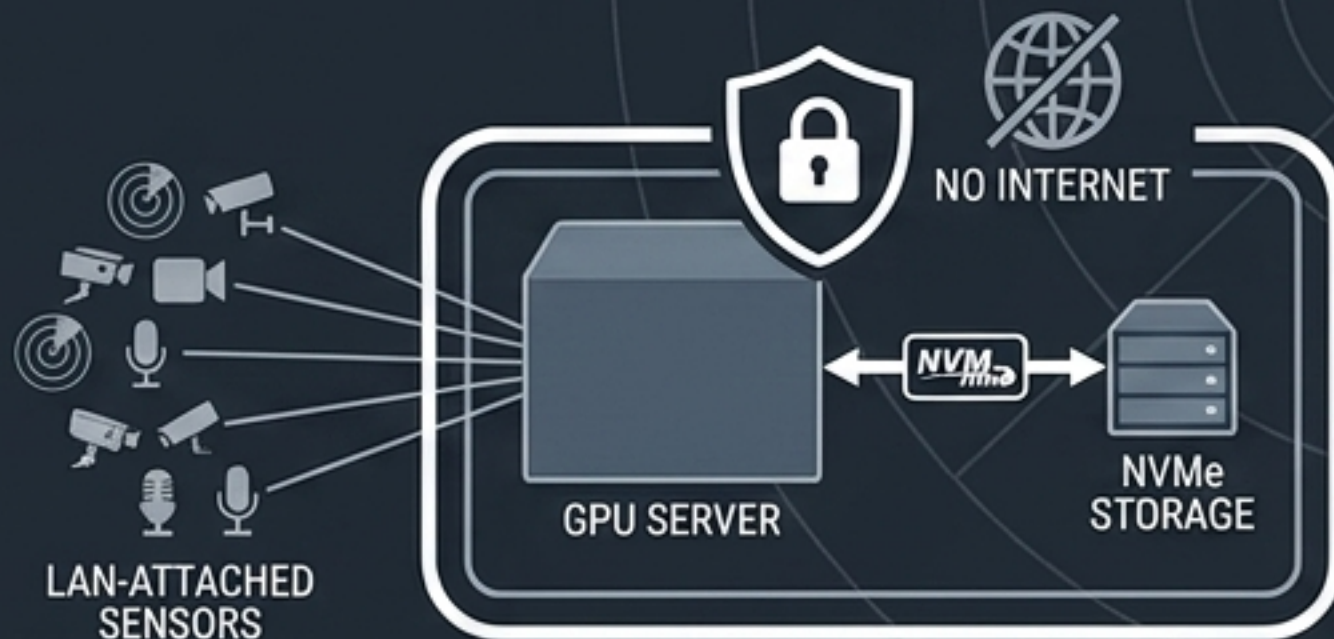
## Immutable Liability Protection

Hashes and local case packs provide perfectly preserved, unalterable chain-of-custody evidence for compliance, insurance, and law enforcement.

# Modular Deployment Architectures

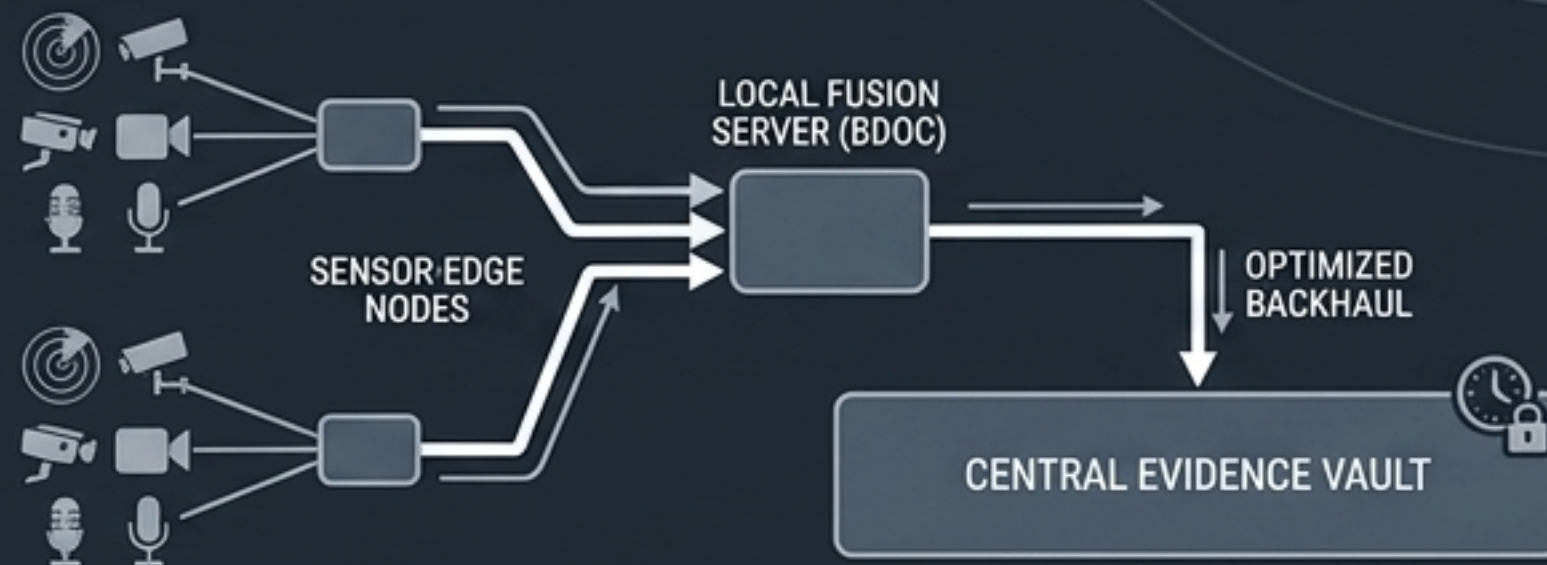
## 1. Single High-End Server (Air-Gapped)

- Setup: GPU server, NVMe storage, LAN-attached sensors.
- Value: Ultimate security. Perfect for critical infrastructure, pilots, and classified sites needing total isolation.



## 2. Distributed Edge

- Setup: Sensor-edge nodes at radar clusters, local fusion server in BDOC, central evidence vault.
- Value: Scales across massive perimeters (airports/ports) while aggressively minimizing backhaul.



## 3. Hybrid / Cloud

- Setup: Local hot path at the edge; cloud handles offline model training, fleet analytics, and long-term evidence search.
- Value: Enterprise scalability and cross-site analytics, bounded by strict policy approvals.



# Governance, Cybersecurity, and Safety Boundaries

## What It IS

- Passive awareness, coordinate-quality fusion, evidence handling, sensor cueing recommendation, buffer governance.
- Governed by: Role-Based Access (RBAC), signed configurations, mTLS trusted links, append-only audit logs.



## What It IS NOT

- No autonomous kinetic engagement.
- No jamming or electronic attack.
- No cyber effects or offensive UAS operations.

The "Dry-Run" Default: Systems deploy with a Champion/Challenger model. New models learn in the background without destabilizing live command operations.

# The Validation Framework: Acceptance Metrics

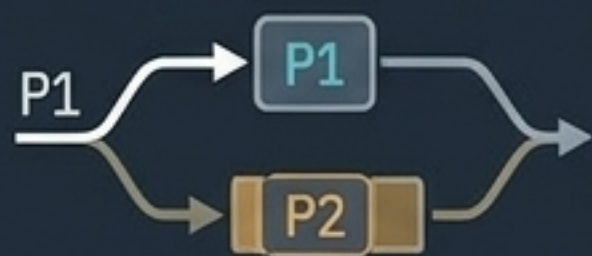
**The Challenge:** We compare PTCP against a standard VMS baseline under controlled swarm load, not nominal demo conditions.



**Hypothesis 1**  
(Predictive Overload)



Queue crosses hard threshold before control action?  
**Failure condition.**



**Hypothesis 2**  
(P1/P2 Reservation)



Inject bulk evidence storm. Does P1 drop or P2 latency exceed envelope? **Failure condition.**



**Hypothesis 3**  
(Evidence Integrity)



Does PoL-TT anomaly quality or confirmation recall degrade beyond accepted threshold? **Failure condition.**



The **Guarantee**: If envelope assumptions fail, the UI visibly warns operators of degraded-state rather than silently displaying stale tracks.

# From Best-Effort Transport to Governed Perception

**Sense:** Extracting physical state with explicit uncertainty.

**Predict:** Forecasting network physics to prevent bufferbloat.

**Govern:** Separating hot-path tracks from slow-path evidence.



PTCP-SARAHAI transforms multi-sensor environments. It guarantees that when the environment and the network are stressed at the exact same time, physical-state awareness remains continuous, prioritized, and undeniable.

Architecture designed for validation. Ready for hardware-attested field deployment.