

From Telemetry Inflation to Decisive Observability

Short strategic brief (board version)

Canonical long paper: [from-telemetry-inflation-to-decisive-observability.md](#)

Illustrative appendix (USD): [appendix-illustrative-cost-model.md](#)

Source references are intentionally centralized in the long paper only (see its [References](#) section).

Third-party product and trademark notice: SolarWinds, Prometheus, Grafana Mimir, Monarch, Google, and other third-party names are used solely for nominative, descriptive, and comparative reference purposes. All trademarks and trade names remain the property of their respective owners. No sponsorship, endorsement, certification, affiliation, or approval by any third-party vendor is implied.

Comparative methodology notice: The quantitative comparisons in this executive brief are illustrative regime-level calculations derived from public documentation, published papers, and explicitly stated modeling assumptions in the canonical long paper. They are not vendor-certified sizing guidance, product benchmarks, procurement recommendations, or claims of functional equivalence.

ESC is a **deviation-triggered explanatory discipline** added upstream of existing observability systems.

It does not replace logs, metrics, packet capture, or time-series infrastructure.

It structures **when deeper explanatory proof should exist**.

It is built on two domain-agnostic invariants — **occurrence** and **observer-time reference** — which shift the dominant scaling variable from raw telemetry volume toward structured occurrence activity and proof-activation timing.

This brief uses an explicit **500,000 supervised interfaces** planning perimeter. Named products and systems are retained as public reference anchors because their documentation or publications provide concrete numeric inputs used in the long paper.

The primary cost driver is not telemetry volume alone.

It is **time-to-explanatory-proof** during operational incidents.

The strategic risk is not telemetry absence.

It is explanatory delay disguised as observability maturity.

ESC reframes observability from “collect more” to “activate proof sooner.”

1.1 The two invariants behind ESC

ESC does not treat telemetry as the primary observable. It treats **structured occurrence behavior** as the primary operational signal.

The discipline rests on two invariants.

Occurrence

An occurrence is a bounded indication that something happened.

It does not require payload content, semantic interpretation, protocol decoding, log parsing, trace reconstruction, or measurement of system state.

For board-level reading: occurrence is the unit that replaces raw telemetry volume as the primary scaling variable in the ESC model.

Observer-time reference

Observer-time reference is the temporal frame in which occurrences are structured for comparison.

ESC does not require the observed system to provide a complete, globally synchronized, producer-authoritative timeline. Instead, occurrences are structured at the point of observation into ticks, cycles, or comparable temporal units.

This allows deviation-relevant behavior to be compared without treating permanent high-fidelity capture or global clock authority as the default operational posture.

Compact reading:

```
Node A: occurrence observed within tick T
Node B: occurrence observed within tick T

ESC/SOBT: compare structured occurrence behavior under observer-time reference
```

The point is not that cost becomes independent of activity. It does not.

The point is that ESC shifts the decision problem away from continuous interpretation of raw telemetry volume and toward bounded occurrence activity, observer-time structuring, and governed proof activation.

1. CAPEX/OPEX decision frame

CAPEX/OPEX BOARD SNAPSHOT		
Decision axis	Legacy default	ESC-oriented discipline
CAPEX policy	Permanent high-fidelity capacity as insurance	Bounded continuous awareness + selective deeper activation
OPEX policy	Bridge-heavy retrospective correlation	Ambiguity-window reduction via earlier explanatory access
SLA posture	Service credits as primary remediation	Faster explanatory access as operational confidence layer
Operating logic	Collect broadly, explain later	Detect deviation, then deepen evidence
Governance requirement	Retention and coverage controls	Activation precision, drift control, backpressure controls

Board reading:

- not a replacement for compliance-grade logs or flow records;
- not a generic archive optimization program;
- not a storage story dressed as innovation.

The CAPEX/OPEX shift is grounded in the same two invariants used in the long paper: **occurrence** provides the bounded operational unit, and **observer-time reference** provides the comparative frame for deciding when proof activation is justified.

2. Raw infrastructure footprint contrast (500,000 interfaces)

This section mirrors the long paper's infrastructure-footprint comparison and keeps a raw baseline posture:


- no HA pairs;
- no spare capacity;
- no multi-region replication;
- no query-acceleration add-ons.

The values below are not vendor-certified bills of materials. They are author-side, regime-level modeled estimates derived from public references and stated assumptions in the long paper.

RAW INFRASTRUCTURE FOOTPRINT AT 500,000 INTERFACES					
Observation regime	CPU cores	RAM	Storage	Footprint	Dominant scaling pressure
SolarWinds class	166-1,334	666 GB-5.3 TB	2.1-20.9 TB	10.4-41.7 polling engines	Polling-engine multiplication and SQL/backend growth
Prometheus / Mimir class	64-256	256 GB-1 TB	1-5 TB / 30 days	Multi-role metrics engine	Scrape fan-out, ingest, index, query, rules, compaction
Monarch-class fabric	Not reducible from public data	Close to 1 PB compressed memory	1-10 PB distributed storage	Planet-scale telemetry fabric	Fleet-class telemetry operation
ESC discipline	16-24	32-64 GB	50-500 GB	1 general-purpose server	Deviation history and selected metadata evidence

Figure G0 - Raw infrastructure footprint contrast at 500,000 interfaces

Raw baseline visual only (no HA, no spare capacity, no multi-region replication).



Board interpretation:

- under the stated assumptions, continuous telemetry regimes can convert observability into infrastructure multiplication;
- ESC keeps baseline capacity compact and shifts pressure toward activation governance quality;
- this compactness is modeled from structured occurrence behavior, not from raw telemetry volume as the primary scaling variable.

This is an architectural and economic comparison, not a product-ranking claim.

3. Capacity snapshot (250 ms planning basis)

Sizing inputs:

- **N** = total interfaces
- **A** = active interface ratio
- **O** = average occurrences per active interface per tick

At 250 ms tick (4 ticks/sec):

$$\text{updates/sec} = N \times A \times O \times 4$$

Here, **O** is deliberately an occurrence count, not a telemetry-volume proxy. The tick term reflects observer-time reference: occurrence behavior is structured in a comparative observation frame.

For the 500,000-interface reference model:

500,000 INTERFACES - 250 MS TICK				
Scenario	A	O	Updates/sec	Traffic
Nominal	5%	1	100,000	8.0-19.2 Mb/s
Elevated	10%	2	400,000	32.0-76.8 Mb/s
Stress	50%	2	2,000,000	160.0-384.0 Mb/s
Theoretical hard cap	100%	15	30,000,000	2.4-5.76 Gb/s

Interpretation:

- nominal/elevated are routine planning zones;
- stress is disturbance-grade planning;
- hard cap is a resilience boundary, not normal operation.

4. What 24 cores means

At 24 cores (3 GHz/core reference), update-loop capacity depends on implementation quality:

- **standard production implementation path (~5k cycles/update):** ~14.4M updates/sec
- **high-performance implementation path (~2k cycles/update):** ~36M updates/sec

So 24 cores is an **operator-grade planning target**, not a guarantee of permanent hard-cap saturation coverage.

Illustrative operator envelope:

- **CPU per node:** 16-24 cores
 - **RAM per node:** 32-64 GB
 - **Topology:** 1-3 general-purpose servers, depending on HA policy
-

5. Economic and leadership implications

ESC is framed as a **bounded, deterministic discipline** at model level, with explicit implementation accountability.

Economic reading:

- bridge OPEX scales with: `concurrency x burdened hourly cost x hours to trusted causal story`;
- CAPEX pressure can come from permanent high-fidelity insurance defaults.

ESC targets both by using structured occurrence behavior to govern when richer proof should exist, rather than making continuous high-fidelity collection the default path to confidence.

Leadership implication:

the organization should not have to buy the same crisis twice - once as infrastructure, and once as calendar time.

Strategic close:

The next advantage is not more telemetry.
It is decisive evidence before decisions become expensive.

Executive takeaway

- Treat proof timing as the missing economic layer in observability strategy.
- Keep existing observability platforms; add ESC as an upstream decision layer.
- The planning perimeter addressed here is **500,000 supervised interfaces**.
- Use explicit assumptions (A, O, tick, disturbance, HA), never single-number slogans.
- Treat **O** as occurrence activity, not raw telemetry volume.
- Treat proof timing as a first-order CAPEX/OPEX variable.
- Govern activation quality (false positives, false negatives, drift, backpressure) as an engineering discipline.

Third-party reference notice

SolarWinds, Prometheus, Grafana Mimir, Monarch, Google, and any other third-party names referenced in this brief are trademarks, trade names, project names, or publication names of their respective owners.

They are used solely to identify public documentation, public capacity guidance, or published research sources used as numeric anchors in the canonical long paper.

No sponsorship, endorsement, affiliation, certification, approval, or commercial relationship is implied.

This executive brief does not claim that ESC is a drop-in replacement for any named product or service, nor does it claim that any named product performs according to the modeled figures in all deployments.

All comparisons are illustrative, architecture-level, and assumption-bound.

License Notice

Copyright (c) 2026 Alain Degreffe.

Except where otherwise noted, this document is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International License (CC BY-ND 4.0).

License deed:

<https://creativecommons.org/licenses/by-nd/4.0/>

Full legal code:

<https://creativecommons.org/licenses/by-nd/4.0/legalcode>

Patent notice:

No patent rights are granted under this license or by this publication.