Appendix: Illustrative Cost Model

Companion to: *From Telemetry Inflation to Decisive Observability* (main strategic paper)

> **Third-party product and trademark notice:** FMADIO, FS, Keysight, Microsoft, AWS, Grafana, Mimir, Sentinel, SolarWinds, Prometheus, Monarch, Google, and all other third-party names referenced in this appendix are trademarks, trade names, project names, or publication names of their respective owners. They are used solely for nominative, descriptive, and comparative reference purposes.
>
> **Methodology notice:** All quantitative ranges in this appendix are illustrative, assumption-bound sensitivity examples derived from public pricing pages, public technical documentation, public research publications, and explicitly stated modeling assumptions. They are not vendor quotations, deployment guarantees, benchmark certifications, procurement recommendations, or claims of product equivalence.

This appendix is a **sensitivity exercise, not a forecast**. It expands the condensed economics note in the main paper with editable placeholders, tables, and meeting prompts. **All monetary amounts are USD.**

Named products and vendors are intentionally retained because they provide publicly accessible pricing, sizing, throughput, or architectural anchors used for the illustrative calculations.

This appendix intentionally starts with a **small operating model**. The objective is not to model hyperscale spend directly, but to expose the **economic unit of delayed explanatory proof** that larger organizations can multiply against their own incident volume, staffing model, and contractual exposure.

Terminology: In the main paper, *high-fidelity proof*, *decisive evidence*, and *richer evidence* refer to the same outcome: trusted explanatory proof delivered on operational time. This appendix uses **proof** and **evidence** in that same sense.

Illustrative Cost Model (Sensitivity Exercise, Not a Forecast)

Leadership conversations stall when observability remains only qualitative. The figures below are a **deliberately simple calculator**: they exist to show **which levers move the budget**—bridge duration, headcount on a bridge, and incremental "insurance-shaped" infrastructure—not to assert audited savings.

The model should be read as a **unit-cost illustration of ambiguity**, not as a full financial statement.

No table in this appendix should be interpreted as a vendor-certified bill of materials, negotiated enterprise quote, or guaranteed deployment cost.

Replace every input with your own (region, labor mix, incident taxonomy, and risk appetite).

Shared assumptions (editable)

BASELINE INPUTS FOR SENSITIVITY RUN		
Input	Example value	What it stands for
Blended fully loaded cost per incident hour	**$200 / hr**	Mixed L2/L3, duty managers, vendor time (order-of-magnitude)
Average concurrency on a major bridge	**8 people**	Parallel streams (NOC, app, network, storage, security, leadership)
Major "ambiguous transient" incidents per year	**12**	Severity-1/2 events where causality is disputed for hours

These inputs are not recommendations. They are **placeholders** so finance can run a one-page sanity check.

They represent a **modest operating entity** (tens of engineers, hundreds of services), not a hyperscale deployment.

OPEX: bridge time dominates

Bridge labor cost per incident is roughly:

> **concurrency × blended hourly cost × hours to trusted explanatory story**

Using the placeholders:

BRIDGE-TIME COST COMPARISON		
Scenario	Hours to trusted story (example)	Labor-only bridge cost (example)
A - Late proof (re-capture, widen windows, debate artifacts)	6 hr	8 x $200 x 6 = **$9,600**
B - Earlier decisive proof (same staffing, shorter ambiguity window)	2 hr	8 x $200 x 2 = **$3,200**

Illustrative **delta per incident**: **$6,400** in bridge labor alone—before customer credits, SLA penalties, opportunity cost, or vendor charges.

At **12** such incidents per year, that single line item is on the order of **~$77k/year** of *bridge-time difference* in this toy model. Change concurrency to **12** people or hourly cost to **$350**, and the same math scales linearly.

> **The point is not the exact dollars. The point is that the function is multiplicative: people × time × incident rate.**
>
> **This is the unit cost of ambiguity before scale.**
>
> A larger organization does not need a different model. It only needs to multiply this unit across more incidents, more teams, and higher labor cost structures.

All figures in this appendix are **USD**. The same formula scales linearly with whatever burdened hourly rate finance accepts (for example **$250/hr** or **$350/hr** instead of $200/hr).

In practical terms, the relevant variable is:

> **time-to-trusted-proof**

because it directly drives how long coordinated human investigation remains active.

Hidden Cost: Context Switching and Delivery Drag

Bridge labor cost captures only the **visible coordination time** during an incident.

In practice, the impact extends beyond the bridge through **context switching and delivery disruption**.

When engineers are pulled into an incident:

- planned work is interrupted;
- cognitive context is lost;
- task resumption requires reloading state and revalidation;
- parallel work streams slow down or stall.

This creates a secondary effect:

> **delivery capacity is reduced beyond the duration of the bridge itself**

A simple way to express this is as a **friction multiplier**:

> **effective delivery impact ~= bridge duration x context-switch multiplier**

Where the multiplier depends on organizational characteristics:

- depth of system complexity;
- number of parallel work streams;
- maturity of engineering processes;
- ability to isolate incident response from delivery teams.

In some environments, one hour of bridge time may translate into **multiple hours of disrupted delivery capacity** across teams.

This appendix does not assign a fixed multiplier.
It is intentionally left as an **operator-defined parameter**.

The key observation is structural:

> **incident cost does not end when the bridge closes - it propagates into delivery capacity.**

CAPEX: “insurance-shaped” observability stacks compound

A separate line—often buried in refresh cycles—is incremental capital spent to reduce uncertainty: more capture, hotter retention, wider analytics clusters.

When delayed proof becomes structurally expected, organizations often compensate by **buying permanent explanatory readiness everywhere**, converting an OPEX ambiguity problem into a CAPEX insurance posture.

Under some deployment assumptions, a three-year program can add **single-digit to low tens of millions** of USD in large environments, depending on scope; amortization and refresh rules determine how it hits the P&L.

Illustrative 3-year TCO shell (always-on capture / wide observability path)

The table below is an **order-of-magnitude TCO skeleton** for a “widen the net” program—not a quote, procurement estimate, or vendor pricing commitment. Ranges intentionally wide; finance should substitute vendor pricing and internal run costs.

THREE-YEAR COST STRUCTURE BY SPEND DRIVER

Line item	Illustrative 3-year TCO band (USD, large enterprise)	What typically drives the spread
Tap / SPAN / packet broker fabric + optics	**$3M-$8M**	Port count, speeds (100G-800G), redundancy
Capture appliances / probes	**$1M-$4M**	Sustained throughput, buffering, form factor
Storage + hot retention for flow / metadata / PCAP-adjacent data	**$2M-$6M**	TB/day ingest, retention policy, replication
Analytics / SIEM / datastore ingest uplift	**$1M-$4M**	EPS, licensing, compute cluster
Implementation + 3-year run (FTE + contractors + training)	**$1M-$3M**	Often under-modeled

Illustrative combined band: about **$8M–$25M** over three years for a heavy “insurance-shaped” posture in a large footprint—**before** you can show a modeled reduction in **hours of ambiguity** per incident.

> **TCO is not the enemy.** The question is whether the next dollar reduces **forensic time** enough to justify itself, or mainly expands **capacity to store uncertainty**.
>
> **Delayed proof is therefore paid twice:** - once as **bridge time** during incidents; - once as **permanent infrastructure** bought in advance of incidents.

An evidence-activation posture does not magically eliminate that baseline. It targets a different failure mode: **spending for universal high fidelity and still paying long bridges** because proof arrives on forensic time.

A useful internal question is therefore:

> **If we spend the next tranche of CAPEX on wider always-on capture, what is our modeled reduction in *hours of ambiguity*—and who owns that assumption?**

SLA credits and contractual exposure (illustrative add-on)

Bridge labor is only one line. Many organizations also carry **contractual** exposure: service credits, rebates, or penalties tied to **minutes** of outage or sustained degradation. This is **not legal advice**; it is a reminder to put **revenue-at-risk** next to **hours on bridge**.

CONTRACT CREDIT TRIGGER EXAMPLE INPUTS	
Input	Example (placeholder)
Monthly recurring revenue (MRR) under an affected contract	**$500,000**
SLA: cumulative **30 minutes** of violation triggers a **10%** monthly service credit (cap varies)	Policy-dependent
One incident drives **45 minutes** of billable violation time	Simplified

If the credit applies as **10% of MRR** for that month, **one** qualifying month exposes on the order of **$50,000** in credits—**before** vendor overtime, customer churn risk, or regulatory reporting costs.

> **A single material SLA month can swamp the bridge-OPEX line.** That is why "faster proof" is not only efficiency—it can be **revenue and relationship protection**.

At larger scale, this exposure compounds across: - multiple concurrent customers, - overlapping incidents, - and cumulative SLA windows.

The same ambiguity unit therefore expands from a labor issue into a **revenue-at-risk mechanism**.

How to use this appendix in a meeting

- Treat the tables as a **spreadsheet shell**, not a claim.
- Force alignment across networking, operations, and finance on **hours to trusted story** for one historical incident; the argument becomes empirical quickly.
- If the organization cannot estimate **hours to trusted story**, that is already a signal: **the cost model is missing the variable that actually invoices.**
- Ask explicitly: **how many such ambiguity units exist per year in our organization?**

Methodology and Public Pricing Anchors

The CAPEX bands in this appendix are not vendor quotes and should not be read as universal deployment prices, certified sizing guidance, or guaranteed implementation outcomes. They are **triangulated order-of-magnitude bands** derived from public component anchors and from the typical cost drivers of large visibility programs: port count, line rate, redundancy, retention, ingest volume, licensing model, and implementation effort.

The method is intentionally illustrative and assumption-driven:

1. identify the major components of an always-on high-fidelity visibility path;
2. use public price or capacity anchors where available;
3. expand from component-level anchors to a multi-site / redundant enterprise posture;
4. treat implementation, support, training, HA, and retention as explicit cost multipliers;
5. require each reader to replace the illustrative bands with quoted internal designs.

Public anchors used

The references below are included solely because they provide public pricing, throughput, or sizing anchors relevant to the illustrative cost ranges used in this appendix.

PUBLIC COMPONENT AND CAPACITY PRICE SIGNALS		
Cost area	Public anchor	How it supports the model
Packet capture appliances	FMADIO Lite Series publicly states pricing starts at **less than $17,000** for a **20 Gbps sustained write-to-disk** system.	Establishes a public lower-bound anchor for packet-capture hardware; higher-throughput, redundant, multi-site deployments scale above this entry point.
High-throughput packet capture	FMADIO publicly describes packet capture systems for **10G, 25G, 40G, 100G**, and high-end capture classes.	Supports the assumption that capture cost scales with sustained throughput and buffering requirements.
Network packet brokers	Public NPB listings include 100G/400G systems, including FS 32-port 100G/400G NPBs and quote-based 400G systems such as NEOX PacketLion 3XL.	Supports the "packet broker fabric + optics" line as a material CAPEX item driven by port count, speed, redundancy, and aggregation policy.
400G visibility class	Keysight Vision Edge 400P and NEOX PacketLion 3XL publicly describe **32 x QSFP-DD 400G-class** packet-broker platforms.	Supports the line-rate and port-density assumptions behind the upper range of visibility-fabric spend.
Analytics / SIEM ingest	Microsoft Sentinel pricing is ingestion-based, with public commitment tiers starting at **100 GB/day** and published daily prices for larger tiers.	Supports the assumption that analytics cost scales with GB/day, EPS, retention, and licensing model.
Hot storage	AWS S3 Standard public pricing provides a commodity cloud-storage anchor in **$/GB-month**, before replication, retrieval, indexing, compute, or analytics overhead.	Supports the storage/retention line as a separately modeled cost driver, not a free side effect.
Time-series backend sizing	Grafana Mimir capacity guidance gives public CPU/RAM sizing anchors, including **1 core and 1 GB RAM per 25,000 samples/sec** for distributors.	Supports the broader claim that ingest-oriented observability paths have measurable compute and memory scaling pressure.

These anchors do not prove a single universal **$8M–$25M** outcome, nor do they imply that any specific vendor deployment necessarily falls within that range. They support only the narrower illustrative claim made here: a heavy always-on visibility posture in a large footprint can plausibly become a **multi-million-dollar three-year program** once packet access, capture, storage, analytics ingest, redundancy, implementation, and run cost are combined.

The scaling mechanism from component anchors to multi-million-dollar programs is driven by:

- **port multiplication** (hundreds to thousands of monitored links);
- **line-rate escalation** (10G -> 100G -> 400G visibility domains);
- **redundancy and HA duplication** (active/active or active/standby paths);
- **multi-site deployment** (regional or datacenter-level replication);
- **retention expansion** (days -> weeks -> months of hot or near-hot data);
- **analytics ingest scaling** (GB/day or EPS-based licensing growth);
- **operational overhead** (integration, training, support, lifecycle management).

In this context, the transition from **component-level pricing (tens of thousands per unit)** to **program-level spend (millions over three years)** is not a single cost jump, but the cumulative effect of scaling across these dimensions.

The illustrative bands in this appendix reflect this **multiplicative expansion model**, not a single-vendor pricing claim, endorsement, or certification.

Comparative and Vendor-Neutrality Notice

This appendix compares economic and architectural scaling behaviors, not vendor quality or product superiority.

Third-party product names, pricing pages, and technical references are cited only because they provide publicly accessible examples of:

- throughput classes,
- ingest-based pricing,
- storage pricing,
- visibility-fabric scale,
- or observability sizing behavior.

No statement in this appendix should be interpreted as:

- a claim that a named vendor requires the illustrated spend levels in all deployments;
- a claim of interoperability or functional equivalence with ESC;
- a benchmark certification;
- a procurement recommendation;
- or a statement of vendor performance guarantees.

All modeled figures remain:

- illustrative,
- assumption-bound,
- deployment-sensitive,
- and subject to substantial variation depending on architecture, retention policy, redundancy posture, operational process, negotiated pricing, and workload profile.

References

- FMADIO Lite Series packet capture appliances: https://www.fmad.io/products-lite
- FMADIO packet capture product family: https://www.fmad.io/products
- FS Network Packet Broker product family: https://www.fs.com/products/152271.html
- FS 400G Network Packet Broker product example: https://www.fs.com/products/346169.html
- NEOX PacketLion 3XL 400G Network Packet Broker: https://www.neox-networks.com/en/shop/neox-networks-en/packetlion-network-packet-broker-en/nx-pbpl-3xl/
- Keysight Vision Edge 400P Network Packet Broker: https://www.keysight.com/us/en/product/SYS-VE400P-BASE-AC/vision-e400p.html
- Microsoft Sentinel pricing: https://www.microsoft.com/en-us/security/pricing/microsoft-sentinel
- Microsoft Sentinel billing documentation: https://learn.microsoft.com/en-us/azure/sentinel/billing
- AWS S3 pricing: https://aws.amazon.com/s3/pricing/
- Grafana Mimir capacity planning: https://grafana.com/docs/mimir/latest/manage/run-production-environment/planning-capacity/

← Back to main paper

License Notice