The Cerbi Insider

Safer Logs and Practical AI for Small Teams



ISSUE 01 / NOVEMBER 2026

A field guide to useful evidence, smaller risks and better experiments.

01 / START SMALL

One workload. One decision. Evidence you can use.

You do not need a new platform to begin. Choose a real engineering question, write down the boundary, and keep enough evidence to know whether the change helped.

> **YOUR FIRST QUESTION**
>
> Can someone diagnose a failed checkout without seeing a password, access token, raw request body or unnecessary personal data?

Three things to leave this guide with

- A **logging contract**: the context you need, the fields you exclude and the person who owns the rules.
- A **small verification record**: what actually reached a destination, what was blocked or altered, and which paths you did not test.
- An **AI test card**: one bounded job, a baseline, representative cases and a clear decision to keep, change or stop.

Inside this issue

03 Write the logging contract
04 Verify the complete path
05 Estimate cost responsibly
06 Test a small AI job
07 Read an honest model result
08 Where Cerbi fits
09 Make the next change
10 Sources and method

02 / THE LOGGING CONTRACT

Keep the context. Reduce the exposure.

Start with fields that answer a specific operational question. Structured logging helps you query named values; it does not make every value safe. [1, 2]

Decision	A fictional checkout example
Keep when justified	event, service, environment, correlationId, outcome and a coarse failure category.
Exclude at the source	Passwords, credentials and raw request/response bodies. Review exceptions and free text for accidental disclosure.
Review before collection	Identifiers, email addresses, IP addresses and business-sensitive fields. Record purpose, access and retention.
Name an owner	The engineering owner maintains the event contract. Security/privacy owners review exposure and exceptions.

A smaller event is often easier to review

```
{
  "event": "checkout.failed",
  "service": "checkout",
  "environment": "staging",
  "correlationId": "fictional-trace-001",
  "outcome": "declined",
  "reasonCategory": "payment_provider"
}
```

Illustrative schema only, not Cerbi configuration or an executed test. Even identifiers and failure categories need a purpose and an access policy. Do not move sensitive values into the message string to evade field rules.

03 / VERIFY THE PATH

A saved policy is only the beginning.

Treat configuration, loaded runtime behavior and destination evidence as separate checks. Use invented canary values in a nonproduction test, never real secrets.

Test case	What to inspect
Expected event	Required context survives and the intended destination retains the event.
Forbidden canary field	The specific rule produces the expected outcome; inspect both structured fields and formatted text.
Alternate output path	Another logger, direct console output, exception stack or exporter does not silently bypass your assumed boundary.
Override or debug mode	Document exactly what can pass downstream and who reviews the exception.
Failure and retry	Observe disabled reporting, unavailable destinations, timeouts and repeated delivery. Do not infer success from a clean HTTP response alone.

KEEP TWO KINDS OF EVIDENCE

1. The application log at its intended destination.
2. The governance decision record, when configured.

One does not automatically prove the other. Record the policy version, runtime version, timestamp and tested route.

Collector-side filtering can help, but it cannot undo earlier copies. OpenTelemetry recommends collecting only necessary data and reviewing instrumentation output. Hashing predictable identifiers is not guaranteed anonymization. [3]

04 / COST WITHOUT HYPE

Count the bytes. Then count the work.

A practical estimate starts with your observed event volume and invoice rules. Treat storage, ingestion, queries, export and engineering time as separate costs.

1.5 GB

less raw event data per day in this fictional example

Assumptions: 1,000,000 events/day, reduced from 2,000 to 500 bytes/event. Difference: 1,000,000 x 1,500 bytes = 1,500,000,000 bytes, or 1.5 decimal GB/day. Over 30 days: 45 GB of raw volume avoided.

> THIS IS NOT A SAVINGS CLAIM
>
> Compression, batching, sampling, pricing tiers and retention rules can change billable volume. No vendor price or customer result is implied. Validate the estimate against your actual invoice and workload.

Record for one comparable window	Before / after
Event count and measured bytes	_____ / _____
Ingestion, storage and query charges	_____ / _____
Time to answer the same incident question	_____ / _____
Required evidence lost or preserved	_____
Engineering and maintenance effort	_____

Do not remove security, legal-hold or incident evidence merely to reduce volume. Agree the purpose and retention requirement with the responsible owner before changing collection.

05 / PRACTICAL AI

Give the model one bounded job.

A useful first pilot might classify a fictional support request, propose a test case or summarize an approved technical note. Keep production writes and customer messages outside the first experiment.

> **EXAMPLE TEST CARD**
>
> **Job:** route synthetic support requests to one of six queues.
> **Input:** approved fictional text; no credentials or customer records.
> **Output:** one label, source reference and an explicit uncertain state.
> **Authority:** suggestions only; a person confirms consequential actions.

Six questions before you expand it

- **Baseline:** can a simple rule or the existing tool do the job well enough?
- **Cases:** include normal, ambiguous, conflicting and out-of-scope inputs. Freeze an untouched evaluation set before selecting a model.
- **Failures:** count harmful error types and abstentions, not just average accuracy.
- **Cost:** measure total latency, tokens or compute, retries and review time. Include maintenance and data preparation.
- **Boundary:** use least-privilege tools; treat retrieved text as data, not authorization. A local model can still leak data through connected tools.
- **Release:** require acceptance criteria, rollback and human review. Preserve the working baseline when the candidate does not clear the gate.

NIST describes its AI Risk Management Framework as voluntary guidance for managing risk across design, use and evaluation. A completed checklist is not a certification. [4]

06 / FROM THE CERBI LAB

A measured gain. An honest hold.

November 6, 2026 / Internal synthetic-data experiment. Laya was evaluated on a six-category bookkeeping choice task. It is a specialized classifier, not a general assistant or tax expert.

Evaluation set	Baseline	Frozen candidate
Development set	59 / 72	64 / 72
Separate reviewed synthetic holdout	55 / 72	59 / 72

What changed

On the paired holdout comparison, four previously incorrect examples became correct and no previously correct examples became incorrect. Ten false-income errors remained. Two loan examples that were already wrong shifted from a transfer label to an income label.

> **RELEASE DECISION: NOT PROMOTED**
>
> The score improved, but the remaining error types were too consequential to treat the candidate as ready. Human review and conservative routing holds remained in place.

How to read these numbers

- These are **correlated synthetic examples from one project**, not independent customer or accounting validation.
- The frozen candidate was compared on a separate holdout. Do not combine development and holdout scores into one headline.
- These counts do not establish production reliability, cost advantage or tax expertise.

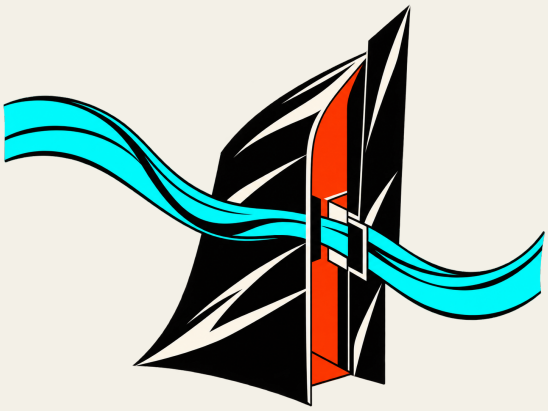
Source: Cerbi internal paired comparison and per-example results, reviewed November 6. Evidence describes this experiment only. [5]

07 / WHERE CERBI FITS

Govern the log path. Keep your stack.

Cerbi focuses on application logging governance. Start with one supported integration and verify its behavior against your agreed policy and destination. [6]

Define. Check. Govern. Review.



Current scope	Boundary to verify
CerbiShield	Versioned policy and governance evidence in a customer-hosted Azure deployment. Reporting depends on configuration.
CerbiStream / Gateway	Selected supported application paths, or explicitly configured OTLP log paths. This is not automatic coverage of every logger or destination.
AI-related logging	Relevant emitted metadata on supported connected logging paths. It does not authorize agent actions or govern model behavior.

Two important limits. Relax is a deliberate override that can allow forbidden fields downstream. The normal reporting path excludes application payload, but evidence/debug behavior depends on configuration. Verify both rather than assuming blanket protection.

The free Scanner can help find risky logging patterns. Review findings against your existing controls, then choose one scoped improvement. Cerbi does not replace your observability stack or guarantee legal compliance.

08 / YOUR NEXT WORK SESSION

Make one small change worth keeping.

Use this as a planning worksheet. The time boxes are a starting suggestion, not a promise about your environment.

01 **Choose the question / 5 minutes**

Name one workload, one incident question and one owner. Choose necessary fields and those that must stay out.

02 **Write the test / 10 minutes**

Build fictional allowed, forbidden, override and failure cases. List tested output paths and coverage gaps.

03 **Compare the evidence / 10 minutes**

Keep the baseline, versions and denominators. Check destination output, useful context and harmful failures before the headline score.

04 **Make the decision / 5 minutes**

Keep, adapt, defer or reject the change. Record why, the next test, rollback and who can approve release.

> THE INSIDER EDITORIAL STANDARD
>
> Practical help, source-linked market and legal developments, AI/model tests and Cerbi progress. Distinguish a proposal, a local test and a deployed result. Publish limitations beside the claim.

The planned newsletter is twice weekly. The Insider page explains current signup availability, consent, privacy and how to unsubscribe. Downloading this booklet alone does not subscribe you.

Visit cerbi.io/insider | Privacy

Educational material, not legal, accounting or security certification. Applicable obligations depend on your data, jurisdiction and use. Seek qualified advice for consequential decisions.

09 / SOURCES & METHOD

Keep the trail with the claim.

Public sources were checked on November 6, 2026. This guide combines source-backed boundaries with original Cerbi worksheets and hypothetical examples.

1 **OWASP Logging Cheat Sheet**

Sensitive-data exclusions, log-injection defenses and protection of logging systems.

2 **Microsoft: Logging in C# and .NET**

Structured logging and destination-provider concepts.

3 **OpenTelemetry: Handling sensitive data**

Data minimization, instrumentation review, processing options and the limits of hashing.

4 **NIST AI Risk Management Framework**

Voluntary risk-management guidance, not a certification or legal-compliance guarantee.

5 **Cerbi lab note / November 6, 2026**

Internal synthetic paired comparison: frozen baseline and candidate results. No customer data. Method and exact counts are on page 7.

6 **Cerbi product scope and Scanner quickstart**

Current public integration boundaries and starting point for a scoped logging review.

About this edition

Issue 01, version 1.0. All events and cost figures are fictional unless explicitly labeled as the internal synthetic experiment. There are no customer testimonials, compliance promises or invented deployment outcomes. Links and product scope can change; verify them before implementation.

Original cover artwork was generated for this edition using Cerbi visual references. Supporting illustrations are conceptual editorial artwork, not evidence of product behavior. Text and tables remain selectable and source links are clickable.

CERBI / Useful context. Deliberate boundaries. Evidence for the next decision.