Observability costs too much: how teams cut telemetry spend without going blind

26 October 2026 · 5 min read · blogs.iictp.org/observability-cost-reduction

Telemetry grows faster than the system it describes unless someone manages it. The goal is not to collect less for its own sake, but to keep the data that helps you make decisions and drop the rest.

TL;DR The short version

- Telemetry grows faster than the system unless someone manages it; keep the data that helps decisions and drop the rest.
- Most money goes on high-volume, low-value logs and high-cardinality metric labels.
- Sample (keep all errors and slow traces), filter at the Collector, and decide first which questions your telemetry must answer.

Read online, listen to it, or change the text size: https://blogs.iictp.org/observability-cost-reduction

Where the money goes

- High-volume, low-value logs.
- Metrics with high-cardinality labels, which multiply the number of series.
- Traces kept at full detail for every request.
- Data kept for longer than anyone looks at it.

Practical levers

- **Sampling:** keep all errors and slow traces, and a share of the rest.
- **Filtering at the Collector:** drop noisy events before they are stored.
- **Cardinality control:** avoid unbounded values, such as user IDs, in metric labels.
- **Tiered retention:** keep recent data hot and older data cheap.

Stay safe

1. Decide what questions you need telemetry to answer.
2. Measure how often each source is actually used.
3. Cut the unused first and watch for gaps.
4. Review costs per team, regularly.

Related credential: Clean Software Telemetry Practitioner. Explore the exam facts and syllabus.

Sources and further reading

1. OpenTelemetry. Sampling. opentelemetry.io/docs/concepts/sampling/.

Published by IICTP, an independent technical credential body, on 26 October 2026. The current version of this post is at https://blogs.iictp.org/observability-cost-reduction. More guides: blogs.iictp.org.