

SOLUTION GUIDE

How to reduce telemetry expenses with Cribl

The rising cost of telemetry and its impact on observability and cybersecurity





CONTENTS

- 03 The rising cost of telemetry and its impact on observability and cybersecurity
- 04 Four areas of telemetry cost management
- 05 How Cribl helps: efficient telemetry with pipelines, data tiering, storage and analysis
- 08 Modern data architecture
- 09 Learn more about cutting telemetry costs with Cribl



INTRODUCTION

The rising cost of telemetry and its impact on observability and cybersecurity

Data is the lifeblood of digital business, powering everything from troubleshooting technical issues to detecting cybersecurity threats.

But as applications, infrastructure, and usage grow, so does telemetry data. Industry analysts are seeing a 29% growth rate per year. At that pace, your daily deluge of telemetry will double every 18 months. Organizations striving for best-in-class observability and security often find themselves struggling to keep up with these costs, which eat into both operational budgets and strategic initiatives.

29%

CAGR annually

2x

telemetry data
every 18 months

But these costs are more than line items. High telemetry expenses limit the amount of data teams can analyze, leading to lower visibility, slower threat detection, reduced compliance, and missed business insights. Fragmented architectures make onboarding new apps painfully slow (sometimes taking years to reach full coverage), while siloed tools and duplicative infrastructure drain both human and computer resources. Ultimately, uncontrolled telemetry costs put both observability and cybersecurity at risk.

How can you gain control?

The answer begins an approach that is focused on regaining choice, flexibility and control over telemetry. And the rewards for gaining control over your growing telemetry are massive.

A global retailer with over 5,000 stores undertook a transformation of its IT and security data management by replacing legacy tools (Logstash and syslog-ng) with Cribl.

What they gained was stellar:

-  41% reduction in data volume
-  Longer data retention
-  Streamlined admin
-  Flexible routing and enrichment



Four areas of telemetry cost management

Successfully managing telemetry cost isn't just about cheaper storage or tossing out legacy tools. It requires your attention across four key domains: **strategic, financial, operational, and risk-related**. Let's break down each area and explore where the biggest wins can be found.



Strategic cost management



Expand



Financial cost management



Expand



Operational cost management



Expand



Risk-related cost management



Expand



Four areas of telemetry cost management

Successfully managing telemetry cost isn't just about cheaper storage or tossing out legacy tools. It requires your attention across four key domains: **strategic, financial, operational, and risk-related**. Let's break down each area and explore where the biggest wins can be found.



Strategic cost management



Financial cost management



Operational cost management



Risk-related cost management



Telemetry strategy must support your business goals. That means getting the right data, to the right people, at the right cost:

- **Increase visibility:** End-to-end visibility is critical for service uptime, compliance, and threat detection. But uncontrolled data growth can reduce visibility—forcing teams to analyze less data, and leave themselves open to downtime and security incidents.
- **Data choice and flexibility:** Avoid vendor lock-in and keep the ability to send data to any tool, now and in the future. Flexible architectures futureproof the stack, reduce the risk of large software license increases, and avoid rearchitecture costs.
- **Unified data access:** A central “front door” for logs and metrics reduces silos and makes onboarding new business units or applications dramatically faster (up to 80%, or even 40x faster in some cases). This is true across your data in motion (ingestion) as well as accessing data at rest.
- **Support for innovation, major business events and customer experience:** Fast, flexible telemetry routing, analysis and storage architectures accelerate time-to-market for new features and support rapid troubleshooting. This is critical in many enterprises, where acquisitions, divestures, and new business initiatives require agility.



Four areas of telemetry cost management

Successfully managing telemetry cost isn't just about cheaper storage or tossing out legacy tools. It requires your attention across four key domains: **strategic, financial, operational, and risk-related**. Let's break down each area and explore where the biggest wins can be found.



Strategic cost management



Financial cost management



Operational cost management



Risk-related cost management



Financial discipline must be applied at every turn, from how you onboard telemetry to how you route, store, and analyze it:

- **Reduce infrastructure costs:** Consolidate duplicative infrastructure like syslog servers, open-source aggregators, and custom pipelines. One institution shaved \$1.8 million (55%) off annual infra costs this way.
- **Optimize storage and licensing:** Data optimization (filtering, shaping, compression) can cut Splunk and Elasticsearch costs by 36–50%, and compress S3/object storage costs by up to 88%. For example, compressing metrics from 92 TB to 10 TB saved \$800k per year for one enterprise company.
- **Tune license spend:** Route only necessary data to high-cost analytics platforms, sending lower-priority data to affordable long-term storage or archive. Further tune telemetry analysis through the implementation of lakehouse technology, where datasets can be accelerated for high-performance querying.
- **Avoid costly toil:** Free up full-time equivalent staff (FTE) from low-value tasks like pipeline management and manual onboarding, allowing them to focus on higher-priority projects and innovation.



Four areas of telemetry cost management

Successfully managing telemetry cost isn't just about cheaper storage or tossing out legacy tools. It requires your attention across four key domains: **strategic, financial, operational, and risk-related**. Let's break down each area and explore where the biggest wins can be found.



Strategic cost management



Financial cost management



Operational cost management



Risk-related cost management



Operational cost reduction is about reducing friction, waste, and delay:

- **Eliminate technology sprawl:** Reduce operations effort of maintaining infrastructure for multiple systems of analysis and storage of logs at scale.
- **Accelerate onboarding:** Onboarding new apps and infrastructure becomes faster and less error-prone — critical when facing backlogs of hundreds or thousands of apps.
- **Simplify compliance:** Efficient routing and tiering make it possible to retain logs for as long as policies require with lower manual effort.
- **Frictionless integration and analysis:** Remove the manual glue code, scripts, and fragmented monitoring that slow down both helpdesk and security teams.



Four areas of telemetry cost management

Successfully managing telemetry cost isn't just about cheaper storage or tossing out legacy tools. It requires your attention across four key domains: **strategic, financial, operational, and risk-related**. Let's break down each area and explore where the biggest wins can be found.



Strategic cost management



Financial cost management



Operational cost management



Risk-related cost management



Telemetry cost and risk are **tightly intertwined**:

- **Improve compliance readiness:** Continually meet retention and data-access requirements for regulations like GDPR, PCI-DSS, and many others, so your organization can rapidly demonstrate full compliance at any time and minimize the risk or impact of regulatory reviews.
- **Reduce error budget and outage risk:** Centralized, lossless pipelines and comprehensive retention shrink mean time to detect (MTTI) and recover (MTTR), directly reducing risk of business-impacting downtime.
- **Vendor independence:** By using pipelines and storage formats that are not tied to a single analytics vendor, you can keep your options open and avoid being trapped by inflexible contracts or price hikes.

A holistic telemetry cost management approach isn't just about technology decisions — it's a business imperative that impacts your operational effectiveness, security, and competitive agility.




How Cribl helps: efficient telemetry with pipelines, data tiering, storage and analysis

Cribl is purpose-built to help organizations take control of their telemetry, balancing observability and security needs with cost efficiency.

Chart a path in a complex digital future

Cribl is your Data Engine tailored for IT and Security. Cribl's product suite — **Stream, Edge, Lake, and Search** — collects, processes, stores, and analyzes your data in a way that's both efficient and cost-effective. Stream lets you route, reduce, and enrich data from any source to any destination, while Edge brings processing closer to where data is collected. Lake provides a simple, secure way to store and manage massive amounts of data, and Search gives you a single, easy-to-use interface to query data across multiple platforms.





The telemetry pipeline solution

At the heart of Cribl's approach is a powerful, universal receiver, connected to intelligent data pipelines. Cribl Stream ingests data from anywhere (logs, metrics, traces), enriches, filters, transforms, and routes it to any required destination — be it Splunk software, Elastic, S3, Datadog, SIEMs, or cheaper bulk storage.

- **Data reduction and optimization:** Cribl Stream enables aggressive filtering (for example, dropping health checks, sampling high-volume logs) and dynamic transformation/routing. This can cut volumes sent to expensive tools by 36–50%, and compress storage by up to 88% via deduplication and advanced compression.
- **Routing and replay:** Need to experiment with new tools or migrate analytics platforms? Stream allows routing of the same data to multiple tools, or replaying archived data for re-analysis.
- **Centralized “front door”:** By onboarding all data — regardless of source — through a single pipeline, your teams get centralized control, faster onboarding (up to 80% faster), and operational simplicity.



Data tiering

Send data to the right place, at the right cost

One-size-fits-all data storage is a recipe for overspending. Cribl enables flexible data tiering:

- **Tier high-value data:** Send critical, high-access logs and metrics to analytics tools for real-time monitoring and detection.
- **Affordable archive:** Route lower-priority or “just-in-case” data into low-cost object storage (for example, S3), freeing up budget while keeping it accessible. Keep medium priority datasets in “warm” data for more frequent investigations, such as analyzing VPC flow logs or empowering developers to search through debug logs.
- **Search-in-place:** Still need to access “hot”, “warm” and “cold” data for compliance, forensics, or ad-hoc investigation? Cribl makes archived data easily searchable, across lakes or clouds, without unnecessary ingestion or egress costs.



Affordable storage for all your telemetry

Cribl Lake gives you a vendor-neutral storage layer for telemetry — optimized for cost, retention, and future flexibility:

- **Store once, use many times:** With a full copy of data stored in a vendor-agnostic format, organizations avoid lock-in, enable re-use, and preserve data value for as long as needed.
- **Search-in-place:** Instead of paying twice (or more) to store and search the same data across multiple analytics platforms, Cribl Lake supports direct search over cost-effective storage. No tedious extract-transform-load, no rehydration delays.
- **Long-term retention:** Lower the burden of compliance by retaining logs for years, without ballooning costs. Cribl Lake gives you a managed layer atop native object storage (cloud or on-prem), making compliance and audit search fast and budget-friendly.



Frictionless, federated query across all data


Cribl Search unlocks fast, accessible search and analytics across heterogeneous storage and analytics platforms:

- **Federated search:** Explore logs, metrics, and traces from both hot and cold storage—across clouds, data lakes, and object stores—from a single UI. No data movement required.
- **Low barriers for investigation:** Security and operations teams can investigate incidents or run forensics on years' worth of data, without bottlenecking expensive analytics tools.
- **Unified experience:** Teams avoid learning (and paying for) multiple query languages or custom data connectors, reducing both training and support costs.



Modern data architecture

Together, **Cribl Stream, Lake, and Search** eliminate data silos, drive operational efficiency, tame infrastructure and storage spend, and de-risk regulatory compliance.





Learn more about cutting telemetry costs with Cribl



Customer insights

One of the fastest ways to learn about cutting telemetry costs comes from our customers.

Dive into our case studies:

SALLYBEAUTY

SpyCloud

YaleNewHaven**Health**



See Cribl in action

To gain a better idea of how Cribl works to reduce log ingestion, tier data, and route data to provide choice, flexibility and control, **check out these interactive demos:**

Routing data



Reducing ingestion



Data tiering





Cribl, the Data Engine for IT and Security, empowers organizations to transform their data strategy. Powered by a data processing engine purpose-built for IT and Security, Cribl's product suite is a vendor-agnostic data management solution capable of collecting data from any source, processing billions of events per second, automatically routing data for optimized storage, and analyzing any data, at any time, in any location. With Cribl, IT and Security teams have the choice, control, and flexibility required to adapt to their ever-changing data needs. Cribl's offerings — [Stream](#), [Edge](#), [Search](#), and [Lake](#) — are available either as discrete products or as a holistic solution.

Learn more: cribl.io | Try now: [Cribl sandboxes](#)

Join us: [Slack community](#) Follow us: [LinkedIn](#) and [X \(Twitter\)](#)

©2025 Cribl, Inc. All Rights Reserved. 'Cribl' and the Cribl Flow Mark are trademarks of Cribl, Inc. in the United States and/or other countries. All third-party trademarks are the property of their respective owners.