Cribl

REPORT

Telemetry trends: Insights unveiled

What last year's source and destination data tells us about future-ready observability



CONTENTS

04 Our top findings

05 How we created this year's report

06 Sources and destinations

10 SIEM

11 Cloud adoption

14 Your telemetry-first roadmap for IT and Security

15 Telemetry resources

This year we saw the rise of agentic AI and its accompanying hype. We saw increasing geopolitical uncertainty, and intensifying cost pressures across IT and Security teams.

Cribl has been right there, tracking the flows of telemetry data that define how these trends affect global enterprises. This year's report sheds light on how IT and Security teams react and adapt to everything from new data tiering architectures to cloud-native security and observability tooling.

We report on an ongoing need to implement a telemetry data strategy built not only on managing costs, but also on realizing the value within telemetry data itself. After all, telemetry data is more than a record of what's happened — it's how enterprises define their future.

Let's get started.

Our top findings



Data strategy, not tool strategy, will determine market leaders

Enterprises deriving the most value from their telemetry data decouple data from tools. They adopt vendor-neutral pipelines, build tiered storage strategies, and enable governed data sharing across AI and security use cases. Leading organizations position telemetry as a business enabler — not just a technical function.



Cost optimization drives pipeline flexibility

Rising data volumes and economic pressure force IT and Security teams to re-evaluate their telemetry strategies through a cost lens. Organizations increasingly use vendor-neutral technologies like Cribl to decouple ingestion from storage and analysis, enabling them to shift high-cost workloads to lower-cost destinations without compromising performance or visibility. This flexibility is especially critical in SIEM migrations, where cost is often the primary motivator for change.



Data tiering has become a strategic imperative

As telemetry data grows in scale and complexity, enterprises are adopting multi-tiered storage strategies to align cost, access, and retention policies with business value. Leading organizations use platforms like Cribl Stream and Cribl Lake to separate high-frequency operational data from long-term archival or governance use cases. This approach enables governed data sharing across teams, supports AI workloads, and ensures that the right data is available to the right tools at the right time.



How we created this year's report

For evidence-based insights, we analyzed anonymized data from Cribl Stream instances within Cribl.Cloud. This data is based on telemetry data collected from Cribl.Cloud users. Our customers use Cribl's data engine for IT and Security to securely route, manage, store, and search their telemetry data. They span every major industry and size, from small enterprises to many of the world's largest organizations.

This data is only representative of Cribl.Cloud customers and the way they use Cribl Stream to connect telemetry data sources with their chosen destinations. The trends described in this report may differ from enterprises not using Cribl.Cloud. Unless otherwise specified, the data included in this report is limited to Cribl.Cloud customers with at least one pipeline deployed.

Unless otherwise noted, this report presents and analyzes data from April 2024, to April 2025, which we refer to as "this year," "today," and "in 2025." Similarly, when we refer to "last year" or "in 2024," we are referring to data from April 2023, to April 2024. "2024" refers to the same period in its respective year. When referring to company size, Cribl uses the term "Developing" to refer to companies with revenue of less than $1B USD, and "Enterprise" to refer to companies with over $1B USD in revenue.

Data on Cribl Packs is collected from the Cribl Packs Dispensary and comprises data from Cribl's entire customer base.

Sources and destinations

Traditional sources remain dominant but are scaling steadily

Top 5 Sources

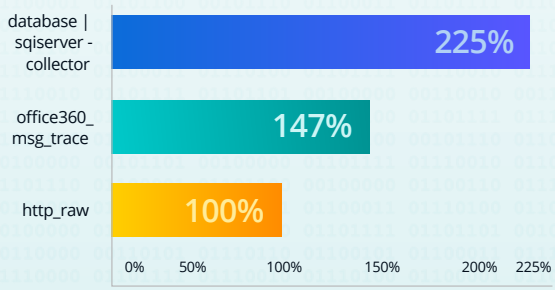
1. syslog
2. TCP
3. TCP JSON
4. REST Collectors
5. Event Hub

Core inputs like **syslog**, **TCP**, and **TCP JSON** continue to form the backbone of telemetry collection. These sources show consistent and substantial growth — each expanding by more than 60%, with **syslog** now contributing the highest volume overall. This steady rise suggests ongoing expansion of infrastructure coverage and a broadening of logging scope, particularly across legacy systems and network devices.

At the same time, ingestion mechanisms such as **Splunk HEC**, **REST Collectors**, and **System State** are also scaling rapidly. These are flexible, API-driven, or vendor-specific methods that allow teams to extract more granular telemetry from a growing number of endpoints and services. Their growth — ranging from 72% to 88% — signals an increasing emphasis on interoperability and real-time stream processing.

The most dramatic trend lies in the long tail

Surging Growth in Lower-Volume Sources Across the Enterprise Stack



Several lower-volume sources have exhibited explosive percentage growth, underscoring a shift toward broader telemetry integration across the enterprise stack. Notably, inputs like the **SQLServer Database Collector** (225% growth), **Office 365 Message Traces** (147%), and **HTTP Raw** (100%) have more than doubled their footprint. These surging sources reflect greater monitoring of SaaS platforms, cloud-hosted databases, and application-layer activity — an indication that organizations are investing in comprehensive, full-stack visibility.

This surge in volume from diverse sources creates both opportunity and strain. On one hand, organizations are gaining richer contextual data to power analytics, threat detection, and performance monitoring. On the other, the expanded telemetry footprint imposes new requirements for pipeline flexibility, scalability, and cost control.

Strategic implications are clear

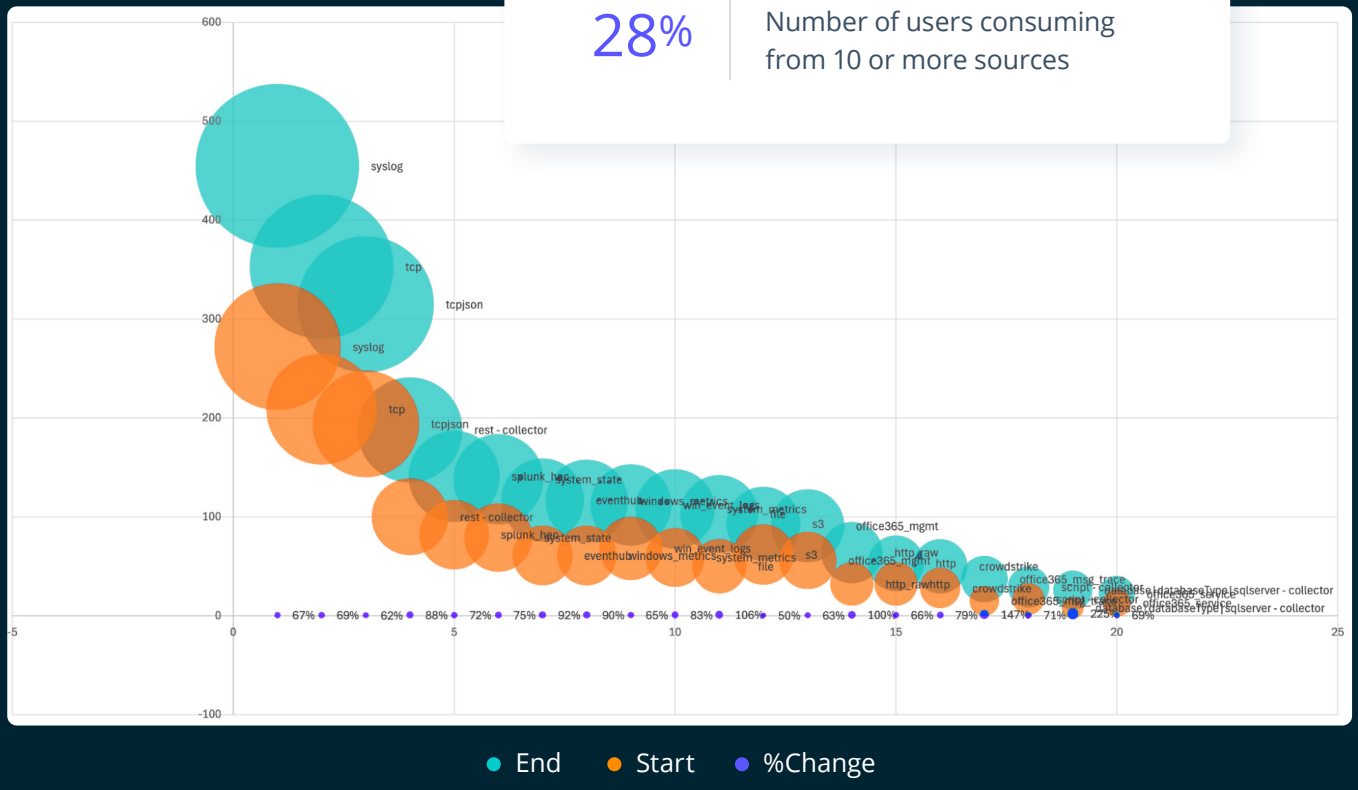
Enterprises must prepare for a bifurcated telemetry environment in which a handful of high-volume sources demand aggressive tuning and cost optimization while a growing array of specialized inputs require agility and extensibility. The ability to normalize, enrich, and route telemetry dynamically — without duplicating data or locking into vendor-native formats — has become a defining feature of modern data pipelines.

CRIBL PERSPECTIVE

Future-ready solutions will need to support not just high-throughput ingestion, but intelligent filtering, routing, and tiering strategies tailored to both technical and business priorities. Organizations that build their data management strategies around data diversity — not just data volume — are best positioned to thrive in this new telemetry economy.

Data sources by the numbers

6.6	Average sources in use
5	Median sources in use
28%	Number of users consuming from 10 or more sources



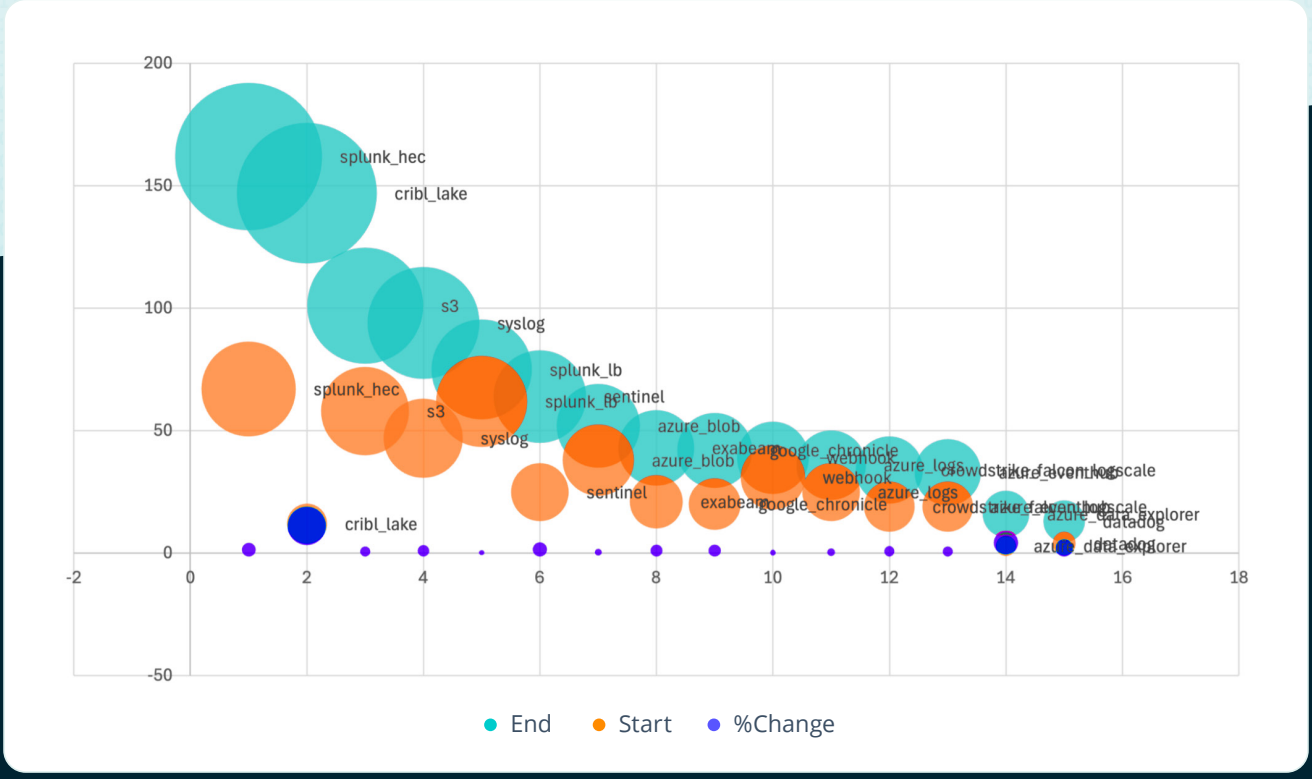
Data destinations by the numbers

As enterprises expand telemetry collection, destination patterns reveal a clear shift toward greater flexibility, cloud adoption, and vendor diversification. Recent data shows a mix of strong growth in traditional destinations and explosive adoption of next-generation platforms.

Splunk remains the most widely used destination, more than doubling since the last measurement period. This reflects its central role in many enterprises' IT and Security teams, particularly in hybrid environments. Similarly, **syslog** and **Amazon S3** have grown substantially (100% and 74% respectively), underscoring their continued importance for log forwarding and long-term storage.

But the most striking trend lies in the **rise of cloud-native and vendor-neutral destinations**. Cribl Lake stands out with a staggering increase in usage, signaling a strong shift toward flexible, schema-less data lakes built for routing, replay, and tiered storage.

Other cloud-forward destinations such as **Azure Data Explorer** (433%), **Datadog** (225%), and **Microsoft Sentinel** (156%) are also seeing rapid adoption. These gains suggest organizations increasingly prioritize tools that allow them to scale analytics without locking into a single observability platform.





SIEM products such as **Exabeam** and **Google Cloud SecOps** also posted over 100% growth, showing growing interest in cybersecurity-specific destinations for enriched threat detection and correlation. This indicates that telemetry pipelines are central not just for IT operations, but also to meet evolving security and compliance demands.

CRIBL PERSPECTIVE

The broader implication is clear: the future of data pipelines is multi-modal and cloud-optimized. Organizations are no longer routing telemetry to a single monolithic platform. Instead, they embrace pipelines that support targeted routing to specialized destinations — balancing real-time visibility, historical context, cost optimization, and tool-specific value.

To remain competitive, modern pipeline architectures must prioritize destination agility: the ability to send the right data to the right tool at the right time based on shifting costs and required value. A modern telemetry data strategy enables this flexibility without introducing complexity, latency, or operational risk.

SIEM

Security teams narrow choices amid budget pressures

The upheaval in the SIEM market we saw in 2024 continues making waves in 2025. Overworked Security teams want modern tools with better automation and detection capabilities, while their bosses must wedge SIEM contract renewals into ever-shrinking budgets.

With those pressures as a backdrop, it's surprising that the number of Cribl.Cloud users routing data to multiple SIEMs fell from 16% to 13% year-over-year. This slight decrease may indicate that some Security teams have selected their new SIEM after testing one or more options in the previous calendar year. Another scenario is that teams are sticking with their incumbent vendor after exploring their choices.

And there are many choices. CrowdStrike Falcon Next-Gen SIEM and Google Cloud SecOps are often paired with Splunk, as is Microsoft Sentinel. Exabeam has also regained some of its popularity after merging with LogRhythm, but this may be a complementary deployment of its UEBA capabilities alongside a SIEM.

Despite these options, one vendor stood out this year: CrowdStrike. Introduced in the middle of our analysis period, CrowdStrike Falcon Next-Gen SIEM quickly became the second most popular SIEM option within Cribl.Cloud, behind Splunk. By combining elements of SIEM, SOAR, and EDR, CrowdStrike delivers modern, integrated security tooling that appeals to users and buyers alike.

HOW CRIBL HELPS

Migrating to a new SIEM means taking on some risk because, without Cribl, it's a one-way door. Once you walk through it, you can't go back. You're committed to deploying new agents, training staff, updating alerts and other content, and, finally, cut over to the new product.

Cribl turns that migration into a two-way door. You can send data to different SIEMs in the format they expect with no loss of fidelity, and without weakening your security posture. Cribl customers migrate their detections and dashboards to compare products, ensuring the new SIEM works as well as the incumbent product. Even better? Keep the legacy product running for a time after the cutover. This takes all the risk out of the migration and accelerates it at the same time.

Cloud Adoption

Multicloud momentum stalls as complexity mounts

Cribl.Cloud customers use a variety of cloud native destinations encompassing SIEM, object stores, message buses, and analytics services. However, Cribl.Cloud users are surprisingly loyal to a single cloud service provider (CSP), with only 8.3% sending data to more than one CSP. This is down from last year, where 11% of Cribl.Cloud customers were sending data to multiple cloud providers.

Why aren't IT and Security teams embracing a multicloud strategy? A multicloud strategy makes sense when you're seeking the latest innovations and when you're trying to maximize flexibility. These advantages also come with increasing complexity, costs, and skills challenges. For IT and Security teams, these downsides outweigh the opportunities of multicloud deployments.

All of the cloud service providers experienced growth year-over-year, with Google Cloud Platform (GCP) experiencing the largest percentage growth. However, it's worth noting that the growth of Amazon Web Services (AWS) is slower because it is already the incumbent CSP for many Cribl.Cloud organizations.

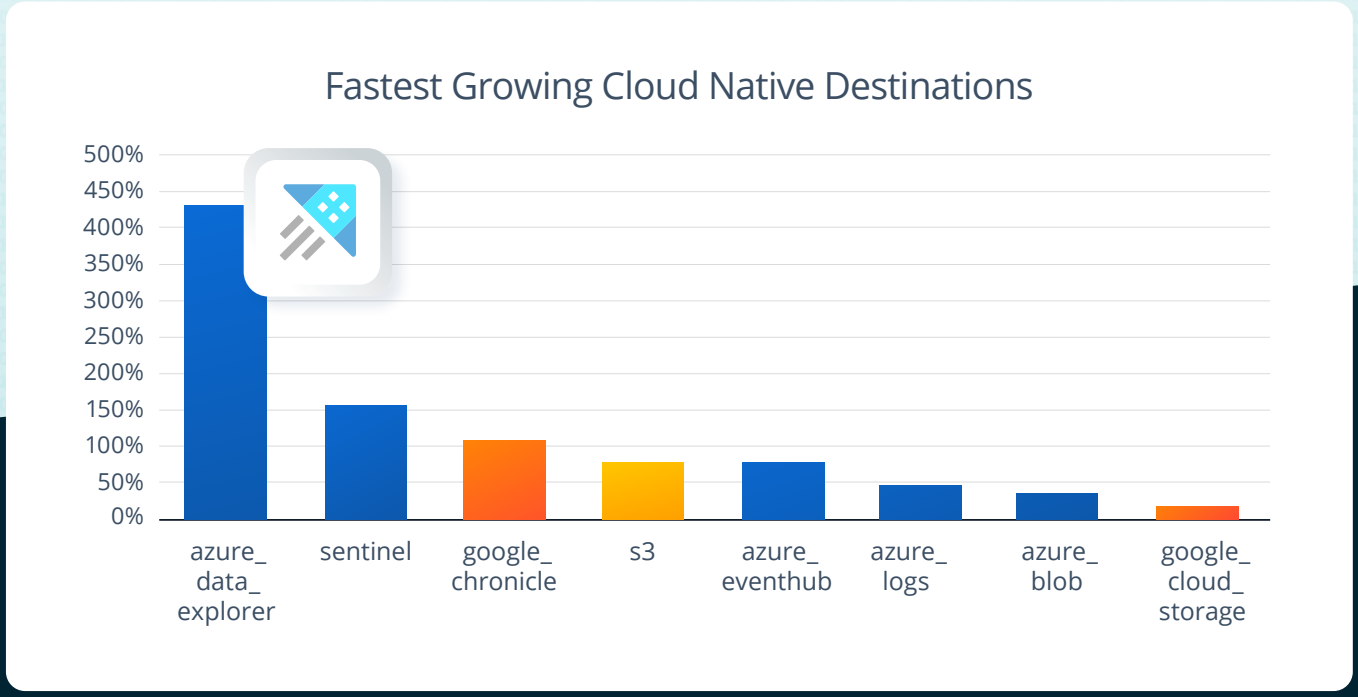


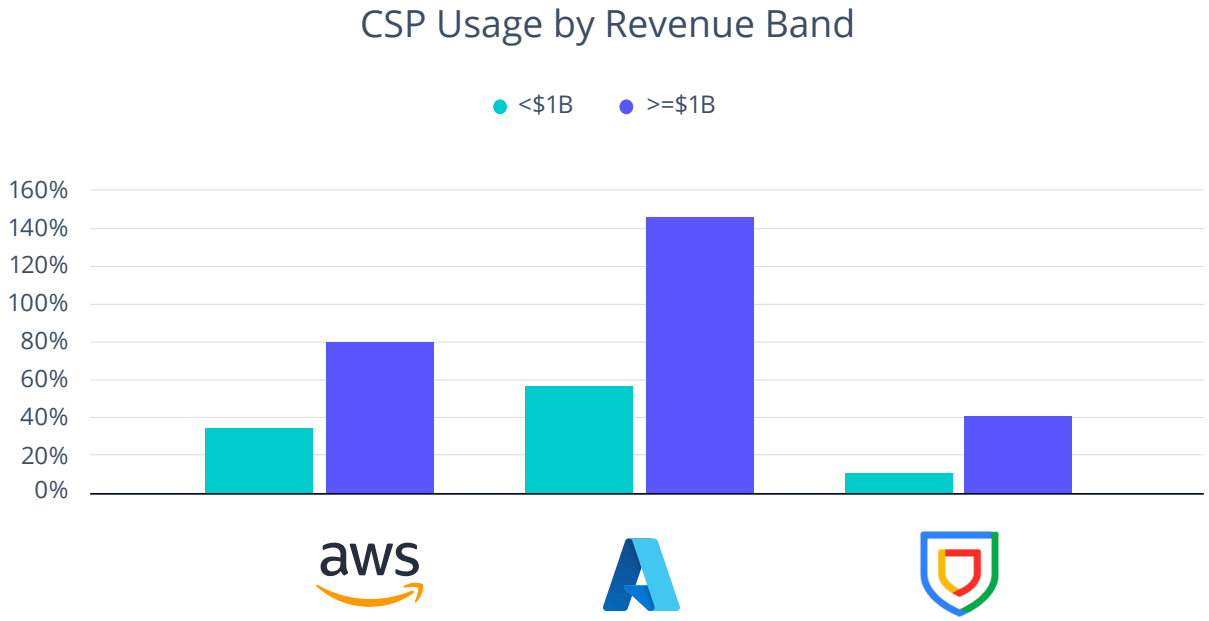
The real growth story over the last twelve months is Azure. While AWS usage is dominated by the ubiquitous Amazon S3 service and GCP's numbers are overwhelmingly represented by Google SecOps users, Azure sees the most service diversity. Unlike competing CSPs, Azure users routinely turn to multiple Azure services, such as Azure Logs, Microsoft Sentinel, Azure Blob, Azure Eventhub, and Azure Data Explorer.

In fact, Azure Data Explorer is our fastest growing cloud native destination year-over-year, with Microsoft Sentinel securing second place ahead of Google SecOps. This is Azure's second year as the leading cloud native destination for Cribl.Cloud users.

Combining these different Azure services allows companies to realize their telemetry data strategies around governed data access and sharing, cost optimization, and data tiering around different use cases, users, and tools.

Azure's growth is especially apparent when we break down CSP usage into Cribl.Cloud revenue cohorts. Companies with over $1B in annual revenue overwhelmingly gravitate toward Azure, owing to its decades-long relationships with Fortune 500 companies, vast partner and integration ecosystem, and entrenched installed base building on Microsoft technologies.





CRIBL PERSPECTIVE

Cloud-native telemetry destinations, such as Microsoft Sentinel, Google Chronicle, and vendor-neutral platforms like Cribl Lake are rapidly outpacing traditional storage and SIEM endpoints. Enterprises are diversifying their data destinations beyond monolithic on-prem destinations to more flexible options in the cloud for better scale, performance, and cost controls. As a result, IT and Security teams must prioritize flexibility, support for diverse CSPs, and intelligent routing to meet the demands of an increasingly distributed, cloud-first telemetry landscape.

Your telemetry-first roadmap for IT and Security

1

Data over tools — your new strategic imperative

As technology choices diversify, IT and Security teams must optimize their data management and technology strategies around telemetry data rather than tools. Enterprises that realize the most value from their telemetry data will lead their industries. That value comes from data rather than tools.

2

Collaborate across functions on telemetry needs

Build a telemetry data tiering strategy for your enterprise that centers on governed sharing, data growth, and cost effectiveness. Work with your peers across the organization to understand how they currently use, or desire to use, telemetry data across their functions, including emerging AI use cases.

3

Insulate your stack from market disruptions

Adopt telemetry pipeline technologies like Cribl Stream to abstract the sources of telemetry data from its native destination. This insulates you from market disruptions like onerous acquisitions and geopolitical shifts, and levels the playing field for contract negotiations with your chosen platform vendors.

4

Avoid captive pipelines to prevent lock-in

Avoid buying your telemetry pipeline product from your cybersecurity or observability platform provider. Without a vendor-neutral position, these captive pipelines will further cement the vendor lock-in you want to avoid. Use your pipeline's data routing and formatting capabilities to test new tools and products without disrupting your incumbent solutions.

5

Streamline your stack

Search for opportunities to consolidate your cybersecurity and observability tools and platforms using the capabilities of a unified data approach that uses open storage and access formats from endpoint, to pipeline, to tiered storage, and search.

Telemetry resources

WHITEPAPER

SIEM migration guide

Read how to build a modern, vendor-agnostic security data strategy while de-risking your SIEM migration and enhancing threat detection.

BLOGS

Why Data Tiering is Critical for Modern Security and Observability Teams

Explore how intelligent data tiering is transforming the way IT and Security teams manage, store, and extract value from telemetry.

Tiered Storage: A Data Strategy for 2025 and Beyond

Learn more about tiered storage and 5 whys for implementing a tiered data strategy.

CASE STUDIES

Autodesk Case Study

Discover how Autodesk broke free of their complex legacy data pipeline and received enterprise-tier support with Cribl Stream.

Sophos Case Study

Cribl Stream data processing capabilities allowed Sophos to quickly identify and filter out unnecessary fields, resulting in a 25% reduction in data volume.

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.

RPT-0003-EN-1-1025

