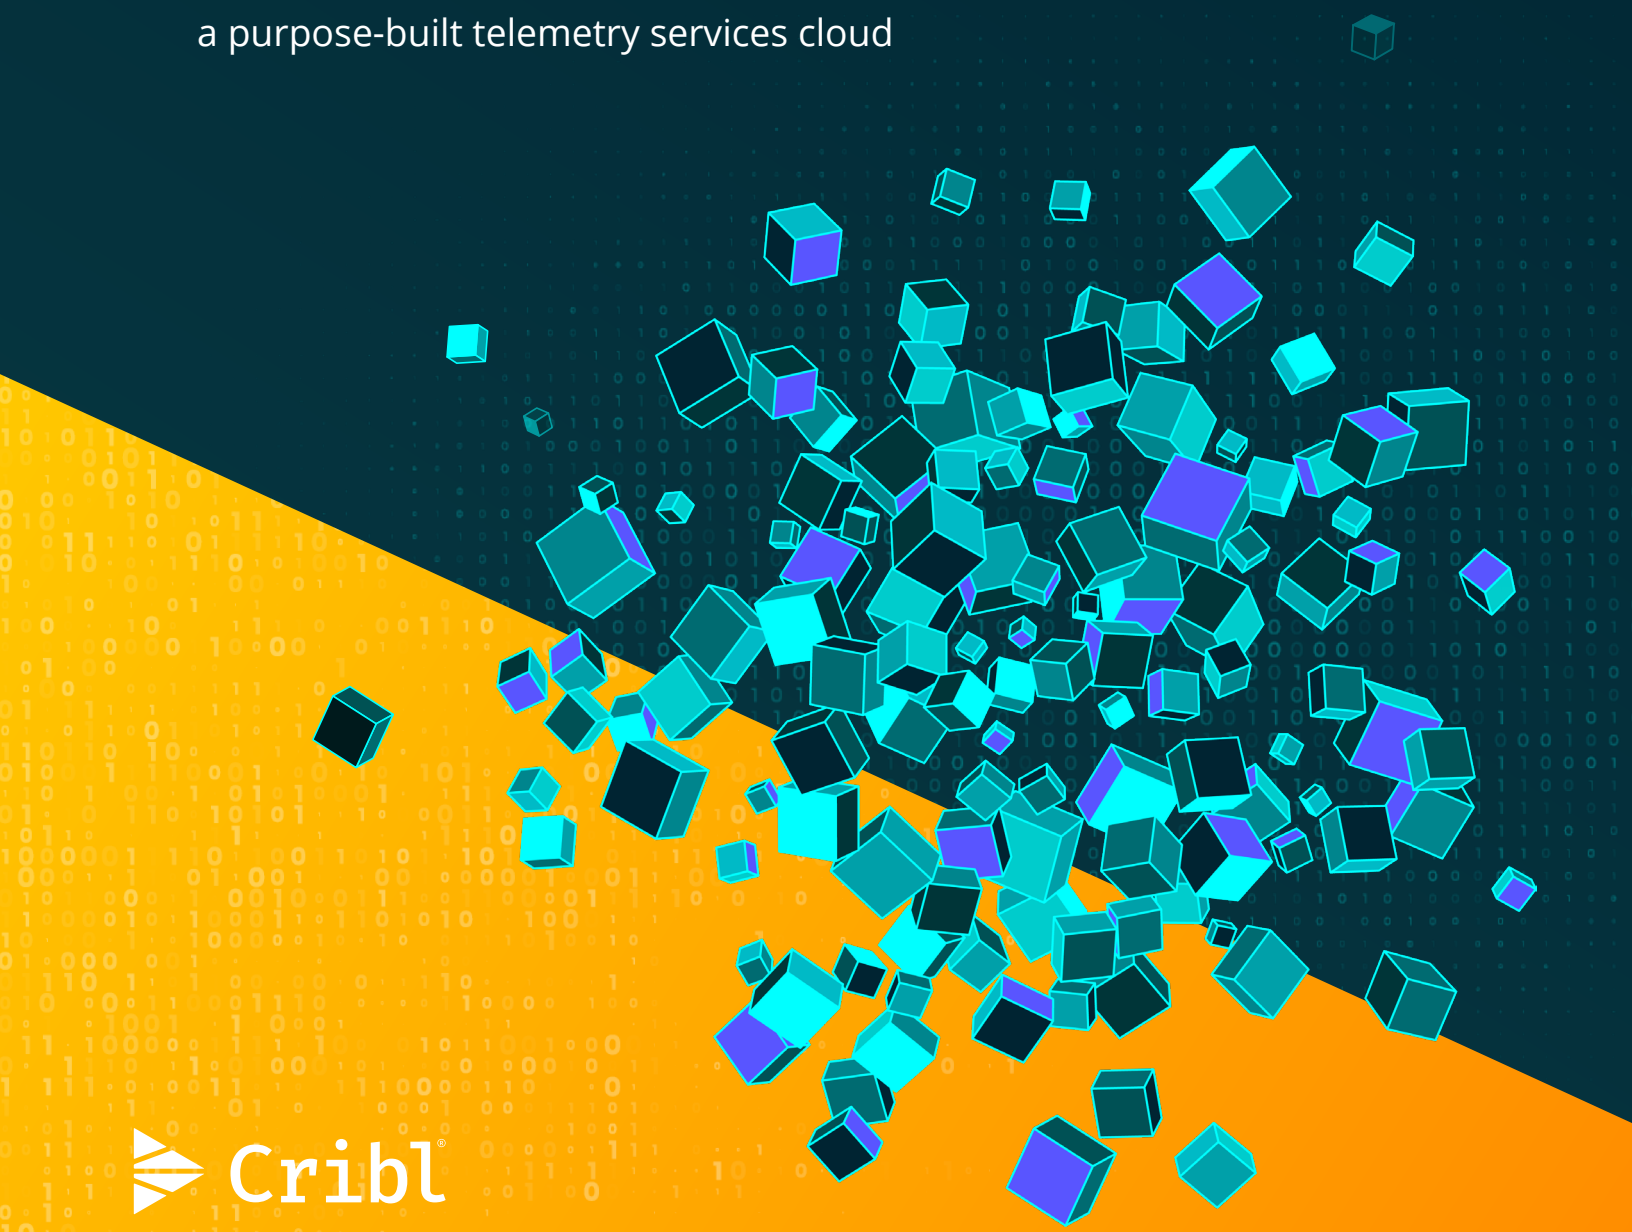
WHITEPAPER

The telemetry time bomb

Why today's telemetry tools are cracking under their own weight — and why the industry needs a purpose-built telemetry services cloud



Cribl®

CONTENTS

03 The explosion of telemetry data: Opportunity or catastrophe?

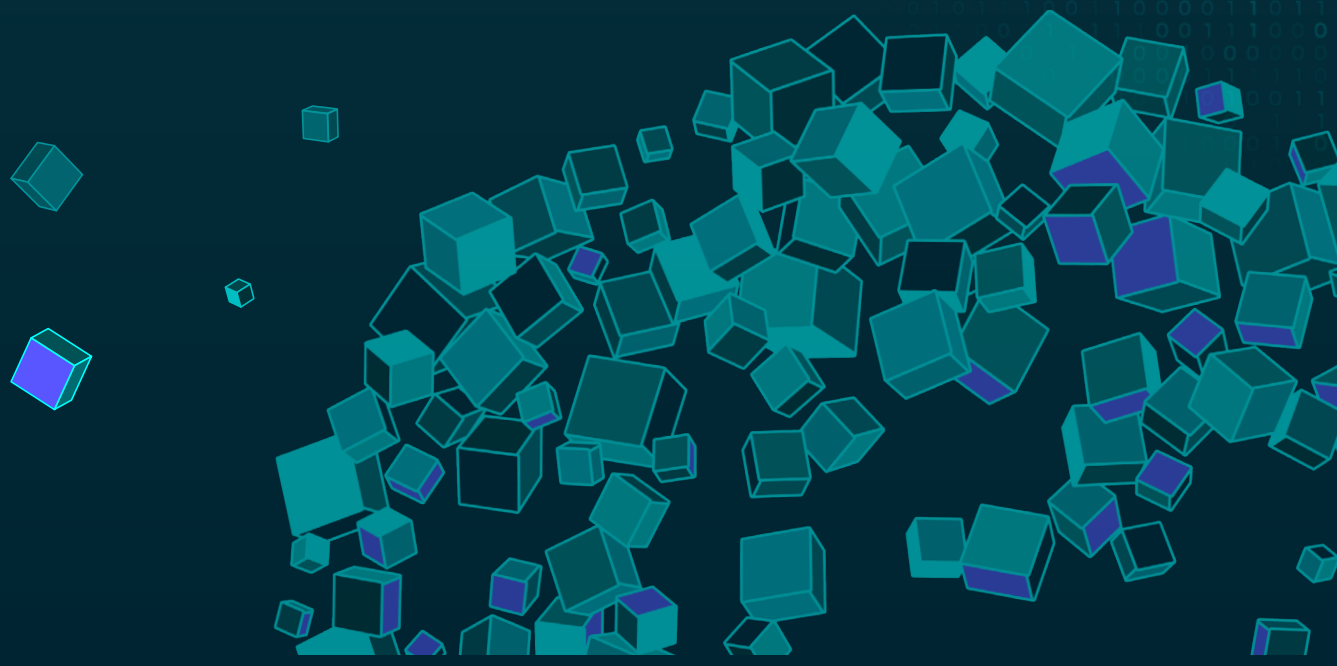
04 Where today's solutions buckle

05 The hidden tax on innovation hits your business

07 Why incremental fixes won't cut it

09 What happens if we do nothing?

10 The road ahead



INTRODUCTION

The explosion of telemetry data: Opportunity or catastrophe?

We are witnessing an unprecedented tidal wave of machine data.
Starting this year, IDC projects that the world will generate 175 zettabytes of data, with nearly 80% being unstructured, machine-generated exhaust — logs, metrics, traces, and sensor streams. These projections don't even account for digital exhaust generated by the sudden explosion of artificial intelligence use cases being deployed, which ultimately pour gas on this telemetry fire.

Every prompt, embedding call, or fine-tuning run emits mountains of additional logs and traces. Analysts now expect AI agents to **triple the volume of network alerts** in the next few years — and every one of those alerts is a log event you have to ingest, route, shape, and retain. At the same time, AI training jobs push **100 Gbps-plus east-west "elephant flows'"** across GPU fabrics, generating far more granular telemetry than traditional workloads.

Here at Cribl we often quote IDC that telemetry is growing at 29% CAGR and budgets simply are not. That number is incredibly high. The reality is that it doesn't account for the sudden surge in AI over the last couple of years. If 29% growth is both conservative and already far outpacing current tool capacity, we're in for a rude awakening if nothing changes.

The world of business data has already begun grappling with this problem. With the rise of cloud stacks built on the likes of Snowflake and Databricks, organizations have the flexibility to deal with the growing mountain of this class of data.

While transactional records live happily in these lakehouses and BI tools, telemetry does not. Machine exhaust, aka telemetry, has significantly higher volume, variety, and velocity compared to business data. Its variety is particularly wide-ranging, coming in raw from numerous sources with uncontrolled formats that cannot be dictated by the organization that is collecting the data. Trying to treat telemetry like business data forces teams into over-provisioned clusters, slow incident response, complex and ever-changing transformation scripting, and region-by-region copy sprawl that can even violate sovereignty rules due to inconsistent policy enforcement.

Most importantly, the value of the vast majority of this data is typically low. Until, that is, something breaks. An audit hits, a system goes down, or an AI tool misbehaves. Suddenly, this data is the most important data in your organization.

That means you have to collect *and* keep it, 24/7.

Where today's solutions buckle

This chart highlights just a few reasons we see legacy systems that operate in a cloud-first world buckle under the pressure:

WHY LEGACY APPROACHES FAIL	
Tool sprawl, data silos, and DIY plumbing	More than a third of enterprises say **"too many monitoring tools and siloed data" (34%)** is the #1 obstacle to full-stack observability. *(New Relic)*
Runaway ingest costs	Popular SIEMs price at **≈US$1,800 per GB per year** — before you pay for storage replicas or the outdated construct of hot/warm tiers. *(Uptrace)*
Capacity trade-offs	Teams delete or sample logs just to stay inside license limits, turning SecOps and SREs into the "Department of No."
Regulatory drag	SOX alone requires **seven years of log retention;** PCI, HIPAA, GDPR and ISO 27001 each add their own clocks, amplifying storage bills, rehydration toil, and e-discovery risk. *(Blumira)*
"Noisy-neighbor" slow-downs	Shared SaaS tiers let one tenant's heavy queries kneecap everyone else — a classic *noisy-neighbor effect* acknowledged by Azure and other cloud providers. *(Microsoft)* *(TechTarget)*
Human toil and burnout	**35% of IT leaders report higher burnout** as ops teams babysit pipelines instead of innovating. *(PagerDuty)*

It's clear the overwhelming presence of too many monitoring tools and fragmented data is holding observability back. High, unpredictable ingest costs exacerbate the issue, with common solutions priced exorbitantly — often upwards of $1,800 per GB annually.

To manage expenses, teams often resort to sampling or deleting logs, inadvertently transforming security operations and reliability engineers into gatekeepers rather than facilitators. And this slows down your entire innovation engine while exposing some of your key systems to risks.

The hidden tax on innovation hits your business

When log ingestion is capped or priced like some rare gold-leaf, data scientists are held back from training AI effectively, product teams throttle experiments, and SecOps misses faint attack signals. The indirect costs — slower releases, longer outages, failed audits, inaccurate root cause analysis — are the brutal hidden tax on innovation that crushes an organizations' ability to compete.

THE SILENT BUDGET DRAIN



Downtime

97% of large enterprises say **one hour of outage now tops $100,000, and 41% put the figure at $1–5M.**[1]



Breaches

The global average breach hit **$4.88M in 2024,** the steepest jump on record.[2]



Dropped logs

When a system of analysis license quota is exceeded, the indexer can stop accepting data, **silently discarding events that exceed the cap.**[3]



Those dollars could bankroll new products, yet they're wasted plugging holes that telemetry should have prevented.

The hidden tax on innovation hits your business *(continued)*

Human drag › **slower releases**

ROLE	WHAT THE DATA SAYS	WHY IT MATTERS
Data scientists	**≈ 80% of their week** goes to finding, cleaning, and organizing data.[6]	Messy, siloed telemetry stymies advanced, revenue enhancing AI projects.
Developers	Amazon engineers average j**ust one hour of real coding per day.**[4] Technical debt consumes **23–42%** of dev time.[5] (See explanation, below.)	Every lost hour delays features.
Security analysts	Log sampling and license ceilings leave analysts "**hunting in the dark.**"[7]	You can't detect what you don't keep.

The compounding opportunity cost

Google's DORA research shows elite performing teams who automate and fully instrument their pipelines:

Deploy
973x
more often

Recover
6570x
faster after failures

Ship with
3x fewer
change-failure rates [8]



Those advantages translate directly into revenue, uptime, and customer trust — yet tool-chain complexity and data fragmentation top the list of blockers to reaching "elite" status.

Why incremental fixes won't cut it

This all should sound familiar. That's because as an industry we've come to terms with this untenable status quo. It's the same story every quarter: You squeeze another discount out of your tool vendor, prune another index, bolt on yet another pipeline sidecar — only to watch **costs, tools, storage, and complexity** roar back a sprint later. Industry surveys prove the cycle is getting worse, not better:

TOP BLOCKER *to full-stack observability*	SURVEY EVIDENCE	WHAT THAT FEELS LIKE *in your day-to-day*
Cost	**61% of practitioners** rank cost or surprise bills as the #1 pain point. (Grafana Labs "Observability Survey 2024")	You keep one eye on license meters, drop "non-critical" logs during spikes, and hope nothing important is hiding in what you throw away.
Tool sprawl	**70% of teams** juggle **4 or more observability tools**; respondents mentioned **62 distinct tools.** (Grafana Labs "Observability Survey 2024")	You spend mornings context-switching between dashboards, exporters, and query syntaxes — then copy-paste screenshots into Slack so the rest of the org can follow along.
Stack complexity	**88% of tech leaders** say their stacks grew more complex last year and **84%** say that complexity makes security harder. (Dynatrace "State of Observability 2024")	Every "simple" production change triggers a war room because nobody trusts the telemetry pipeline and data storage to survive another microservice or region.

Why incremental fixes won't cut it *(continued)*

You could renegotiate ingest pricing, share another Elasticsearch cluster, or write one more Lambda to filter logs. **But the pattern is familiar:**



Cost Gravity

Data grows ~29% YoY; but tool and storage pricing doesn't shrink 25%! It's typically quite the opposite. Each quarter the same red-line reviews return, only bigger.



Integration tarpit

Every new point tool buys visibility into one corner, silos another copy of data, and adds two new failure modes everywhere else.



Cognitive overload

More dashboards mean more "Is this metric even correct?" debates, delaying every root cause analysis and audit request.



That's why, when surveys ask what really blocks full-stack observability, **cost, tool sprawl, data fragmentation, and complexity** leap off the charts. The Band-Aids — license tier shuffles, query sampling, cheaper "cold" storage — treat symptoms, not the structural rot.

What happens if we do nothing?

The price of doing nothing has never been steeper, more diverse, or pervasive. **Organizations that settle face:**

Exploding spend

Ingest fees rise faster than revenue.

Audit failures

Fragmented retention and masking policies break compliance.

Slower product cycles

Teams ration data instead of iterating.

Talent drain

Top engineers leave rather than babysit brittle pipelines.

No organization — startup or Fortune 50 — can absorb that trajectory.

But what happens when you remove these blockers? **You unlock:**

Released velocity on elite footing

Ship daily, not monthly.

Shorter MTTR & fewer rollbacks

Full-fidelity signals speed root-cause.

Security that scales with threat volume

Hunt first, then license math later.

Happier engineers and scientists

More time building, less time babysitting systems, and fewer resignations.



In other words, capped telemetry isn't just today's pain — **it's tomorrow's missed launch.** The goal has to be an elastic resource that can handle all data and make it accessible.

The road ahead

This post mapped the crisis of volume, cost, variety, fragmentation, risk, and human burnout that make a telemetry services cloud inevitable.

Cribl.Cloud was built for this exact moment. Unlike legacy tools that treat telemetry as an afterthought to their analytics stack, Cribl.Cloud is purpose-built as a telemetry services cloud. That means elastic scale without runaway costs, centralized governance without silos, and open pipelines that let you route any data to any destination.

The result: You get full-fidelity visibility without forcing engineers into painful trade-offs between cost, compliance, and speed.

Most importantly, Cribl.Cloud puts the control back in your hands. Instead of ballooning ingest bills and brittle pipelines dictating what data you can keep, you decide what's valuable and where it lives. That shift — from telemetry controlling your budget to you controlling telemetry — is what makes a telemetry services could not just a better solution, but a necessary foundation for competing in the age of AI, compliance scrutiny, and relentless digital velocity.

Don't just take our word for it — try Cribl.Cloud yourself. Process up to 1TB per day for free, and see how easy it is to take back control of your telemetry.

Cribl Stream™

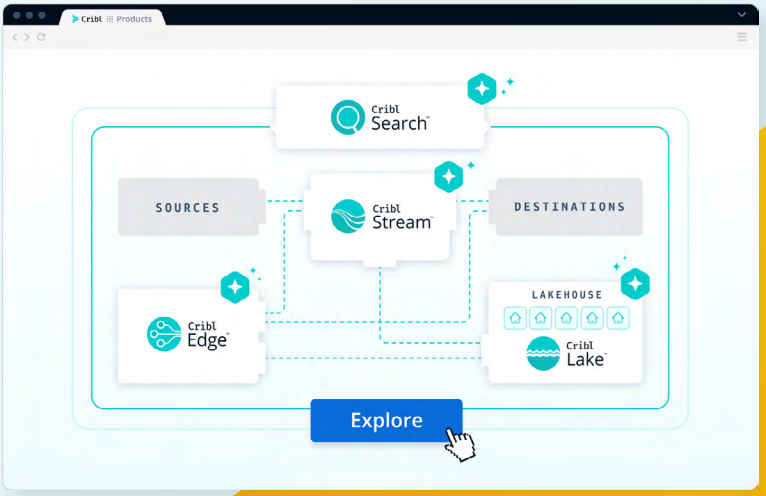
Cribl Stream filters and shapes noisy data at ingest to slash unnecessary volume.

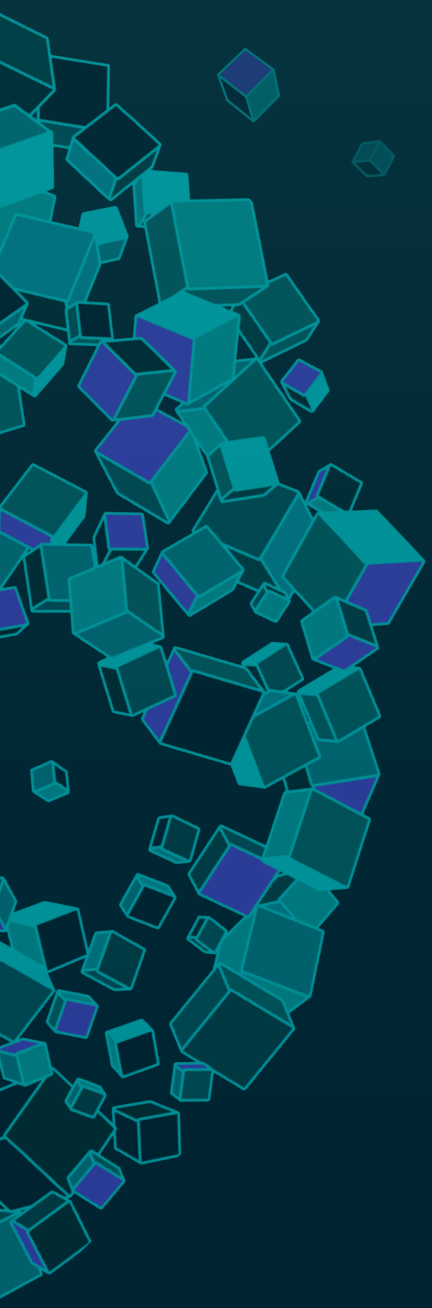
Cribl Lake™

Cribl Lake provides vendor-neutral long-term retention with search-in-place.

Cribl Search™

Cribl Search gives you on-demand investigations without dragging every byte through an expensive indexer.





REFERENCES

1. ITIC, 2024 Hourly Cost of Downtime
2. IBM, Cost of a Data Breach Report 2024
3. Splunk Docs, About license violations
4. Fortune, "Amazon developers spend just one hour per day coding"
5. vFunction, "Technical Debt: Who's Responsible?"
6. Pragmatic Institute, "Overcoming the 80/20 Rule in Data Science"
7. Edge Delta, "What is Log Sampling (and When to Use It)"
8. Google Cloud, Accelerate State of DevOps Report 2021
9. Stack Overflow, 2024 Developer Survey – AI Section

Cribl®

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.