chronosphere

Six Reasons Your Logs Are Out of Control

Why engineering teams are drowning in data.

Your team is generating more log data than ever before. But what's happening with all that data? Fact: most of it rarely provides value during incidents or investigations. This guide breaks down what's driving the log data overload and how to get back in control of your logging.

CHALLENGE

As log data growth explodes, challenges scale with it.

Inflated costs

Teams pay a premium to ingest their log data, while a vast majority of that data goes unused.

Performance bottlenecks

Many incumbent platforms can't deliver a responsive experience when processing massive data volumes.

Low signal-to-noise

Engineers waste valuable time sifting through meaningless entries to find the few logs that matter.

DATA BREAKDOWN

Recent data gives us a glimpse into the log data growth explosion.

1 Year-over-year log growth [1]

Logs are growing ***faster than anyone anticipated***. In the past year alone:

- **Log volume grew 250% on average across engineering teams**
- **22% of orgs generate 1TB+ of logs per day**
- **12% produce 10TB or more daily**

Without any way to control log data growth, these volumes can quickly outgrow your budget.

2 Log data supports a wider range of use cases than ever before.

- Logs are **the most helpful in troubleshooting** production systems
- Incident response is **the next most common use case**
- Logs are **frequently used to monitor application performance** and load testing

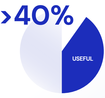

3 But, most logs are not providing value

Over 90% of recently surveyed companies said that less than 40% of their log data was actually useful. This means:

- The sheer volume of logs **created by modern distributed systems** is overwhelming observability teams
- Logs are **contributing to cost, not clarity**
- Engineers are **spending too much time** looking for the "needle in the haystack"

4 The #1 log challenge: extracting insight

According to recent research:

- **38%** of teams say their biggest pain is getting useful insights
- That's **12%** higher than the next most cited issue
- With rising scale and noise, **clarity is becoming harder** — and more critical — to achieve.

38% of teams say their biggest pain is getting useful insights.

5 Legacy tools weren't built for this

Most log management tools were designed in a previous era:

- Built **before Kubernetes and AI workloads**
- **Architected and priced for companies generating gigabytes of data each day** — not terabytes or petabytes
- Can't keep up with the **volume of data you create**

6 Three common — but flawed — approaches to solving the log data challenge

To solve data growth, teams typically choose one of three painful paths:

- **Keep it all** and pay for **100%** of the data, despite only using a small fraction of it
- **Drop potentially critical logs** and risk creating blind spots
- **Manually optimize** and waste engineering time while adding toil

SOLUTION

To fix the log data problem, you need visibility into which data your team actually uses and *how*.

Is your log solution ready for scale?

Use this checklist to assess where your current logging tool stands:

Runaway data volume and cost

- ☐ Is your provider enforcing pricing models that make scaling prohibitively expensive?
- ☐ Are teams able to understand log utility?
- ☐ Do teams get recommendations on how to reduce noise?

Performance and reliability at scale

- ☐ Are teams forced to route logs to low-cost storage solutions that cut costs while slowing investigations?
- ☐ Can teams investigate issues across any time range without compromising speed or reliability?

Signal-to-noise

- ☐ Are teams ingesting excessive low-value logs that create noise while masking critical signals?
- ☐ Are teams wasting valuable engineering time during incident response and slowing resolution due to sifting through mountains of data?

The new approach: log management that offers control

Instead of storing everything or guessing what to drop, teams are starting with one simple question:

What do we actually use?

That shift enables:

- **Smarter routing** — Only store what matters
- **Signal-focused ingest** — Less noise, faster insight
- **Sustainable cost** — Aligned to business value

What success looks like

Teams adopting log management have reported:

- **52%+ reduction** in logging costs
- **Faster troubleshooting** during high-pressure incidents
- **Happier engineers**

How Chronosphere can help

Chronosphere Logs transforms log management in containerized, microservices environments. By gaining clear insight into log volume and how your team uses its data, you quickly and easily identify the data that matters.

These insights power our recommendations to help you remove noise while preserving data value — all just in a matter of clicks. Our platform unifies all MELT data in one place and delivers fast performance no matter how much data you generate.

Control your log data volume and costs.

Join a demo to see how Chronosphere Logs help you control log data volume and costs in containerized, microservices environments

Book your spot!

[1] The state of log data: 6 trends impacting observability | [2] Logs: Love 'em, don't leave 'em