



Cribl Pricing Guide

Exploding data volumes shouldn't mean exploding data costs. Get the choice, control, and flexibility you need to manage all your telemetry data while keeping costs in check.



CONTENTS

03 Overview

05 Cribl Credits

08 Pricing by Product

08 Cribl Stream

09 Cribl Edge

10 Cribl Search

11 Cribl Lake

13 Cribl Insights

14 Configuration Options

16 Billing

17 Additional Resources

Overview



At Cribl, we pride ourselves on giving you the choice, control, and flexibility you need to manage your IT and security data—and that same approach applies to how we price our product suite and offerings. Our flexible pricing structure ensures you only pay for what you use, empowering you to minimize costs while maximizing the value of your existing tools and investments.

This guide will walk you through our simple and transparent pricing model. We offer flexible, usage-based pricing, whether you're collecting data from on-premises environments, or cloud infrastructure. The list price for a Cribl Credit is always one dollar, and our products burn down credits at different rates. Cribl's portfolio includes the following four products:



**Cribl
Stream™**

An observability pipeline



**Cribl
Edge™**

A scalable vendor-neutral agent



**Cribl
Search™**

A search and investigation solution



**Cribl
Lake™**

A turnkey data lake

We also offer a volume-based licensing model for Stream and Edge only, for customers who need to manage and run their own infrastructure entirely. Learn more about [Cribl Credits](#) or our [configuration options](#) below.



Cloud

PRODUCT	PRICING (1 credit = \$1)
Stream	Ingest pricing <ul style="list-style-type: none"> 0.32 Credits/GB per Cloud Worker 0.26 Credits/GB per Hybrid Worker <i>*Infrastructure pricing based on region</i>
Guard	<ul style="list-style-type: none"> 0.15 Credits/GB scanned
Edge	Ingest pricing <ul style="list-style-type: none"> 0.21 Credits/GB per Edge Node
Search <small>See cribl.io/pricing for details</small>	Federated Engine Lakehouse Engine Lakehouse Storage <ul style="list-style-type: none"> 0.05 Credits/GB per month
Lake <small>Lakehouse pricing see cribl.io/pricing</small>	Storage pricing based on 30-day retention <ul style="list-style-type: none"> 0.05 Credits/GB if data lake managed by Cribl 0.02 Credits/GB if self-managed
Insights	<ul style="list-style-type: none"> 450 Credits/month for 7 days rolling retention 900 Credits/month for 1 month rolling retention 1980 Credits/month for 3 months rolling retention 4500 Credits/month for 1 year rolling retention

Licensed

Per day volume of data measured in GBs

PRODUCT	DAILY INGEST	PRICING
Stream	Up to 500GB / day	\$50,000
	Up to 1TB / day	\$84,000
	Over 5TB / day	Contact us
Edge	Up to 500GB / day	\$50,000
	Up to 1TB / day	\$84,000
	Over 5TB / day	Contact us
Guard	Up to 500GB / day	\$29,000
	Up to 1TB / day	\$48,000
	Up to 2TB/ day	Contact us
	Over 2TB / day	Contact us

For the latest pricing details, visit www.cribl.io/pricing.

Cribl Credits

WHAT IS A CRIBL CREDIT?

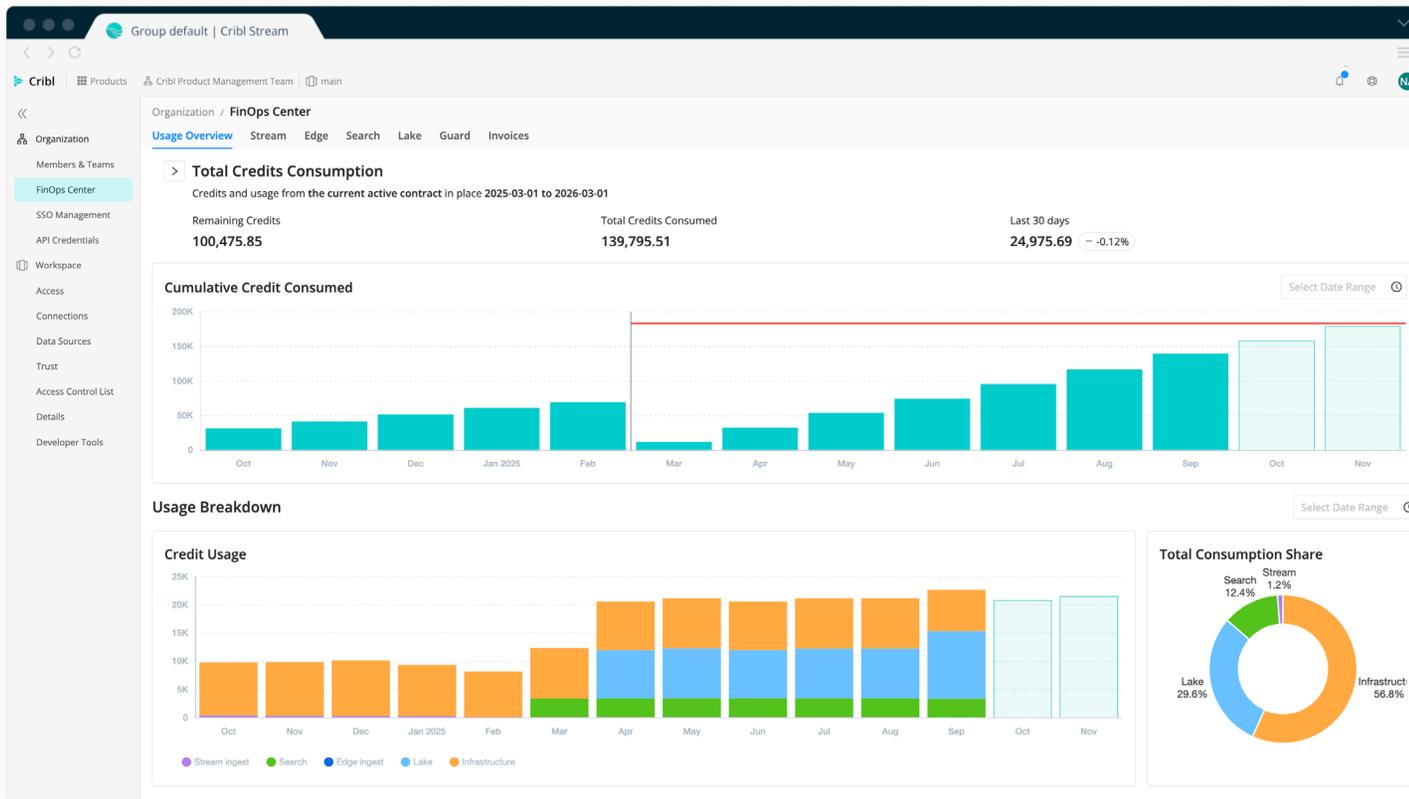
A Cribl Credit is a unit of measurement for Cribl billing across all Cribl products — Stream, Edge, Search, and Lake. The rate at which credits are consumed depends on the specific product being used. \$1 = 1 Credit.

Cribl Credits offer flexibility, allowing you to allocate your spending as needed. This means you can easily adjust for changes in data volume or seasonal usage fluctuations without exceeding your committed spend or having to initiate a new sales process.

You purchase your pool of Cribl Credits upfront — your account team will do the math to convert data volumes to credits to help you estimate how many credits you should purchase — and as you use Cribl products, you draw down from that pool of credits every month to use it towards any Cribl product or service.

As Cribl continues to innovate, you will be able to use your credits to try new products and offerings as they become available. **Talk about future-proofing!**

Within the **Cribl.Cloud** platform, you can keep tabs on your credits and check usage: from the sidebar, Go to **Account > Organization > FinOps Center**. New and renewing Cloud customers will be subject to minimum commitment and credit rollover terms. At the end of each year, no more than 20% of the credits purchased for that year may be rolled over into a new subscription term. Any unused credits beyond the 20% rollover amount for that year will be forfeited. Note that credits only rollover if you renew for at least the same amount, and all unused credits will be forfeited if you choose not to renew.





Your Cribl.Cloud spend will vary depending on your product usage.

- For Stream and Edge, your total spend equates to data ingest plus the infrastructure resources needed for processing your data. We charge on ingress, but you never pay for egress, so you can send your processed data to as many destinations as you'd like.
- For Lake, it equates to the amount of compressed data stored.
- For Search, three pricing options are available:
 1. **Flat rate:** Opt for a fixed monthly fee based on the size of your dataset, independent of usage.
 2. **Usage-based:** The amount of compute resources used to deliver search results. You'll be charged in compute hours, while most searches take less than a second.
 3. **Lakehouse:** Only searches time-sensitive Lakehouse datasets at high speeds. Flat fee, priced by tier (Small to 2XL), for given data rate (GB/day).

Universal Subscription

Customers who choose to manage their own workers for Cribl Stream and Edge can leverage Universal Subscription to still benefit from Cribl's consumption-based licensing model. With Universal Subscription, existing on-prem Edge and Stream licenses can be converted into an equivalent amount of credits, allowing you to use any Cribl product including Cloud. This replaces your Cribl license while allowing you to maintain your self-managed data collection and worker processing without multiple plans, or subscription types.

Universal Subscription simplifies operations by eliminating the need for separate credit pools for each org, enabling you to balance heavy-use orgs with lighter-use ones seamlessly. It also lets you run Stream and Edge deployments closer to your data, providing the flexibility to manage unlimited self-managed deployments without requiring separate licenses or sharing licenses across environments.

Additionally, with Universal Subscription, you can easily track spend and usage across multiple environments and across Cribl products through the FinOps Center — all under one billing plan. [Learn more.](#)

Pricing by Product



[Cribl Stream](#) is an observability pipeline for parsing and processing of data. Cribl Stream helps you get the data you need, in the formats you want, routed to the right destinations.

Things to consider when using Stream, there are two cost components:

INGEST COST

Two types of Cribl Workers can process data:

- a. **Hybrid Workers** are managed and maintained by you. It leverages the credits you purchase, and consumes **0.26** of a Cribl Credit per GB of data processed.
- b. **Cloud Workers** are managed and maintained by Cribl. It leverages the credits you purchase, and consumes **0.32** of a Cribl Credit per GB of data processed.

Your Cribl account team can help you determine which type of Workers are best to suit your specific needs.

INFRASTRUCTURE COST

Cribl handles the infrastructure (networking, compute, memory, etc.) needed for your Worker Group. Set your expected ingestion in MB/sec, and based on the configured ingest rate, Cribl will automatically provide and manage the necessary infrastructure.

- a. Credits are consumed by Cribl-managed Cloud Workers hourly once they are provisioned. Credits are consumed even when data is not passing through
- b. Infrastructure costs vary depending on the region your Workers are deployed.
- c. See cribl.io/pricing for more details on infrastructure pricing.

With this pricing model, you pay for ingest and for the Workers processing your data. You hold the reins of your Cribl deployment and can control your Stream usage.

[Cribl Guard](#) uses AI with human-in-the-loop controls to detect and protect sensitive data—like payment cards, passports, and Social Security numbers—as it flows through Cribl Stream. With 200+ policy rules, teams can review or override actions in real time, while data is masked or blocked before storage to ensure PCI DSS, HIPAA, and GDPR compliance and reduce exposure risk. Guard is priced at **0.15** credits per GB scanned.



[Cribl Edge](#) is a vendor-neutral, intelligent agent designed for the variety and scale of today's modern architectures. With a unified telemetry collection system, you can have hundreds of thousands of agents at your fingertips to automatically discover and collect endpoint data from your Windows, Linux, and Kubernetes environments.

Any data that Cribl Edge collects will burn down your Cribl credits at a rate of **0.21** credits per GB.

NOTE

Data sent from Cribl Edge to Cribl Stream is only counted once—at the initial point of collection. There is no double counting. Additionally, data dropped before routing to any single destination, including DevNull is not included in your ingest rate. Read more about [dropping events in this blog](#).



[Cribl Search](#) delivers a high-performance search experience across all data — hot or cold, in-place or directly ingested — so IT and security teams can investigate faster, streamline workflows, and use AI-guided exploration to get from questions to answers fast.

Cribl Search is built around three components:

- 1. Lakehouse Engine** is a tiered configuration that bundles ingest, compute, and high-speed search into a single unit. Each tier provides a fixed allocation of resources, with sizing recommendations based on daily ingest volume and expected investigation workloads.
- 2. Lakehouse Storage** is used to retain telemetry stored in Cribl Search. Storage scales independently from compute and is billed at standard Cribl Lake rates (0.05 credits per GB per month), allowing you to adjust retention based on your needs.
- 3. Federated Engine** enables search across external data sources such as Amazon S3, Azure Blob Storage, and other supported environments without moving the data. It is also offered in tiers and can be resized at any time, giving teams flexibility to scale search capacity and control spend.

In your Cribl environment, you can always reference the dashboard at [Monitoring > System > Licensing](#) to review your actual license or billed usage consumption.



[Cribl Lake](#) is a turnkey data lake solution that makes it easy to store, access, and retrieve data when you need it. With unified retention, security, and access control policies, Cribl Lake streamlines workflows and stores data in open formats to allow a schema-on-need approach.

Cribl-managed Lake is priced based on the volume of data stored at **0.05** credits per GB. The total amount you're charged depends on the volume of data stored and how long you have stored the data during the month. You are only charged for the amount of compressed data that is stored per month, not the raw ingested data.

Self-managed Bring Your Own Storage (BYOS) is priced at **0.02** credits per GB. Keep your data under your control while leveraging Cribl's powerful processing, management, and analytics. Get a unified view of federated data with seamless access and querying—no matter where it resides.

In addition to using Cribl Stream to route data to Cribl Lake, administrators can now use the new [Direct Access](#) option *if Stream data shaping services are not required*. Direct Access ingest is charged at **0.10** credits per GB for ingest, plus the standard [Lake storage charges](#).



Cribl measures your storage based on your usage daily. We also include two Cribl-specific datasets, [cribl_metrics](#) and [cribl_logs](#) for free for 30 days.

There are three components to Cribl Lake pricing:

VOLUME

The amount of data being stored.

RETENTION

How long data is stored for.

COMPRESSION

The amount of data zipped up to save on costs.

To figure out how many credits you need:

1. Determine the daily ingest or amount of data that needs to be stored in GB (volume)
2. How long you need to store the data for (retention). Note: Data compounds monthly until it reaches its retention period, then stabilizes in usage monthly.
3. Use the [sizing calculator](#) to determine how many credits you need.

Quickly spin up dedicated datasets in Lakehouse to run fast, flexible searches on your most recent data for real-time analysis. Easily deprecate them when no longer needed to save on costs.

Lakehouses are priced by tier as a flat fee for the infrastructure provided. Tiers range from Small to 2XL — similar to t-shirt sizes — based on Capacity GB/day. Each tier has a fixed monthly price, which can be modified anytime. See cribl.io/pricing for more details about Lakehouse pricing.



[Cribl Insights](#) provides centralized monitoring and alerting across your Cribl environment, with up to a year worth of context retained. Instead of reacting to broken pipelines or missing data, Insights defines what 'normal' looks like and flags deviations instantly. From silent log drops to pipeline slowdowns, actionable alerts flow straight into your workflow—helping teams and agents stay ahead of issues, optimize costs and resources, and keep data moving.

Pay only for the retention you select. Usage is calculated daily on a rolling basis and billed monthly against your annual commitment.

Enterprise Cloud customers get **48 hours of metrics and alerts included** automatically in their workspaces.

Extend monitoring retention from 7 days to 1 year for robust auditing and root cause analysis.

7 DAYS

Rolling retention

450

credits per month

1 MONTH

Rolling retention

900

credits per month

3 MONTHS

Rolling retention

1980

credits per month

1 YEAR

Rolling retention

4500

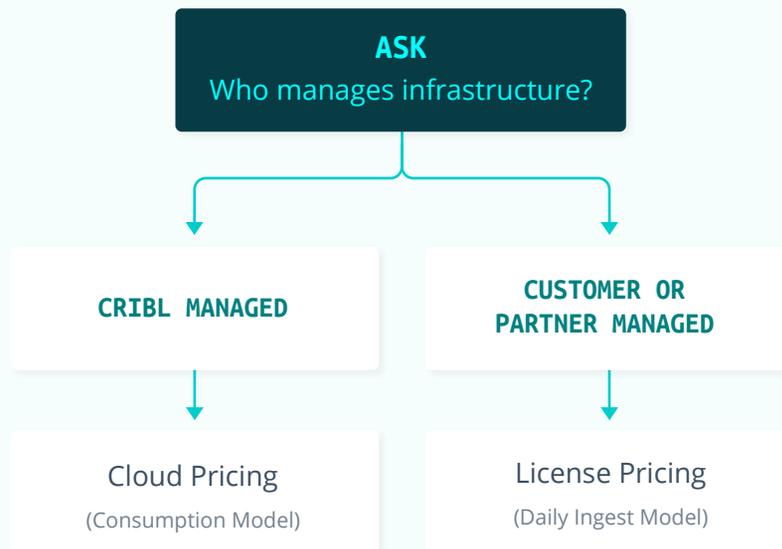
credits per month

Configuration Options



There are two ways to collect and manage your data with Cribl, either through Cribl-managed Workers or self-managed. You can rely on Cribl-managed workers for processing data in Cribl.Cloud, or with hybrid-workers they can collect data from on-premise environments sent to your Cribl.Cloud environment. Additionally, Lake and Search are only available in Cribl.Cloud, there is no self-managed option. If you need to manage your Cribl Leader and workers on your own premises, the software licensing model is available.

Which Pricing Model Should I Use?





Cloud pricing

Customers with either Cribl-managed or hybrid workers operate on a consumption-based model. Customers purchase credits and consume them flexibly when they use Cribl products. This approach aligns with industry-standard consumption models, offering a “pay for what you use” experience. Cribl Search and Cribl Lake products are exclusively available in Cribl’s cloud platform. This model makes it easy to scale up or down as needed, and track usage directly on the platform. No surprises or getting in trouble with procurement.



Self-managed software licensing

For customers who are exclusively using Stream and Edge with self-managed workers, license-based pricing based on daily ingestion rates is available. Your Cribl rep will help you figure out what your estimated daily ingestion rate is in a Proof of Value (POV), per day volume of data measured in GBs or TBs.

Self-managed Cribl Stream is licensed based on the peak amount of data processed daily, measured in gigabytes (GB) or terabytes (TB).

IMPORTANT

- If your data volume consistently exceeds the licensed amount, it is your responsibility to remain compliant by notifying your Cribl account team of the overage and purchasing a license for the new daily volume. But don’t sweat it — there are no hard stops to data if license volume is exceeded, we won’t turn off your data!

You can easily monitor your license usage within your Cribl environment by navigating to [Monitoring](#) > [System](#) > [Licensing](#) in the dashboard.



EXAMPLE

If you choose a 3TB per day license for a 12-month period, you can process up to 3TB of data through Cribl Stream each day during that time.

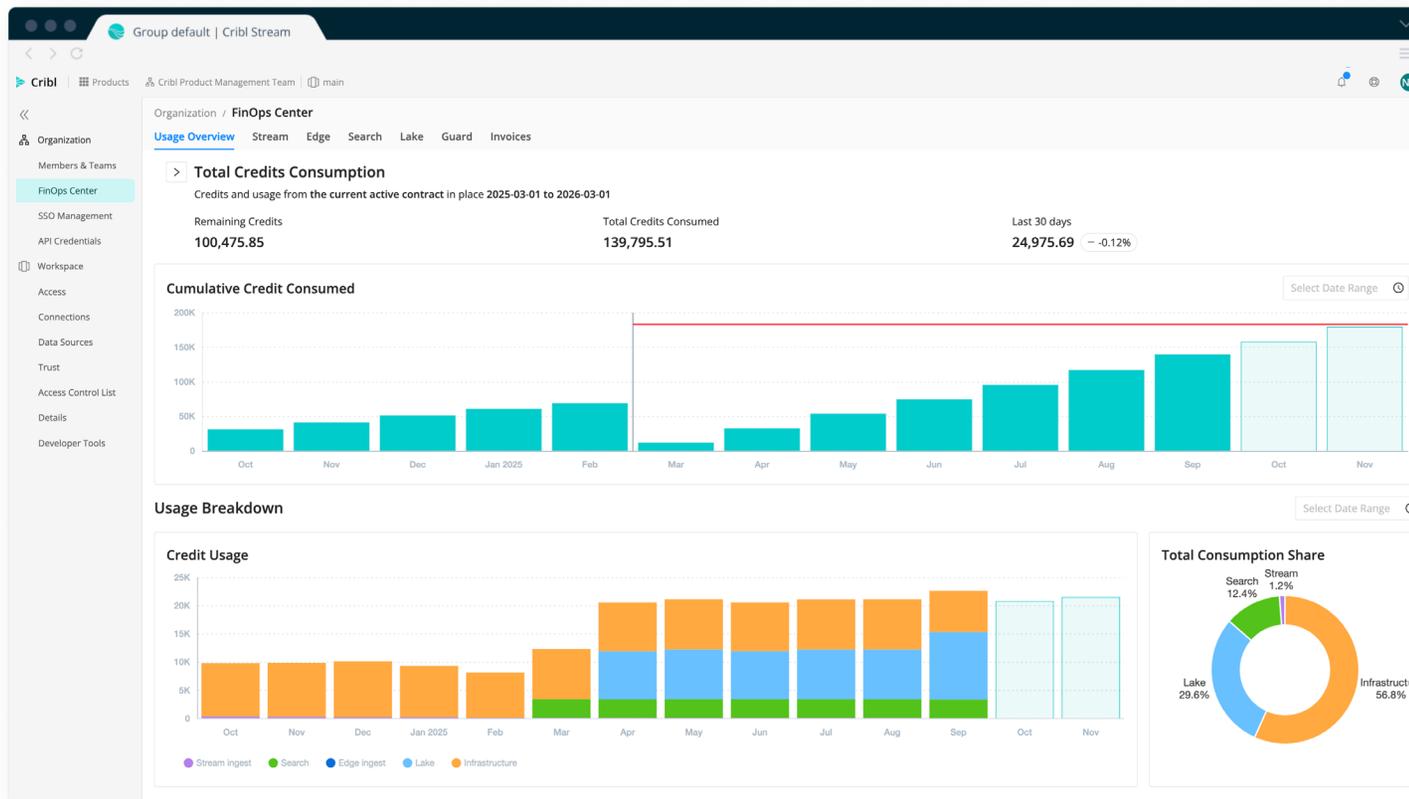
To determine the appropriate volume-based license, we recommend discussing your daily data ingestion volume with your system of analysis (SOA) team.

We meter ingestion starting from when data is first collected by a Cribl Stream Worker. There is no cost to send data between Edge Agents and Stream Workers, and no cost to send data between Stream Workers. Edge and Stream share the same license.

Billing



From the **FinOps Center** tab, under your Cribl Organization, you can track your credit usage throughout the month, and compare your monthly invoices. Usage accrues throughout the month and is drawn down automatically from your credits; your invoice will post monthly with the total credits used across all products. Draft invoices are also available for mid-month usage tracking.





Additional Resources

- [Cribl Pricing Page](#)
- [Cribl Sizing Calculator](#)
- [Blog: How Cribl's Cloud Pricing Works](#)
- [Free Training with Cribl University](#)



Cribl, the AI Platform for Telemetry, empowers enterprises to manage and analyze telemetry for both humans and agents. Trusted by organizations worldwide, including half of the Fortune 100, Cribl bridges the gap between AI ambition and infrastructure reality. No lock-in. No data loss. No compromises. Cribl's vendor-agnostic platform ensures data remains portable and interoperable. By cost-effectively handling increasing data volume and variety without delay, Cribl gives enterprises the choice, control, and flexibility to build what's next.

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

Follow us: [LinkedIn](#) and [X](#)

©2026 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. BGDE-0002-EN-03-09-26