

Multicloud: the good, the bad, and the unmanageabl.

How the right data management platform can help optimize performance, cost, and security in your multicloud environment.



Table of contents.

03 Multicloud adoption is on the rise.

17 What is a data engine?

30 In your multicloud era? Get equipped with a Data Engine.

07 Challenges with multicloud.

19 Cribl for multicloud success.

31 Get 1TB/Day for free with Cribl.Cloud.

12 A data engine for multicloud success.

27 Customer story: Autodesk.

Multicloud adoption is on the rise.



What is multicloud?

Let's start with what isn't multicloud. That's when organizations only use one cloud computing service for hosting and storage. Now, when organizations use at least two public cloud providers, it means they are taking a multicloud approach. **Multicloud simply means more than one cloud.**


A multicloud environment can be made up of private, public, and hybrid clouds, and can be combined with on-premises operations along with apps and services running on various cloud providers.



Public cloud offerings include third-party cloud service providers; the most popular public clouds are Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP), Alibaba Cloud, and Oracle Cloud. These companies manage everything from the hardware, storage, and network, which are all used and accessed by many cloud tenants through a web browser.



Private cloud means the service is solely dedicated and controlled by a single organization. That organization is responsible for maintaining infrastructure and ensuring compliance requirements are met. Private clouds offer the ability to adjust security posture to the specific needs of the business. Since fewer people have access to private clouds, data becomes less vulnerable to be compromised.



Hybrid cloud combines on-premises infrastructure with a public cloud, and allows data to move between the two. Many organizations take advantage of a hybrid cloud approach because they get the best of both worlds — keeping sensitive data secure on-premises and maximizing their existing investments while innovating fast by building newer technologies in the cloud.



On-premises data center
(Infrastructure is hosted locally and owned by an organization)



Private cloud
(Owned and controlled by a single organization)



Public cloud



+



Hybrid cloud




+



Multicloud
(Any combination of public clouds)

Why the shift to multicloud?

Organizations are realizing it's risky to put all their eggs (and data!) in one basket. According to Flexera's 2022 State of the Cloud Report, 89% of organizations have a multicloud strategy, with 80% taking a hybrid cloud approach.¹ And in another report by ERG, most use between 2 and 4 public clouds.²



of organizations have a multicloud strategy, with 80% taking a hybrid cloud approach.

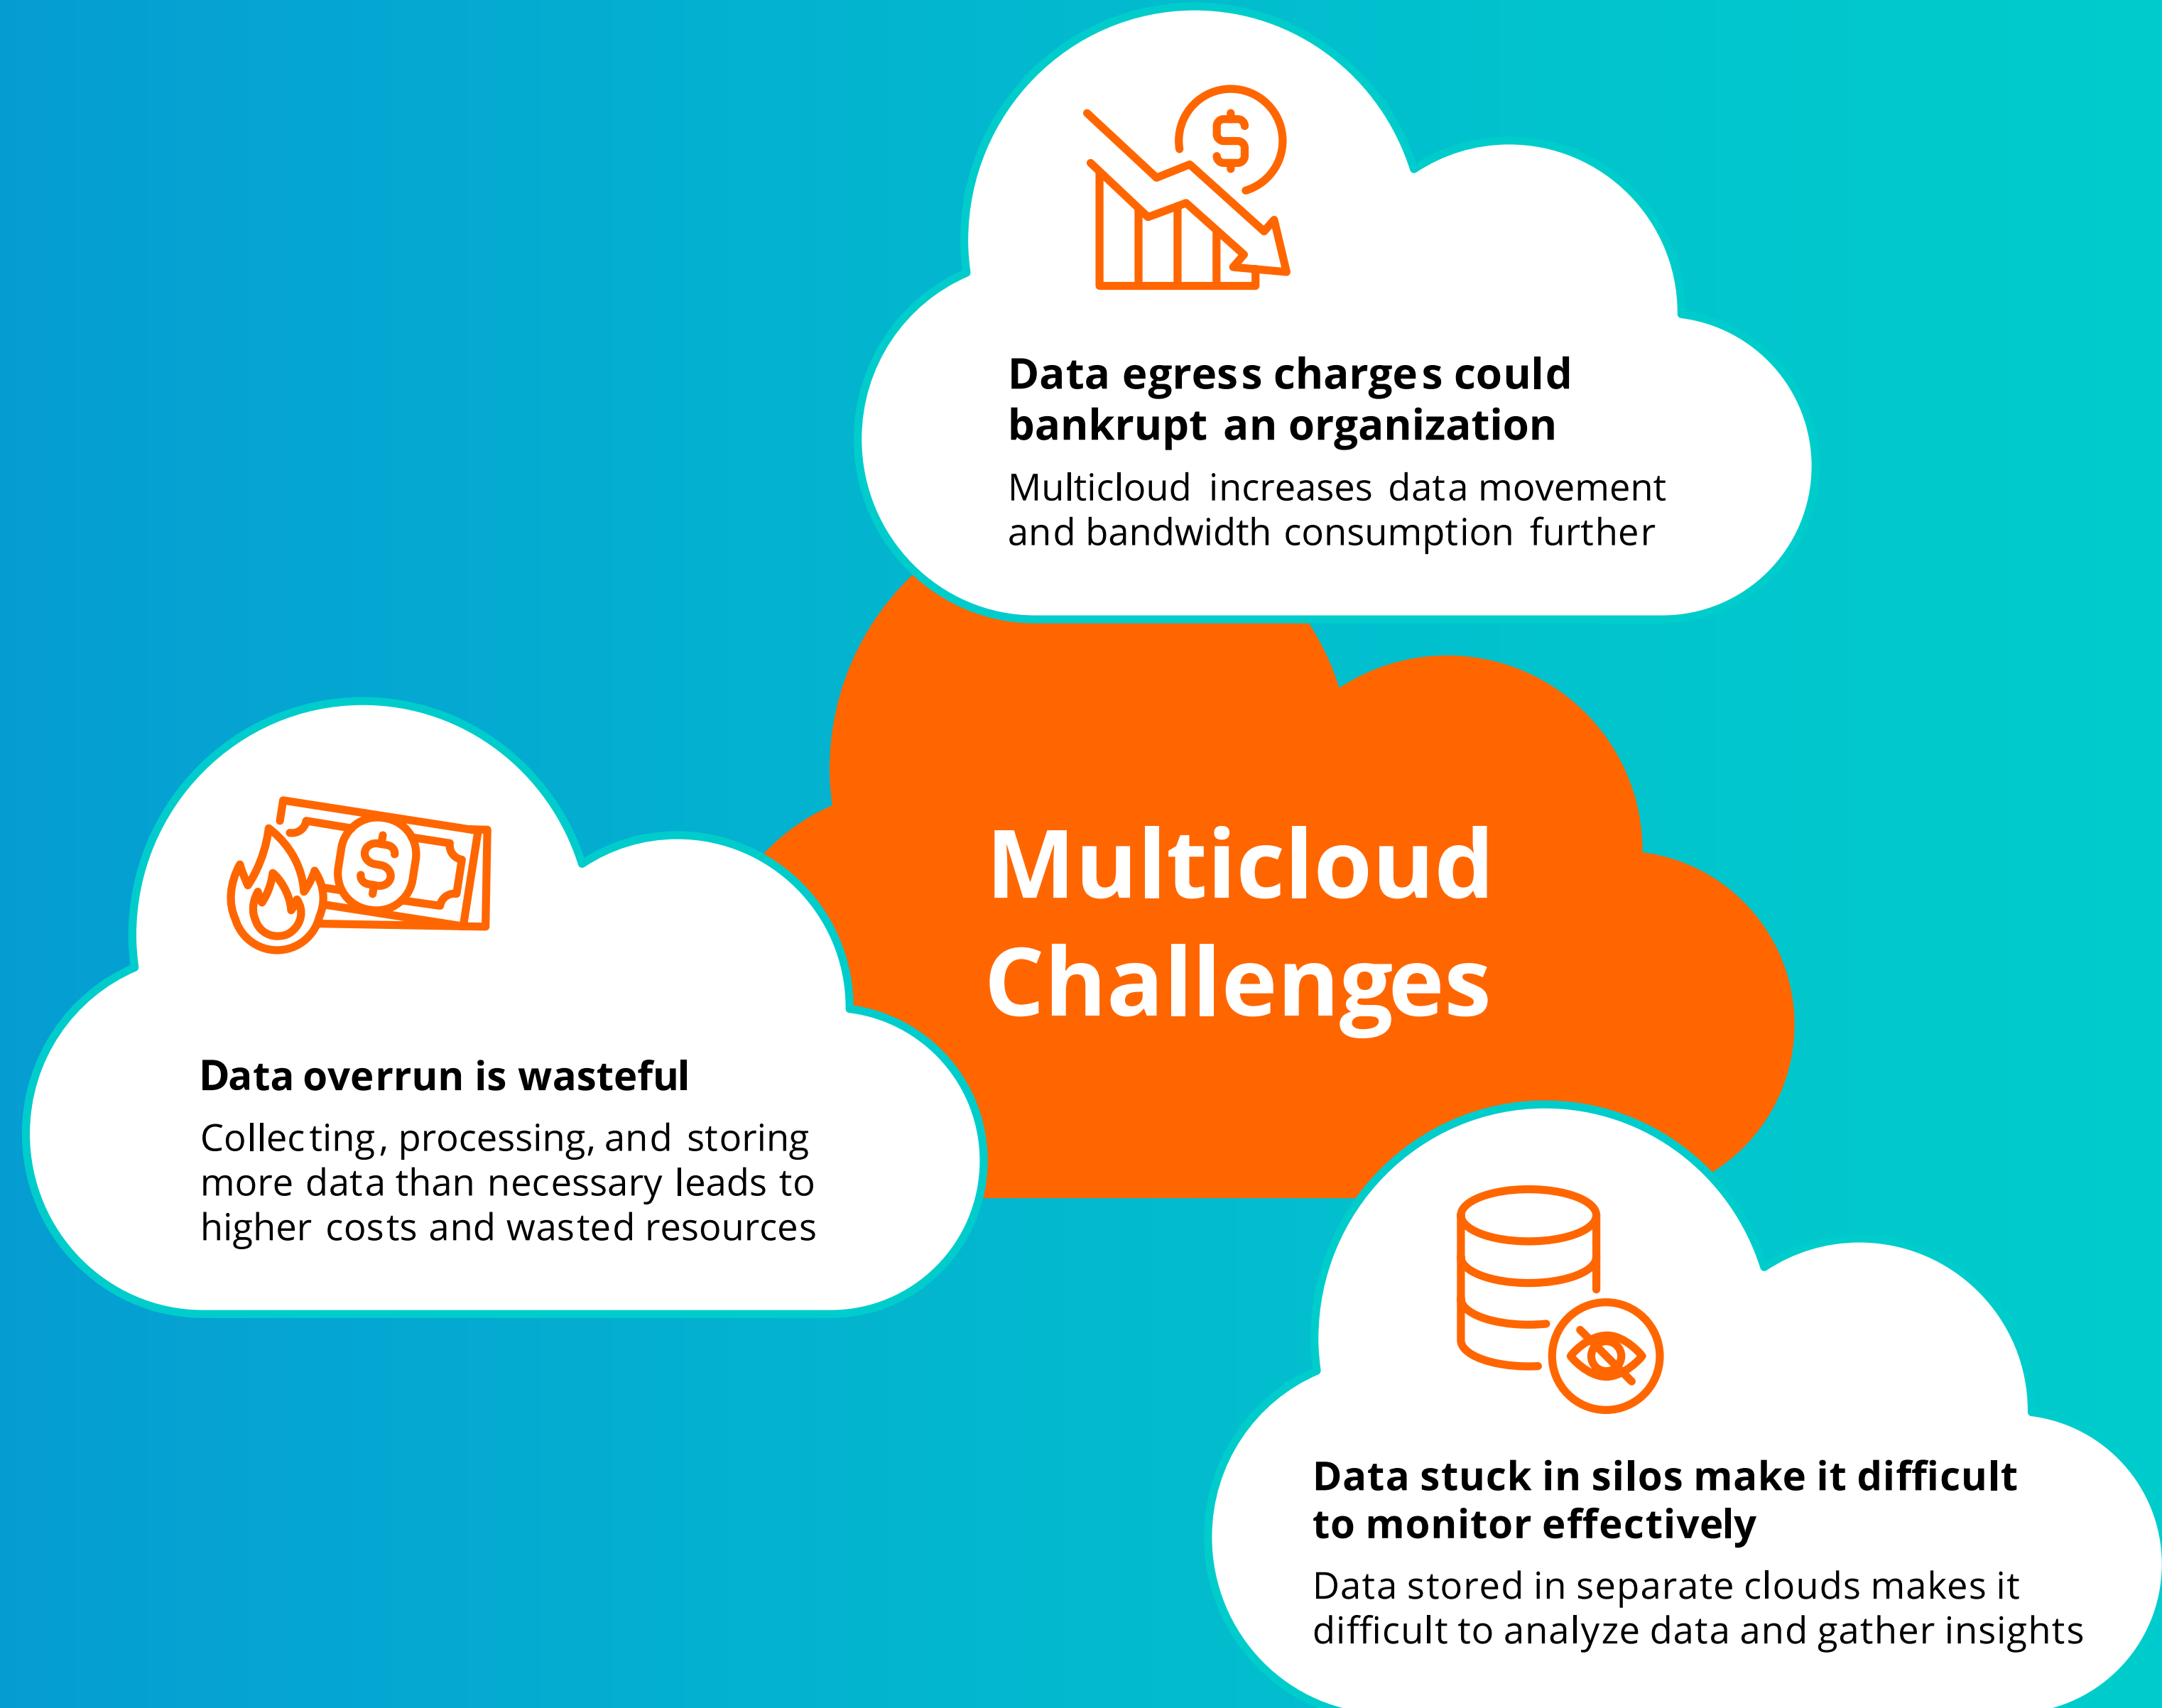
The ongoing trend to adopt multicloud has become more popular recently because it allows IT and engineering teams to choose the best cloud platform for each of their workloads, allowing them to take advantage of all the great features and capabilities, resiliency, performance, and price. It also prevents vendor lock-in and prevents them from a situation where they're strong-armed into paying higher costs because they've become vulnerable and reliant on just one cloud computing service.

Organizations may also adopt multicloud unintentionally through acquisitions, mergers, and migrations. It's also not uncommon for different teams to choose which cloud provider they want to use, without running it by the larger team or getting formal approval. When this happens, IT leaders then need to figure out a proper multicloud strategy that will set their teams up for future success.

¹ [Flexera's 2022 State of the Cloud Report](#)

² [ESG Global Blog - Multi-cloud is Sweet](#)

Managing a multicloud environment can be a daunting task, it requires dealing with different cloud providers each with its own set of tools and interfaces. Having to integrate and coordinate many different systems together is anything but simple.



Challenge 1:

High data egress charges.

81% of businesses cite managing cloud spend as the top cloud challenge they face. This will only get more difficult as cloud consumption grows.³

One of the greatest challenges organizations have is dedicating the time and energy to move apps and data between data centers and across multiple public cloud services.⁴ Transferring data in and out of cloud accounts can quickly cause an organization to go broke. Compared to on-premises solutions, cloud services actually require more bandwidth to transfer data. More egress means more costs and you can't get around it — you have to transfer data to work in the cloud.

There are ways around keeping data transfer costs reasonable. On AWS, you can cut down on having to send outbound data, since transfer costs are higher when sending data out vs in. You can also keep data transfers confined within a single region. But these approaches can only go so far in a multicloud environment.

Multicloud will increase data movement and bandwidth consumption further. And for companies with a multicloud strategy, managing and optimizing costs across providers is nearly impossible due to proprietary and dissimilar pricing models.⁵

How a data management strategy can help



Reducing cloud spend is a top priority for businesses of all sizes and industries. A data engine can help tame bandwidth charges by reducing data volume. Easily eliminate duplicate fields, null values, and any elements that provide little analytical value. That way, only necessary data gets transferred.

³ [Flexera's 2022 State of the Cloud Report](#)

⁴ [ESG Global Blog - Multi-cloud is Sweet](#)

⁵ [2022 Anodot State of Cloud Cost Report](#)

Challenge 2:

Immense data overrun.

A vast majority of data ends up becoming unusable, and this happens for many reasons. Not only is it expensive to collect all the data that gets generated, but it's also impractical to then move and store it all somewhere. And then there's the need to query data so it can be searchable.

94% of organizations are overspending in the cloud, with a lot of this avoidable. One of the top reasons for this overspend is overprovisioned resources.⁶

Data overrun occurs when organizations collect, process, and store more data than necessary, resulting in higher costs and

wasted resources. This is often overlooked, and organizations end up paying more for data than their production systems. Cloud platform vendors charge by data volume, which means that observability plans may create additional egress charges from the cloud provider.

According to the Anodot 2022 report, one-third of cloud spending is wasted, with cloud resource expenditures going unused or underused, resulting in increased operational expenses without providing any business value. Additionally, 44% of IT executives agree that at least a third or more of their spend is wasted.

How a data management strategy can help



Organizations can use a data engine to reduce log volume, automate cloud cost management, and get visibility into cloud spend across multicloud environments. A data engine can help collect, process, and store only the data that is necessary.

⁶ [HashiCorp 2022 State of Cloud Strategy Survey](#)

Challenge 3:

Data silos and monitoring woes.

Data silos occur when data is stored in separate clouds that aren't integrated. This can lead to inconsistencies in the data and also hinder teams from being able to access and use the data they need to do their jobs. It also makes it difficult to analyze data and gain cohesive insights from it.


When data is siloed, gathering insights become a challenge. Each cloud provider offers its own dedicated monitoring and analysis tool, but none are able to work across multiple clouds and services. This makes it difficult to gain a holistic view of the entire infrastructure and track performance and

issues. Without proper visibility across cloud platforms, teams have a more difficult time collaborating with each other. This makes it significantly more difficult to detect and resolve issues, and leaves them stagnant when trying to find potential areas for optimization and improvement.


How a data management strategy can help




Managing a multicloud environment requires careful planning, coordination, and intention to prevent data silos. A data engine can ensure effective and efficient operations by providing one unified platform for end-to-end visibility, full control of data, and easier cloud operations.



of businesses cite managing cloud spend as the top cloud challenge they face.⁷



of IT executives agree that at least a third or more of their spend is wasted.⁸



of organizations say multicloud is working. Out of those who have already adopted a multicloud approach, the vast majority say it is already helping their organization advance or achieve their business goals.⁹

⁷ Flexera's 2022 State of the Cloud Report

⁸ 2022 Anodot State of Cloud Cost Report

⁹ HashiCorp 2022 State of Cloud Strategy


A data engine for multicloud success.

Behind every successful multicloud strategy is a robust data management plan. Managing data requires a way to watch and understand the health and performance of various systems, applications, and services being used. This is especially crucial in a multicloud environment, where cloud resources and services are distributed across different cloud platforms and are accessed and managed through different tools and interfaces.



Benefits

A data engine in a multcloud environment can help to:



— Gain a holistic view of the entire multcloud landscape.

A data engine offers panoramic views into all data across all cloud providers so that anyone within an organization — with the right access controls and data governance and compliance rules in place — can view.

When teams have easy access and real-time visibility to the data they need, they're able streamline release cycles and optimize for better user experiences.

Benefits

A data engine in a multicloud environment can help to:



— **Understand the health of systems by tracking and measuring performance and cloud resources.**

A data engine makes it quick and easy to pinpoint exactly where service degradations are occurring. Especially in a multicloud environment, this is crucial because performance bottlenecks and resource capacity issues can be extremely detrimental.

A data engine provides the insight needed to speed up resolution, and optimize resources — such as CPU, storage, memory — so teams can make smarter capacity planning decisions and save on costs.

Benefits

A data engine in a multicloud environment can help to:



Discover trends to predict and prevent issues from happening in the future.

When an issue occurs, a monitoring tool will tell you where the issue is. A data engine takes it a step further by monitoring trends, getting ahead of issues, and letting you know before something goes down.

The power of the data engine is that it proactively tracks how systems perform, and can predict and prevent similar issues from occurring again in the future. A data pipeline, which is part of a data engine, opens doors to new possibilities by allowing you to ask questions of your data that you didn't even know you needed to ask beforehand.

Benefits

A data engine in a multicloud environment can help to:

— Enhance security and compliance posture.



Data governance and security are integral in any organization, especially when working in the cloud. There's a huge liability keeping sensitive data in various cloud platforms, and security teams need to ensure everything is properly secured.

A data engine allows you to easily centralize data governance principles. You only need to put the guidelines in once, and those automatically apply to multiple checkpoints. A data pipeline also configures data streams for maximum protection. Redaction and masking keep sensitive data private. You want to protect your customers and limit liability, and it's easy to do just that with a data engine.

What is a data engine?



IT and Security teams face a deluge of uniquely challenging combinations of structured, unstructured, and semi-structured data. A data engine offers a clear path to make sure the right data is in the right place, and can be accessed when it's needed. Each organization adopts its own set of analytics tools and enterprise architectures to effectively deal with its distinct data challenges. A data engine should be designed to provide users with the ability to move processing intelligence to the optimal point in the enterprise architecture, abstracting the complexity of data exploration, retrieval, and processing to a simple, cost-effective user experience.

An ideal data engine should be:

Open – Own your data in open formats, in your storage, with an open ecosystem. No lock-in, ever.

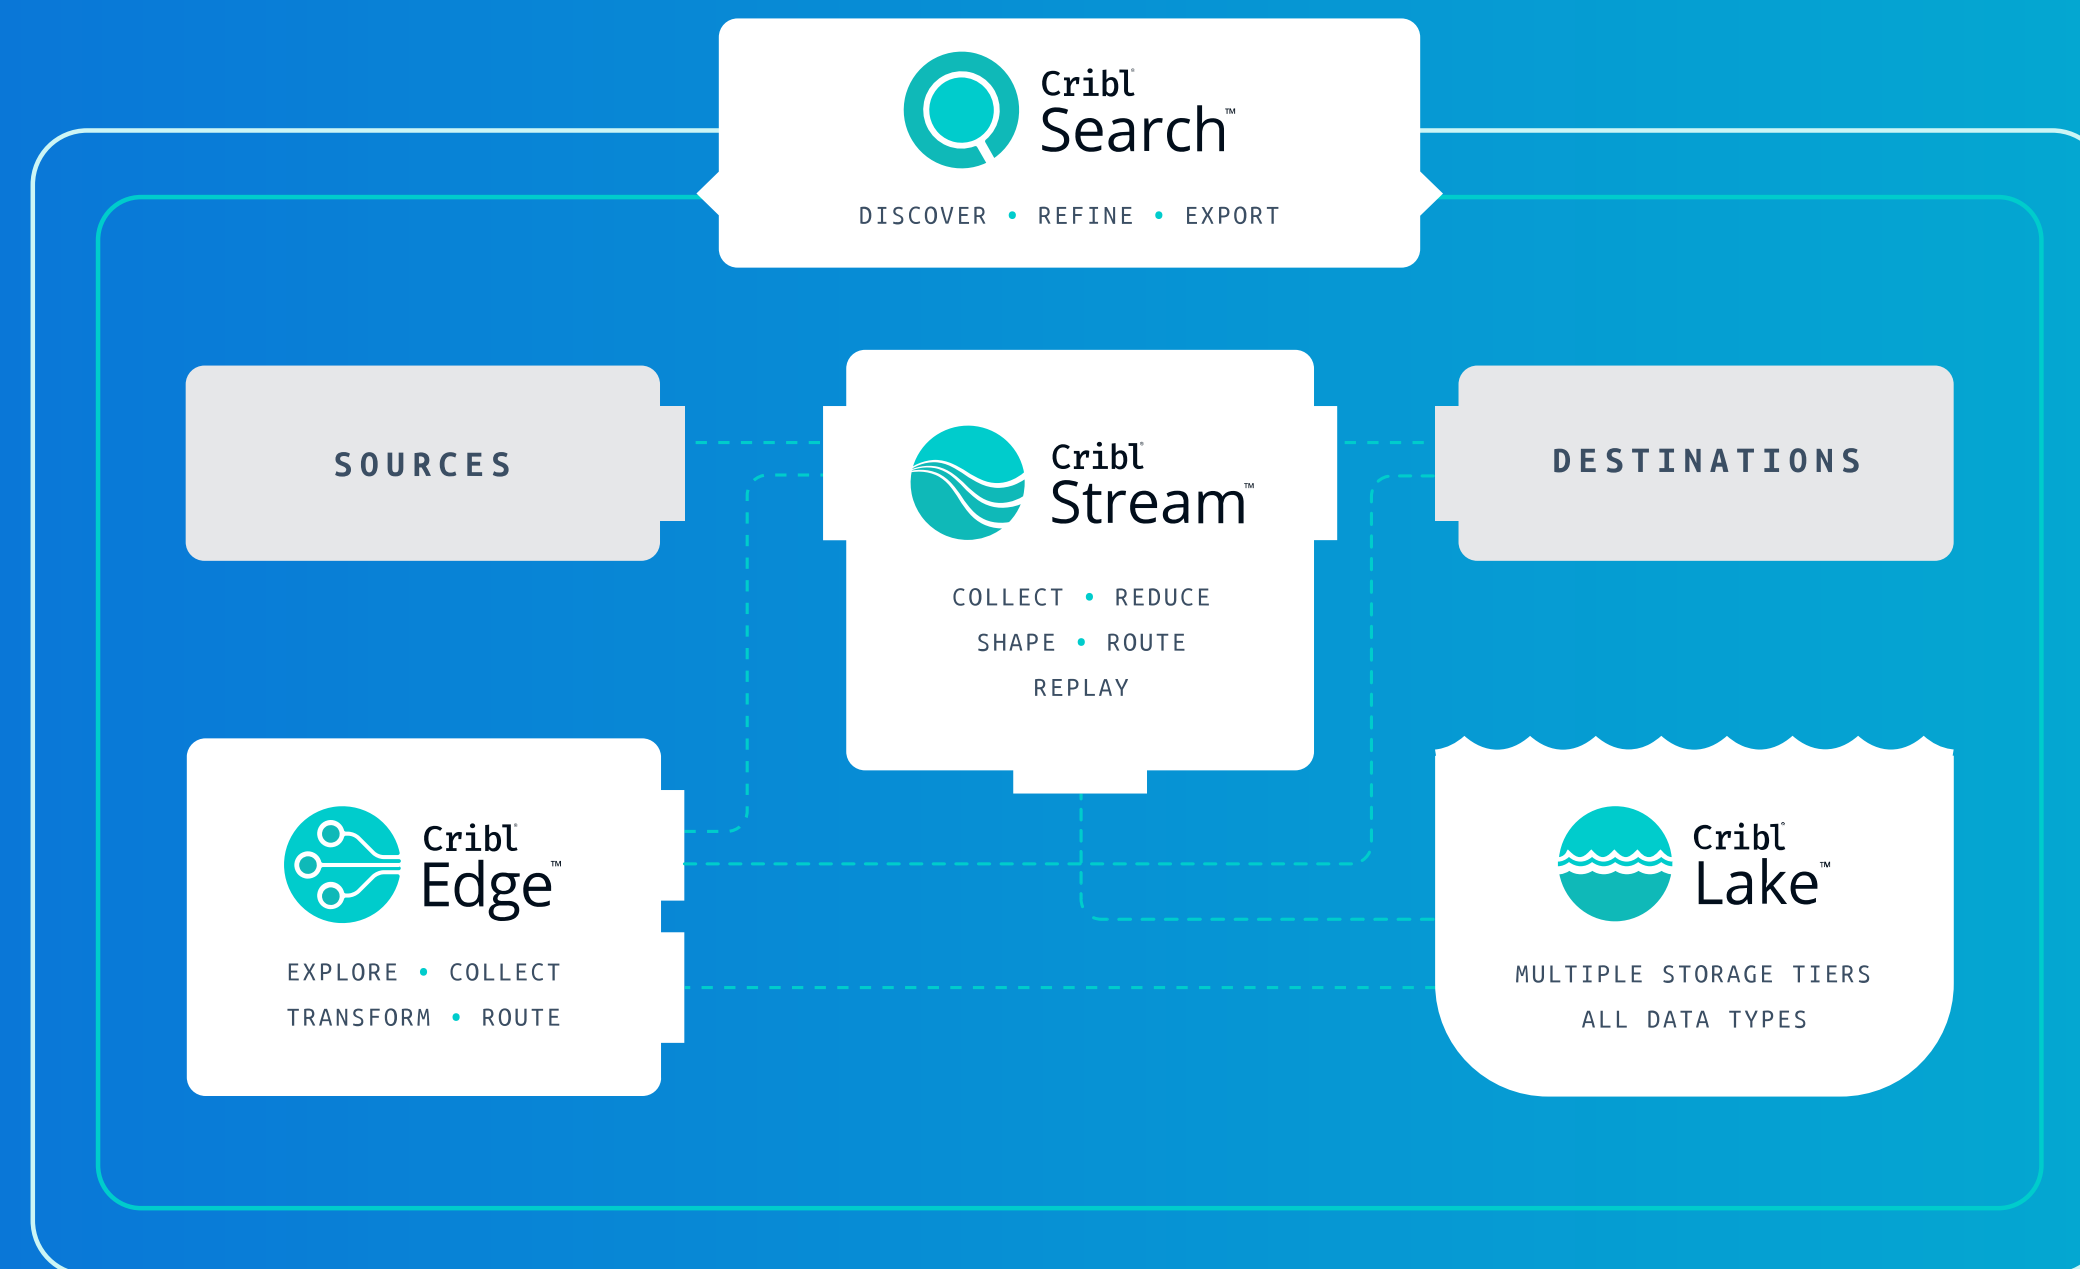
Focused – Software built for you. Focused on meeting your needs and adapting to your environment.

Performant – Performs at global scale for petabytes of data. Designed for the scale and complexity of your data.

Versatile – Highly customizable and adaptable. Can be deployed independently or as a unified engine.

Cribl for multicloud success.





Organizations are struggling with never-ending data growth, making it difficult for IT and security teams to gain control over routing, control over costs, and control over their multicloud environments. Moving data from all different sources into the right tools, and into the right cloud services, can mean sacrificing control and flexibility. Teams are having to reconfigure architectures and data flows to ensure parity and visibility, all while keeping a handle on ingress and egress charges.

Cribl puts your IT and Security data at the center of your data management strategy and provides a one-stop shop for analyzing, collecting, processing, and routing it all at any scale.

Cribl's suite of products, all built on a unified data processing engine, includes **Stream**, **Edge**, **Search**, and **Lake**. Gain the flexibility to build a data engine that's right for your IT and security infrastructure and the control to let you analyze, collect, process, and route telemetry data at any scale with total freedom. Mix and match products or go full beast mode with the whole enchilada — the choice is yours.

Collect, process, and route data from various sources to the right destinations.



Cribl Stream, Cribl's flagship product, is the leading observability pipeline in the cloud. Stream helps gather, process, and distribute data from various sources to the right destinations in a flexible and efficient way. Teams can create and apply intelligent routing logic to reduce data transfer within the cloud or between cloud providers. This helps reduce log volume to control costs and improve system performance.

With Cribl support of PrivateLink in AWS, teams can save significantly on bandwidth costs. Cribl helps transit data over less expensive methods, and makes it possible to look at how all systems work together. Teams can set rules for how data travels between different clouds and systems, and get a better understanding of how all data is related.

Not only that, Stream can help reduce log volume and lower egress charges. Easily eliminate duplicate fields, null values, and any elements that provide little analytical value. Filter and screen events with dynamic sampling, or aggregate log

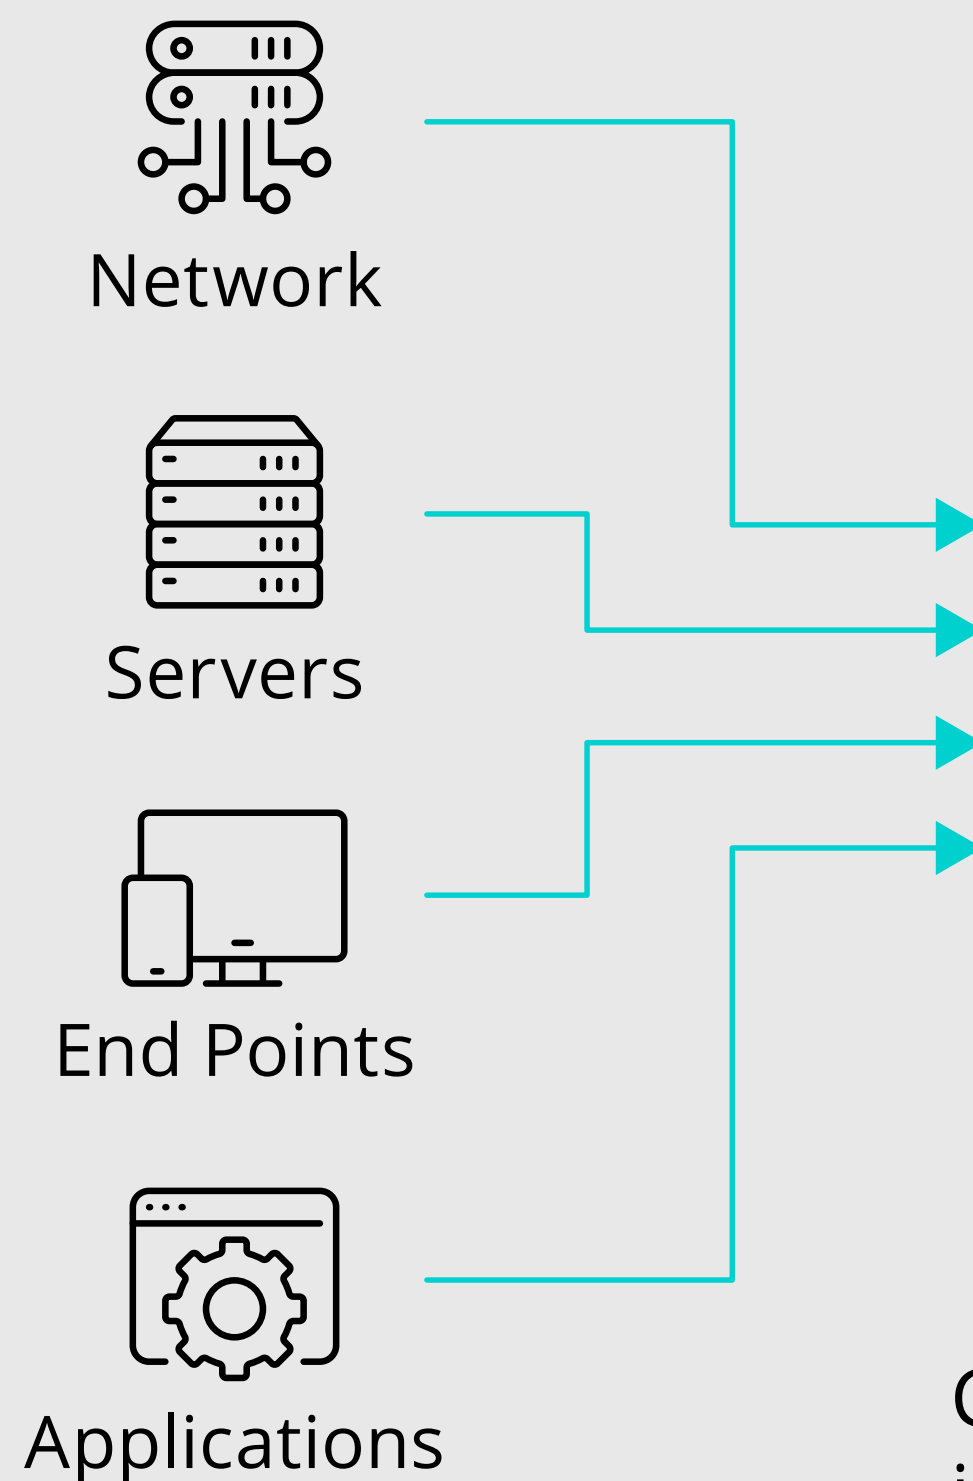
data into metrics for massive volume reduction. Reduce without worry: you can keep a full-fidelity copy in a low-cost destination and replay it back if needed.

One of the greatest concerns enterprises have about working in the cloud is security. Moving data between systems and cloud platforms leaves sensitive data vulnerable to attacks. But with Stream, you can get peace of mind with data masking, governance, and compliance — ensuring you have full control and protection over data routing.


Stream makes working in multicloud environments so much simpler, secure, and cost-efficient. Automate cloud cost management, gain visibility into cloud spend across your entire multicloud landscape and collect, process, and store only the observability data that is necessary.

What is a data pipeline?

Sources



Destinations



Cribl Stream allows you to ingest data and get value from that data in **any format**, from **any source**, and then direct it to **any destination**.

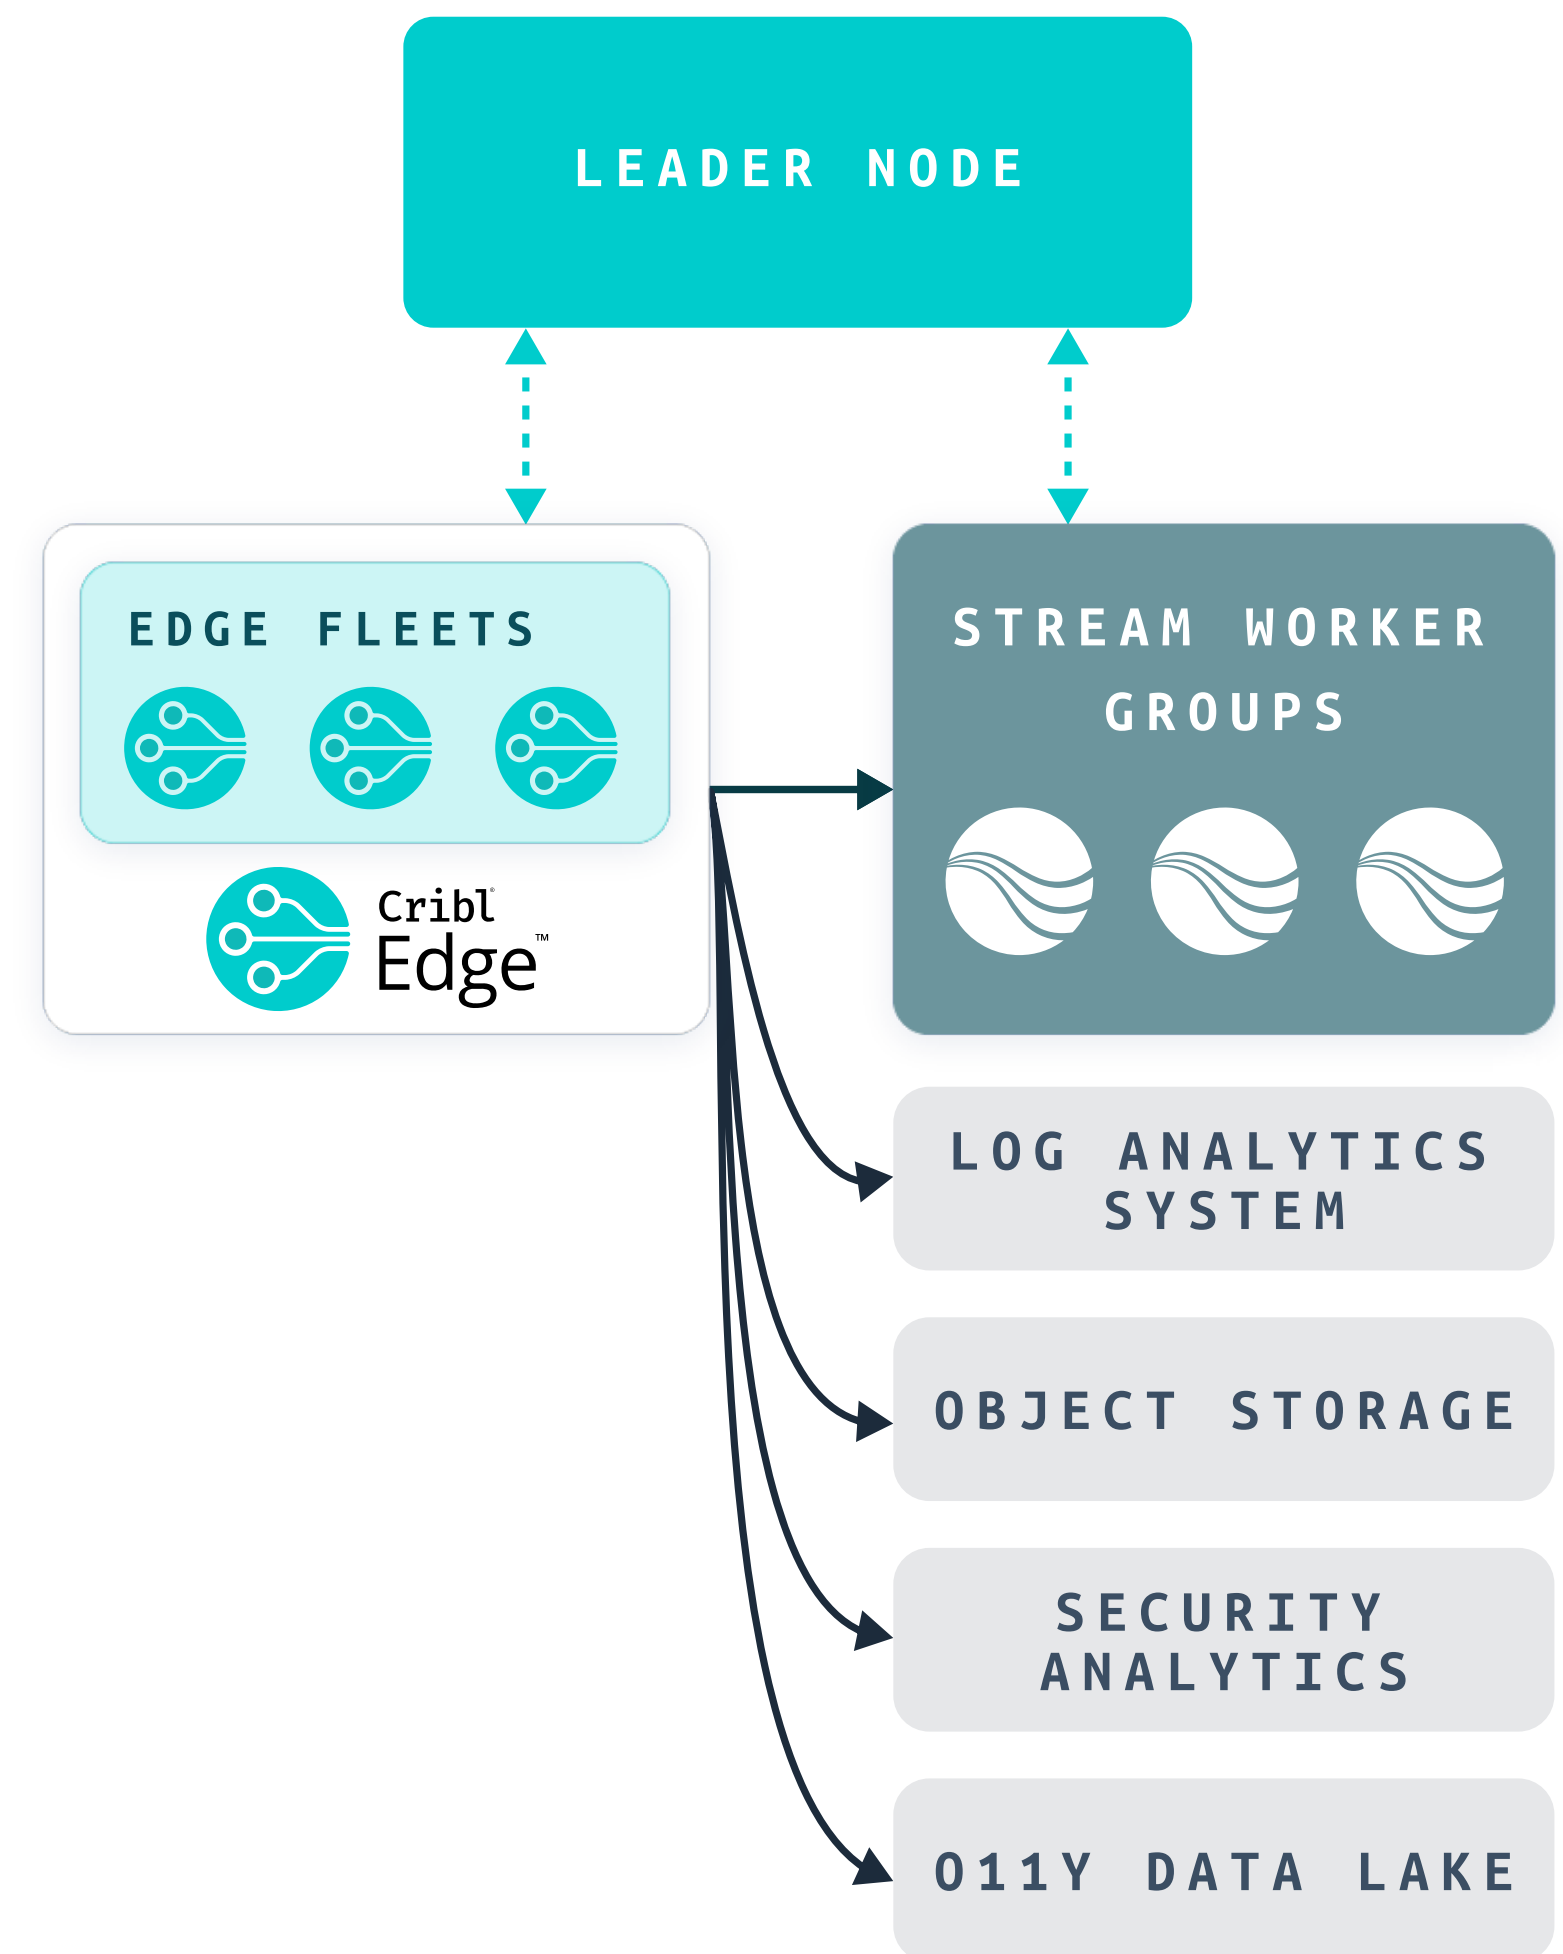
Auto-discover IT and Security Data at its egress point. Collect all the data, at the edge, at scale.



Cribl Edge is an intelligent, scalable edge-based data collection system for logs, metrics, and application data and was designed to support today's modern multicloud architectures. Cribl Edge is centrally managed, auto-discovers IT and Security Data at its egress point, and opens up additional, cost-effective options for data collection and processing.

With Edge's built-in Fleet Management, you can effortlessly manage tens of thousands of Edge nodes while lowering data collection total cost of ownership. Collect all the data you need, at the edge, at scale.

Cribl Edge is centrally managed, configured, and version controlled for easy expansion and low cost of ownership.



Perform federated “search-in-place” queries on any data, in any format, at any location, and across all cloud services.



Cribl Search performs federated “search-in-place” queries on any data in any format at any location, and across all your cloud services. Increase the scope of analysis without requiring the cost or complexity of first shipping, ingesting, and storing the data. This gives teams relevant, valuable data that’s only routed for further analysis if necessary.

With Cribl Stream and Edge, you’re able to cost-effectively ingest and process data, and route to low-cost storage. Further reduce storage costs by storing less frequently accessed data in low-cost object storage like Cribl Lake, Amazon S3, Google Cloud Storage, and Microsoft Azure Blob Storage, and with Search, you still have the ability to search across them all.

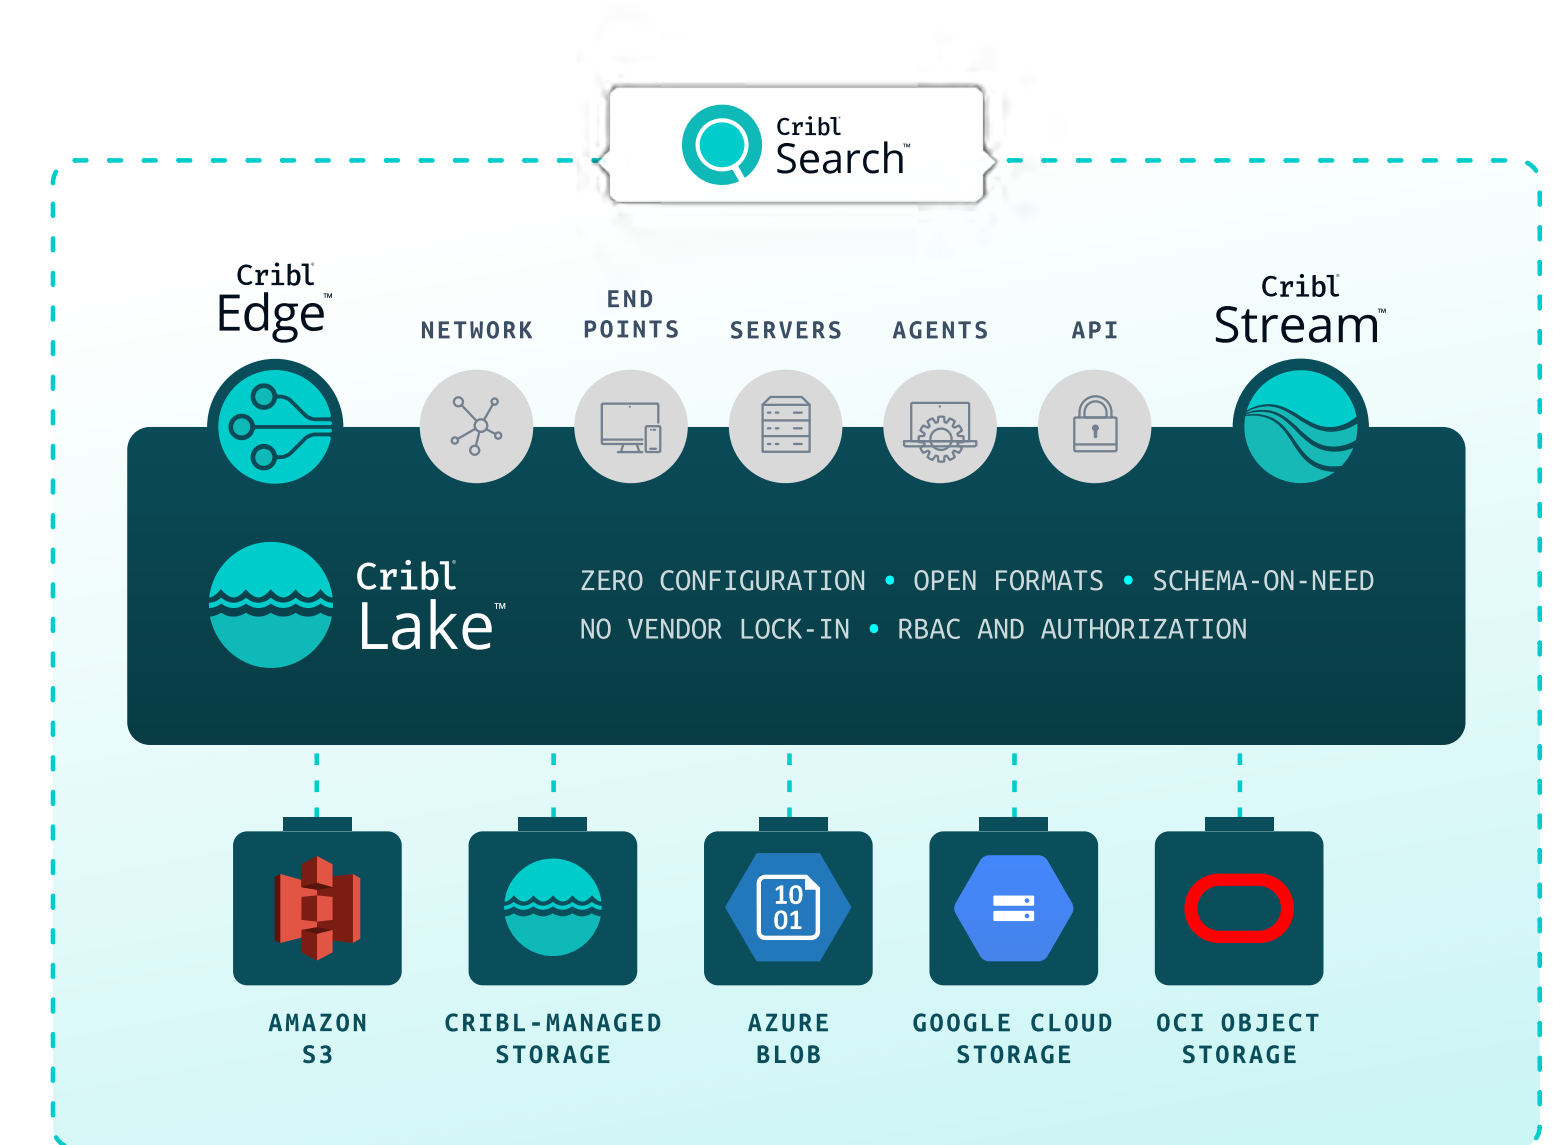


Designed to
make setting up
and managing
a data lake
fast and easy,
with no cloud
or data-specific
skills required.



Cribl Lake was designed to make setting up and managing a data lake fast and easy, with no cloud or data-specific skills required. IT and security teams can leverage low-cost object storage — whether managed by Cribl or owned by you — to onboard huge volumes of data and make it easy to store, access, and retrieve data.

Cribl Lake offers a single view into federated data. Organizations face challenges of managing data across multiple clouds and systems due to acquisitions, mergers, migrations. Cribl Lake gives access across all clouds, sources, and regions. This helps understand what data you have and where, realizing the full value of hybrid cloud deployments regardless of where data resides.

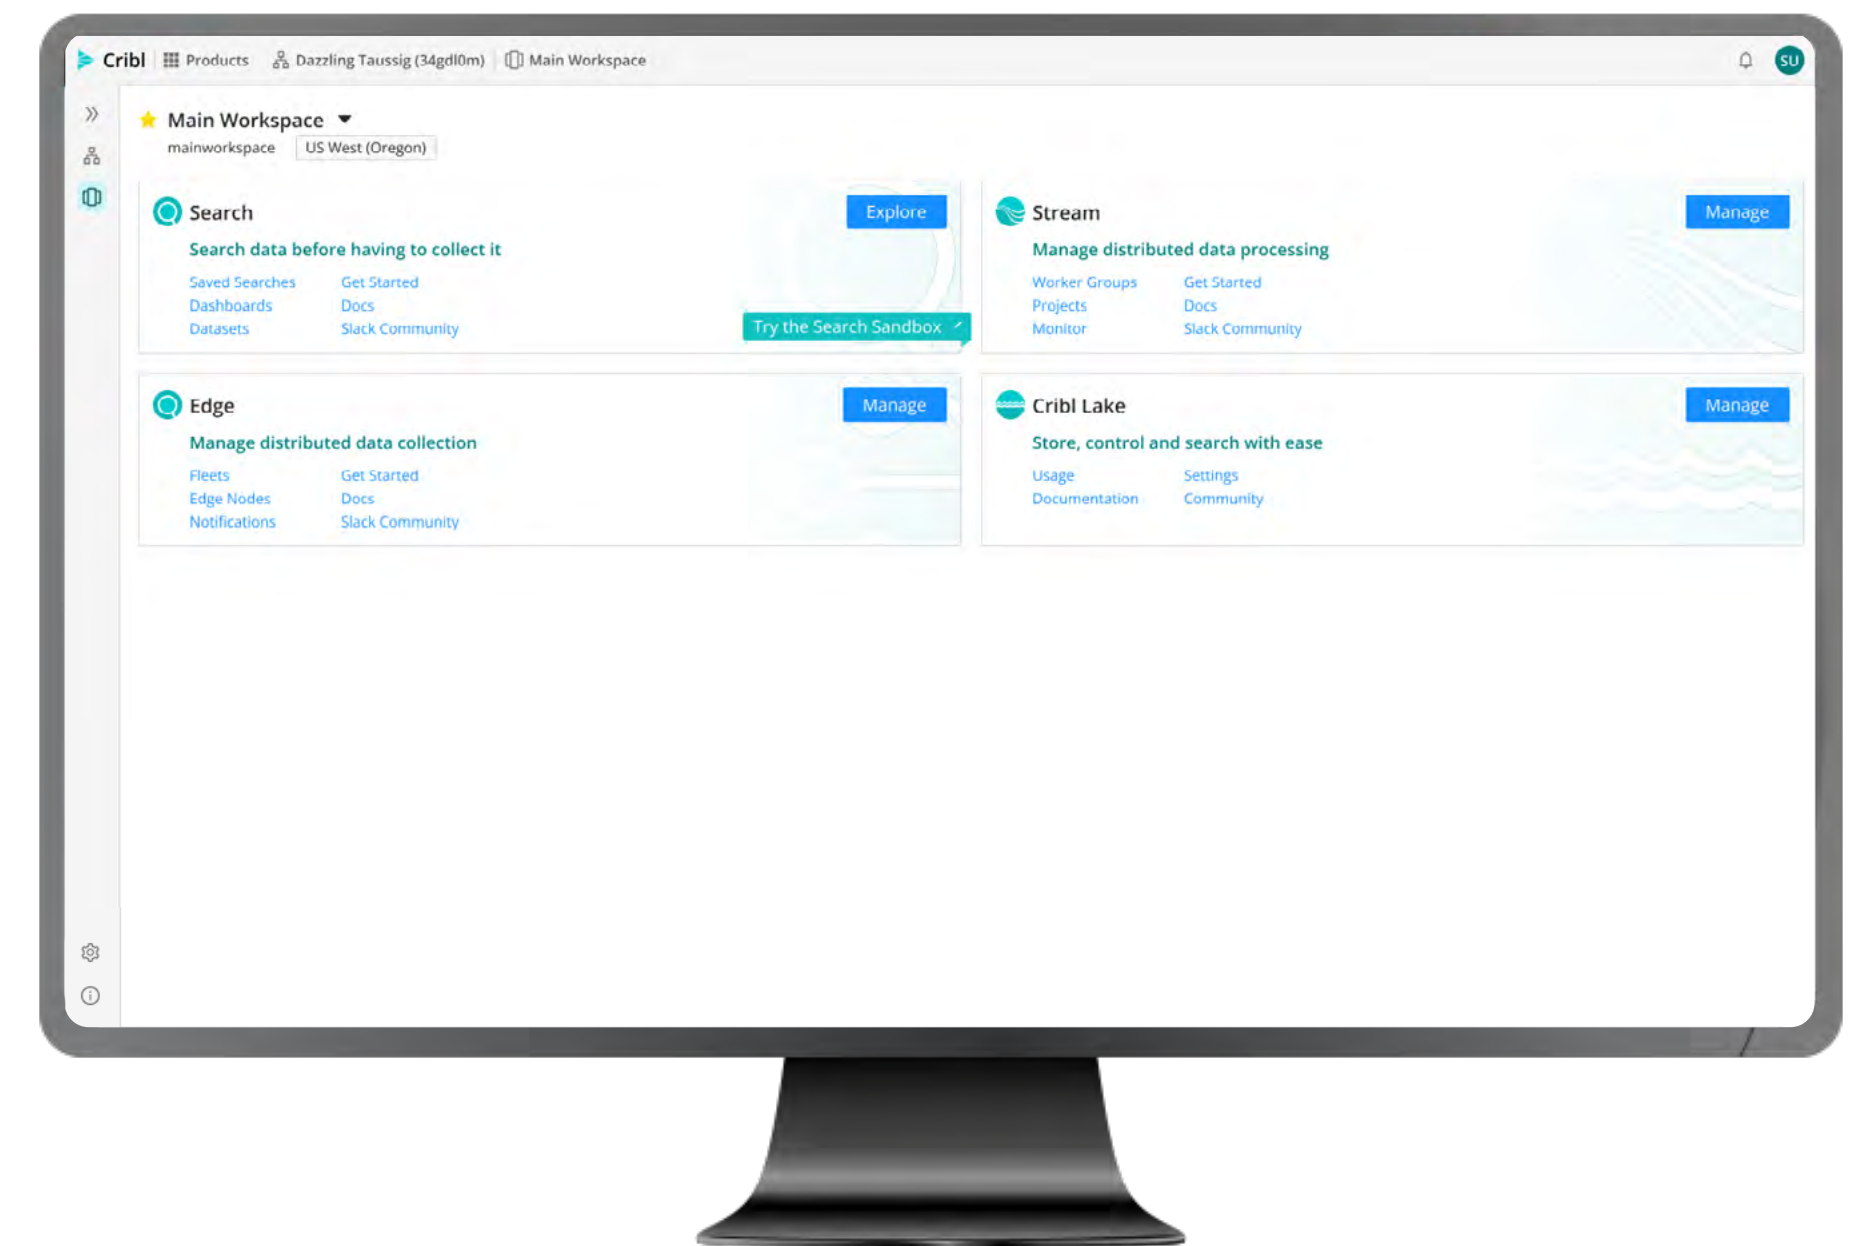


A unified SaaS platform to quickly spin up Cribl products without the hassle or costs of running infrastructure.

Cribl[®].Cloud

Managing a multicloud environment requires careful planning, coordination, and intention. Cribl.Cloud is a unified platform that can be used to quickly spin up Cribl products without the hassle or costs of running infrastructure.

With Cribl.Cloud, now all your data becomes usable and valuable, so you can enjoy the benefit of fast deployment at scale, as your data engine in the cloud manages the security risk of interfacing with different cloud platforms. Lower risk, faster results.



Customer Story



AUTODESK



Challenge

The Cloud Architecture and Insights team at Autodesk was responsible for logging practice, general data governance and oversights, policies, architecture, and spend for AWS. They were challenged with complex legacy data pipelines, and needed to consolidate monitoring tools and also better optimize their log data. They were on a quest for a solution that could onboard new data sources faster and easier.

It was at this time that the technical leadership team also made the decision to steadily expand its use of cloud-based services and decrease its datacenter footprint. This was an opportunity for a multicloud migration initiative, involving business data sources that could be used to break free from the complexity of accumulated technical debt in associated data pipelines.

How Cribl Helped

Cribl helped Autodesk's data pipeline infrastructure become more modern, uncomplicated, and resilient. Costs associated with analytics tools were significantly lower as well. There was a lot of data being duplicated in their logging platforms, and Cribl Stream helped easily detect when this happened. The team was able to reduce extra ingestion by 166%.

Cribl didn't just help Autodesk save money, but also helped better prioritize more valuable data. Autodesk teams can now use Cribl to refine data ingestion to focus on the events that matter most for the business

[Read more about how Autodesk simplifies enterprise-tier data pipeline with Cribl](#)



The goal is to optimize the data. Our priority is to make the data more valuable, mark it up, add more - it's not just about cost cutting, it's about getting more context to make better decisions.

Jacob Gorney, Cloud Architect at Autodesk



“Cribl is transforming our relationships across the business. It gives us the flexibility to achieve our goals, get data onboarded and routed to the appropriate systems more easily. We are more agile and provide a better experience to our customers for a lower cost.

Samy Senthival, Director of Observability at Autodesk



In your Multicloud Era? Get equipped with a Data Engine.

In December of 2022, the Pentagon awarded contracts to Amazon, Google, Microsoft, and Oracle — combined worth of \$9 billion through 2028.¹⁰ Organizations are quickly realizing it's risky to rely on just one cloud. And working with more than one cloud provider means they're able to better withstand service disruptions brought on by outages.

According to Hashicorp, 81% of organizations surveyed are choosing to adopt multicloud — 60% are already using multicloud infrastructures, with an additional 21% saying they will be within the next 12 months. Of those that are already embracing a multicloud environment, 90% are saying it's helping their organization advance and achieve their business goals.¹¹


As multicloud continues to gain more traction, a robust data engine is required. Set your cloud environment up for success with Cribl and enjoy greater flexibility, instant scalability, and more choice.

¹⁰ [CNBC Google, Oracle, Amazon and Microsoft awarded Pentagon cloud deal of up to \\$9 billion combined](#)

¹¹ [HashiCorp 2022 State of Cloud Strategy Survey](#)



Get 1TB/day for free with Cribl.Cloud!



Cribl.Cloud is the fastest and easiest way to try Cribl products and start building your data engine today! Nothing to install, no infrastructure to manage, no license required, no payment collected. You can also access Cribl.Cloud through the **[AWS Marketplace!](#)**

About Cribl

Cribl, the Data Engine for IT and Security, empowers organizations to transform their data strategy. Customers use Cribl's vendor-agnostic solutions to analyze, collect, process, and route all IT and security data from any source or in any destination, delivering the choice, control, and flexibility required to adapt to their ever-changing needs. Cribl's product suite, which is used by Fortune 1000 companies globally, is purpose-built for IT and Security, including Cribl Stream, the industry's leading observability pipeline, Cribl Edge, an intelligent vendor-neutral agent, Cribl Search, the industry's first search-in-place solution, and Cribl Lake, a turnkey data lake. Founded in 2018, Cribl is a remote-first workforce with an office in San Francisco, CA.

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

©2024 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.

EB-0003-EN-3-1024