Cribl

CHANNEL FAQ

Cribl Stream™

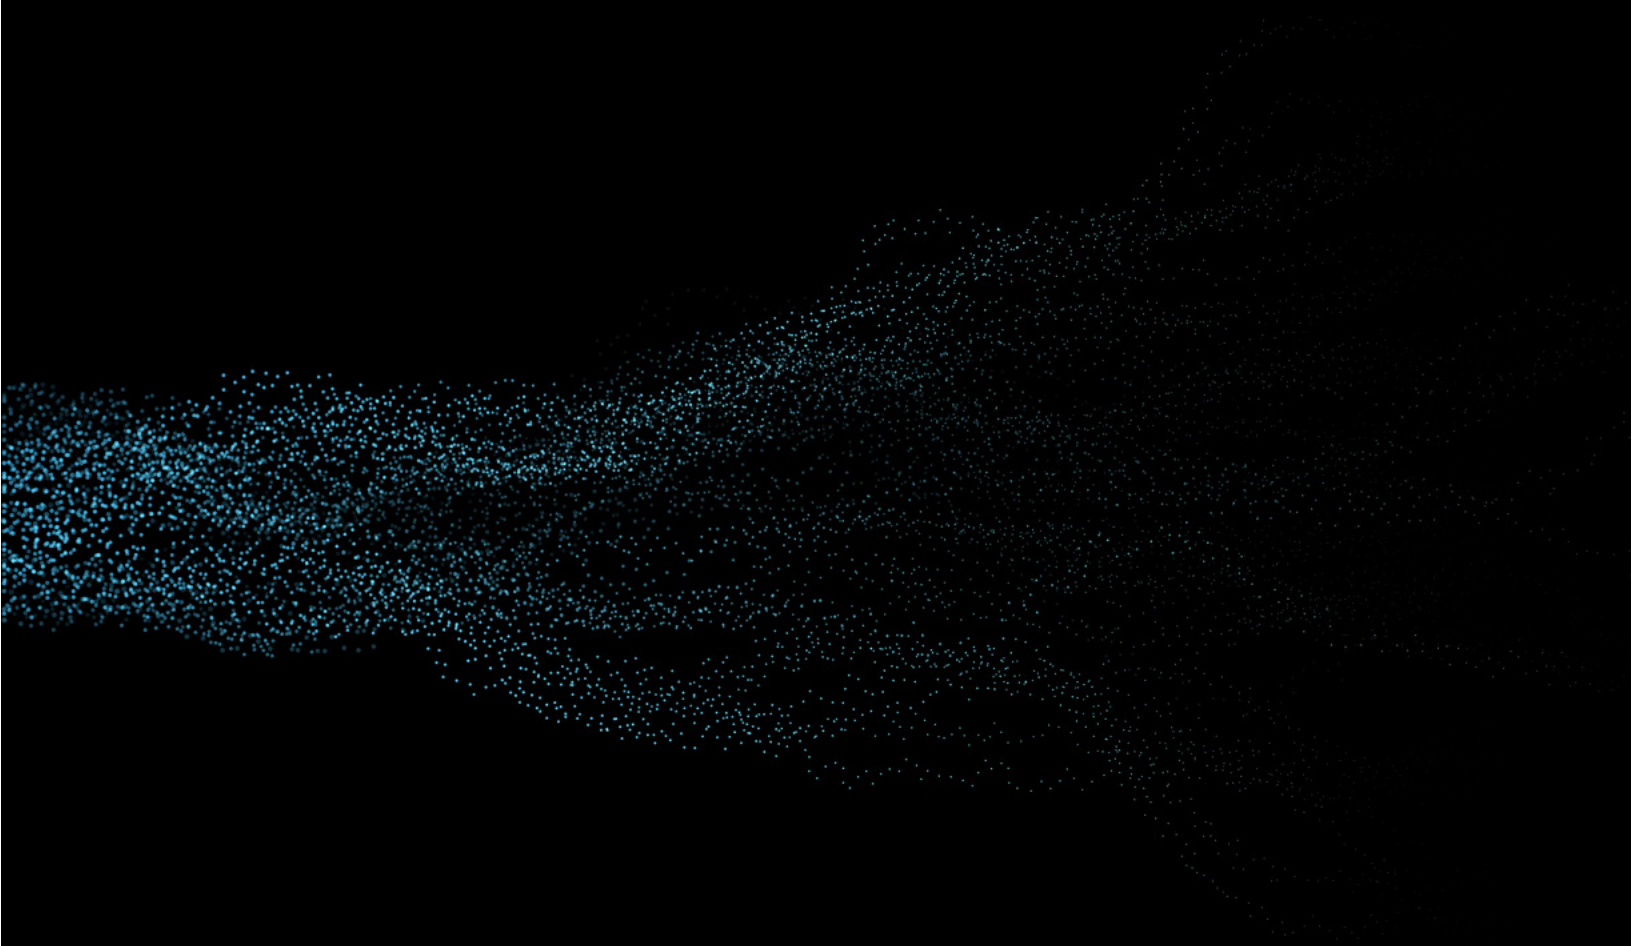

CHANNEL FAQ

Cribl Stream™

The goal of this document is to preemptively address expected questions associated with the Cribl Stream product updates for the March 2022 release, including technical capabilities and sales positioning.

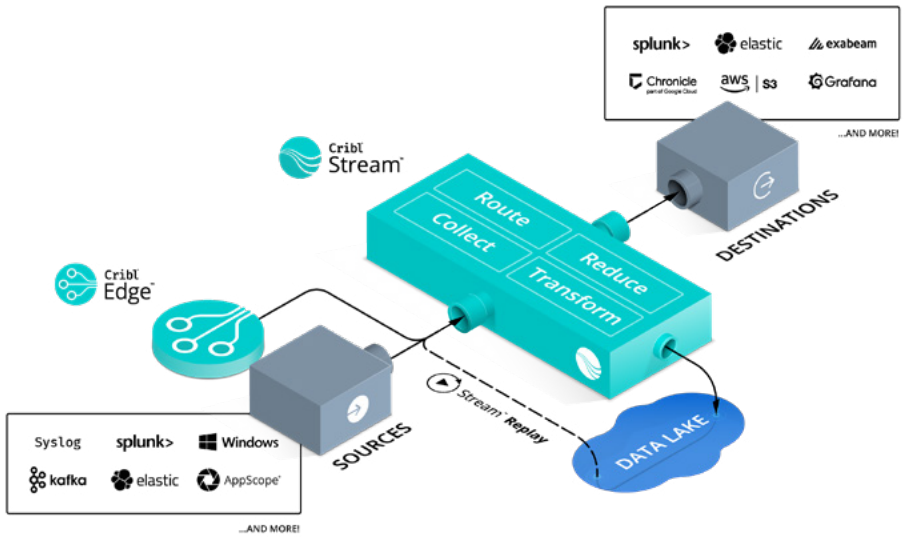

Chapter 1: General FAQ

1. WHAT IS CRIBL STREAM?

- Cribl Stream is the new name for Cribl LogStream. The name change aligns with our new naming conventions and product branding. Stream is a vendor-agnostic data collection, reduction, enrichment, and routing system for observability data. Stream enables you to give your IT, Security, and DevOps customers the flexibility to route, shape, restructure, and enrich data from any source to any destination without adding new agents.

2. IS THERE A DIFFERENCE BETWEEN STREAM AND LOGSTREAM?

- There is no difference beyond branding. However, we timed the name change with the Spring 2022 release where we introduce new capabilities as part of the normal release program.

3. WHAT'S NEW WITH STREAM'S SPRING RELEASE?

- Packs Registry: A Cribl-supported repository of Packs
- Source-side Queueing: Reduces data loss by queueing data coming from the sources onto the disk
- Value Dashboards: a quick view of Stream usage for customers who want to get more details about their consumption patterns
- Windows Event Forwarding: Native support for Windows Event forwarding
- HTTP Load Balancing: Ideal for customers who don't want to use external load balancers for HEC/ Elastic kind of destinations

TABLE OF CONTENTS

4. IS THERE ANYTHING ELSE NEW BESIDES THE STREAM NAME CHANGE AND UPDATES?

- Yes. We also released a new product: Cribl Edge. Cribl Edge is a highly-scalable edge-based data collection system for logs, metrics, and application data. Edge is the easiest way to get data out of any Linux system, enabling you to give your IT and SecOps customers the ability to collect, process, and forward observability data at scale while keeping costs down. Manage fleets of Edge nodes without relying on external tools. Process and collect at the edge – at scale – and increase options for your customers' data.

5. CAN CRIBL EDGE WORK WITH CRIBL STREAM?

- Yes, Edge is designed as a companion to Stream, depending on the network design and the customer's requirements. Stream and Edge can be used independently or in combination to offer greater distribution of processing. A single Leader can manage Stream Workers, Stream Worker Groups, and Edge Fleets in a combined deployment.
- Cribl Edge, combined with Cribl Stream, creates the first fully-distributed observability pipeline, allowing organizations to cost effectively collect and recall even more full-fidelity event data.

6. HOW DOES AN ADMINISTRATOR MANAGE CRIBL STREAM AND EDGE?

- When you click the Stream Icon 

at top of the left- hand navigation, a user can access either the Stream or Edge Admin page.

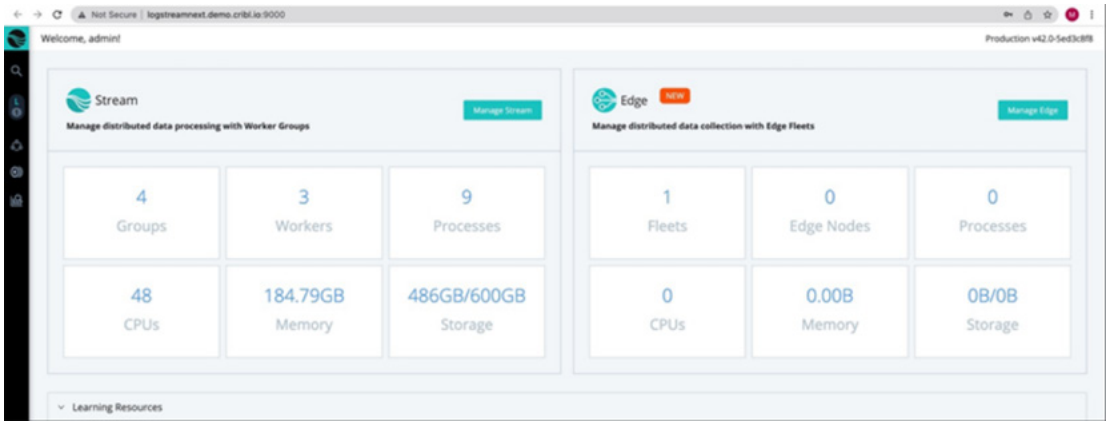

Chapter 2: Technical Specifications

7. WHICH AWS REGIONS ARE SUPPORTED TODAY? WHAT'S ON THE ROADMAP?

- We currently support US-West with plans to expand to US-East in Spring 2022.

8. WHAT NEW INTEGRATIONS ARE AVAILABLE IN THIS RELEASE?

- Elastic Destination – Added support for Elastic Datastreams and in-product load balancing of traffic to Elastic nodes.
- DataDog Agent Source – Introduced support for the DataDog agent and DogStatsd protocol.
- Windows Event Forwarder (WEF) Source – Native support for Windows Event Forwarding. Our Windows Event Forwarder (WEF) Source provides Windows admins with a safe and reliable way to collect, process, and route Windows logs without extra servers or agents. The WEF Source for Cribl Stream (and Edge) supports TLS authentication, and enables you to add custom WEF subscriptions directly through its UI.

9. WHAT ARE STREAM VALUE DASHBOARDS?

- Stream Value Dashboards uplevel Stream's existing Monitoring views to provide more information about your data, including daily usage stats, top talkers, license utilization, and change markers. Stream Value Dashboards also provide details about top Sources, Destinations, Pipelines, Routes, and Packs. In all, you'll get a better understanding of how Stream is being used in your environment.

10. DOES STREAM SUPPORT TCP/HTTP LOAD BALANCING?

- Yes. We've enhanced our existing destinations to support TCP and HTTP load balancing. Admins have the ability to add as many destinations as they wish, distributing data across their endpoints. This adds support for ElasticSearch and Splunk HEC, to go along with the existing Stream and Splunk destination load balancing already available.

11. TELL ME MORE ABOUT SOURCE-SIDE QUEUING.

- Source-side queueing is now available to alleviate data loss issues. This can be especially useful for UDP sources which don't respond to backpressure. Source-side queueing allows sources to write data to disk when the downstream destination is unavailable or is exerting backpressure. This is typically not a problem with TCP sources, but for UDP sources like syslog, source-side queueing allows you to queue that data until the downstream destination recovers and is ready to receive data again.

12. WHAT NEW ENTERPRISE FEATURES CAN OUR CUSTOMERS LEVERAGE?

- **ACLs** – Support for customer – or partner – managed allowlists for workspace Load Balancers.
- **Billing Updates** – An updated in-cloud UI for tracking usage and billing for paid plans.
- **AWS Trust** - Self-service, cross-account trust for your customers' AWS sources and destinations.

13. IS STREAM AVAILABLE AS SOFTWARE, CLOUD, OR BOTH?

- Stream is deployable as software or in the Cloud and can also be deployed in hybrid environments.

14. CAN A CRIBL CUSTOMER WITH AN EXISTING SOFTWARE STREAM DEPLOYMENT MIGRATE TO A CLOUD- OR HYBRID-HOSTED DEPLOYMENT WHILE RETAINING ALL CONFIGURATIONS?

- Yes. Loop in your Cribl sales rep to figure out the best next steps.

15. WHAT ARE THE TECHNICAL DIFFERENCES BETWEEN THE DIFFERENT DEPLOYMENT OPTIONS?

Criteria	On-prem	SaaS	Hybrid
Infrastructure Ownership & Management	Customer-hosted	Cribl-hosted (AWS US WEST)	Cribl hosts leader node + 1 worker group Customer adds self-hosted worker groups
Leader Backup	Local or remote GIT	Managed by Cribl	Managed by Cribl
Security & Compliance	RBAC, KMS Integration, Secret Store, UI Managed credentials, CLI, API	Secret Store, UI Managed credentials, API, SOC2 Type 2 coming	Secret Store, UI Managed credentials, API
Other Considerations	N/A	No customer defined scripts No tee function No persistent queuing for destinations No shell/CLI access	No shell/CLI access

Chapter 3: Positioning and Target Audiences

16. CAN I GET A QUICK OVERVIEW?

- Stream is a vendor-agnostic data collection, reduction, enrichment, and routing system for observability data. Stream enables you to give your IT, Security, and DevOps customers the flexibility to route, shape, restructure, and enrich data from any source to any destination without adding new agents.

17. WHAT CUSTOMER PROBLEMS IS CRIBL STREAM SOLVING?

- Data is growing year over year, and at the same time, companies are trying to analyze new sources of data to get a complete picture of their IT and Security environments. They need flexibility to get data into multiple tools, from multiple sources, but don't want to add a lot of new infrastructure and agents. These companies need a better strategy for retaining data long-term that is also cost-effective.
- Cribl Stream is a vendor-agnostic data collection, reduction, enrichment, and routing system for observability data that allows you to support your customers in instrumenting everywhere, gaining more insights from analytics tools, and retaining more data for longer periods of time, all while controlling costs and increasing performance.
  - Give customers more control over their data and help them send it to the right teams and the right tools
  - Flexibility to add any analytics tool without adding new agents
  - Simplify observability, with a universal log router for any tooling environment
  - Enable customers to instrument everything, analyze more, and control costs

18. WHAT IS A TYPICAL USE CASE FOR CRIBL STREAM?

- For IT and Security professionals struggling to cost-effectively gain control of their data, Cribl Stream gives them the flexibility to route, shape, restructure, and enrich data from any source to any destination without adding new agents. With Stream, they gain control over their data and simplify their observability efforts, enabling them to instrument everything, analyze more data, and pay less.
- Key Use Cases:
  - **Collect** – Vendor-agnostic integrations allow users to collect from and send to any observability tool
  - **Reduce** – Control costs by removing unneeded fields and dropping unimportant events
  - **Transform** – Convert data to common formats and reshape data as needed to make it usable
  - **Enrich** – Add context to data using external sources like GeoIP, Cloud metadata, or custom tables and databases
  - **Replay** – Stash low-value data in cost-effective object storage and replay back through your customer's data pipelines and to their destination of choice for no additional cost
  - **Route** – Route the same data to multiple places to break down data silos and increase customer choice and flexibility

19. ARE THERE ANY COMPETITORS TO STREAM?

- Competitors to Stream include Splunk DSP and Elastic's Logstash. Please reference the Stream Quick Sell Guide in the LMS or reach out to your Cribl sales rep for direct comparisons and tips on objection handling.

20. WHAT TYPE OF CUSTOMER OR DEPARTMENT PERSONAS WOULD BE PRIMARY TARGETS?

- CIO/CISO/CSO, VP CyberThreat/Security Engineering, Security/SOAR/CyberThreat Engineer, IT Director, System Administrator, Tooling Director/Manager
- Looking to get data to and from multiple sources while keeping costs down and getting valuable insights from data
- Needs to secure internal and external company data with added investigation capabilities in case of a breach

21. WHAT IS THE SALES WORKFLOW FOR STREAM AND EDGE?

- Stream is our "land" product. Edge is the "expand" product after a customer has been successful with Stream and wants to improve their data collection experience and/or move some of the data processing and routing off of their central Stream deployment.

22. WHAT ARE THE KEY BENEFITS OF STREAM?

- **Routing:** Send data to the most effective destinations, including low-cost storage locations like S3 for long-term retention. Route data to the best tool for the job – or all the tools for the job – by translating and formatting data into any tooling schema you require. Let different departments choose different analytics environments without having to deploy new agents or forwarders.
- **Reduction:** Reduce as much as 50% of your customer's ingested log volume to control costs and improve system performance. Eliminate duplicate fields, null values, and any elements that provide little analytical value. Filter and screen events for dynamic sampling, or aggregate log data into metrics for massive volume reduction. Do all of this without worry: You can keep a full-fidelity copy in a low-cost destination, and replay it later if needed.
- **Data Collection and Replay:** Stream is the best way to get multiple data formats into your customer's analytics tools. Use the Stream universal receiver to collect from any observability data source – and even to schedule batch collection from multiple APIs. In addition, use ad-hoc data collection to recall data from low-cost storage, when your customer needs to replay logs to analytics tools to assist them in later investigations.
- **Data Shaping:** Stream collects data from all of your customer's sources, and shapes it into actionable logs and metrics for analysis. Your customers can shape all of their data to drive decisions about their environment. Translate and transform data from any source to the tools they choose. Help them get a more complete picture of their data by enriching logs with third-party information.
- **Control:** Reduce management overhead, with robust and easy-to-use GUI-based configuration and testing. Capture live data and monitor observability pipelines in real time. Gain further insights with graphical data flow mapping. Take advantage of automated upgrades to easily and efficiently keep a whole fleet of Workers up to date without tedious CLI-based management. Easily share functionality between Leader nodes and Worker groups to reduce time-to-value. Use role-based activity control to ensure that only the right teams are seeing what they need to do their jobs. Manage Stream on-premises, use our cloud implementation, or go hybrid.

Chapter 4: Additional Information

23. WHERE CAN I FIND ADDITIONAL SALES AND TECHNICAL RESOURCES AND MATERIALS ON CRIBL STREAM?

- Please visit Cribl's LMS for the latest resources and sales materials. Both sales and technical resources are available.

24. WHERE DO I FIND THE LOGO/ICONS FOR CRIBL STREAM?

- You can visit our brand repository for the latest Cribl logos and assets. If you need additional assistance, please reach out to us at **partners@cribl.io**.

ABOUT CRIBL

Cribl makes open observability a reality for today's tech professionals. The Cribl product suite defies data gravity with radical levels of choice and control. Wherever the data comes from, wherever it needs to go, Cribl delivers the freedom and flexibility to make choices, not compromises. It's enterprise software that doesn't suck, enables tech professionals to do what they need to do, and gives them the ability to say "Yes." With Cribl, companies have the power to control their data, get more out of existing investments, and shape the observability future.Founded in 2017, Cribl is a remote-first company with an office in San Francisco, CA. For more information, visit **www.cribl.io** or our **LinkedIn**, **Twitter**, or **Slack** community.