


CHANNEL FAQ

Cribl Edge™



CHANNEL FAQ

Cribl Edge™

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

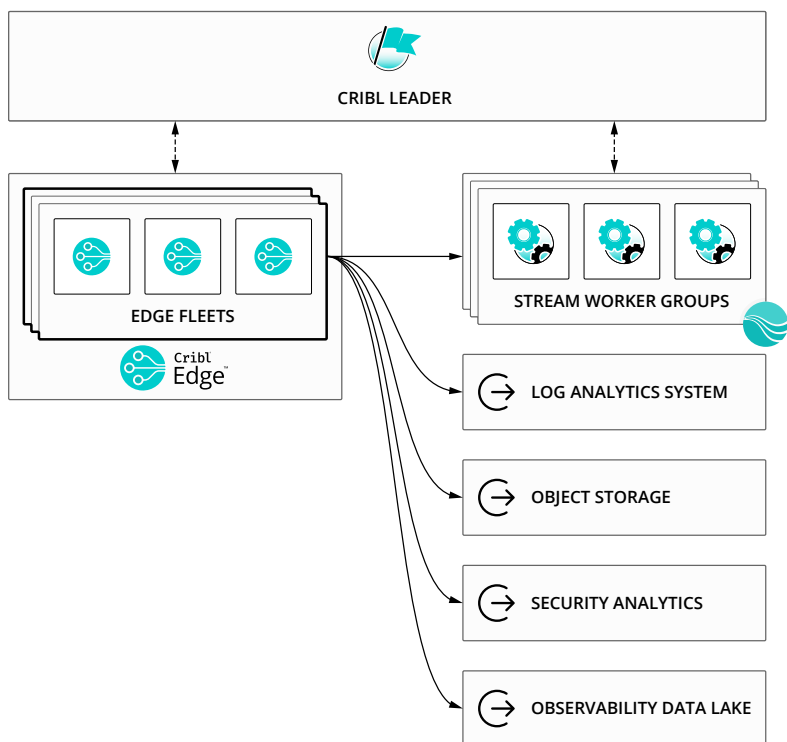


TABLE OF CONTENTS

CHAPTER 1

General FAQ

CHAPTER 2

Technical Specifications

CHAPTER 3

Positioning

CHAPTER 4

Additional Information

Chapter 1: General FAQ

1. WHAT IS CRIBL EDGE?

Cribl Edge allows you to collect and process observability data – logs, metrics, and application data – in real time, from your Linux machines, apps, microservices etc., and deliver them to Cribl Stream or any supported destination. Cribl Edge is a next generation, intelligent agent that efficiently gathers and auto-discovers observability data at its egress point. Edge enhances Cribl's observability pipeline, giving users the choice to reduce, enrich, and route data, in any format, to any destination. In addition, Edge provides a central way to manage your fleets, simplifying how you install, enforce policies, and configure thousands of Edge nodes. Cribl Edge applies auto-discovery to enhance search queries by automatically discovering useful data. Customers can choose to incorporate Edge data into analysis without adding cost, complexity, or additional agents.

2. WHAT IS MEANT BY FLEET MANAGEMENT?

- A Fleet is a collection of Edge Nodes that share the same configuration. Fleets simplify how you author, manage, and configure settings for a particular set of Edge Nodes. The Leader node is responsible for syncing configurations and tracking edge node status and monitoring metrics.
- You can configure multiple Fleets to reflect your organizational or geographic constraints. E.g., you might have a U.S. Fleet with certain TLS certificates and output settings, versus an APAC Fleet and an EMEA Fleet, each with their own distinct certs and settings. Multiple Fleets requires an Edge Enterprise or Standard license.
- An Edge node can only exist in one Fleet at a time, with only one active configuration rule set.

3. WHAT TYPES OF USERS IS CRIBL EDGE TARGETED AT?

- Cribl Edge is targeted at ITops and SecOps Engineers and users of operational and security intelligence products and services.

4. 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.

5. WHAT ARE THE DIFFERENCES BETWEEN CRIBL STREAM AND CRIBL EDGE?

- Cribl Stream and Cribl Edge are essentially built on the same core code, but each is optimized for the specific use case in how they discover, gather, process and route data.
- 99% of Stream capabilities are supported by Edge. One key difference is Edge does not work with Pull or Collector sources. Edge also supports discovering and collecting log files and metrics directly from the host system.

6. WHAT IS THE ADVANTAGE OF CRIBL EDGE OVER A CUSTOMER'S EXISTING AGENTS?

- Most existing agents only do basic forwarding; multiple agents are often required for full visibility. Every one of those agents introduces a potential compatibility issue. Solutions end up fragmented and siloed. You may have one tool for logs, a different tool for traces, another tool for metrics, and so on. Agents are also required for full visibility, and every one of those agents introduces a potential compatibility issue. Finally, adding or switching vendors typically requires an agent overhaul.
- Edge's key differentiation compared to existing agents:
 - Edge can collect, process, and forward data.
 - Edge can discover what is emitting logs and auto-tail those files.
 - Edge unifies collection of logs, metrics, and traces in a single agent.
 - Edge requires fewer resources to shape and route to a destination.
 - Edge provides Fleet management.
 - Edge supports distributed processing, so you can leverage capacity at the source that typically goes unused.
 - Edge gives you access to on-demand instrumentation of any process or executable via AppScope.

7. CAN CRIBL EDGE WORK IN PARALLEL WITH A CUSTOMER'S EXISTING AGENTS FROM OTHER PRODUCTS?

- Yes. Please reach out to your Cribl sales rep for additional guidance.

8. IS CRIBL EDGE READY FOR USE IN PRODUCTION ENVIRONMENTS?

- Yes, as of March 2022.

9. IS EDGE AVAILABLE AS SOFTWARE, CLOUD, OR BOTH?

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

10. CAN EDGE SECURE THE DATA TO A DESTINATION?

- Yes, Edge supports TLS.

11. WILL EDGE HAVE A CENTRAL MANAGEMENT FOR DEPLOYMENT?

- Yes, this is on the roadmap for Fall 2022.

12. HOW MANY EDGE NODES CAN BE CONNECTED TO A LEADER?

- We've tested a single leader and can confirm support for up to 1000 Edge nodes. We'll be working on adding additional scale in the Summer 2022 timeframe.

13. CAN A CUSTOMER MIGRATE FROM A STREAM ONLY DEPLOYMENT TO A MIXED STREAM/EDGE DEPLOYMENT?

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

14. CAN AN EXISTING STREAM CLOUD CUSTOMER ADD EDGE FROM THE CLOUD?

- Yes, existing Stream customers with Cloud deployments can add Edge directly from their cloud environment.

15. CAN A CUSTOMER GET EDGE WITH ANY EXISTING STREAM DEPLOYMENT (FREE, STANDARD, OR ENTERPRISE)?

- Yes. Edge is available for existing Stream customers regardless of licensing tier.

16. ARE THERE ANY CAPACITY OR FEATURE LIMITS BETWEEN EDGE FREE AND THE LICENSED OFFERINGS?

- There are no functionality limits between the Free and licensed versions, everything is available.
- Edge Free is bound by the same 1TB limit as Stream. Standard and Enterprise are not capacity limited.
- Edge Free entitles users to manage up to a 100 Edge nodes, in a single Fleet, and ingest up to 1TB/day.
- Edge Enterprise entitles users to manage unlimited Edge nodes, across unlimited Fleets and ingest up to its rated volume.

17. WILL WE OFFER A SINGLE UI FOR STREAM AND EDGE?

- Yes, you can access Stream and Edge in the same UI.

18. WHO WOULD BE A GOOD CANDIDATE FOR EDGE?

- Existing Stream customers who are experiencing significant pain keeping up with collecting data from their endpoints, customers looking to replace their existing agents, and customers that simply want to ship or forward data to a destination are all great prospects for Cribl Edge.

19. WHAT TYPE OF DATA WILL CRIBL EDGE GATHER UPON THIS FIRST RELEASE?


- Edge will gather system- and container-level metrics and logs, provide detailed information on processes and containers running on machines or hosts, and support data that can be collected using custom scripts.

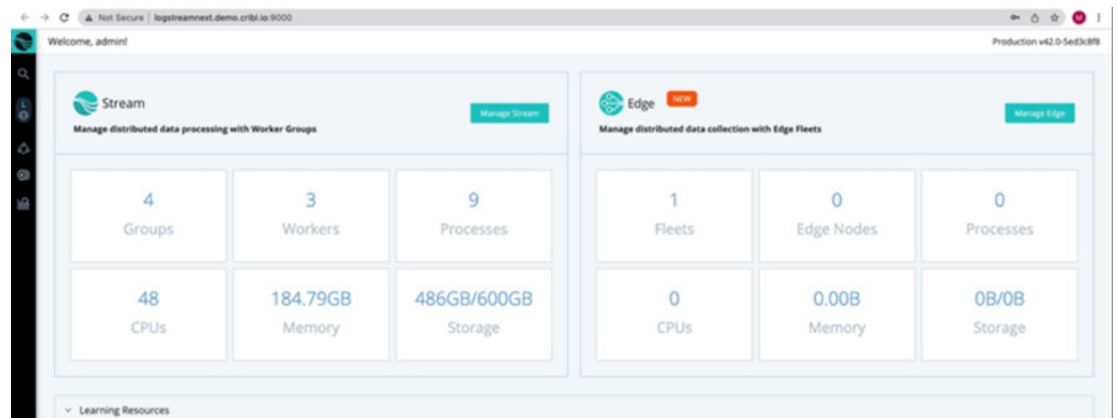
20. CAN YOU ROUTE EDGE DATA TO DIFFERENT DESTINATIONS?

- Yes, Cribl Edge has the same routing and forwarding capabilities available in Cribl Stream.

Chapter 2: Technical Specifications

21. HOW DOES A USER ACCESS EDGE FROM THE UI?

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



22. WHAT TYPES OF PLATFORMS DOES EDGE CURRENTLY SUPPORT?

- At this time, Edge supports Linux. We will be adding additional support in the future.

23. WHAT ARE THE 'HOST' REQUIREMENTS FOR CRIBL EDGE?

- The minimum requirements for running Edge are the same as for a Stream instance.
- Edge node (minimum):
 - OS: Linux: RedHat, CentOS, Ubuntu, AWS Linux, Suse (64 bit)
 - System: ~1Ghz processor, 512MB RAM, 5GB of free disk space (more if persistent queuing is enabled on Workers)
- Leader Node (recommended):
 - OS: Linux: RedHat, CentOS, Ubuntu, AWS Linux, Suse (64 bit)
 - System: +4 physical cores, +8GB RAM, 5GB free disk space
- Browsers supported include Firefox 65+, Chrome 70+, Safari 12+, and Microsoft Edge.

24. WHAT ARE THE NETWORK PORT REQUIREMENTS?

- Cribl Edge leverages the following ports:
 - UI default • 9420
 - HTTP Inbound • 10080
- We can support additional user options and other ports as needed.
- Please reach out to your Cribl sales rep for additional guidance.

25. WHAT DEPLOYMENT OPTIONS DOES EDGE OFFER?

- Just like Cribl Stream, Edge supports Single Instance, Distributed, Cloud, and Hybrid deployments.

26. WHAT DATA TYPES ARE SUPPORTED?

- Same as Stream, plus host system Metrics and Logs

27. WILL EDGE ALLOW ADMINISTRATORS TO CONNECT INTO A SPECIFIC SOURCE/HOST?

- Yes, users can connect directly into the host from the UX, eliminating the need to SSH into the host.
- **Note:** Any changes made directly to the host in this mode are not saved on a reset or when a new update from the Leader occurs.

28. DOES EDGE SUPPORT PIPELINES AND PACKS?

- Yes. Routes, QuickConnect, Pipelines, Packs, Knowledge Objects, QuickConnect, and monitoring capabilities all operate the same in Edge as they do in Stream. As with Stream, some data may need to be transformed via Pipelines/Functions to support destination requirements.

29. DOES EDGE NEED TO BE CONFIGURED AS ROOT ON THE HOST?

- No, however it is a recommended practice if the customer organization permits them to do so. Without root rights, some additional configuration may be required.

30. WHAT STREAM FUNCTIONALITY IS NOT AVAILABLE IN EDGE?

- Generally speaking, Collectors and Pull sources are the only functionality that isn't available in Edge.

31. HOW DOES EDGE INTEGRATE WITH A CUSTOMER'S EXISTING AGENT/OTHER SOLUTION?

- Depending on the vendor, Edge may replace an existing agent. In other cases, Edge can work in parallel with a customer's existing agent. Contact your Cribl sales rep for additional guidance.

32. FOR CRIBL EDGE CLOUD, WHO IS THE CLOUD PROVIDER?

- Amazon Web Services (AWS) is the cloud provider for Cribl Edge Cloud.

33. WHAT OPTIONS DO I HAVE TO MANAGE EDGE AND EDGE FLEETS?

- Edge can be managed from your existing Stream Leader. Fleets are groups of systems with the same configuration – similar to worker groups – and can be managed from the UI, in the command line, or via API.

34. 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.

37. HOW IS EDGE DELIVERED?

- Edge is available as a single binary downloaded from our website or as software as a service (SaaS).

36. ARE THERE ANY SECURITY OR CONCERNS WITH INSTALLING EDGE ON THE SOURCE?

- Some customers will not want their hosts/devices communicating directly to the Internet via an Edge instance. In that case, we recommend leveraging Stream here. Your Cribl sales rep will be able to offer additional guidance and best practices.

37. CAN EDGE REPLACE EXISTING LOG SHIPPERS?

- Yes.

38. CAN EDGE SEND DATA DIRECTLY TO NON-CRIBL DESTINATIONS?

- Yes, Edge is vendor-agnostic and can send data to any destination(s).

39. CAN EDGE SEND DATA TO MULTIPLE DESTINATIONS?

- Yes, Edge is vendor-agnostic and can send data to any destination(s).

40. CAN EDGE SEND SOME DATA TO STREAM AND OTHER DATA DIRECTLY TO A DESTINATION?

- Yes, Edge can send data to any destination(s), including Stream.

41. WHAT DESTINATIONS ARE CURRENTLY SUPPORTED BY EDGE?

- Edge supports all of the same destinations as Stream. [Please reference our docs](#) for more detailed information and examples.

42. WHICH SOURCES WILL BE SUPPORTED BY EDGE AT GA?

- Edges supports most of the same sources as Stream, except for Collectors and Pull sources. [Please reference our docs](#) for more detailed information and examples.

43. WHAT RELATIONSHIP/SUPPORT EXISTS BETWEEN EDGE AND APPSCOPE? IS APPSCOPE SHIPPED WITH EDGE OR DOES IT NEED TO BE DOWNLOADED SEPARATELY?

- AppScope ships with both Edge and Stream. With both Edge and AppScope operating within a host, AppScope can send data to Edge which will then process and forward that data as needed. AppScope can be deployed independently and is not dependent on Edge for operation. Please note that AppScope is open source and is not for sale at this time.

44. DOES THE EDGE DATA TRANSIT THROUGH THE MANAGER/LEADER (IN PATH)?

- No, data flow is independent of management traffic Leader and Edge node.

45. ARE THERE ANY PROCESSING ADVANTAGES OF DEPLOYING EDGE WITH STREAM?

- Yes. By deploying Edge instances on the host and performing processing and routing locally, less data will need to flow to a central Stream instance. This reduces Stream processing requirements, thus reducing the number of workers and worker processes.

46. WILL WE HAVE A CRIBL EDGE DEMO ENVIRONMENT?

Yes, the Cribl Edge Demo environment is available [here](#).

47. WILL WE ADD AN EDGE SANDBOX?

- Yes, our existing Sandboxes will be updated to include Edge.

48. WHAT FILE MONITORING CAPABILITIES DOES EDGE PROVIDE?

- Via the File tab, Edge automatically identifies files with actively running processes that are open for writing logs. This essentially allows you to install Edge on almost any system and immediately start gathering data. You can then select that file and view events. From there, filtering and search capabilities can assist in locating specific information.
- You can also configure Manual mode, allowing you to identify a top-level directory where Edge will search for logs in that location. Note: These files do not need to be 'open', as with auto mode in order to access data. You can also apply filters, create an Allowlist to only locate specific files or types (i.e., files ending in .log), or search sub-directories.

49. HOW MANY VCPUS WILL EDGE CONSUME?

- Currently Edge is limited to using a single vCPU.

50. HOW DO YOU INSTALL EDGE?

- You install Edge using the same process as Stream. Please visit our documentation for more detailed information, or consult with your Cribl sales rep.

51. WHAT HAPPENS TO EDGE NODES WHEN THE LEADER IS DOWN?

- Data processing continues uninterrupted when the Leader goes offline. Management and monitoring are impacted until the connection is reestablished.

Chapter 3: Positioning and Target Audiences

52. CAN I GET A QUICK OVERVIEW?

- 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.

53. WHAT CUSTOMER PROBLEMS IS CRIBL EDGE SOLVING?

- Data volumes – and the number of data sources – are growing year over year. Companies need to collect, process, and forward that data, but data collection at the edge in modern app and microservice environments just can't keep up with the scale and types of observability data. They need to process data closer to the source to keep costs down. They also need to collect all of the necessary data without installing, configuring, and managing multiple, independent agents. On top of all that, they need better visibility into edge data collection.
- Cribl Edge is the easiest way to get observability data out of any Linux system. Edge allows you to support your customers in reliably collecting and processing logs, metrics, and application data in real time from their Linux machines, apps, and microservices and deliver them to Cribl Stream or any supported destination. Edge's intelligent agent efficiently gathers and auto-discovers observability data at its egress point, so you open up additional, cost-effective options for data collection and processing for your customers.

With Edge's built-in Fleet Management, you can effortlessly manage hundreds to thousands of Cribl Edge nodes while lowering data collection TCO. Collect all of the data you need, at the edge, at scale.

54. WHAT IS A TYPICAL USE CASE FOR CRIBL EDGE?

- For system administrators struggling to collect data at scale, Cribl Edge efficiently gathers and processes observability data at its egress point, in any format, and delivers it to any analytics tool or destination(s) the customer chooses. With Edge's built-in Fleet Management, you can manage thousands of edge nodes from a central location.
- Key Use Cases:
 - **Explore** – Explore logs, metrics, and application data at the edge before deciding to collect and process
 - **Collect** – Collect logs, metrics and application data at scale in any Linux environment
 - **Process** – Reduce, transform, and route data at the edge giving more flexibility and lower costs for forwarding and storing data
 - **Verify** – Visually author configurations for data collection and transformation at the edge
 - **Forward** – Forward data to Cribl Stream or to any of the supported Destinations
 - **Manage** – Manage hundreds to thousands of Edge nodes with Fleet Management, lowering time and cost of ownership of data collection

55. ARE THERE ANY DIRECT COMPETITORS TO EDGE?

- Not at this time. Edge is the only vendor-agnostic, smart agent.

56. SHOULD I POSITION EDGE INSTEAD OF STREAM?

- Consider Edge an upsell to Stream. Your Cribl sales rep can provide more guidance here.

57. HOW DO WE POSITION EDGE IN COMPARISON TO A STREAM DEPLOYMENT?

- 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.

58. MY CUSTOMER IS ALREADY USING OTHER DATA COLLECTORS AND AGENTS.

- Collection of all the data a customer needs often requires installation, configuration, and management of multiple agents. To control costs and increase options for your data, processing has to be moved closer to the source. With Edge's built-in Fleet Management, you can help them manage hundreds to thousands of Cribl Edge nodes without relying on external tools for updates, upgrades, and management. Cribl Edge opens up more collection, processing, and forwarding options for your customer's data, so you can scale these use cases up or down with their data volumes and sources. When you pair Edge with Cribl Stream, your customer always has access to top-tier data processing and transformation capabilities and can get data moving in minutes.

59. IS THERE A USE CASE FOR JUST DEPLOYING EDGE, WITHOUT ANY OTHER CRIBL PRODUCT?

- Yes. Cribl Edge provides a single, data-agnostic agent with the ability to process data at the edge and forward it to analytics processors and/or storage systems.

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

- System administrators and existing Stream customers
- Anyone looking to collect data from multiple sources while keeping costs down and getting valuable insights from data
- Need to increase visibility into data collection and edge node configurations
- Currently experiencing pain points managing several agents at scale or across teams

61. 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.

62. WHAT ARE THE KEY BENEFITS OF EDGE?

- **Scalable Data Exploration at the Source:** Cribl Edge enables admins to explore logs, metrics, and application data at their egress points. When you do data exploration closer to the source, you get the option to dive into the data before deciding to collect and process it.
- **Better Data Collection:** Collect logs, metrics, and application data in any Linux environment at unprecedented scale. Easily manage hundreds to thousands of Edge nodes with Edge’s built-in Fleet Management, lowering time and cost of ownership (TCO) of data collection.
- **Cost-Effective Data Processing:** From there, you can reduce, transform, and route data at the edge, giving you more flexibility and lowering costs for forwarding and storing data. Forward data to Cribl Stream or the destination of your choice.
- **Interactive Troubleshooting:** Visually author configurations for data collection and transformation at the source. Cribl Edge’s rich, interactive agent allows an admin to teleport into any Edge node and get a detailed look at system performance metrics, running processes, containers, open log files, and more.

Chapter 4: Additional Information

63. WHERE CAN I FIND ADDITIONAL SALES AND TECHNICAL RESOURCES AND MATERIALS ON CRIBL EDGE?

Please visit [Cribl’s LMS](#) for the latest resources and sales materials. Both sales and technical resources are available.

64. WHERE DO I FIND THE LOGO/ICONS FOR CRIBL EDGE?

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.