Cribl

Simplify Kubernetes Instrumentation with **Cribl Edge for AWS EKS Add-on**

HOW-TO GUIDE

CONTENTS

INTRODUCTION

Simplify Kubernetes Instrumentation with Cribl Edge for AWS EKS Add-on

Kubernetes (K8s) has revolutionized application deployment, but it has simultaneously introduced a labyrinth of monitoring challenges that traditional IT approaches struggle to address. Modern Kubernetes environments require efficient data collection and security without adding operational complexity.

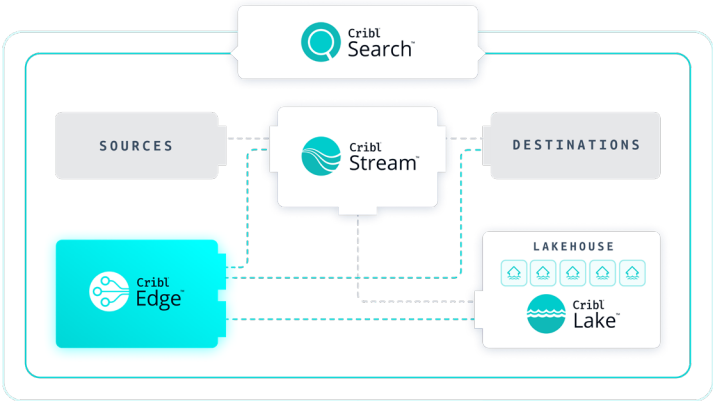

Cribl Edge provides that solution by seamlessly deploying on Amazon Elastic Kubernetes Service (EKS) clusters, allowing organizations to simplify Kubernetes instrumentation. With a single platform to collect, enrich, normalize, and route telemetry data, teams gain precise control over logs, metrics, and traces at the source.

Kubernetes environments generate massive volumes of telemetry data due to their dynamic and ephemeral workloads. Without effective data management, organizations face operational blind spots, overwhelmed observability stacks, and escalating costs. Running as a scalable agent, Cribl Edge ensures that only the most relevant data reaches downstream analytics tools, helping teams improve efficiency, manage data volumes, and strengthen security without disrupting performance.

Understanding the Problem: the complexity of Kubernetes

Kubernetes is an open-source container orchestration platform for automating the deployment, scaling, and management of containerized applications. It can be used in on-premises, hybrid, or public cloud infrastructures. Despite automating many processes, it's often difficult to bring them to fruition because traditional IT infrastructure tools weren't designed to handle our present-day challenges.

Architectural complexity

K8s is a complex system with numerous interdependent components, making it challenging to determine where to start instrumenting, what data to collect, and how to monitor performance effectively. Traditional monitoring tools often struggle to correlate these interconnected data points, leading to blind spots that hinder system understanding and operational visibility.

Growing data volume

K8s generates massive volumes of data due to its dynamic, ephemeral workloads and the sheer scale of unique metrics it produces. As containers spin up and down rapidly, logs and metrics flood monitoring systems, while high-cardinality data adds further strain, such as unique pod names and labels. Without efficient data collection and routing, organizations risk overwhelming their observability stack, driving up costs and slowing down insights.

Variety in instrumentation methods

Kubernetes has many different ways to instrument it, including using native K8s metrics and logs or third-party tools. While this variety is beneficial, too many options can make it challenging to choose the right approach for your goals.

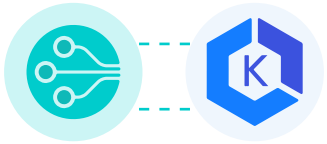

Cribl Edge for Amazon EKS Add-on: an overview

Amazon EKS is a fully managed service for running Kubernetes across the AWS Cloud and on-premises, providing scalability, reliability, and deep integration with AWS services. To simplify integrating third-party software, Amazon EKS includes an add-on functionality that enables customers to seamlessly install, manage, and deploy solutions directly from AWS Marketplace. This eliminates the need for manual deployment and configuration, reducing operational overhead and accelerating time to value.

The EKS add-on functionality streamlines Cribl Edge deployment onto Kubernetes clusters. Once deployed, Cribl Edge provides robust capabilities for collecting, transforming, enriching, and routing telemetry data. Instead of deploying multiple collectors or agents to route data to different destinations, Cribl Edge centralizes data management at the source, eliminating complexity while ensuring that only relevant data reaches downstream tools. This allows teams to control data volumes, improve operational efficiency, and maintain high performance across their Kubernetes environments.

Cribl Edge is purpose-built for dynamic environments like Kubernetes, offering intelligent data management directly within containerized applications. Unlike traditional logging agents that require complex configurations and sidecar deployments, Cribl Edge simplifies data collection by automatically gathering logs, metrics, and application data at scale. Teams can route this data to multiple destinations without the need for redundant agents, reducing management complexity and minimizing the risk of misconfigurations. Plus, developers can easily share EKS data between security and operations.

By using Cribl Edge with Amazon EKS, organizations gain precise control over their Kubernetes telemetry. Built-in fleet management capabilities allow teams to efficiently oversee thousands of distributed nodes, while a centrally managed, version-controlled approach reduces vendor lock-in and ensures flexibility as infrastructure needs evolve. The result is a streamlined, efficient, and scalable data management strategy that empowers teams to gain deeper insights into their Kubernetes environments for faster troubleshooting and enhanced security without compromising performance.

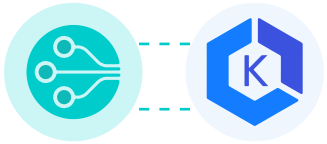

Cribl Edge for Amazon EKS Add-on: step by step

Follow these steps to integrate Cribl Edge into an Amazon EKS environment using the EKS add-on. For advanced configuration, refer to Cribl Edge for EKS Add-On Documentation or this Setup Video.

Prerequisites

1. Cribl requirements:
   - A Cribl.Cloud deployment or a self-hosted Cribl Leader instance.
2. AWS requirements:
   - Subscribe to Cribl Edge on AWS Marketplace.
   - Install `kubectl`, AWS CLI, and (optional) `eksctl`.
   - Access to an EKS cluster with permissions to create add-ons.

Step 1 | Retrieve Cribl Leader URI and token

1. Log into your Cribl Deployment in a self-hosted deployment, or log in to Cribl.Cloud, and select **Manage Edge**.
2. In Cribl Edge, select your Fleet and navigate to your Fleet configuration.
3. Select **Add/Update Edge Node** › **Kubernetes**.
4. Copy the `cribl.leader` URI from the Script field (starts with `tls://`). Example URI: `tls://<auth_token>@<cribl_org_id>.cribl.cloud?group=<default_fleet>`

Step 2 | Deploy via EKS Console

1. Open the **Amazon EKS Console** and select your cluster.
2. Go to **Add-ons** › **Get more add-ons**.
3. Search for "Cribl Edge" and enable the add-on.
4. Configure settings:
   - **Version**: Select the latest version.
   - **IAM role**: Choose "Inherit from node".
   - **Optional configuration settings** › **Configuration values**: Insert this JSON (replace placeholders with your URI):

```
{
 "cribl": {
     "leader": "tls://<token>@<cribl_cloud_leader>.criblcloud?
      group=<default_fleet>"
 }
}
```

5. Replace the `leader` value with the URI for your Edge Leader gathered in Access the Cribl URI and Token.
   - For self-hosted Cribl deployments, enter the publicly available URI and token for your Leader Node.
6. Set **Conflict resolution method** to "None".

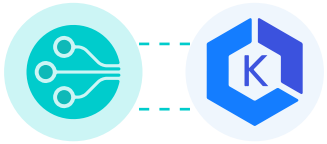

Cribl Edge for Amazon EKS Add-on: step by step *(continued)*

Step 3 | Enable Cribl Edge Add-on using AWS CLI

1. Run this command to install the add-on:

   ```
   aws eks create-addon --addon-name cribl_cribledge --cluster-name $YOUR_CLUSTER_NAME --region $AWS_REGION --configuration-values file://path_to_cribl_leader.json
   ```

1. Replace:
   - `$YOUR_CLUSTER_NAME` with your EKS cluster name
   - `$AWS_REGION` with your AWS region
   - `path_to_cribl_leader.json` with your configuration file containing the Leader URI

Step 4 | Verify installation

1. Check status with:

   ```
   aws eks describe-addon --addon-name cribl_cribledge --cluster-name $YOUR_CLUSTER_NAME --region $AWS_REGION
   ```

2. Look for `"status": "ACTIVE"` in the output.

Post-deployment configuration

1. Add Edge nodes:
   - Use the CLI command `./cribl mode-managed-edge` on worker nodes to connect to your Leader.
2. Configure data routing:
   - Set up Sources and Destinations in Cribl Edge UI.
   - Commit and deploy configuration changes.

Conclusion

Modern Kubernetes environments aren't monolithic. They're complex, interconnected processes that need to work together. Cribl Edge with Amazon EKS transforms Kubernetes monitoring from a reactive process to an intelligent, proactive data management strategy.

By consolidating data at the source and enabling flexible routing, organizations can ensure that only the most valuable information reaches their analytics tools. With built-in fleet management and a centrally managed approach, teams can efficiently oversee thousands of distributed nodes while maintaining control over their data strategy. This integration empowers security and operations teams with deeper insights, faster troubleshooting, and enhanced performance, all without adding unnecessary operational overhead.

Additional resources

- Setup video
- Cribl Docs on EKS Add-on
- Simplify Kubernetes with Cribl Edge on EKS Add-on Blog
- Direct Deployment to Amazon EKS Clusters Technical Blog
- Learn more about Amazon EKS
- Amazon EKS overview video

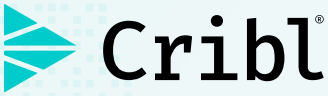

Cribl, the Data Engine for IT and Security, empowers organizations to transform their data strategy. Powered by a data processing engine purpose-built for IT and Security, Cribl's product suite is a vendor-agnostic data management solution capable of collecting data from any source, processing billions of events per second, automatically routing data for optimized storage, and analyzing any data, at any time, in any location. With Cribl, IT and Security teams have the choice, control, and flexibility required to adapt to their ever-changing data needs. Cribl's offerings — Stream, Edge, Search, and Lake— are available either as discrete products or as a holistic solution.

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

EB-0011-EN-1-0525