



A vendor-neutral, centralized agent management system to simplify endpoint data collection at any scale.



The **AI Platform**
for Telemetry

Cribl Edge lets you discover, collect, and process telemetry in real time from servers, laptops, applications, and microservices running on Linux, Windows, macOS, or Kubernetes. You can send that data to Cribl Stream for deeper processing or to any other destination you choose.

Explore

- Investigate logs, metrics, and application data right at their egress points.

Collect


- Automatically collect logs, metrics, and application data at scale.

Monitor

- Quickly deploy thousands of nodes to monitor your growing infrastructure.

Forward

- Shape endpoint data and send it to the best tool for the job—or all the tools you need.



A modern agent

Use a rich UI with visual change authoring, fast upgrades, and autodiscovery to explore your Edge nodes and what they're collecting.



Centralized management

Use Fleet Management to deploy and manage hundreds of thousands of nodes from one place. Centralize configurations, upgrades, and reboots in just a few clicks.



Efficient collection

Analyze data at the source before you decide to collect it. Take advantage of unused endpoint capacity to cut infrastructure costs instead of driving up the cost of centralized processing.



Vendor-neutral

Route data dynamically with a vendor-agnostic agent. Use out-of-the-box integrations to send data to Cribl Stream for heavier processing or directly to any destination you want.

Product Features

AGENT FEATURES

- An intelligent, highly-scalable edge based data collection system for logs, metrics, and application data.
- Supports Linux or Windows machines, apps, microservices, and Kubernetes environments.
- Automatic discovery of host, container, and application metrics, logs, etc. on endpoints.
- Collect, process, and forward data with low resource overhead.
- Host, container, and cloud metadata discovery and event enrichment.
- Centrally managed, configured, and version controlled for easy expansion and low cost of ownership.
- Experience designed to allow users to see what is being collected and how it can be optimized from the UI.

ARCHITECTURE

- Centralized configuration, monitoring and software upgrades for hundreds of thousands of nodes.
- Collection of independent (un-managed) nodes.
- Sub-millisecond latency.
- Tested to upwards of 20PB/day.
- Deployment options include:
 - SW including Linux or Windows binaries, docker containers, and helm charts for easy deployment in any K8s environment and runtimes.
 - Cloud provides SaaS experience via Cribl.Cloud; entirely Cribl managed, no infrastructure overhead, and scales as needed.
 - Hybrid-leader / control plane in the cloud and Edge Node performing local processing.

SYSTEM MANAGEMENT

- Enterprise grade authentication support (LDAP, SSO, etc).
- Policy-based RBAC for fine-grained permissioning.
- Configuration version control via Git.
- Intuitive, rich user interface for distributed system management.
- Built-in, real-time configuration change validation.
- Centralized support for certificate and key management, ability to leverage external key management services for managing secrets / tokens across all nodes.
- Centralized upgrade of Edge nodes on Windows endpoints without stepping out of Cribl UI with just a single click.
- Built-in synchronization with external code repositories for CI / CD integrations and disaster recovery.

INTEGRATIONS

- 60+ out-of-the-box integrated sources / destinations.
- Sources include logs, metrics, and application data, as well as system and container level metrics and logs.
- New Prometheus Edge Scraper purpose built for K8s allows discovery and metrics collection at scale and route to destination of choice.
- Support for monitoring and displaying compressed and uncompressed files and one-time ingestion of files.

- Support for Kubernetes logs and metrics spooling allows for searching at the edge using Cribl Search without feeding data into an analysis system.
- Native protocol support for leading sources and destinations of logs and metrics.
- Out-of-the-box TLS support for all integrations that support it.
- Out-of-the-box support for IAM and assume roles (AWS specific).
- Rich logging, metrics, and real-time status for each integration.
- Baked-in connectivity tests and results for each integration.
- Support for sending and collecting from all major Cloud PaaS storage services.
- Support for arbitrary Script based data collection.

FLEET MANAGEMENT

- Fleet: a collection of Edge nodes that share the same configuration.
- Fleets facilitate authoring and management of configuration settings for a particular set of Edge nodes.
- Centralized management for up to 1000s of fleets / nodes.
- Support for multiple fleets / sub-fleets within the enterprise.

MONITORING

- Kubernetes control plane events.
- Built-in monitoring covering all aspects of a distributed deployment.
- Built-in centralized log search across 100s of groups / nodes / fleets.
- Rich, visually dense, dashboards built for admins / operators.
- Contextual monitoring for all sources and destinations.
- Notification system alerts operators when data flows have stopped.
- Dataflow visualizations provide birds eye view of all sources, routes, pipelines, and destinations.

WORKING WITH DATA

- Interactive, user-friendly UI for working with streaming data.
- Visual authoring, validation, and troubleshooting of data pipelines.
- Data preview with instant feedback for visual inspection of events as they're being transformed.
- Built-in data generators for pipeline and destination testing.
- Built-in documentation and contextual help on every screen.
- Live capture on multiple points as events travel from source to destination for inspection and / or troubleshooting.
- Ability to forward full-fidelity data to external systems.

- Over 30 out of the box functions that support arbitrary data transformations, securing, and enrichment.
- Over 40 built-in C. function methods for finer processing capabilities.
- ...plus all the power of JavaScript for almost and - arbitrary data transformations.
- Automatic byte-stream to events conversion / breaking using intelligent rules with optional user overrides.
- Timezone recognition and / or correction.
- Built-in JavaScript expression editor with live result preview.
- Built-in Regex editor with live match and capturing group preview.
- Built-in Regex Library for most common regex, extensible.
- Out-of-the-box parsing support for many well known data sources.
- Regex-based field extractions and native Grok pattern support.
- Event schema validation support using JSON Schema standard.
- Support for Global Variables – Re-usable and composable JS expressions that can be referenced by any function.
- Real-time data enrichment via lookup tables. Exact, Regex, and CIDR support out of the box.
- Support for geoup enrichment using Maxmind binary databases.
- Access to a growing community of Packs with pre-built pipelines, and custom functions to accelerate time to value of Edge.
- Global search makes finding data easy and fast.
- Gather live data samples to aid in development of pipelines or to share with teammates working on similar projects.

TECHNICAL REQUIREMENTS

Edge leader

- OS: Linux: RedHat, CentOS 7+, Ubuntu, AWS Linux, Suse (64 bit).
- System: ~ +4 physical cores, +8GB RAM, 5GB free disk space.
- Also available as a SaaS solution through Cribl.Cloud.

Edge nodes (minimum requirements)

- OS: Windows 10, 11 (always-connected); Windows Server 2016, 2019, 2022, 10. Linux: RedHat, CentOS 7+, Ubuntu, AWS Linux, Suse (64 bit).
- System: ~1Ghz processor, 512MB RAM, 5GB of free disk space (more if persistent queuing is enabled on Edge Node).

Browsers Supported:

- Firefox 65+, Chrome 70+, Safari 12+, Microsoft Edge.



The AI Platform for Telemetry

Learn more at cribl.io | Join our [Slack community](#)
Try [Cribl Sandboxes](#) | Follow us on [LinkedIn](#) and [X](#)

© COPYRIGHT 2026 CRIBL, INC. ALL RIGHTS RESERVED

DS-0009-EN-1-0726