



A vendor-neutral, centralized agent management system to simplify the collection of endpoint data - no matter the scale.

Cribl Edge enables real-time discovery, collection, and processing of telemetry data from servers, laptops, applications, and microservices in Linux, Windows, macOS, or Kubernetes environments, seamlessly delivering it to your destination of choice.

A modern agent

Experience a rich UI, visual change authoring, rapid upgrades, and autodiscovery capabilities for exploring Edge nodes.

Centralized management

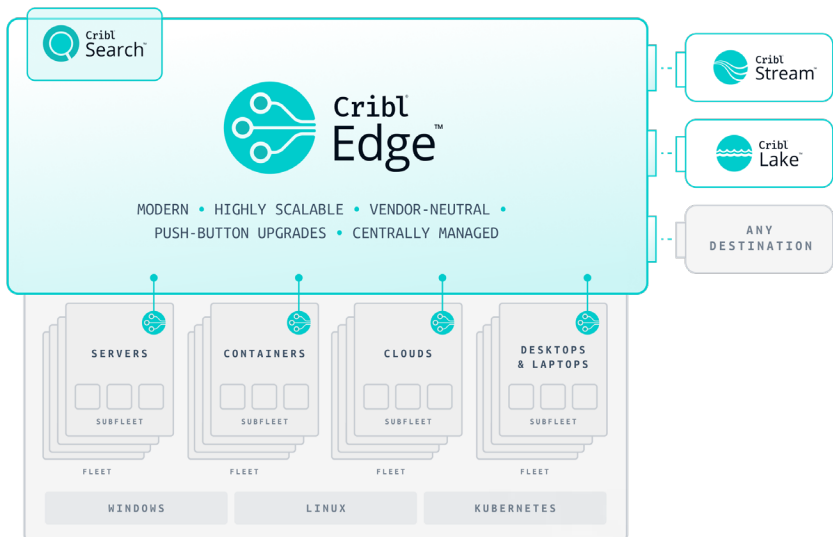
Effortlessly deploy and manage hundreds of thousands of nodes with Fleet Management. Centralize configurations, upgrades, and reboots with just a few clicks.

Efficient collection

Analyze data at the source, before deciding to collect it. Leverage unused endpoint capacity to reduce infrastructure costs rather than contributing to the costs of centralizing it.

Vendor-neutral

Dynamically route data using an agnostic agent with out-of-the-box integrations to send to Cribl Stream for heavier processing or any destination of your choice.



EXPLORE

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

COLLECT

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

MONITOR

Easily deploy thousands of nodes to monitor your growing infrastructure.

FORWARD

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

Product Features

AGENT FEATURES

- An intelligent, highly-scalable edge based data collection system for logs, metrics, and application data
- Supports Linux, Windows, and macOS machines, laptops and desktops, 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, Windows, and macOS 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 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 geoip 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; Windows Server 2016, 2019, 2022, 2025; Linux: RedHat, CentOS, Ubuntu, AWS Linux, Suse; macOS Tahoe, Sequoia, and Sonoma
- 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

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.

DS-0002-EN-6-1025