DATA SHEET

Cribl Edge™

Intelligent, highly scalable edge-based data collection system for logs, metrics, and application data. Designed to support today’s modern microservice architectures and sprawling environments.

Cribl Edge provides the ability to discover, collect, and process observability data – logs, metrics, application data – in real time, from your Linux machines, apps, microservices etc., and deliver them to Cribl Stream or any supported destination.

Benefits

INTELLIGENCE

Intelligent agent that efficiently auto discovers and gathers observability data at its egress point

EXPLORE

Explore logs, metrics, and applications data at the edge before deciding to collect and process

COLLECT

Automatically collect logs, metrics, and application data at scale in any Linux environment

MANAGEMENT

Manage hundreds to thousands of Edge nodes with Fleet Management, reducing time and effort of ownership.

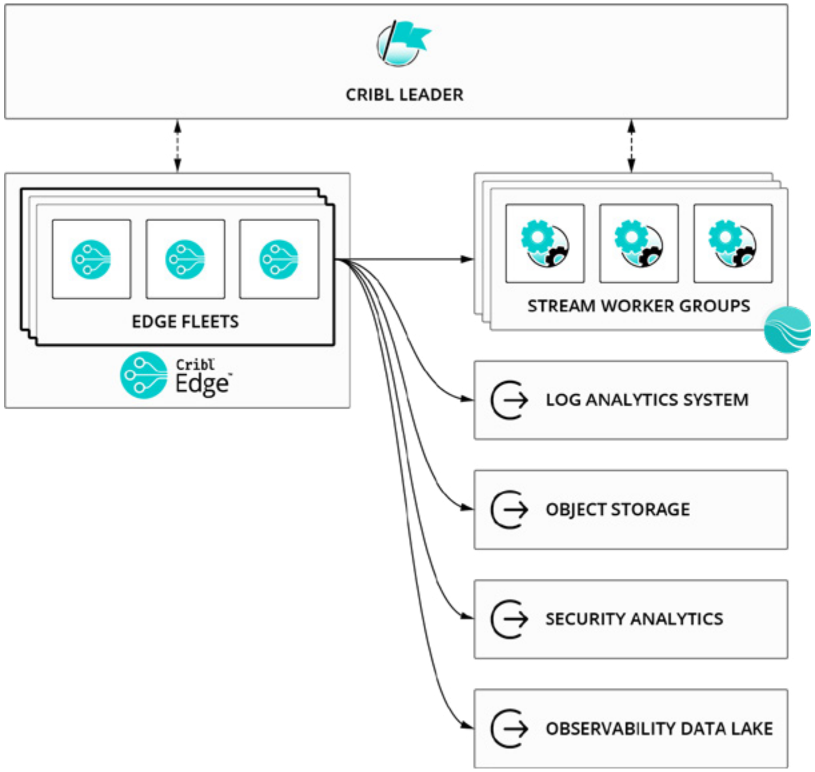

COLLECT THE RIGHT DATA
Get data from any source to any tool in the right format

SHAPE YOUR DATA
Take data as it comes, shape it into what you need

ROUTE YOUR DATA
Put data where it has the most value

REDUCE YOUR DATA
Eliminate uninteresting data to control costs

Product Features

AGENT FEATURES

- *An intelligent, highly-scalable edge based data collection system for logs, metrics, and application data*
- *Supports Linux machines, apps, microservices 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*
- *When used in conjunction with AppScope, allows on-demand instrumentation of any process or executable operating on host*

ARCHITECTURE

- *Leader/Edge Node deployment supporting 1000s of nodes comprising 1 or more fleets*
  - *Collection of independent (un-managed) nodes*
- *Sub-millisecond latency*
- *Tested to upwards of 20PB/day*
- *Deployment options include:*
  - *SW including Linux binaries, docker containers, and helm charts for easy deployment in any K8s environment*
  - *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*

INTEGRATIONS

- *Out of the Box integrated Sources/Destination*
  - *Sources include Logs, Metrics and Application data, as well as system and container level metrics and logs*
  - *40+ Destinations including Streaming, Non-Streaming and special purpose (Output router, DevNull, ...)*
- *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 arbitrary Script based data collection*
- *Support for sending and collecting from all major Cloud PaaS storage services*

SYSTEM MANAGEMENT

- *Enterprise grade authentication support (LDAP, SSO etc)*
- *Policy-based RBAC for fine-grained permissioning*
- *Intuitive, rich user interface for distributed system management*
- *Configuration version control via Git*
- *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*
- *Built-in synchronization with external code repositories for CI/CD integrations and disaster recovery*

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 within the enterprise*

MONITORING

- *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*
- *Ability to forward full-fidelity data to external systems*
- *Built-in data generators for pipeline and destination testing*
- *Live capture on multiple points as events travel from source to destination for inspection and /or troubleshooting*
- *Built-in documentation and contextual help on every screen*
- *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, Ubuntu, AWS Linux, Suse (64 bit)*
- *System: ~ +4 physical cores, +8GB RAM, 5GB free disk space*

Edge Node

- *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 Edge Node)*

Browsers Supported:

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