

Cribl Stream™

Take Control and Shape Your Data.

All of Your Observability Data. All Under Control. All Under Budget.

Stream allows you to implement an observability pipeline helping you parse, restructure, and enrich data in flight – ensuring that you get the right data, where you want, in the formats you need.

Benefits

FLEXIBILITY


Add new analytics tools and other destinations without adding agents or collectors.

CONTROL YOUR DATA

Route data to the right teams and the right tools.
Deploy enterprise-grade security.

SIMPLIFY OBSERVABILITY

Use a universal log and metrics router for any tooling environment.
Easily visualize observability data flows and port functionality across worker nodes.



Instrument everything. Observe more. Pay less.

ROUTE YOUR DATA

Put data where it has the most value

REDUCE YOUR DATA

Eliminate uninteresting data to control costs

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

Product Features

ARCHITECTURE

- Single binary distribution with zero dependencies
- Shared-nothing, super scalable distributed architecture
- Scales from laptop to 100s of nodes and 10000s of cores
- Highly parallelizable, highly performant platform built for extensibility
- 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 through Cribl.Cloud, entirely Cribl managed, no infrastructure overhead and scales as needed
 - Hybrid-Leader/Control plane in the cloud and workers performing local processing Integrations
- Over 80 source/destination integrations available out of the box
- Native protocol support for leading sources and destinations of logs, metrics, and traces
- Out-of-the-box TLS support for all integrations that support it
- Out of the box support for IAM and Assume roles (AWS specific)
- Live data capture for integration for troubleshooting and inspection
- Rich logging, metrics and real-time status for each integration
- Baked-in connectivity tests & results for each integration
- Support for arbitrary REST endpoint data collection
- Support for arbitrary Script based data collection
- Support for sending and collecting from all major Cloud PaaS storage services

MANAGEMENT

- Full control of all your observability data from a central control plane
- Enterprise grade authentication support (LDAP, SSO etc)
- Policy-based RBAC for fine-grained permissioning
- Intuitive, rich user interface for distributed system management
- Single, centralized management via cloud or self-hosted software for 100s of groups/nodes

- Dependable configuration version control with ability to revert changes
- Built-in, real-time configuration change validation
- Centralized support for certificate and key management
- Built-in data generators for pipeline and destination testing
- Fully automated and distributed upgrades of all Stream workers
- 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

MONITORING

- Notification system alerts operators when data flows have stopped
- Built-in Monitoring covering all aspects of a distributed deployment
- Built-in centralized log search across 100s of groups/nodes
- Rich, visually dense, dashboards built for admins/operators
- Contextual monitoring for all sources and destinations
- Ability to forward full-fidelity internal logs & metrics to external solutions
- Dataflow visualizations provide Birds Eye View of all sources, routes, pipelines and destinations

WORKING WITH DATA

- Interactive, user-friendly, efficient 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
- Live capture on multiple points as events travel from source to destination
- Built-in documentation and contextual tooltips 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-arbitrary data transformations
- IDE-like experience with auto-complete and typeahead assist
- Automatic byte-stream to events conversion/breaking using intelligent rules with optional user overrides

- Automatic timestamp format recognition 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 & capturing group preview
- Built-in Regex Library for most common regex, extensible
- Out-of-the-box parsing support for many well known data sources
- User-defined data parsers for K=V, CSV, ELFF, CLF, JSON and delimiter based values
- 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
- Packs support for building, packaging and sharing routes, pipelines, data samples, functions internally or with members of the community
- Access to a growing community of Stream Packs with pre-built pipelines, custom functions and other out-of-the-box features for speeding up and improving the value of Stream
- Global search makes finding anything in Stream easy and fast
- Gather live data samples to aid in development of pipelines or to share with teammates working on similar projects.

TECHNICAL REQUIREMENTS

System

- +4 physical cores
- +8GB RAM
- 5GB free disk space (more if PQ enabled)
- Also available as a SaaS solution through Cribl.Cloud

Sizing Guidance

- 1 physical core for each 400GB/day of IN+OUT throughput
E.g., 4 TB IN -> full 4TB to Destination A, plus 2 TB to Destination B = 10TB total = 25 physical cores.