

Cribl LogStream™

All of your Observability Data. All under Control. All under Budget.

LogStream 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

COST REDUCTION

Reduce log data by 30% or more to dramatically reduce your license, storage, and infrastructure costs.

COMPLETE CONTROL OF YOUR DATA

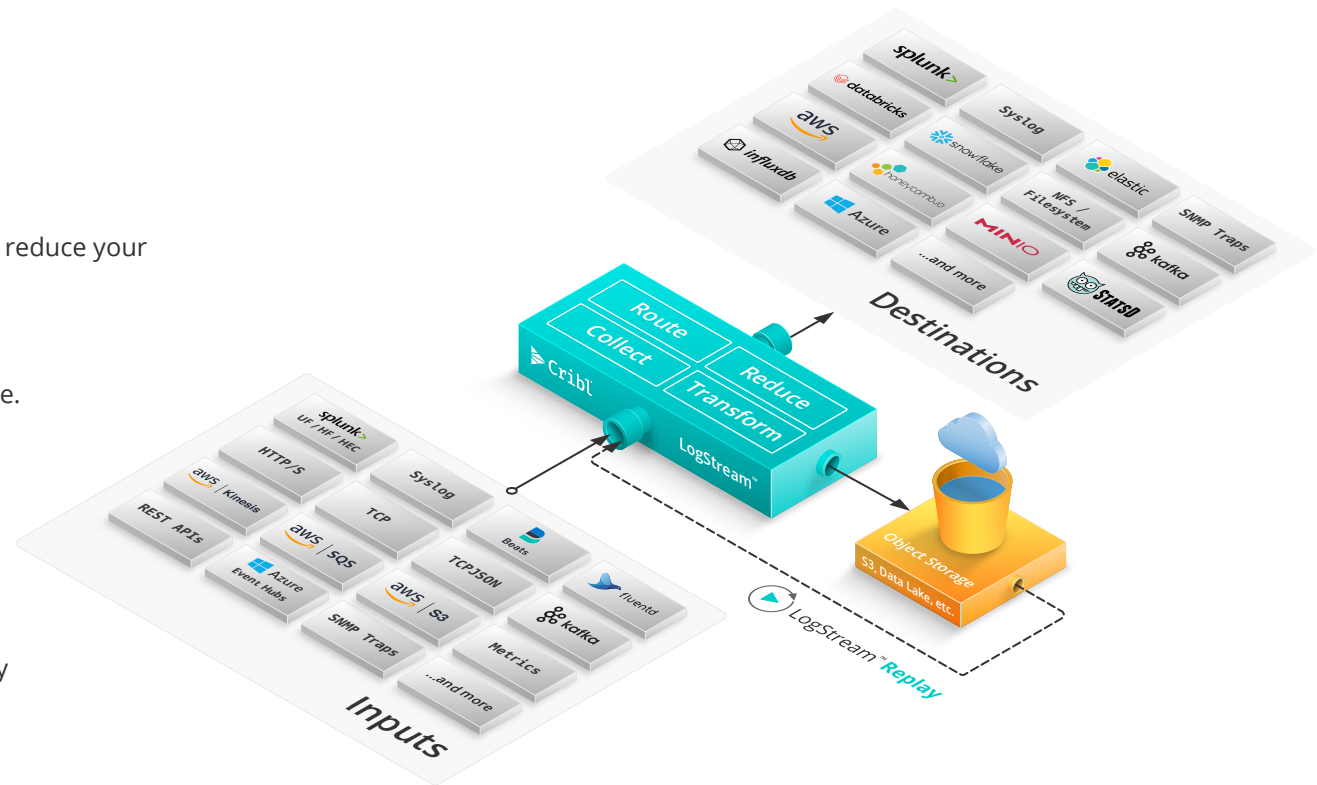
All machine data, in any format, to any data store. As data volumes continue to grow, we give you control to choose the right data sources and the best tools to get the answers you need.

PLATFORM INDEPENDENCE

Route to multiple destinations without having to install new agents or collectors with a vendor-agnostic solution that gives you flexibility with the vendors that are right for you.

IMPROVED TIME TO INSIGHT

Centralize forwarding of all machine data to be shaped, enriched, and ready for analysis. Filter out data with little analytical value so you focus on insightful data you can use.



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

INTEGRATIONS

- Over 50 integrations out of the box (more added on every release)
- Native protocol support for leading sources and destinations of logs & 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)
- 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

MANAGEMENT

- 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 console for 100s of groups/nodes
- Dependable config 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

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 (more added with every release) that support arbitrary data transformations, securing and enrichment.
- Over 40 built-in C.* function methods for finer processing capabilities ... plus powerful 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 geoiip enrichment using Maxmind binary databases.

MONITORING

- 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

MINIMUM TECHNICAL REQUIREMENTS

System

- 4 physical cores
- 8GB RAM
- 5GB free disk space

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.