



REPORT ON CONFLUENT'S SERVICE RELEVANT TO SECURITY AND AVAILABILITY (SOC 3 REPORT)

FOR THE PERIOD JANUARY 1, 2019 TO DECEMBER 31, 2019



Section I – Report of Independent Service Auditors

To: Confluent, Inc.

Scope

We have examined Confluent’s accompanying assertion, titled “Confluent’s Assertion” (assertion), that the controls within Confluent’s system were effective throughout the period January 1, 2019 to December 31, 2019, to provide reasonable assurance that Confluent’s service commitments and system requirements were achieved based on the trust services criteria relevant to security and availability (applicable trust services criteria) set forth in TSP section 100, *2017 Trust Services Criteria for Security, Availability, Processing Integrity, Confidentiality, and Privacy* (AICPA, *Trust Services Criteria*).

Service Organization’s Responsibilities

Confluent is responsible for its service commitments and system requirements and for designing, implementing, and operating effective controls within the system to provide reasonable assurance that Confluent’s service commitments and system requirements were achieved. Confluent has provided the accompanying assertion about the effectiveness of controls within the system. When preparing its assertion, Confluent is responsible for selecting, and identifying in its assertion, the applicable trust services criteria, and for having a reasonable basis for its assertion by performing an assessment of the controls within the system.

Service Auditor’s Responsibilities

Our responsibility is to express an opinion, based on our examination, on whether management’s assertion that controls within the system were effective throughout the period to provide reasonable assurance that the service organization’s service commitments and system requirements were achieved based on the applicable trust services criteria. Our examination was conducted in accordance with attestation standards established by the American Institute of Certified Public Accountants. Those standards require that we plan and perform our examination to obtain reasonable assurance about whether management’s assertion is fairly stated, in all material respects. We believe that the evidence we obtained is sufficient and appropriate to provide a reasonable basis for our opinion.

Our examination included:

- Obtaining an understanding of the system and the service organization’s service commitments and system requirements
- Assessing the risks that controls were not effective to achieve Confluent’s service commitments and system requirements based on the applicable trust services criteria
- Performing procedures to obtain evidence about whether controls within the system were effective to achieve Confluent’s service commitments and system requirements based on the applicable trust services criteria

Our examination also included performing such other procedures as we considered necessary in the circumstances.

Service Auditor's Independence and Quality Control

We have complied with the independence and other ethical requirements of the Code of Professional Conduct established by the AICPA. We applied the Statements on Quality Control Standards established by the AICPA and, accordingly, maintain a comprehensive system of quality control.

Inherent Limitations

There are inherent limitations in the effectiveness of any system of internal control, including the possibility of human error and the circumvention of controls. Because of their nature, controls may not always operate effectively to provide reasonable assurance that the service organization's service commitments and system requirements are achieved based on the applicable trust services criteria. Also, the projection to the future of any conclusions about the effectiveness of controls is subject to the risk that controls may become inadequate because of changes in conditions or that the degree of compliance with policies or procedures may deteriorate.

Opinion

In our opinion, management's assertion that the controls within Confluent's system were effective throughout the period January 1, 2019 to December 31, 2019, to provide reasonable assurance that Confluent's service commitments and system requirements were achieved based on the applicable trust services criteria, is fairly stated, in all material respects.

Cadence Assurance LLC

March 27, 2020
Salt Lake City, Utah



Section II – Confluent’s Assertion

We are responsible for designing, implementing, operating, and maintaining effective controls within Confluent’s system throughout the period January 1, 2019 to December 31, 2019, to provide reasonable assurance that Confluent’s service commitments and system requirements relevant to security and availability were achieved.

We have performed an evaluation of the effectiveness of the controls within the system throughout the period January 1, 2019 to December 31, 2019, to provide reasonable assurance that Confluent’s service commitments and system requirements were achieved based on the trust services criteria relevant to security and availability (applicable trust services criteria) set forth in TSP section 100, *2017 Trust Services Criteria for Security, Availability, Processing Integrity, Confidentiality, and Privacy* (AICPA, Trust Services Criteria). Confluent’s objectives for the system, in applying the applicable trust services criteria, are embodied in its service commitments and system requirements relevant to the applicable trust services criteria.

There are inherent limitations in any system of internal control, including the possibility of human error and the circumvention of controls. Because of these inherent limitations, a service organization may achieve reasonable, but not absolute, assurance that its service commitments and system requirements are achieved.

We assert that the controls within the system were effective throughout the period January 1, 2019 to December 31, 2019, to provide reasonable assurance that Confluent’s service commitments and system requirements were achieved based on the applicable trust services criteria.

Confluent, Inc.
March 27, 2020



Section III – Confluent’s Description of its Confluent Service

Company Overview

Confluent was founded by the team that built Apache Kafka. Apache Kafka is a community-distributed, event-streaming platform capable of handling trillions of events a day. Initially conceived as a messaging queue, Kafka is based on an abstraction of a distributed commit log. Since being created and open sourced by LinkedIn in 2011, Kafka has quickly evolved from messaging queue to a full-fledged event streaming platform. Confluent delivers the most complete distribution of Kafka with its Confluent Cloud System (Confluent Cloud). Confluent Cloud improves Kafka with additional community and commercial features designed to enhance the streaming experience of both operators and developers in production, at massive scale. Confluent provides a streaming platform that enables companies to access data as real-time streams. Confluent is headquartered in Mountain View, CA with additional offices in San Francisco, CA and London, UK, and remote workers at various locations. Currently, Confluent employs approximately 1,000 employees across these locations.

Confluent believes every byte of data has a story to tell, something of significance that informs the next thing to be done. In a data-driven enterprise, how data moves is nearly as important as the data itself. With greater speed and agility, data’s value increases exponentially.

System Description

Confluent Cloud is comprised of the following components:

- *Auto Data Balancer* – Allows organizations and users of the system to:
 - Easily add and remove nodes from Kafka clusters
 - Rebalance partitions across a cluster via rack aware algorithm
 - Throttle traffic from balancing when data transfer occurs
- *Confluent Control Center™* – A comprehensive management and monitoring system for Apache Kafka. Control Center provides:
 - Users the ability to monitor and manage clusters from a rich user interface
 - Users the ability to quickly scan through clusters for anomalies and track down messages to their sources
 - Full integration with connectors, allowing users to manage data pipelines without a line of code
 - The delivery of real-time analysis of the performance of Kafka
 - The ability to drill into topics, producers, consumers, and more to understand the activity within their data pipelines enabling organizations to govern a growing ecosystem of stream data applications



- *Java Messaging Service (JMS) Client* – An adaptor client, which allows the use of the standard JMS 1.1 Application Programming Interface (API) backed by Kafka. Because the JMS client is a drop-in implementation, it allows users to migrate applications from legacy message queues and gain the modern design, implementation, and scalability of Kafka. Key features allow users to:
 - Develop Java applications leveraging the standard JMS APIs
 - Transparently swap in the JMS client and Kafka without recompiling to migrate applications to Kafka
 - Utilize both producer and consumer APIs
 - Be compatible with most JMS 1.1 features
- *Kafka Connect* – Provides organizations and users with a framework that integrates Kafka with other systems to make it easy to add new systems to scalable and secure stream data pipelines. Connectors translate data between Kafka and other systems, while supporting a variety of data formats and lightweight inline transformations. The following connectors are developed, tested, documented, and fully supported by Confluent Cloud:
 - Active MQ Connector (Source)
 - Amazon S3 (Sink)
 - Confluent Replicator (Source & Sink)
 - Elasticsearch (Sink)
 - Filestream Connector (Source & Sink)
 - IBM MQ Connector (Source)
 - HDFS (Sink)
 - JDBC (Source & Sink)
 - JMS (Source)

Additional supported connectors may be installed separately via Confluent Hub.

- *KSQL* – The streaming SQL engine that enables real-time data processing against Apache Kafka. It provides an easy-to-use, yet powerful interactive SQL interface for stream processing on Kafka without the need to write code in a programming language such as Java or Python. KSQL is scalable, elastic, fault-tolerant, and it supports a wide range of streaming operations, including data filtering, transformations, aggregations, joins, and windowing.
- *Replicator* – A tool for managing streaming pipelines spanning multiple data centers. It enables users to define topologies spanning multiple data centers and perform multi-datacenter replication to execute these topologies. Replication can be performed in both active/passive and active/active modes between multiple data centers, or to aggregate data into a global data center. Features of Replicator include:
 - Increased reliability by giving users the ability to easily configure and maintain cross-cluster replication and real-time monitoring of replication lag
 - Easy management of multi-cluster deployments
 - Centralized configuration and monitoring



- The ability to replicate an entire cluster or a subset of topics
- Automatic replication of topic configuration
- Automated security using Kafka's Simple Authentication and Security Layer (SASL) for Kerberos and Active Directory
- Transport Layer Security (TLS) encryption between data centers

Confluent Cloud Products and Services

The Confluent Cloud product provides customers with the distribution of the Confluent technology, inclusive of Apache Kafka, as a service in the public cloud, simplifying engineering operations and administration of Kafka clusters and related services such as KSQL, Connect, and Schema Registry. Deployed in minutes, it is a streaming data service for the cloud-first developer on a mission or the operations-starved organization. It complements Apache Kafka with administration, monitoring, and management tools. The Confluent Cloud system is comprised of the following components:

- *Kafka* – Confluent Cloud provides an API-based service for the latest, stable Apache Kafka version. Confluent handles the upgrades on behalf of its customers and provides it in a seamless fashion.
- *Customer Support (optional)* – Full support for the range of products offered by Confluent, including Java, Python, C/C++, Go, .NET, as well as the Kafka Streams API.
- *Managed Service* – Remove the operations burden with a fully Confluent-managed cloud service.
 - Clusters can be created and destroyed on-demand, in any cloud region the service is offered, with any configuration of throughput available; and,
 - API keys for access to each cluster are self-managed and completely under customers' control.

The Confluent Cloud system offers the following features:

- *Performance* – Highest throughput rate of any streaming data service. Standard plans and custom plans are available for scales up to 1GBs ingest and 2GBs egress.
- *Reliability* – Optional support for high availability across multiple Availability Zones is available. The service comes with a 99.95 percent service level agreement.
- *Flexibility* – Configurable retention period, storage, and throughput rate to suit customer workloads.
- *Kafka Expertise* – Support provided by the team that created Kafka and that has the most extensive experience operating it at scale.

System Boundaries

Included within the scope of this report are the production systems, infrastructure, software, people, procedures, and data supporting Confluent Cloud and the hosted Confluent Cloud environment. This report is specific to Confluent Cloud and does not include Confluent Platform.



Subservice Organizations

Confluent utilizes cloud service providers Amazon Web Services (AWS), Microsoft Azure (Azure), and Google Cloud Platform (GCP) for Confluent's data center, infrastructure, software, and hosting services. AWS, Azure, and GCP are excluded from the scope of this report; the controls for which they are responsible are found in the subsequent section entitled *Complementary Subservice Organization Controls (CSOCs)*.

Principle Service Commitments and System Requirements

Confluent communicates operational requirements to support the achievement of security and availability through its policies and in its contracts with customers. Confluent's commitments are documented and communicated to customers through the following:

- Terms of Service (<https://www.confluent.io/marketplace-terms-of-service/>)
- Data Processing Addendum (<https://www.confluent.io/cloud-customer-dpa/>)
- Confluent Cloud Security Addendum (<https://confluent.io/cloud-enterprise-security-addendum>)



System Components

The components of Confluent Cloud include the following infrastructure, software, people, procedures, and data elements.

Infrastructure

Confluent Cloud is designed, implemented, and operated to achieve specific business objectives in accordance with management-specified requirements. The purpose of the system description is to delineate the boundaries of the system, which include the services outlined above and the components described below.

Confluent utilizes cloud service providers AWS, Azure, and GCP for its computing platforms, network technologies, and internal global infrastructure. Confluent Cloud primarily uses Alpine Linux and Debian operating systems in AWS, Azure, and GCP to support the infrastructure systems supporting Confluent Cloud.

Software

Confluent has various software programs and tools used to support Confluent Cloud. These programs and tools assist with monitoring, authentication, automation of software development, issue tracking, incident response, customer relationship management, and encryption.

People

Confluent teams and functions who support the Confluent environment include Security, Product Engineering, Growth and Marketing, Customer Success, Business Systems, People Operations and Recruiting, Finance, Legal, Information Technology, and Information Security Compliance.

Procedures

Confluent maintains a set of policies that are published and communicated to Confluent personnel. Policies are updated as necessary and are reviewed and approved.

The following policies are relevant to the scope of this report:

- Acceptable Use Policy
- Information Security Policy
- Access Management Standard
- Asset Management Standard
- Business Continuity Plan
- Cloud Security Standard
- Configuration Management Standard
- Cryptography Standard
- Data Classification and Handling Standard



- Incident Management Standard
- Mobile Device Management Standard
- Risk Management Standard
- Vendor Management Standard
- Vulnerability Management Standard

Data

The Confluent Cloud system processes and stores only data associated with the Kafka clusters each customer elects to deploy. Confluent may also record, store, and access metadata associated with each customer's Kafka clusters to further improve performance and make feature updates to future releases. Users are required to provide an email address and password to provision their account and operate the web front-end. Customer data is stored in AWS, Azure, or GCP virtual instances.



Internal Control Framework

Confluent has adopted the following control framework to meet its security and availability commitments. This framework includes the following aspects: control environment, risk assessment, control activities, information and communication, and monitoring.

Control Environment

An organization's control environment represents the attitude, awareness, and actions of the board of directors, executive management, and other key stakeholders concerning the importance of controls and the emphasis given to controls in the company's policies, procedures, operations, and organizational structure.

Risk Assessment

Confluent maintains an ongoing risk management process intended to proactively identify vulnerabilities within Confluent systems, and assess new and emerging threats to company operations. Processes to assess and identify risk include risk assessments, vulnerability scans, penetration tests, and an annual review of risk mitigation plans.

Control Activities

Controls have been implemented to address system and data risks. Controls have been designed and implemented in the following areas:

- System inventory
- Physical security
- Perimeter controls
- Network access and logical security
- User access
- Anti-malware protection
- Remote access
- Hardware security
- Vulnerability assessment
- System Monitoring
- Incident management
- Change management
- Backup and disaster recovery

Information and Communication

To help align Confluent's business strategies and goals with operating performance, management is committed to maintaining effective communication both with employees and customers.



External Communications

Contact information for the Customer Success team is available via the company website for customers to make requests, ask questions, and report security incidents or any additional concerns. Customer Success tracks issues to resolution. Customers are specifically notified via email if there are changes to the Terms of Service, Privacy Policy, or Subscriber Agreement.

Internal Communications

Confluent defines job descriptions outlining roles and responsibilities, including those related to designing, developing, implementing, operating, monitoring, and maintaining the Confluent systems. Job descriptions are made available to enable employee awareness of their responsibilities.

The Security Management Plan and Information Security Policy, which are communicated to internal personnel, define the information security roles and responsibilities. The Security Steering Committee and Head of Information Security approve these documents annually. Employees complete annual security awareness training. Confluent publishes written policies and procedures to its employees related to the following areas: acceptable use, access management, asset management, cloud security, configuration management, cryptography standard, data classification and handling, mobile device management standard, incident management, information security, risk management, vendor management, and vulnerability management.

Monitoring

Confluent has developed a suite of controls to monitor the compliance of its control environment. These controls are designed to be complimentary to Confluent's existing suite of controls. Monitoring control activities include, annual vendor assessments, and an annual internal control evaluation.



Complementary User Entity Controls

Confluent's controls were designed under the assumption that certain controls would be implemented by user organizations, the application of which is necessary to meet certain trust services criteria identified in this report. This section highlights those internal control responsibilities Confluent believes should be present at each customer, and has considered in developing its controls reported herein. Confluent customers should evaluate their own control environment to assess if the following controls are implemented and operating effectively. These complementary user entity controls do not represent a comprehensive list of controls that should be employed by Confluent customers, but provide a summary of controls necessary to meet the stated trust services criteria presented in this report. These controls include the following:

- User entities are responsible for provisioning, deprovisioning and reviewing user access (CC6.1, CC6.2, CC6.3).
- User entities are responsible for sending data to Confluent via a secure connection and/or the data should be encrypted (CC6.7).



Complementary Subservice Organization Controls

Confluent contracts with Amazon Web Services (AWS), Google Cloud Platform, and Microsoft Azure to provide data center, infrastructure, software, and managed hosting services. Controls managed by these third-party subservice providers are not included in the scope of this report. Expected subservice provider controls that have an effect on specific criteria are included below.

Criteria	Expected Controls
<p>CC6.1 – The entity implements logical access security software, infrastructure, and architectures over protected information assets to protect them from security events to meet the entity’s objectives.</p>	<p>Access to hosted systems requires users to use a secure method to authenticate.</p> <p>User content is segregated and made viewable only to authorized individuals.</p> <p>Network security mechanisms restrict external access to the production environment.</p>
<p>CC6.2 – Prior to issuing system credentials and granting system access, the entity registers and authorizes new internal and external users whose access is administered by the entity. For those users whose access is administered by the entity, user system credentials are removed when user access is no longer authorized.</p>	<p>New user accounts are approved by appropriate individuals prior to being provisioned.</p> <p>User accounts are removed when access is no longer needed.</p> <p>User accounts are reviewed on a regular basis by appropriate personnel.</p>
<p>CC6.3 – The entity authorizes, modifies, or removes access to data, software, functions, and other protected information assets based on roles, responsibilities, or the system design and changes, giving consideration to the concepts of least privilege and segregation of duties, to meet the entity’s objectives.</p>	<p>Access modifications to hosted systems are approved by appropriate individuals prior to being provisioned.</p> <p>User accounts are removed when access is no longer needed.</p> <p>User accounts are reviewed on a regular basis by appropriate personnel.</p>
<p>CC6.4 – The entity restricts physical access to facilities and protected information assets (for example, data center facilities, back-up media storage, and other sensitive locations) to authorized personnel to meet the entity’s objectives.</p>	<p>Access to physical facilities is restricted to authorized users.</p>
<p>CC6.5 – The entity discontinues logical and physical protections over physical assets only after the ability to read or recover data and software from those assets has been diminished and is no longer required to meet the entity’s objectives.</p>	<p>Production media is securely decommissioned and physically destroyed prior to being removed from the data center.</p>

Criteria	Expected Controls
<p>CC6.6 – The entity implements logical access security measures to protect against threats from sources outside its system boundaries.</p>	<p>Network security mechanisms restrict external access to the production environment.</p> <p>Encrypted communication is required for connections to the production system.</p>
<p>CC6.7 – The entity restricts the transmission, movement, and removal of information to authorized internal and external users and processes, and protects it during transmission, movement, or removal to meet the entity’s objectives.</p>	<p>Access to hosted data is restricted to appropriate users.</p> <p>Hosted data is protected during transmission through encryption and secure protocols.</p>
<p>CC6.8 – The entity implements controls to prevent or detect and act upon the introduction of unauthorized or malicious software to meet the entity’s objectives.</p>	<p>Antivirus or antimalware solutions are installed to detect or prevent unauthorized or malicious software.</p>
<p>CC7.1 – To meet its objectives, the entity uses detection and monitoring procedures to identify (1) changes to configurations that result in the introduction of new vulnerabilities, and (2) susceptibilities to newly discovered vulnerabilities.</p>	<p>System configurations changes are logged and monitored.</p> <p>Vulnerabilities are identified and tracked to resolution.</p>
<p>CC7.2 – The entity monitors system components and the operation of those components for anomalies that are indicative of malicious acts, natural disasters, and errors affecting the entity’s ability to meet its objectives; anomalies are analyzed to determine whether they represent security events.</p>	<p>Security events are monitored and evaluated to determine potential impact per policy.</p>
<p>CC7.3 – The entity evaluates security events to determine whether they could or have resulted in a failure of the entity to meet its objectives (security incidents) and, if so, takes actions to prevent or address such failures.</p>	<p>Operations personnel log, monitor and evaluate to incident events identified by monitoring systems</p>
<p>CC7.4 – The entity responds to identified security incidents by executing a defined incident response program to understand, contain, remediate, and communicate security incidents, as appropriate.</p>	<p>Operations personnel respond, contain and remediate incident events, and update stakeholders, as needed.</p>
<p>CC8.1 – The entity authorizes, designs, develops or acquires, configures, documents, tests, approves, and implements changes to infrastructure, data, software, and procedures to meet its objectives.</p>	<p>System changes are documented, tested, and approved prior to migration into production.</p> <p>Access to make system changes is restricted to appropriate personnel.</p>

Criteria	Expected Controls
<p>A1.1 – The entity maintains, monitors, and evaluates current processing capacity and use of system components (infrastructure, data, and software) to manage capacity demand and to enable the implementation of additional capacity to help meet its objectives.</p>	<p>Operations personnel monitor processing and system capacity.</p>
<p>A1.2 – The entity authorizes, designs, develops or acquires, implements, operates, approves, maintains, and monitors environmental protections, software, data back-up processes, and recovery infrastructure to meet its objectives.</p>	<p>Environmental controls protect the physical devices supporting the production environment.</p>
<p>A1.3 – The entity tests recovery plan procedures supporting system recovery to meet its objectives.</p>	<p>System failover and backup procedures are tested.</p>