Proprietary & Confidential

Confluent Cloud

System Description of the Confluent Cloud System

SOC 3

Relevant to Security, Availability, and Confidentiality

JANUARY 1, 2022 TO DECEMBER 31, 2022

MOSSADAMS

Table of Contents

I. Independent Service Auditor's Report

Confluent, Inc.
899 West Evelyn
Mountain View, CA 94041

To the Management of Confluent:

Scope

We have examined Confluent's accompanying assertion in Section II titled "Confluent's Assertion" (assertion) that the controls within Confluent's Cloud System (system) were effective throughout the period January 1, 2022 to December 31, 2022, to provide reasonable assurance that Confluent's service commitments and system requirements were achieved based on the trust services criteria relevant to Security, Availability, and Confidentiality (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 uses the following subservice organizations:

- Amazon Web Services for managed hosting services of its infrastructure, software, and data
- Microsoft Azure for managed hosting services of its infrastructure, software, and data
- Google Cloud Platform for managed hosting services of its infrastructure, software, and data

The description indicates that complementary subservice organization controls that are suitably designed and operating effectively are necessary, along with controls at Confluent, to achieve Confluent's service commitments and system requirements based on the applicable trust services criteria. The description presents the types of complementary subservice organization controls assumed in the design of Confluent's controls. The description does not disclose the actual controls at the subservice organization. Our examination did not include the services provided by the subservice organizations, and we have not evaluated the suitability of the design or operating effectiveness of such complementary subservice organization controls.

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 also 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 effectiveness 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 (AICPA). 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 the applicable trust services criteria

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

We are required to be independent and to meet our other ethical responsibilities in accordance with relevant ethical requirements relating to the engagement.

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 were 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 the policies or procedures may deteriorate.

Opinion

In our opinion, management's assertion that the controls within Confluent's Cloud System were effective throughout the period January 1, 2022 to December 31, 2022, 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.

Moss Adams LLP

San Francisco, California
March 28, 2023

II. Confluent's Assertion

We are responsible for designing, implementing, operating, and maintaining effective controls within Confluent's Cloud System (system) throughout the period January 1, 2022 to December 31, 2022 to provide reasonable assurance that Confluent's service commitments and system requirements relevant to Security, Availability, and Confidentiality were achieved. Our description of the boundaries of the system is presented in Section III titled "Confluent's Description of the Boundaries of the Confluent Cloud System" and identifies the aspects of the system covered by our assertion.

We have performed an evaluation of the effectiveness of the controls within the system throughout the period January 1, 2022 to December 31, 2022, to provide reasonable assurance that Confluent's service commitments and system requirements were achieved based on the trust services criteria relevant to Security, Availability, and Confidentiality (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. The principal service commitments and system requirements related to the applicable trust services criteria are presented in Section III titled "Confluent's Description of the Boundaries of the Confluent Cloud System".

Confluent uses the following subservice organizations:

- Amazon Web Services for managed hosting services of its infrastructure, software, and data
- Microsoft Azure for managed hosting services of its infrastructure, software, and data
- Google Cloud Platform for managed hosting services of its infrastructure, software, and data

The description indicates that complementary subservice organization controls that are suitably designed and operating effectively are necessary, along with controls at Confluent, to achieve Confluent's service commitments and system requirements based on the applicable trust services criteria. The description presents the types of complementary subservice organization controls assumed in the design of Confluent's controls. The description does not disclose the actual controls at the subservice organization(s).

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, 2022 to December 31, 2022, to provide reasonable assurance that Confluent's service commitments and system requirements were achieved based on the applicable trust services criteria.

III. Confluent's Description of the Boundaries of the Confluent Cloud System

A. System Overview

1. Services Provided

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 a messaging queue to a full-fledged event streaming platform. Confluent delivers the most complete distribution of Kafka with the 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 2,200 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:

- *Confluent Cloud Console*™ – A comprehensive management and monitoring system for Apache Kafka. Control Console provides:
  - The ability to monitor and manage clusters from a rich user interface
  - 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
- *Kafka* – Apache Kafka® is a distributed streaming platform that:
  - Publishes and subscribes to streams of records, similar to a message queue or enterprise messaging system
  - Stores streams of records in a fault-tolerant durable way
  - Processes streams of records as they occur

- *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.

- *ksqlDB* – ksqlDB is the streaming SQL engine for 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. ksqlDB is scalable, elastic, fault-tolerant, and real-time. It supports a wide range of streaming operations, including data filtering, transformations, aggregations, joins, windowing, and sessionization.
- *Schema Registry* – The Confluent Schema Registry provides a serving layer for metadata. Specifically, it provides a RESTful interface for storing and retrieving Avro®, JSON Schema, and Protobuf schemas. It stores a versioned history of schemas based on a specified subject name strategy, provides multiple compatibility settings and allows evolution of schemas according to the configured compatibility settings and expanded support for these schema types. It also provides serializers that plug into Apache Kafka® clients to handle schema storage and retrieval for Kafka messages sent in any of the supported formats.
- *Confluent Health+* - Confluent Health+ allows users to monitor and manage the Confluent Cloud environment, ensuring the health of clusters and minimizing business disruption with intelligent alerts, monitoring, and proactive support based on best practices created by the inventors Apache Kafka®.

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. Confluent Cloud 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.
    - Role-Based Access Control protects Confluent Cloud resources and data by authorizing and restricting access of user and service accounts and by delegating access authorization to the appropriate business units and teams in an organization
    - Audit Events capture event records from auditable event methods for Kafka cluster event categories and organization event categories
    - Metrics API supports a diverse set of querying patterns to support usage and performance analysis over time
    - Public and Private Networking allows access to clusters through secure internet endpoints, Private Link connections, VPC/VNet peering, or AWS Transit Gateway, according to cluster type. All connections to Confluent Cloud are encrypted with TLS and require authentication using API keys.

Confluent Cloud 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 the Confluent Cloud System, and does not include the Confluent Platform.

2. 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.

Infrastructure for Confluent Cloud is unique and delimited. People, policies, procedures, and data are shared across Confluent Cloud's boundaries and described below:

CONFLUENT CLOUD INFRASTRUCTURE: OVERVIEW

Confluent utilizes cloud service providers AWS, Azure, and GCP for its computing platforms, network technologies, and internal global infrastructure. The system components of Confluent Cloud are hosted within the cloud service provider and region selected by customers. Confluent Cloud primarily uses Alpine Linux and Debian operating systems in AWS, Azure, and GCP to support the infrastructure systems supporting Confluent Cloud. See **Error! Reference source not found.** for the Technical Overview Diagram.

To facilitate operation of Confluent Cloud, Confluent uses the following AWS services:

AWS Service	Function
Amazon Application Load Balancer (ALB)	Amazon ALB operates at the request level (layer 7), routing traffic to targets (EC2 instances, containers, IP addresses, and Lambda functions) based on the content of the request.
Amazon CloudFront	Amazon CloudFront is a fast content delivery network (CDN) service that delivers data, videos, applications, and APIs to customers globally with low latency and high transfer speeds.
Amazon CloudWatch	Amazon CloudWatch provides monitoring for AWS cloud resources and applications. Amazon CloudWatch provides visibility into resource utilization, operational performance, and overall demand patterns—including metrics such as utilization, disk reads and writes, and network traffic.
Amazon DynamoDB	Amazon DynamoDB is a fully managed proprietary NoSQL database service that supports key–value and document data structures. DynamoDB uses synchronous replication across multiple data centers for high durability and availability.
Amazon Elastic Block Storage (EBS)	Amazon EBS provides block level storage volumes for use with EC2 instances. EBS volumes that are attached to an instance are exposed as storage volumes that persist independently from the life of the instance.
Amazon ElastiCache	Amazon ElastiCache allows users to seamlessly set up, run, and scale popular open-source compatible in-memory data stores in the cloud.
Amazon Elastic Container Registry (ECR)	Amazon ECR is a container registry used to store, manage, share, and deploy container images and artifacts.
Amazon Elastic Compute Cloud (EC2)	Amazon EC2 provides a virtual computing environment that uses web service interfaces to perform the following functions: • Launch instances of operating systems. • Create Amazon Machine Images (AMI) containing applications, libraries, data, and associated configuration settings. • Configure security and network access on the EC2 instances.

AWS Service	Function
Amazon Elastic Kubernetes Service (Amazon EKS)	Amazon EKS provides the flexibility to start, run, and scale Kubernetes applications in the AWS cloud. Amazon EKS provides highly available and secure clusters and automates key tasks such as patching, node provisioning, and updates.
Amazon Elastic Load Balancing (ELB)	Amazon ELB automatically distributes incoming application traffic across multiple targets, such as Amazon EC2 instances, containers, IP addresses, Lambda functions, and virtual appliances.
Amazon Elasticsearch Service	Amazon Elasticsearch Service is a fully managed service used to deploy, secure, and run Elasticsearch.
AWS Global Accelerator	AWS Global Accelerator is a networking service that improves the performance of users' traffic using Amazon Web Services' global network infrastructure.
Amazon Network Load Balancer (NLB)	Amazon NLB is optimized to handle sudden and volatile traffic patterns while using a single static IP address per Availability Zone.
Amazon OpenSearch	Amazon OpenSearch provides a highly scalable system for providing fast access and response to large volumes of data with an integrated visualization tool that makes it easy for users to explore their data.
AWS PrivateLink (feature of EC2)	AWS PrivateLink provides private connectivity between VPCs, AWS services, and the on-premises networks, without exposing traffic to the public internet.
Amazon Relational Database Service (RDS)	Amazon RDS is a web service used to operate relational databases in the AWS cloud.
Amazon Route 53	Amazon Route 53 is a Domain Name System (DNS) web service.
Amazon Simple Notification Service (SNS)	Amazon SNS is a notification service that provides infrastructure for the mass delivery of messages.
Amazon Simple Queue Service (SQS)	Amazon SQS is a message queue service used to exchange messages through a polling model, and can be used to decouple sending and receiving components.
Amazon Simple Storage Service (S3)	Amazon S3 is virtual storage used in conjunction with Amazon EC2 to store object data. Amazon S3 is also used to automatically replicate data across AWS regions.

AWS Service	Function
Amazon Virtual Private Cloud (VPC)	Amazon VPC is used to provision logically isolated virtual networks in the AWS cloud. Amazon VPC is used to manage the virtual networking environment, including selection of IP address ranges, creation of subnets, and configuration of route tables and network gateways.
AWS CloudTrail	AWS CloudTrail provides event history of AWS account activity, including actions taken through the AWS Management Console, AWS SDKs, command line tools, and other AWS services.
AWS GuardDuty	Amazon GuardDuty is a threat detection service that continuously monitors for malicious activity and unauthorized behaviors to protect AWS accounts, workloads, and data stored in Amazon S3.
AWS Identity and Access Management (IAM)	AWS IAM enables users to manage access to AWS services and resources securely. Using IAM, users can create and manage AWS users and groups, and use permissions to allow and deny their access to AWS resources.
AWS Key Management Service (KMS)	AWS KMS makes it easy to create and manage cryptographic keys and control their use across a wide range of AWS services.
AWS Shield Advanced	AWS Shield is a managed DDoS protection service that safeguards applications running on AWS.
AWS Web Application Firewall (WAF)	AWS WAF is a web application firewall that helps protect web applications or APIs against common web exploits and bots that may affect availability, compromise security, or consume excessive resources.

Refer to **Error! Reference source not found.** for the AWS Data Flow Diagram below for an overview of AWS services within the Confluent environment. Similar data flows exist for Azure and GCP.

To facilitate operation of Confluent Cloud, Confluent uses the following Microsoft Azure services:

Azure Service	Function
Azure Active Directory (AD)	Azure AD enterprise identity service provides single sign-on and multi-factor authentication to help protect users from cybersecurity attacks.
Azure Application Gateway	Azure Application Gateway is a web traffic load balancer that enables traffic management on web applications.
Azure Container Registry	Azure Container Registry provides a managed, private Docker registry service for building, storing, and managing container images and artifacts for all types of container deployments

Azure Service	Function
Azure Database for PostgreSQL	Azure Database for PostgreSQL provides a relational database service based on the open-source Postgres database engine.
Azure DNS	Azure DNS provides the scale and redundancy to give users ultra-high availability for over domains.
Azure Kubernetes Service (AKS)	Azure AKS is a container orchestration service used to deploy, scale, and manage Docker containers and container-based applications in a cluster environment.
Azure Load Balancer	Azure Load Balancer delivers high availability and network performance to applications.
Azure Monitor	Azure Monitor is a monitoring solution that provides visibility into application, infrastructure and network performance to identify problems in real time.
Azure Blob Storage	Azure Blob Storage is an object storage solution for the cloud.
Azure Container Instances	Azure Container Instances run Docker containers on-demand in a managed, serverless Azure environment.
Azure Content Delivery Network (CDN)	Azure CDN is a global CDN solution for delivering high-bandwidth content.
Azure Database for MySQL	Azure Database for MySQL is a relational database service powered by the MySQL community edition.
Azure Event Grid	Azure Event Grid allows you to easily build applications with event-based architectures
Azure SQL Database	Azure SQL Database is a fully managed platform as a service (PaaS) database engine that handles most of the database management functions such as upgrading, patching, backups, and monitoring without user involvement.
Azure Traffic Manager	Azure Traffic Manager is a DNS-based traffic load balancer.
Azure Security Center	Azure Security Center provides unified security management and advanced threat protection across hybrid cloud workloads.
Azure Virtual Machines	Azure Virtual Machines allow users to provision Windows or Linux in seconds, operate seamlessly in hybrid environments, and protect data and monitor cloud health.

Azure Service	Function
Azure Virtual Network	Azure Virtual Network provides an isolated and highly secure environment to run virtual machines and applications.

To facilitate operation of Confluent Cloud, Confluent uses the following GCP services:

GCP Service	Function
Google App Engine	Google App Engine is a cloud computing platform as a service for developing and hosting web applications in Google-managed data centers.
Google BigQuery	Google BigQuery is a serverless, highly scalable, and cost-effective multi-cloud data warehouse designed for business agility.
Google Cloud Audit Logs	Google Cloud Audit Logs provide audit logs for each Cloud project, folder, and organization.
Google Cloud CDN	Google Cloud CDN enables customers to deliver content over Google's high-performance distributed edge caching infrastructure.
Google Cloud DNS	Google Cloud DNS is a reliable, resilient, low-latency DNS serving from Google's worldwide network.
Google Cloud Identity and Access Management (IAM)	Google IAM is a fine-grained access control and visibility for centrally managing cloud resources.
Google Cloud Load Balancing	Google Cloud Load Balancing is a high performance, scalable load balancing on Google Cloud Platform.
Google Cloud Operations Suite	Google Cloud Operations Suite provides metrics for migration components
Google Cloud SQL/Cloud Spanner	Google Cloud SQL/Cloud Spanner is a fully managed relational database service for MySQL, PostgreSQL, and SQL Server
Google Cloud Storage	Google Cloud Storage provides object storage in the cloud.
Google Container Registry	Google Container Register is used to store, manage, and secure Docker container images.
Google Compute Engine	Google Compute Engine lets users create and run virtual machines on Google infrastructure.

GCP Service	Function
Google Event Threat Detection	Google Event Threat Detection monitors organization's Cloud Logging stream and detects threats in near-real time.
Google Kubernetes Engine (GKE)	GKE provides a managed environment for deploying, managing, and scaling containerized applications using Google infrastructure.
Google Pub/Sub	Google Pub/Sub is an asynchronous and scalable messaging service that decouples services producing messages from services processing those messages to reduce latencies.
Google Security Command Center	Google Security Command Center is a Security and risk management platform for Google Cloud.
Google Virtual Private Cloud	Google Virtual Private Cloud is a managed networking functionality for Google Cloud resources.

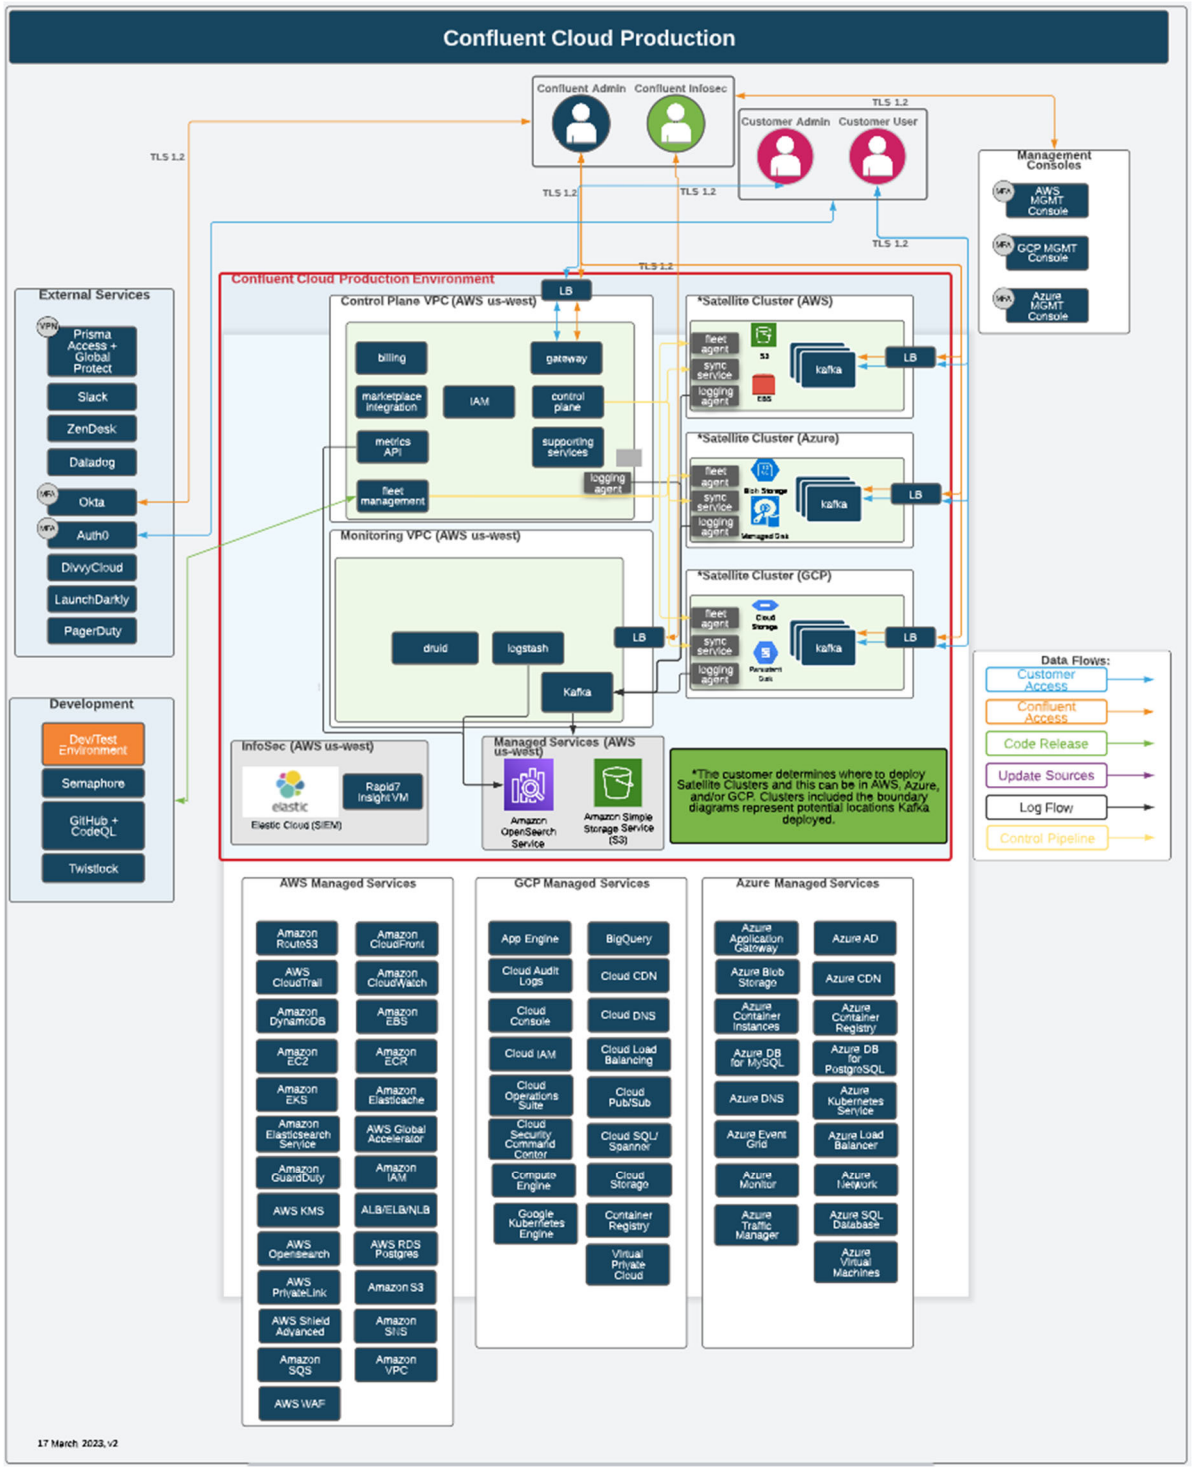

FIGURE 1: TECHNICAL OVERVIEW DIAGRAM

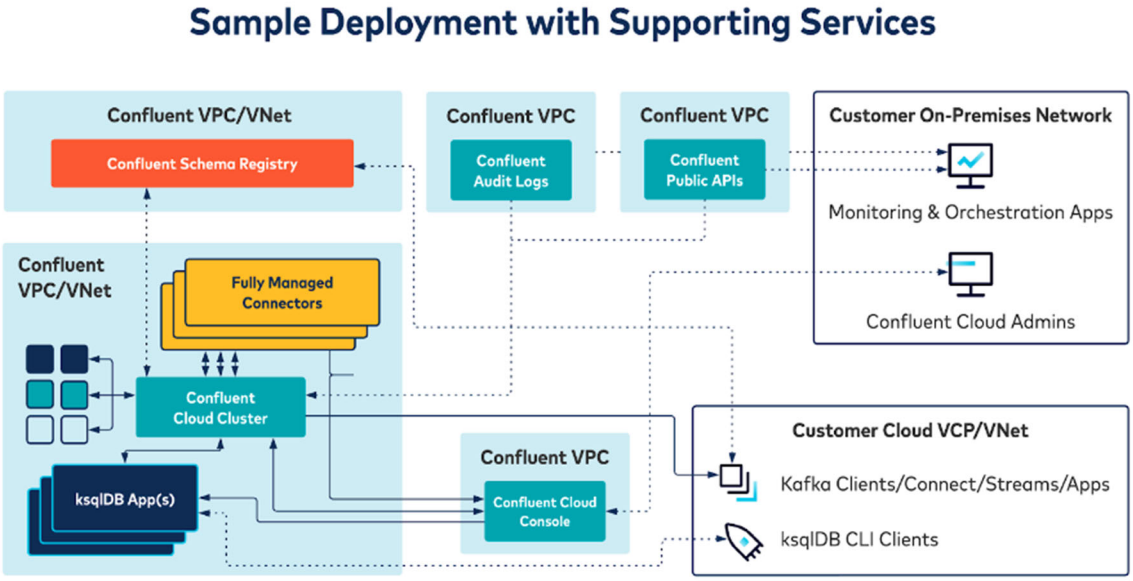

FIGURE 2: AWS DATA FLOW DIAGRAM

3. Software

Confluent has various software programs and tools used to support Confluent Cloud, with key programs listed below:

Vendor Software	Function
Auth0	The Auth0 Customer Identity Management service is a third-party solution to add authentication and authorization services to web applications.
Chef	Chef is a configuration management tool used to streamline the task of configuring and maintaining servers.
Datadog	Datadog is a monitoring service for cloud-scale applications, providing monitoring of servers, databases, tools, and services, through a SaaS-based data analytics platform.
GitHub	GitHub provides hosting for software development version control using Git. GitHub provides access control and several other collaboration features such as bug tracking, feature requests, task management, and wikis for all projects.
HashiCorp Vault	HashiCorp Vault is an encryption service with centralized key management.
Jenkins	Jenkins is an open-source automation server that enables developers to build, test, and deploy their software.

Vendor Software	Function
Jira	Jira is an issue tracking and project management tool.
Kubernetes	An open-source system for automating deployment, scaling, and management of containerized applications.
Okta	Okta is an identity and access management platform providing cloud software that helps companies manage and secure user authentication into modern applications.
PagerDuty	PagerDuty is a cloud-based alerting and on-call monitoring software.
Salesforce	Salesforce is a customer relationship management tool and integrated CRM platform.
Terraform	A cloud agnostic tool that allows a single configuration to be used to manage multiple providers, and to even handle cross-cloud dependencies.

4. People

Confluent is organized in the following functional areas:

- *Business Systems* – Responsible for helping Confluent operate as a data-driven, strategic, and efficient company. This group also includes the internal IT group responsible for IT support at a corporate level.
- *Engineering* – Responsible for the engineering functions across Confluent products. This group is also responsible for maintaining and developing infrastructure components supporting the Confluent suite of products.
- *Finance* – Responsible for the oversight of financial processes for Confluent including fees, internal finance, and budgeting.
- *Global Technical Support* – Responsible for providing customer service, support, and training to Confluent customers.
- *Growth and Marketing* – Responsible for the promotion and demand creation for Confluent products and services.
- *Information Technology (IT)* – Responsible for managing and supporting IT systems, onboarding and offboarding processes and asset management functions.
- *Legal* – Responsible for providing legal advice to foster sound decision-making for Confluent's operations.
- *People Operations and Recruiting* – Responsible for ensuring that Confluent attracts, hires, develops, and retains the best talent to fuel business growth.
- *Product* – Responsible for designing the systems and planning feature updates, enhancements, and bug fixes based on customer and internal requests.

- *Security Steering Committee* – Responsible for the oversight of internal control and includes members independent from control operators.
- *Trust, Security, and Reliability* – Responsible for security and compliance efforts, including production, client-facing, corporate information assets, and data.

5. Data

Confluent Cloud 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. Further, customers are required to provide either an email address and password to provision their account and operate the web front-end or utilize a single sign-on integration. Customer data is stored in AWS, Azure, or GCP block and cloud object storage instances.

6. Processes and Procedures

Confluent has developed and communicated to its personnel procedures to protect service data and the company's assets. Teams are expected to adhere to Confluent policies and procedures that define data is protected through rules and requirements. These are located on the company's intranet and shared drive and can be accessed by any Confluent employee.

The policies and standards used to safeguard Confluent Cloud include:

- Acceptable Use Policy
- Access Management Standard
- Asset Management Standard
- Business Continuity and Disaster Recovery Plan
- Cloud Permissions Procedure
- Cloud Security Standard
- Configuration Management Standard
- Cryptography Standard
- Data Classification and Handling Standard
- Information Security Policy
- Issue Management Standard
- Logging and Monitoring Standard
- Mobile Device Management Standard
- Risk Management Standard
- Security Incident Response Standard
- Secure System Development Standard
- Vendor Management Standard
- Vulnerability Management Standard

B. Principal Service Commitments and System Requirements

Confluent communicates operational requirements to support the achievement of security, availability, and confidentiality 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)

C. Complementary Subservice Organization Controls

Confluent's controls related to the Confluent Cloud System cover only a portion of overall internal control for each user entity of Confluent It is not feasible for the criteria related to the Confluent Cloud System to be achieved solely by Confluent Therefore, each user entity's internal controls must be evaluated in conjunction with Confluent's controls, taking into account the types of controls expected to be implemented by the subservice organization as described below.

Complementary Subservice Organization Controls	
1	Access to hosted systems requires users to use a secure method to authenticate.
2	User content is segregated and made viewable only to authorized individuals.
3	Network security mechanisms restrict external access to the production environment.
4	New user accounts are approved by appropriate individuals prior to being provisioned.
5	User accounts are removed when access is no longer needed.
6	User accounts are reviewed on a regular basis by appropriate personnel.
7	Access modifications to hosted systems are approved by appropriate individuals prior to being provisioned.
8	Access to physical facilities is restricted to authorized users.
9	Production media is securely decommissioned and physically destroyed prior to being removed from the data center.
10	Encrypted communication is required for connections to the production system.
11	Access to hosted data is restricted to appropriate users.
12	Hosted data is protected during transmission through encryption and secure protocols.
13	Antivirus or antimalware solutions are installed to detect or prevent unauthorized or malicious software.
14	System configuration changes are logged and monitored.
15	Vulnerabilities are identified and tracked to resolution.

Complementary Subservice Organization Controls	
16	Security events are monitored and evaluated to determine their potential impact.
17	Operations personnel log, monitor, and evaluate incident events identified by monitoring systems.
18	Operations personnel respond to, contain, and remediate incident events, and update stakeholders, as needed.
19	System changes are documented, tested, and approved prior to migration to production.
20	Access to make system changes is restricted to appropriate personnel.
21	Operations personnel monitor processing and system capacity.
22	Environmental protections, software, data back-up processes, and recovery infrastructure are implemented.
23	System failover and backup procedures are tested.

MOSSADAMS