CONFLUENT

Considerations when starting with Confluent Cloud

Welcome to Confluent

As you journey from day 1 to production, we know you're going to have many questions along the way. Using our experience of helping other customers be successful, this session is intended to provide suggestions of questions to consider as you start out with Confluent Cloud

What do I need to think about before starting with Confluent Kafka?

Roles & Responsibilities

Who owns what? Who needs to get out of bed when an alarm goes off?!

Infrastructure

What business rules or restrictions are we working to and how can we meet those with Confluent Kafka?

Enablement

How do I enable myself & my teams? Where do I go for resources?

Operations

What are my priorities and SLAs? How can we meet them with Confluent Cloud and best practices?

Use Cases

Has the business identified priorities and SLAs for aspects like retention, uptime, etc?

Enablement

What do we offer?

Hands On

- Demos and tutorials at developer.io
- Training

Interactive

- Support
- Forums
- Our professional services team

If you prefer self serve enablement:

- Documentation
- Resources such as
  - White Papers
  - Blog
  - Podcasts
- Support KB

Identify Roles and Responsibilities

At the outset we need to identify both who is responsible for building or managing components *and* who has responsibility for providing answers. We recommend you consider the following questions to help determine responsibilities.

Operations

Who will own the responsibility of deploying and monitoring Confluent Kafka environments? Will this be centrally managed if there are multiple teams?

Business Factors

Who can provide answers to questions about data retention, uptime SLAs, throughput & latency objectives, security needs and so forth?

Infrastructure

Who can provide answers to questions about security? How can your external system admins be contacted? What are your networking setup limitations and requirements?

Operations

“A company can seize extraordinary opportunities only if it’s very good at the ordinary operations”

In order to design, plan, and implement your project, there are important questions to which you need answers at the outset:

- **Outages**: Should we be able to cope with an outage at any part of the flow? If so, for how long?
- **Monitoring**: What are our thresholds and who needs to be notified?
- **Support**: What are the SLAs for the various support levels and how do we adjust?
- **Process & Procedures**: Do we need runbooks and if so, who will produce and maintain them?
- **Management:** How will new use-cases / teams be onboarded? How will new teams be educated with best practices, procedures, etc?
- **Automation**: Does our Kafka deployment need to be automated? Does the deployment of clusters or other components need to be automated (e.g. Terraform)?

Infrastructure

When thinking about your infrastructure, we recommend considering the following questions:

- How do you plan to **connect external systems** to Confluent Kafka? Have you checked out **Kafka Connect**?
- Do you have organizational restrictions regarding:
  - Security
  - Connectivity
  - Networking
- Are you aware of **Confluent Cloud's Networking** options and how these can map to your business needs?

Security

What **security** do you need? Ensure you have answers to the following questions before you begin:

- **Encryption** (In flight & at rest)
- **Authentication**
- **Authorization**

Define your use-cases and prove the functionality

Functionality

- **Retention**
- **Throughput / Latency**
- **Durability / Availability**
- **Connectivity to external systems**

Data

- **Performance**: How much data to expect, how long must it be kept, what are the measures of success?
- **Scaling:** Is there a requirement to dynamically scale up / down in relation to demand?
- **Monitoring:** Capture expectations for monitoring & alerting (i.e. Thresholds, tools etc)

SLAs

- Identify business SLAs and map them to Confluent offerings (cluster types, configurations, etc)


CONFLUENT