

# The Past, Present, and Future of Cloud Native API Gateways


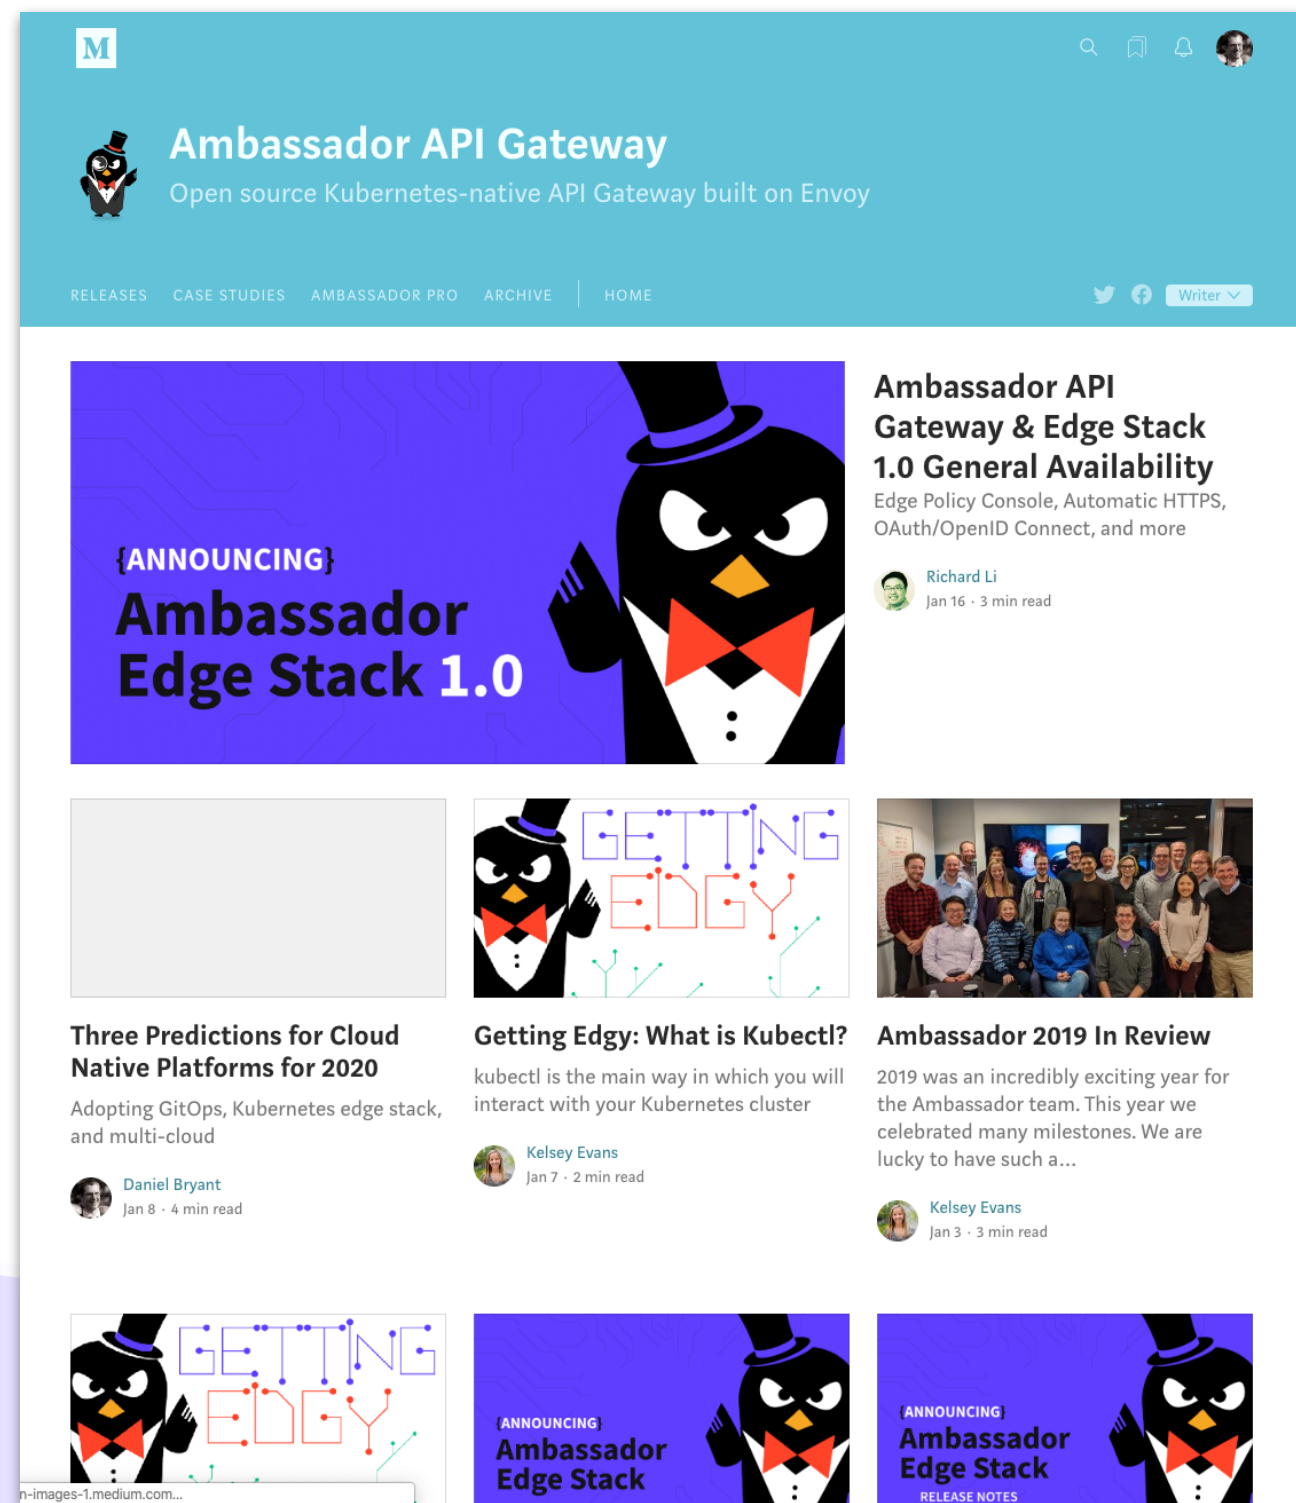
 DATAWIRE

Daniel Bryant

# tl;dr

- Edge gateways have undergone a series of evolutions, **driven by architecture**
- Adopting microservices/Kubernetes changes architecture **and workflow**
- Chose your cloud API gateway solution **intentionally**

# @danielbryantuk






**Edge: The boundary between your data center and your user(s)**

**Thesis: The evolution of the edge has been driven by application architecture**




**~1995**

# Application Architecture in the '90s



# Hardware Load Balancer

<b>User</b>	Systems administrators
<b>Purpose</b>	High availability / scalability
<b>Key Features</b>	Load balancing (round robin, sticky sessions) Health checks

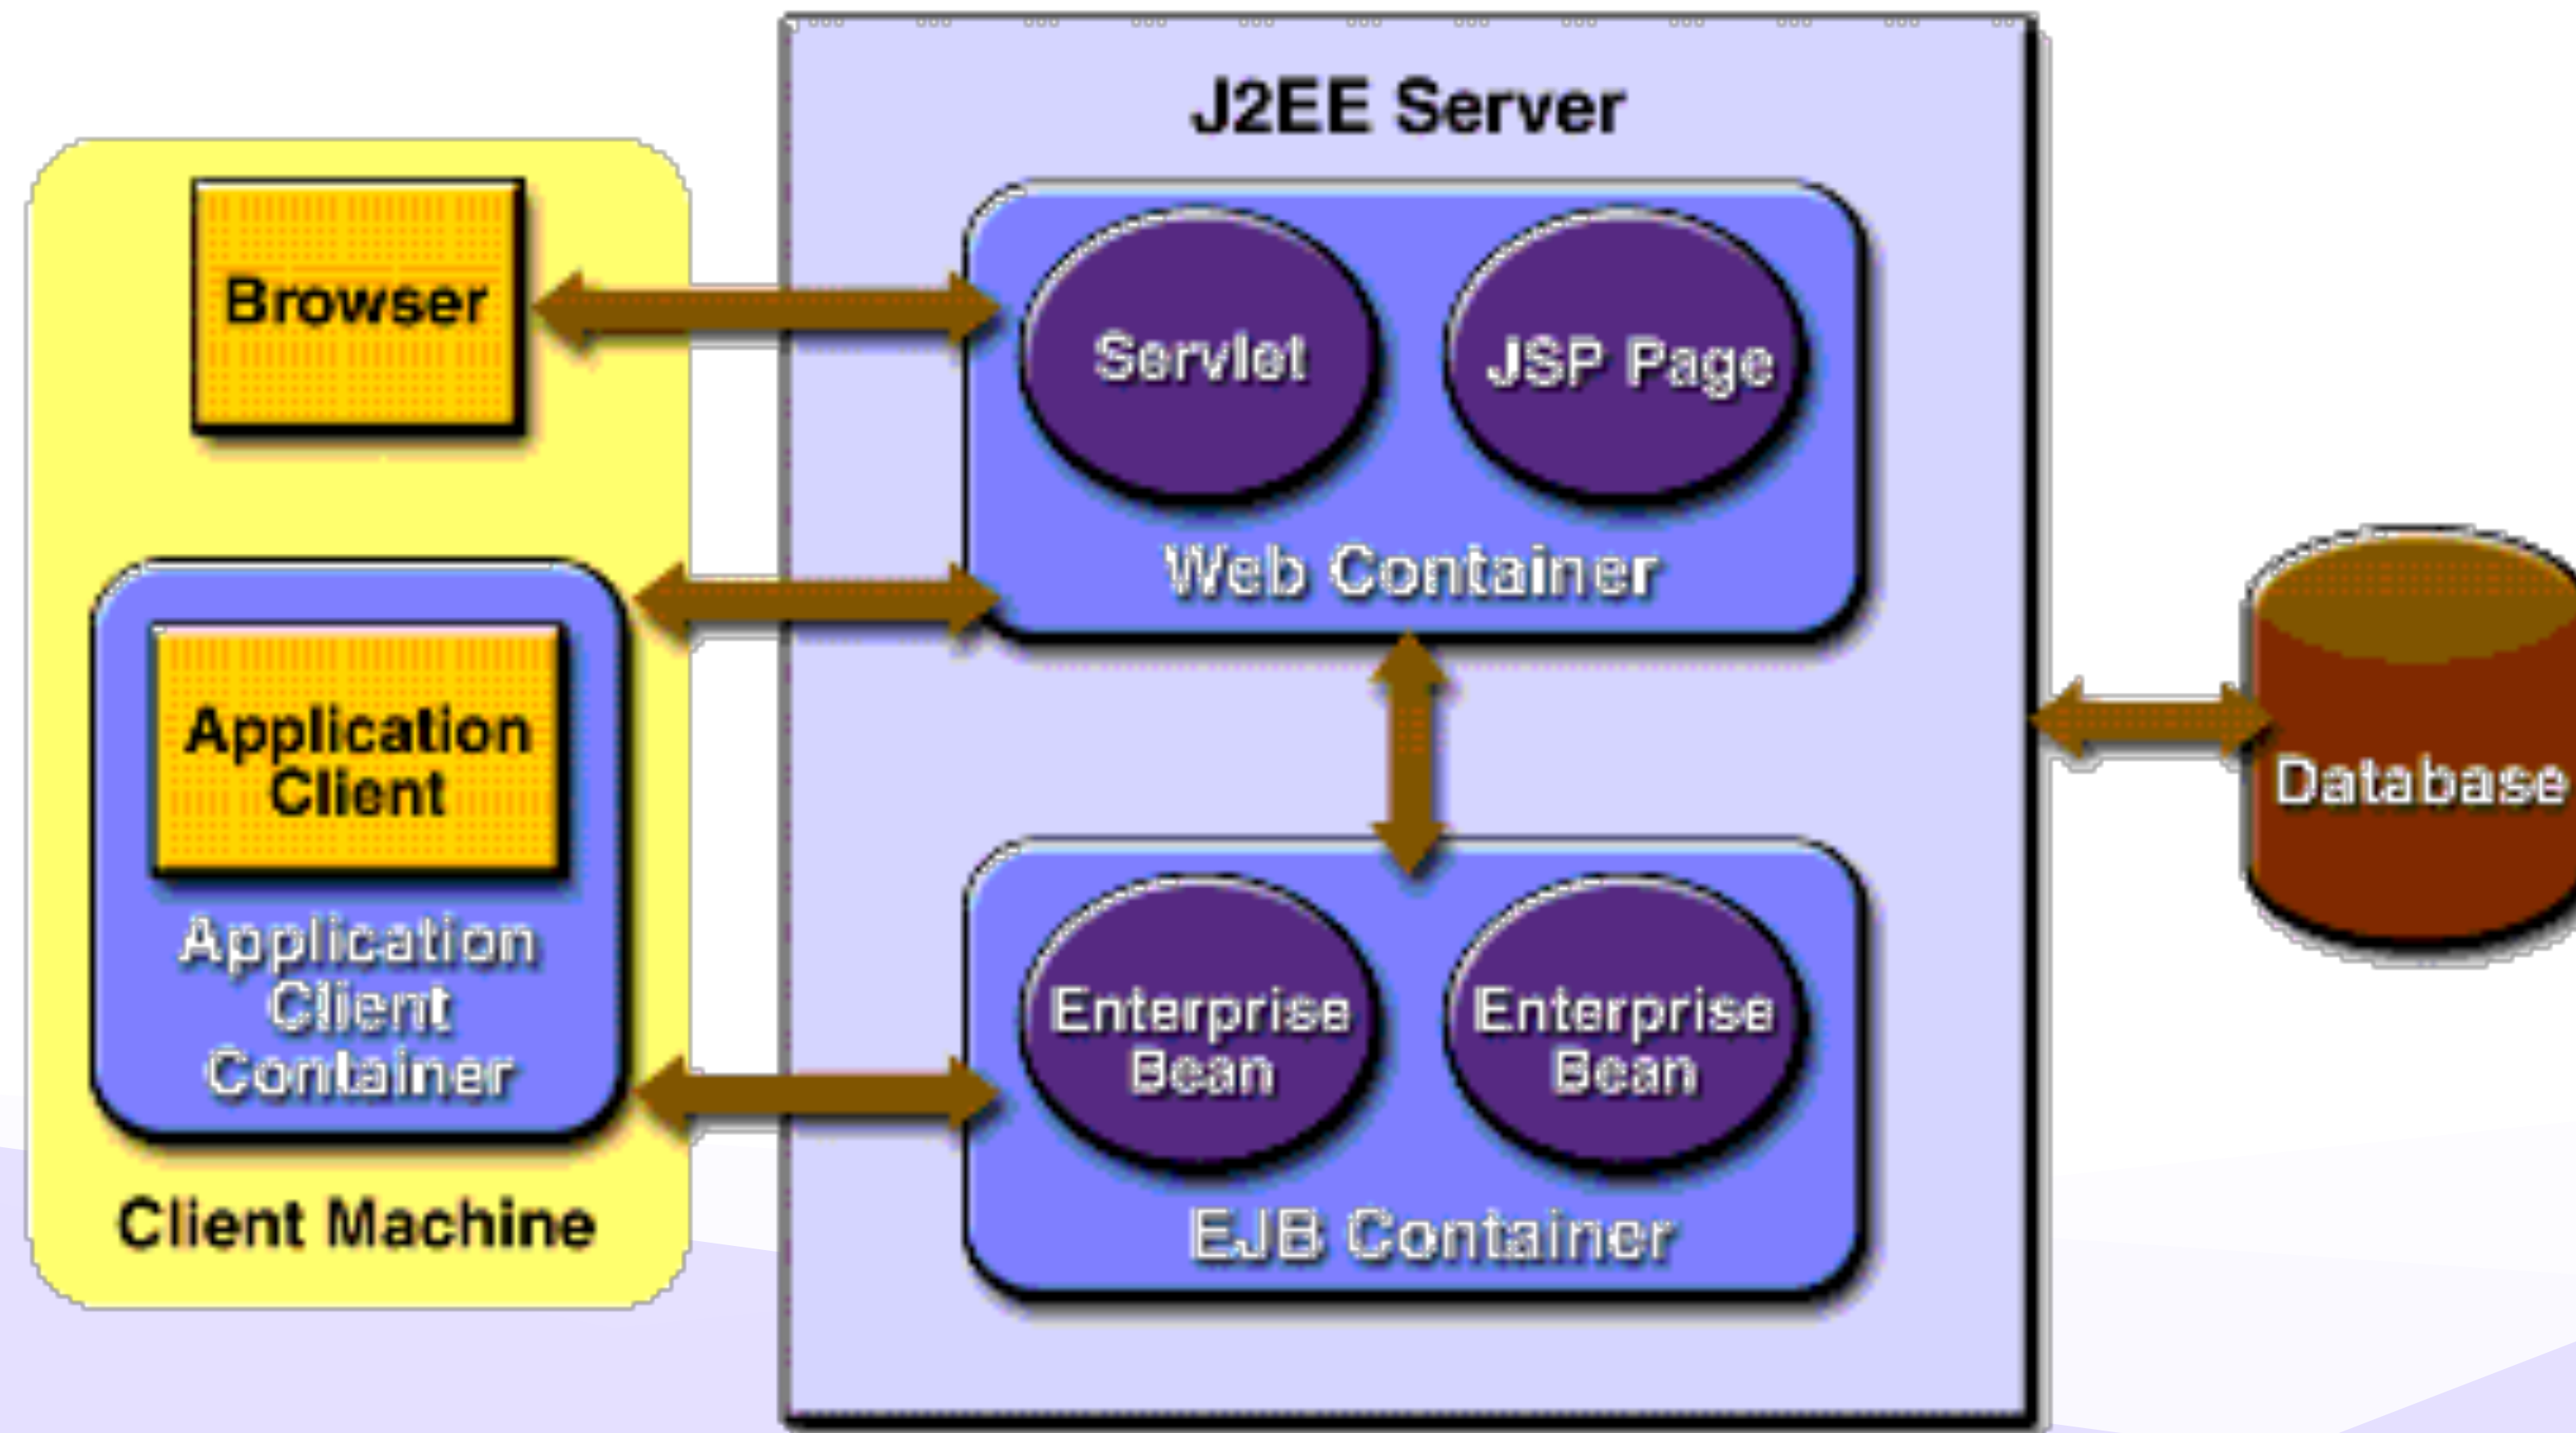







**~2000**

# Similar application architecture





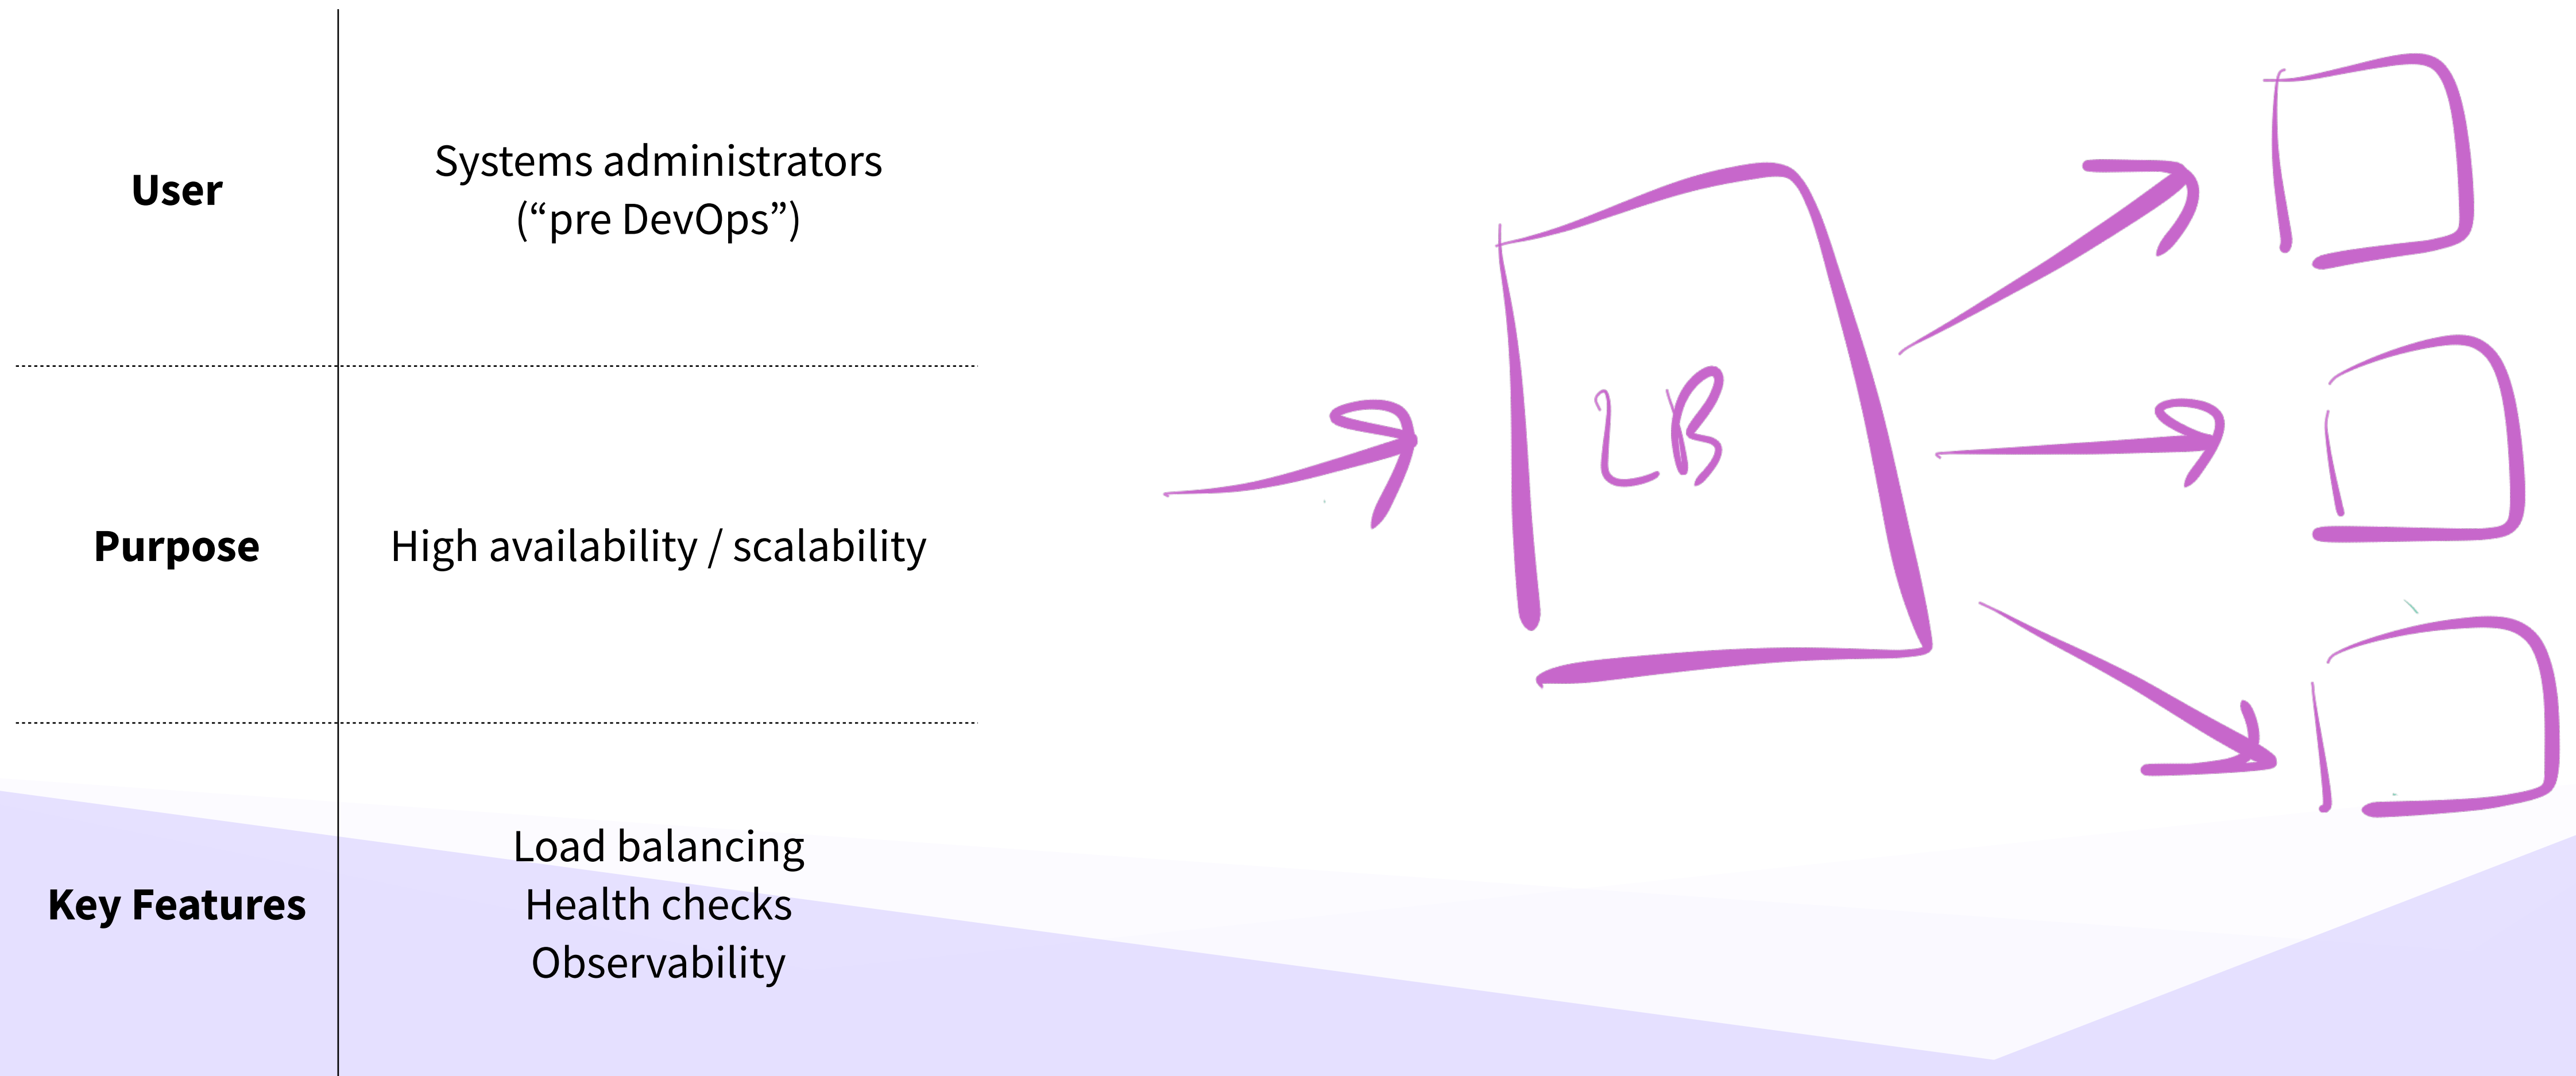
**HAPROXY**

2001

**NGINX**

2002

# Software Load Balancer

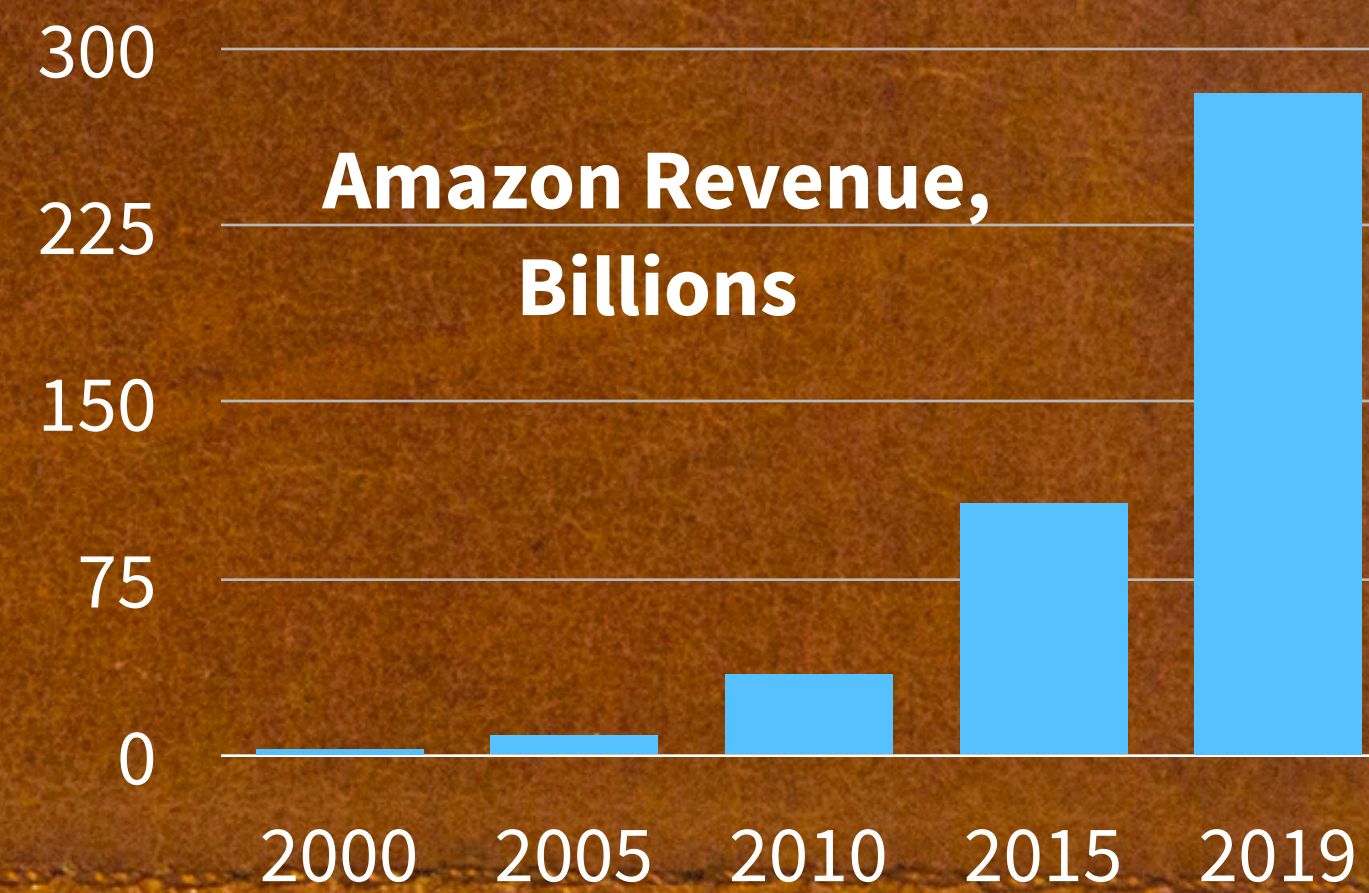




~2005



# Ecommerce




# AJAX

Asynchronous Javascript And XML




# The Application Delivery Controller (ADC)



# Application Delivery Controllers

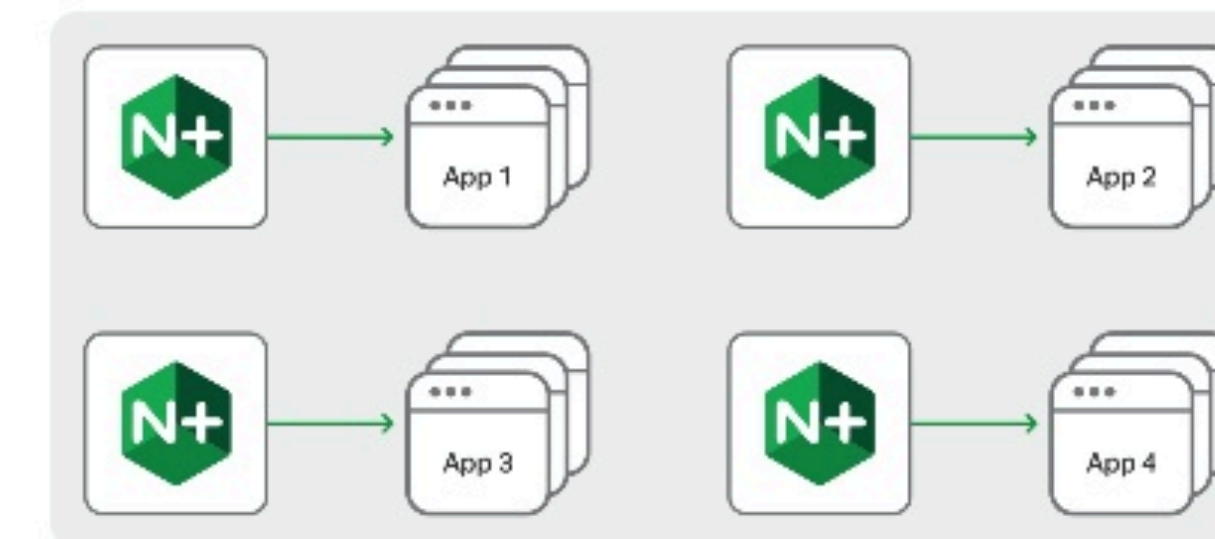
<b>User</b>	Systems administrators
<b>Purpose</b>	High availability and application acceleration
<b>Key Features</b>	SSL offload, caching, compression + load balancing





### 3. Micro Load Balancers/Gateways

Legacy Hardware ADC replace to a application centric architecture



- Load balancer per application
- Load balancer per customer for SaaS providers
- Configuration stored along with application in GitHub
- Fully portable

13





**~2010**

# The proliferation of APIs



2005: API launched

stripe

2008

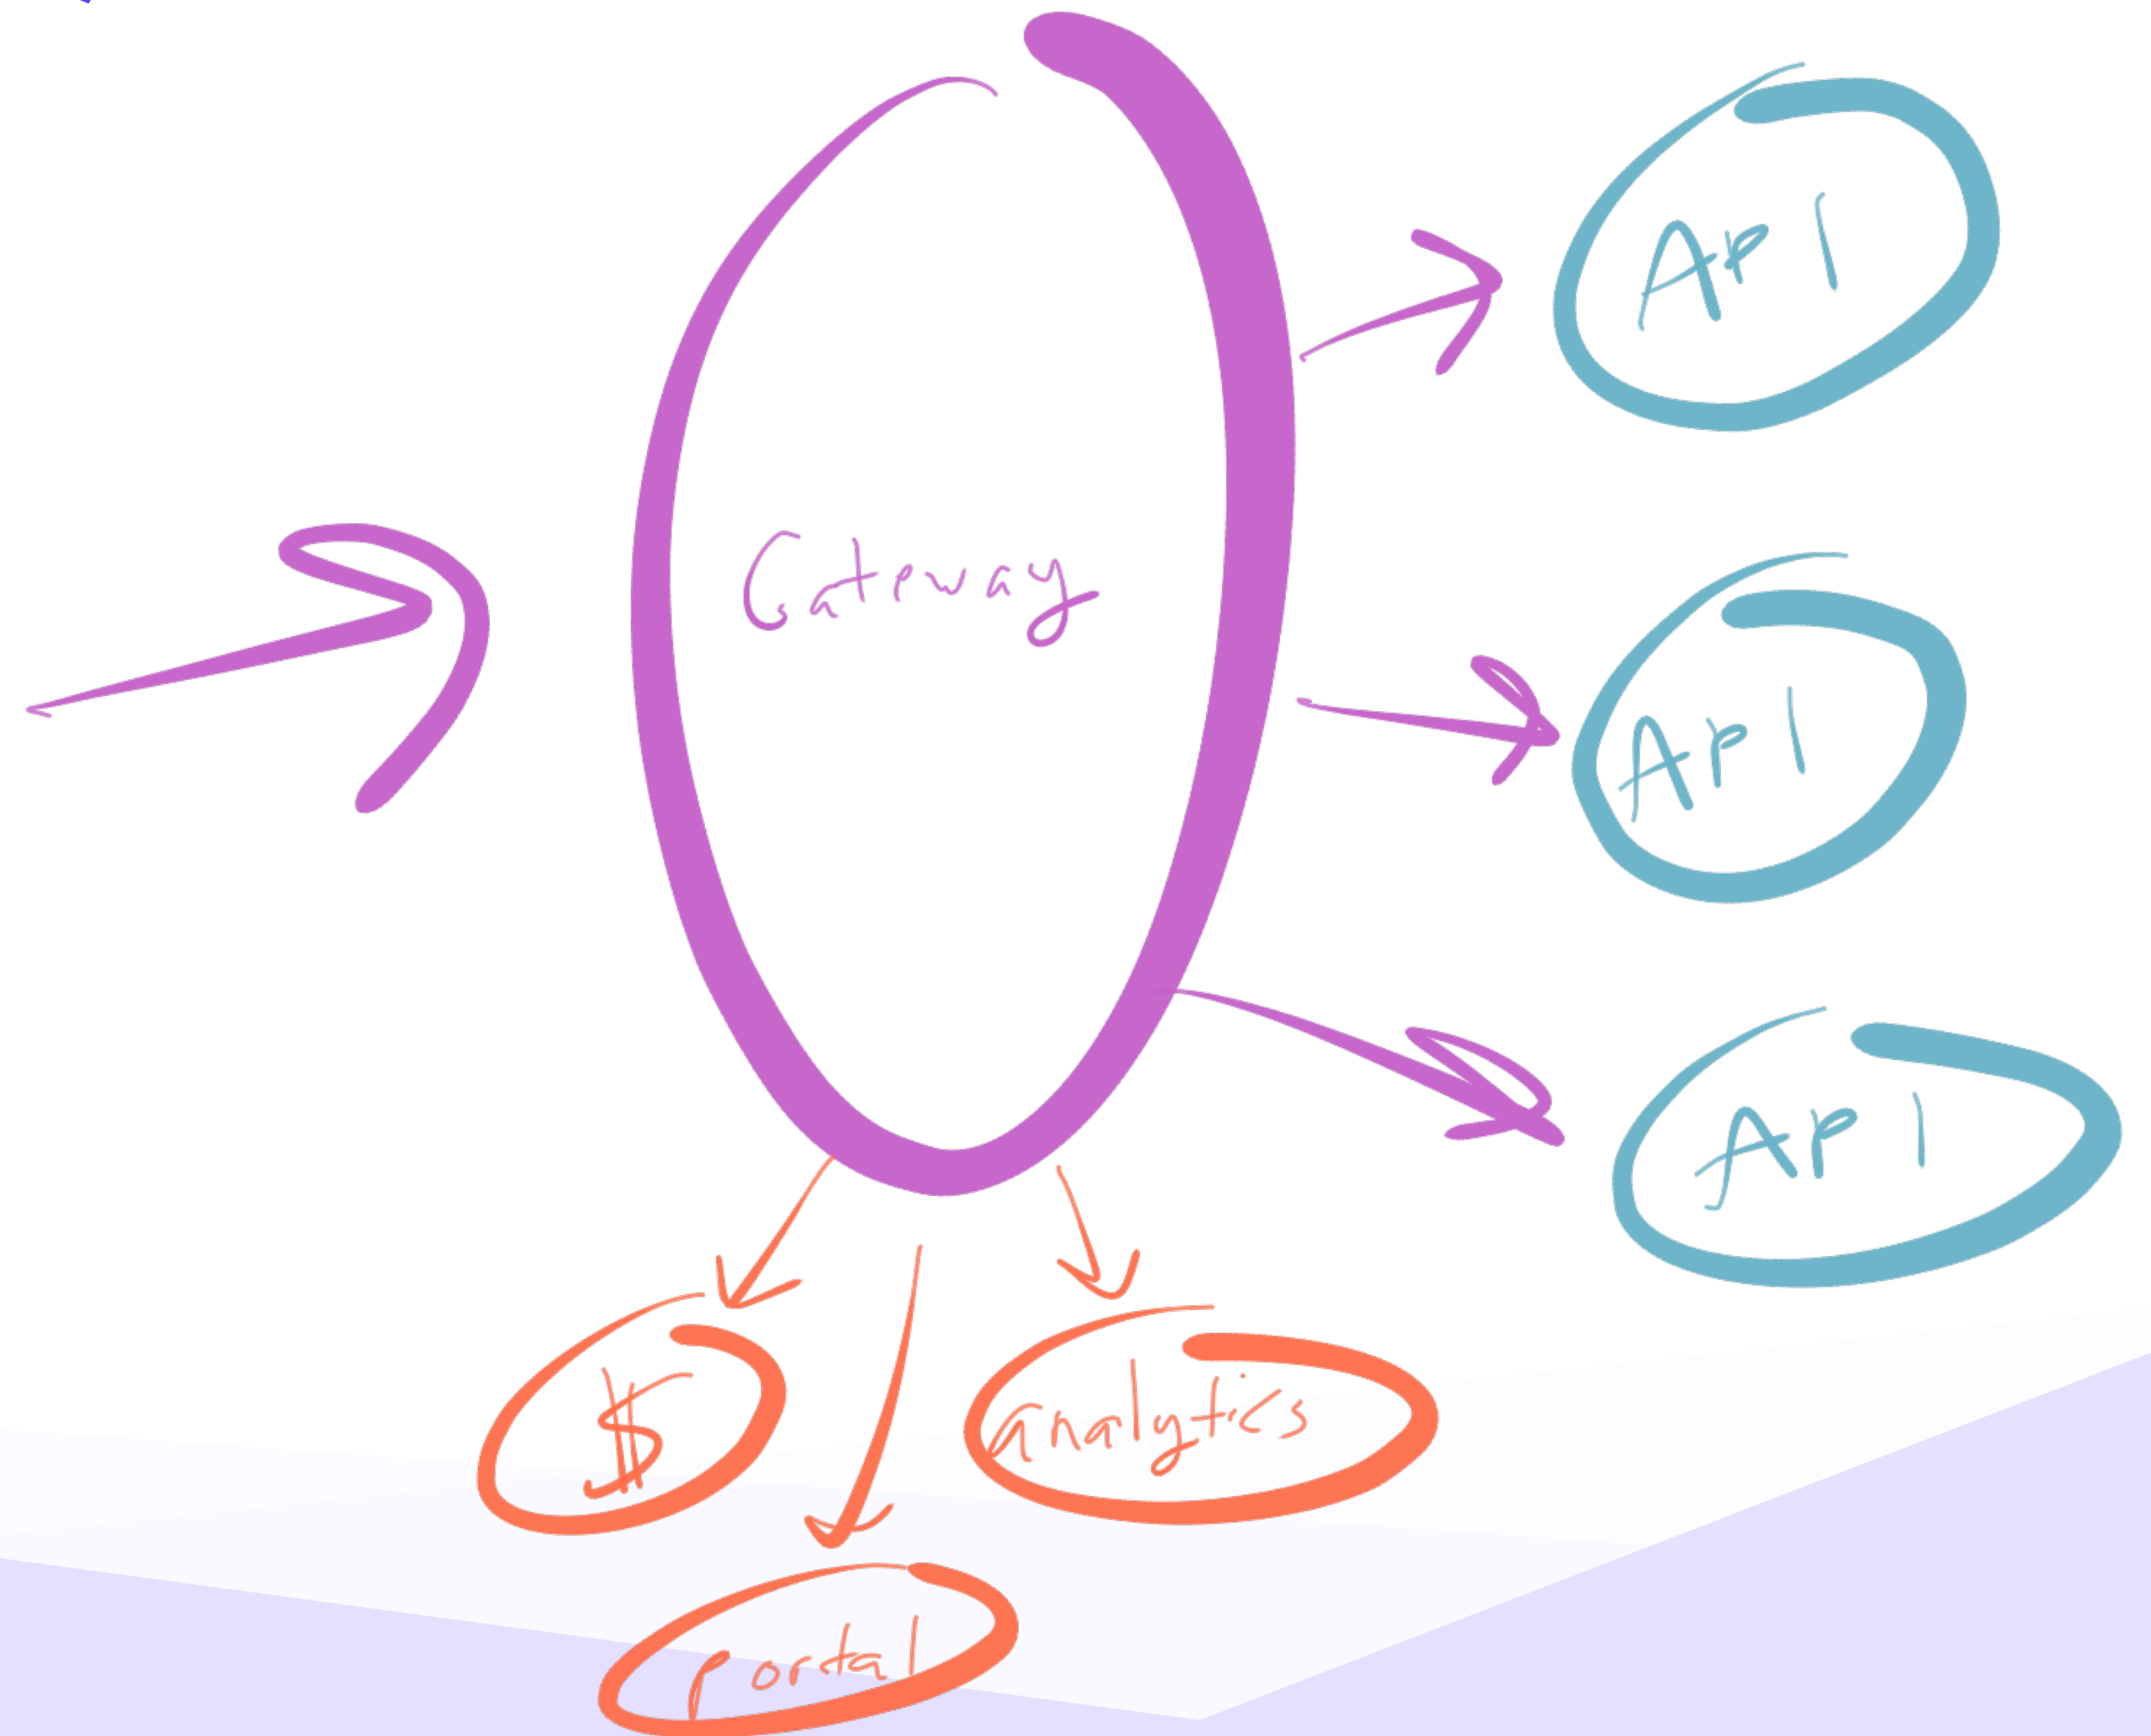
 SendGrid

 twilio

2009


# API Gateway (1st Gen)

<b>User</b>	Systems administrators & API developers
<b>Purpose</b>	Expose business APIs to broader ecosystem (“API management”)
<b>Key Features</b>	L7 routing (e.g., throttling), Publishing, Dev Portal, Analytics, Monetization





**~2015**



# Twitter is over capacity.

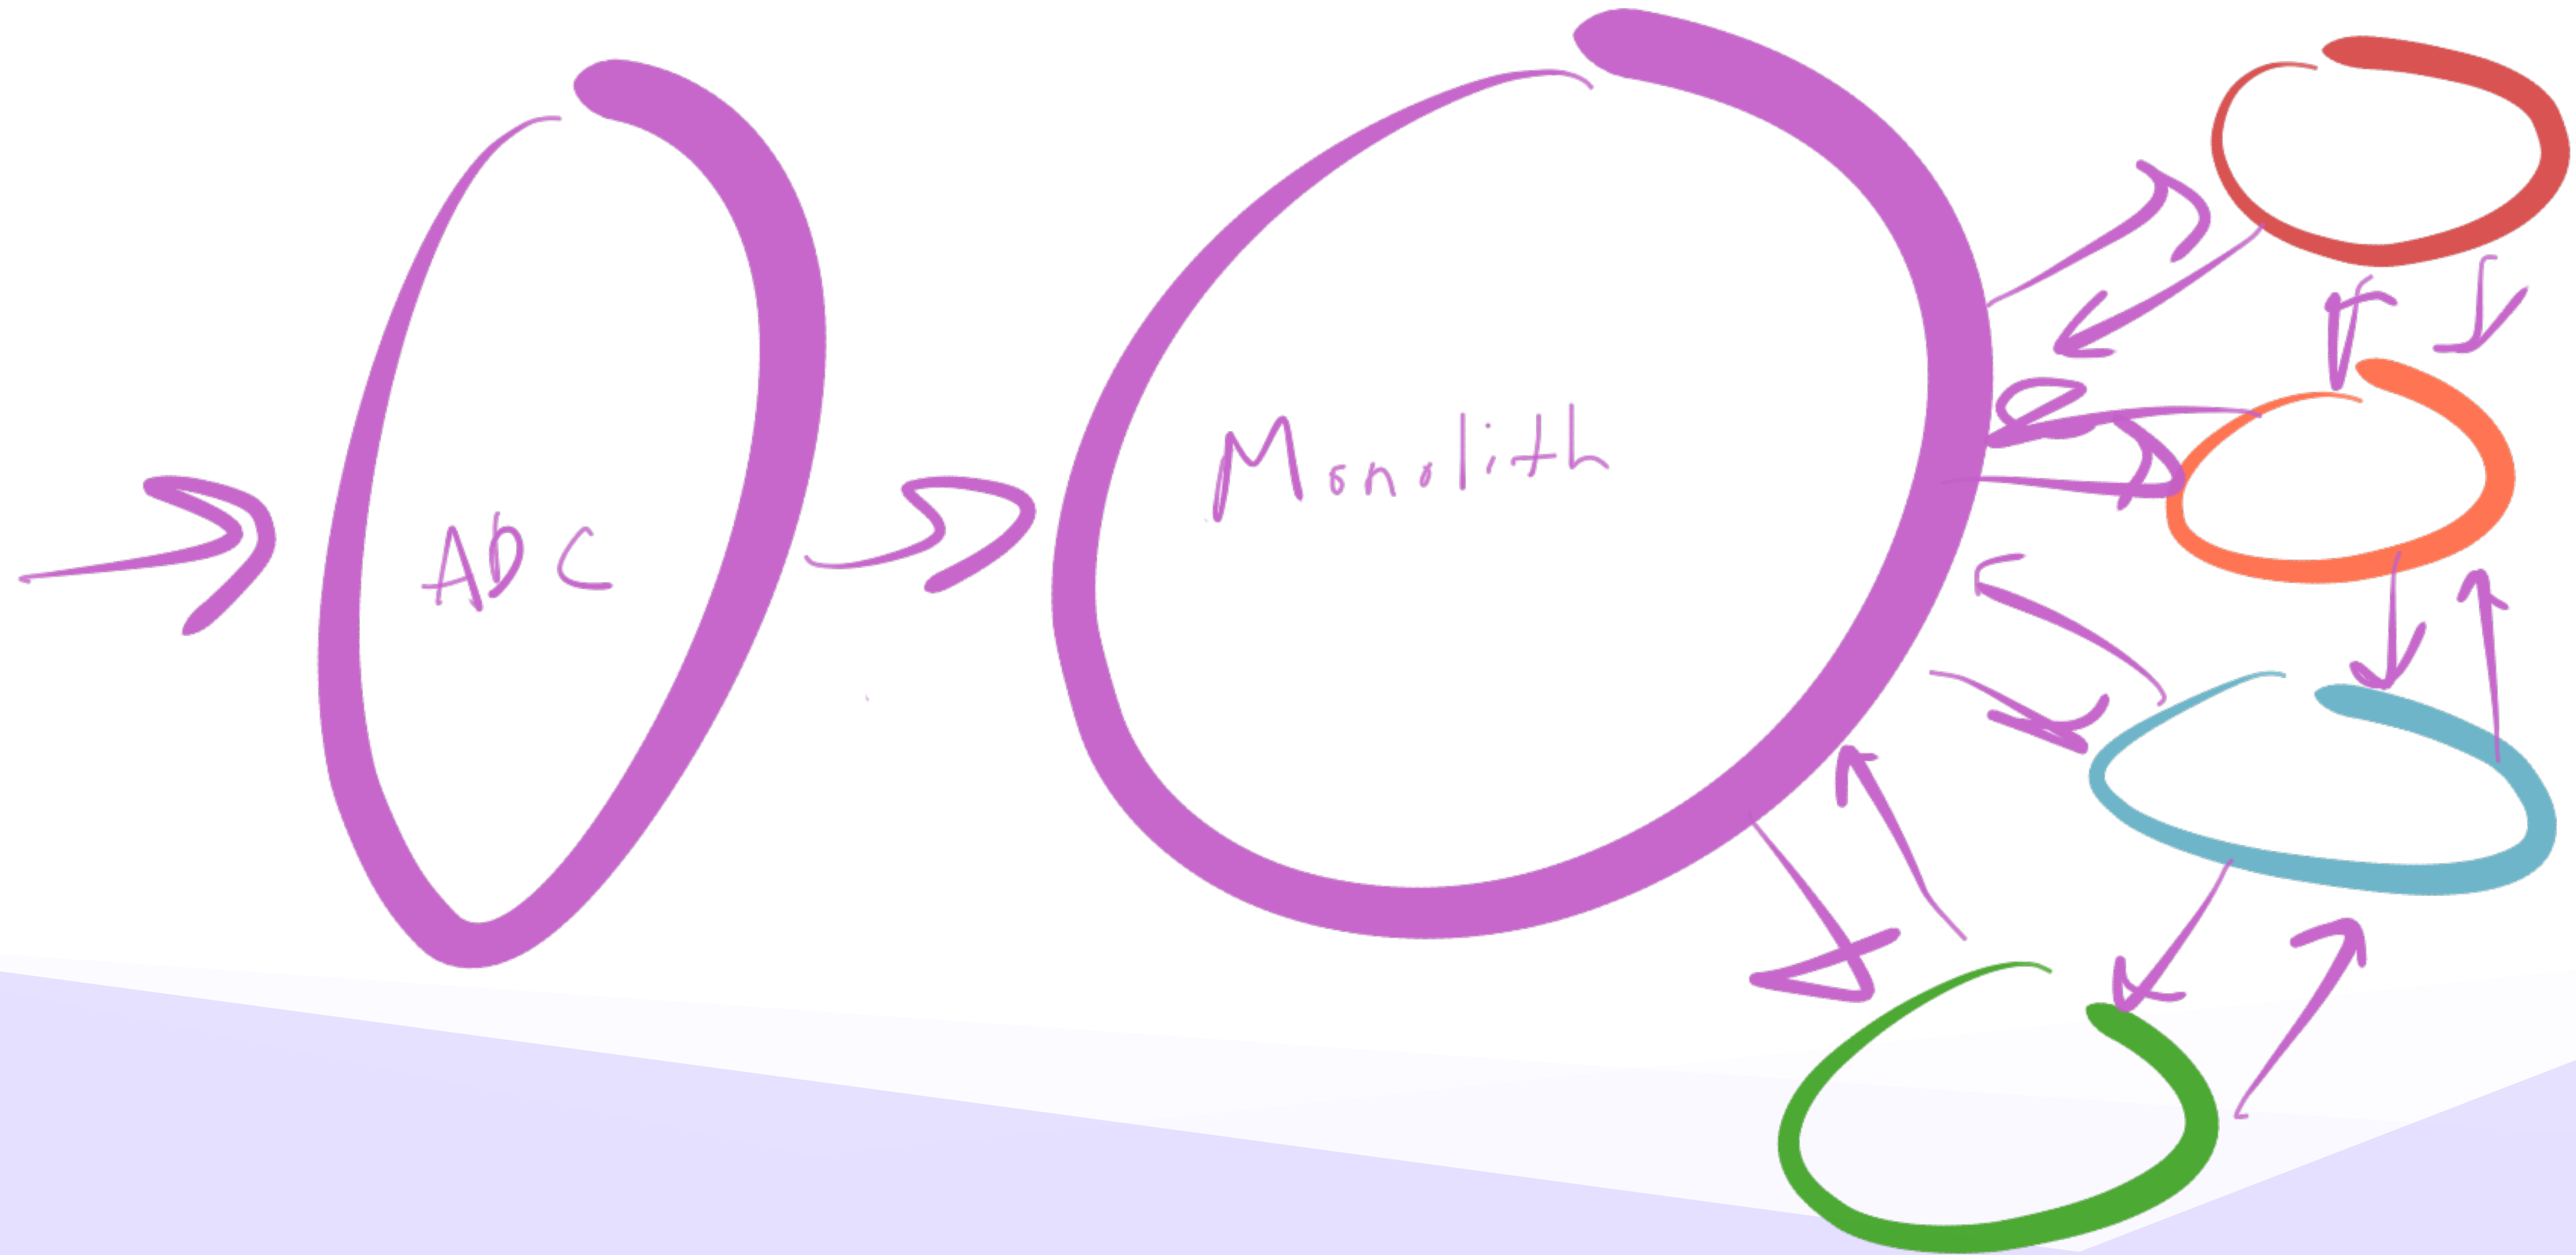
Please wait a moment and try again. For more information, check out **Twitter Status**.

[Bahasa Indonesia](#) [Bahasa Melayu](#) [Deutsch](#) [English](#) [Español](#) [Filipino](#) [Français](#) [Italiano](#) [Nederlands](#) [Português](#) [Türkçe](#)  
[Русский](#) [हिन्दी](#) [日本語](#) [简体中文](#) [繁體中文](#) [한국어](#)


© 2012 Twitter [About](#) [Help](#) [Status](#)



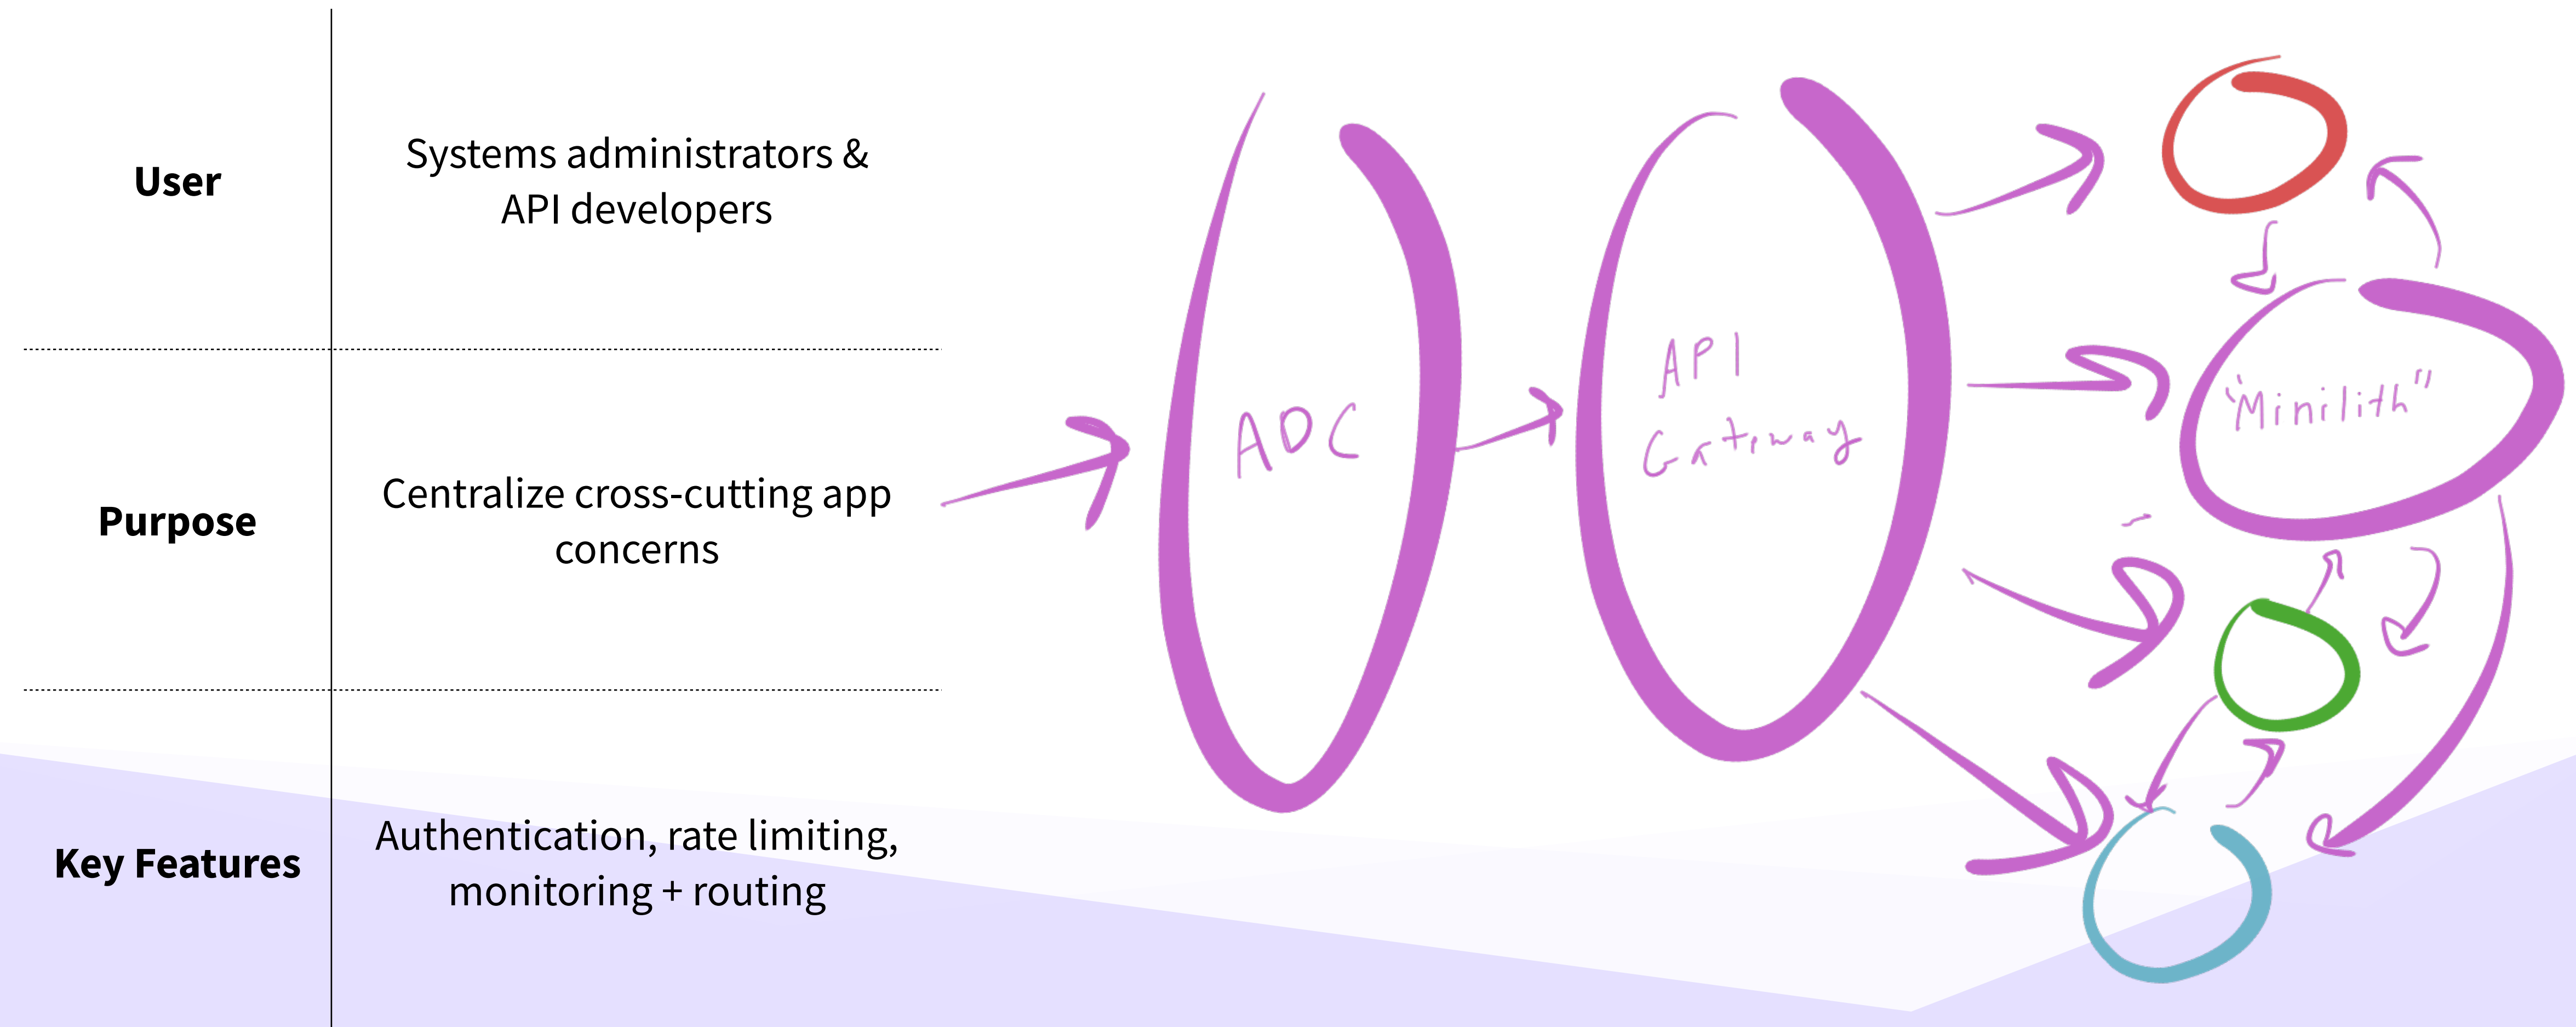
# Mini-services



# API Gateway (2nd Generation)




# API Gateway (2nd Generation)





# Cloud-native applications

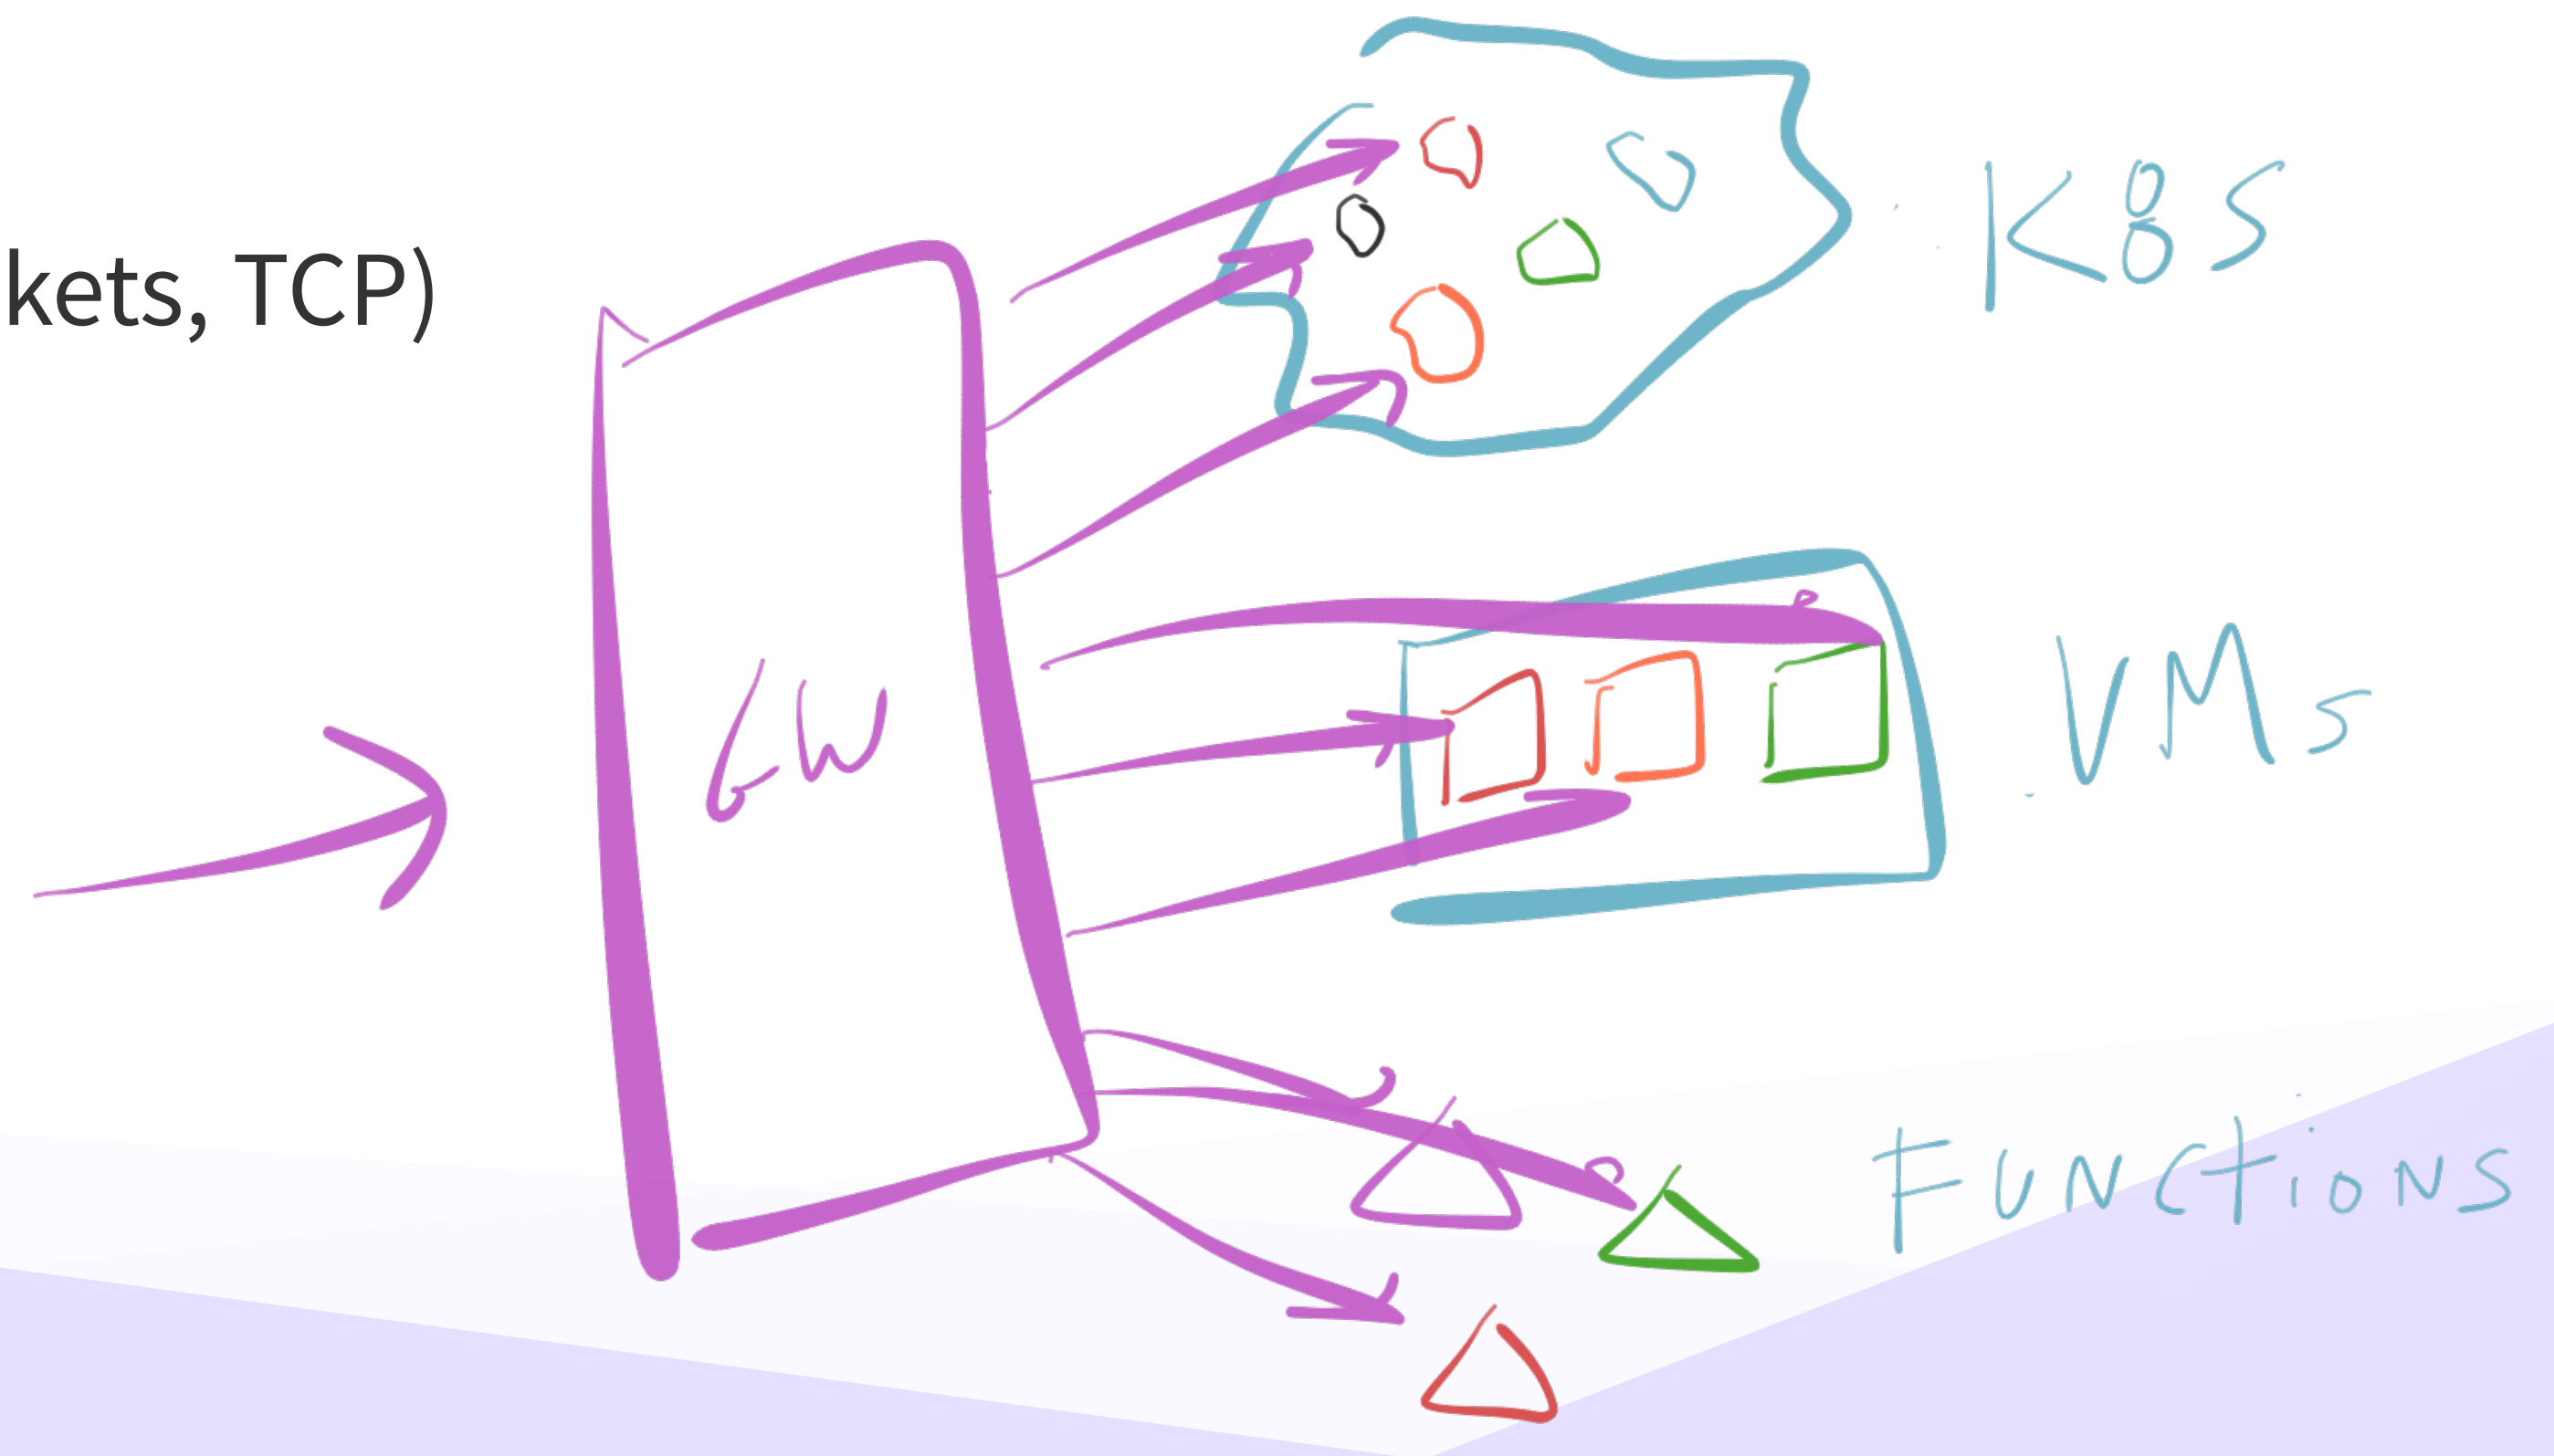
# Cloud-Native Microservices



- Modularisation (“microservices”)
- Built, released, & operated by independent application teams
- Scaled independently

# App Architecture: A Spectrum of Services

- Different locations (K8s, VMs, FaaS)
- Different protocols (gRPC, HTTP, WebSockets, TCP)
- Different load balancing requirements (sticky sessions, round robin)
- Different authentication requirements



# Cloud Gateway

1

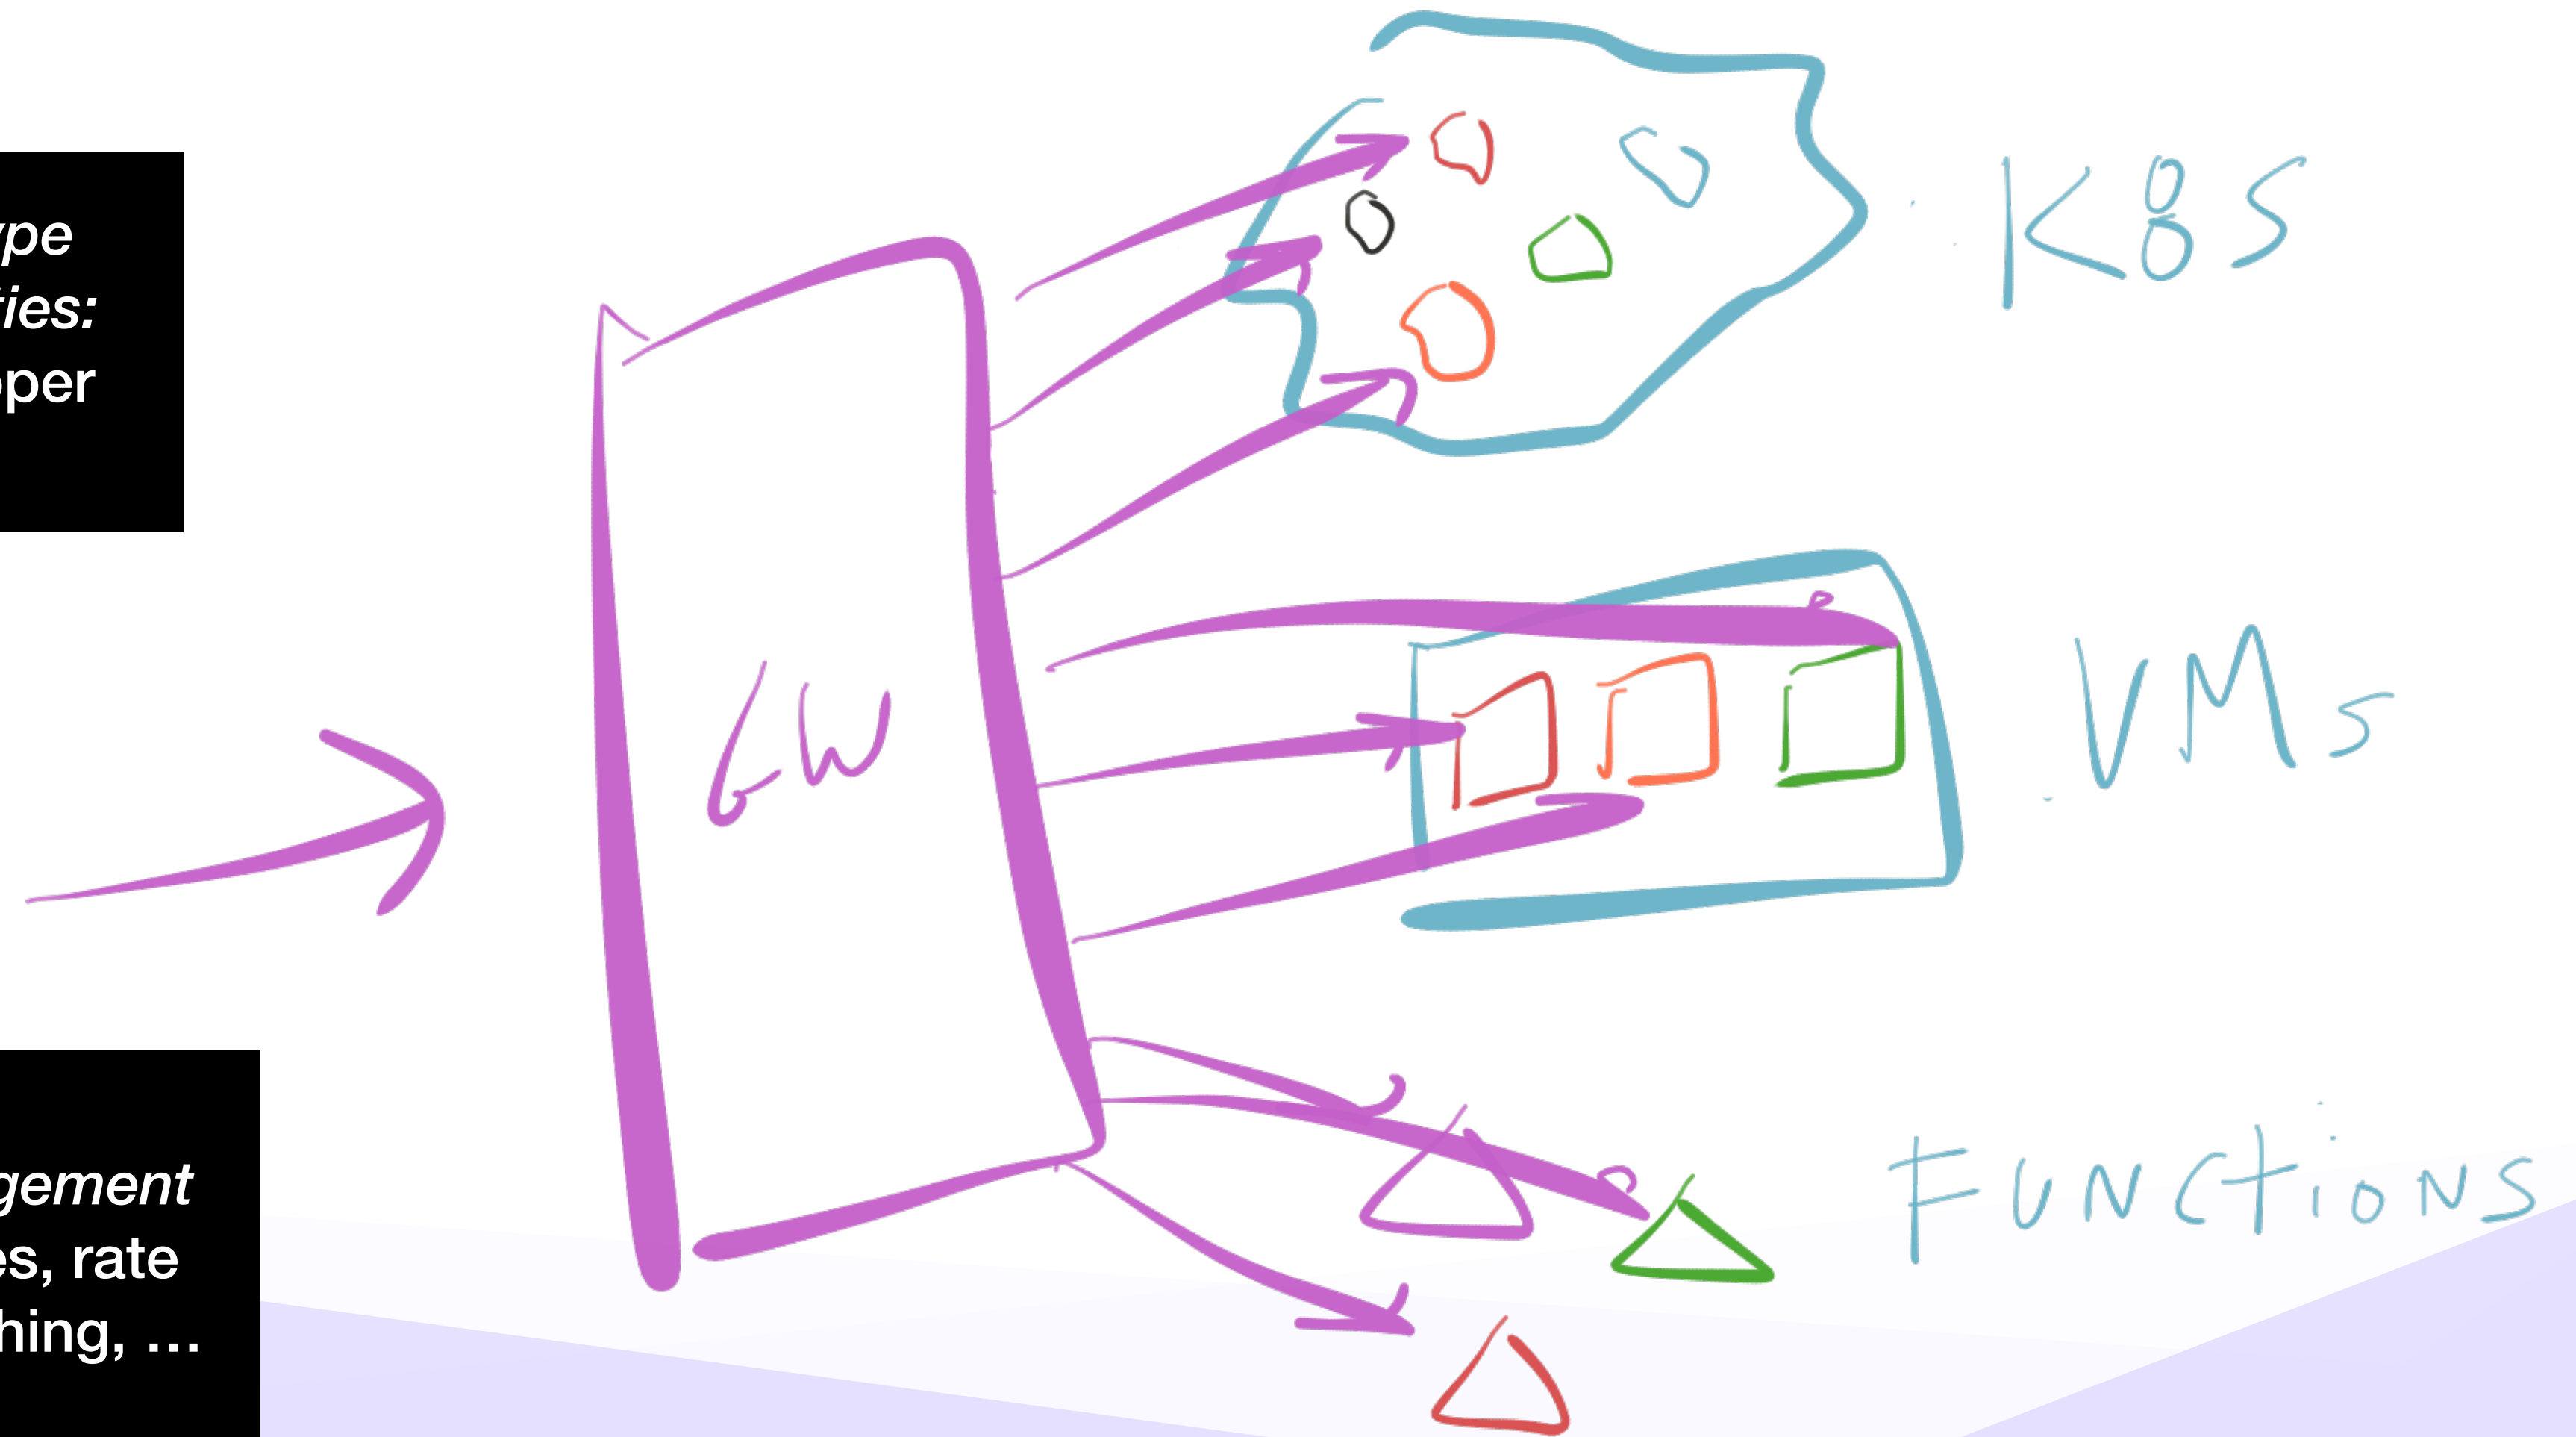
*Need API Gateway-type management capabilities: authentication, developer portal, metrics, ...*

2

*Need ADC-like traffic management capabilities: timeouts, retries, rate limiting, load balancing, caching, ...*

3

**Real-time Service Discovery**



# A spectrum of services means Cloud Gateways merge:

Load balancers / ADC functionality +

API management +

Service discovery



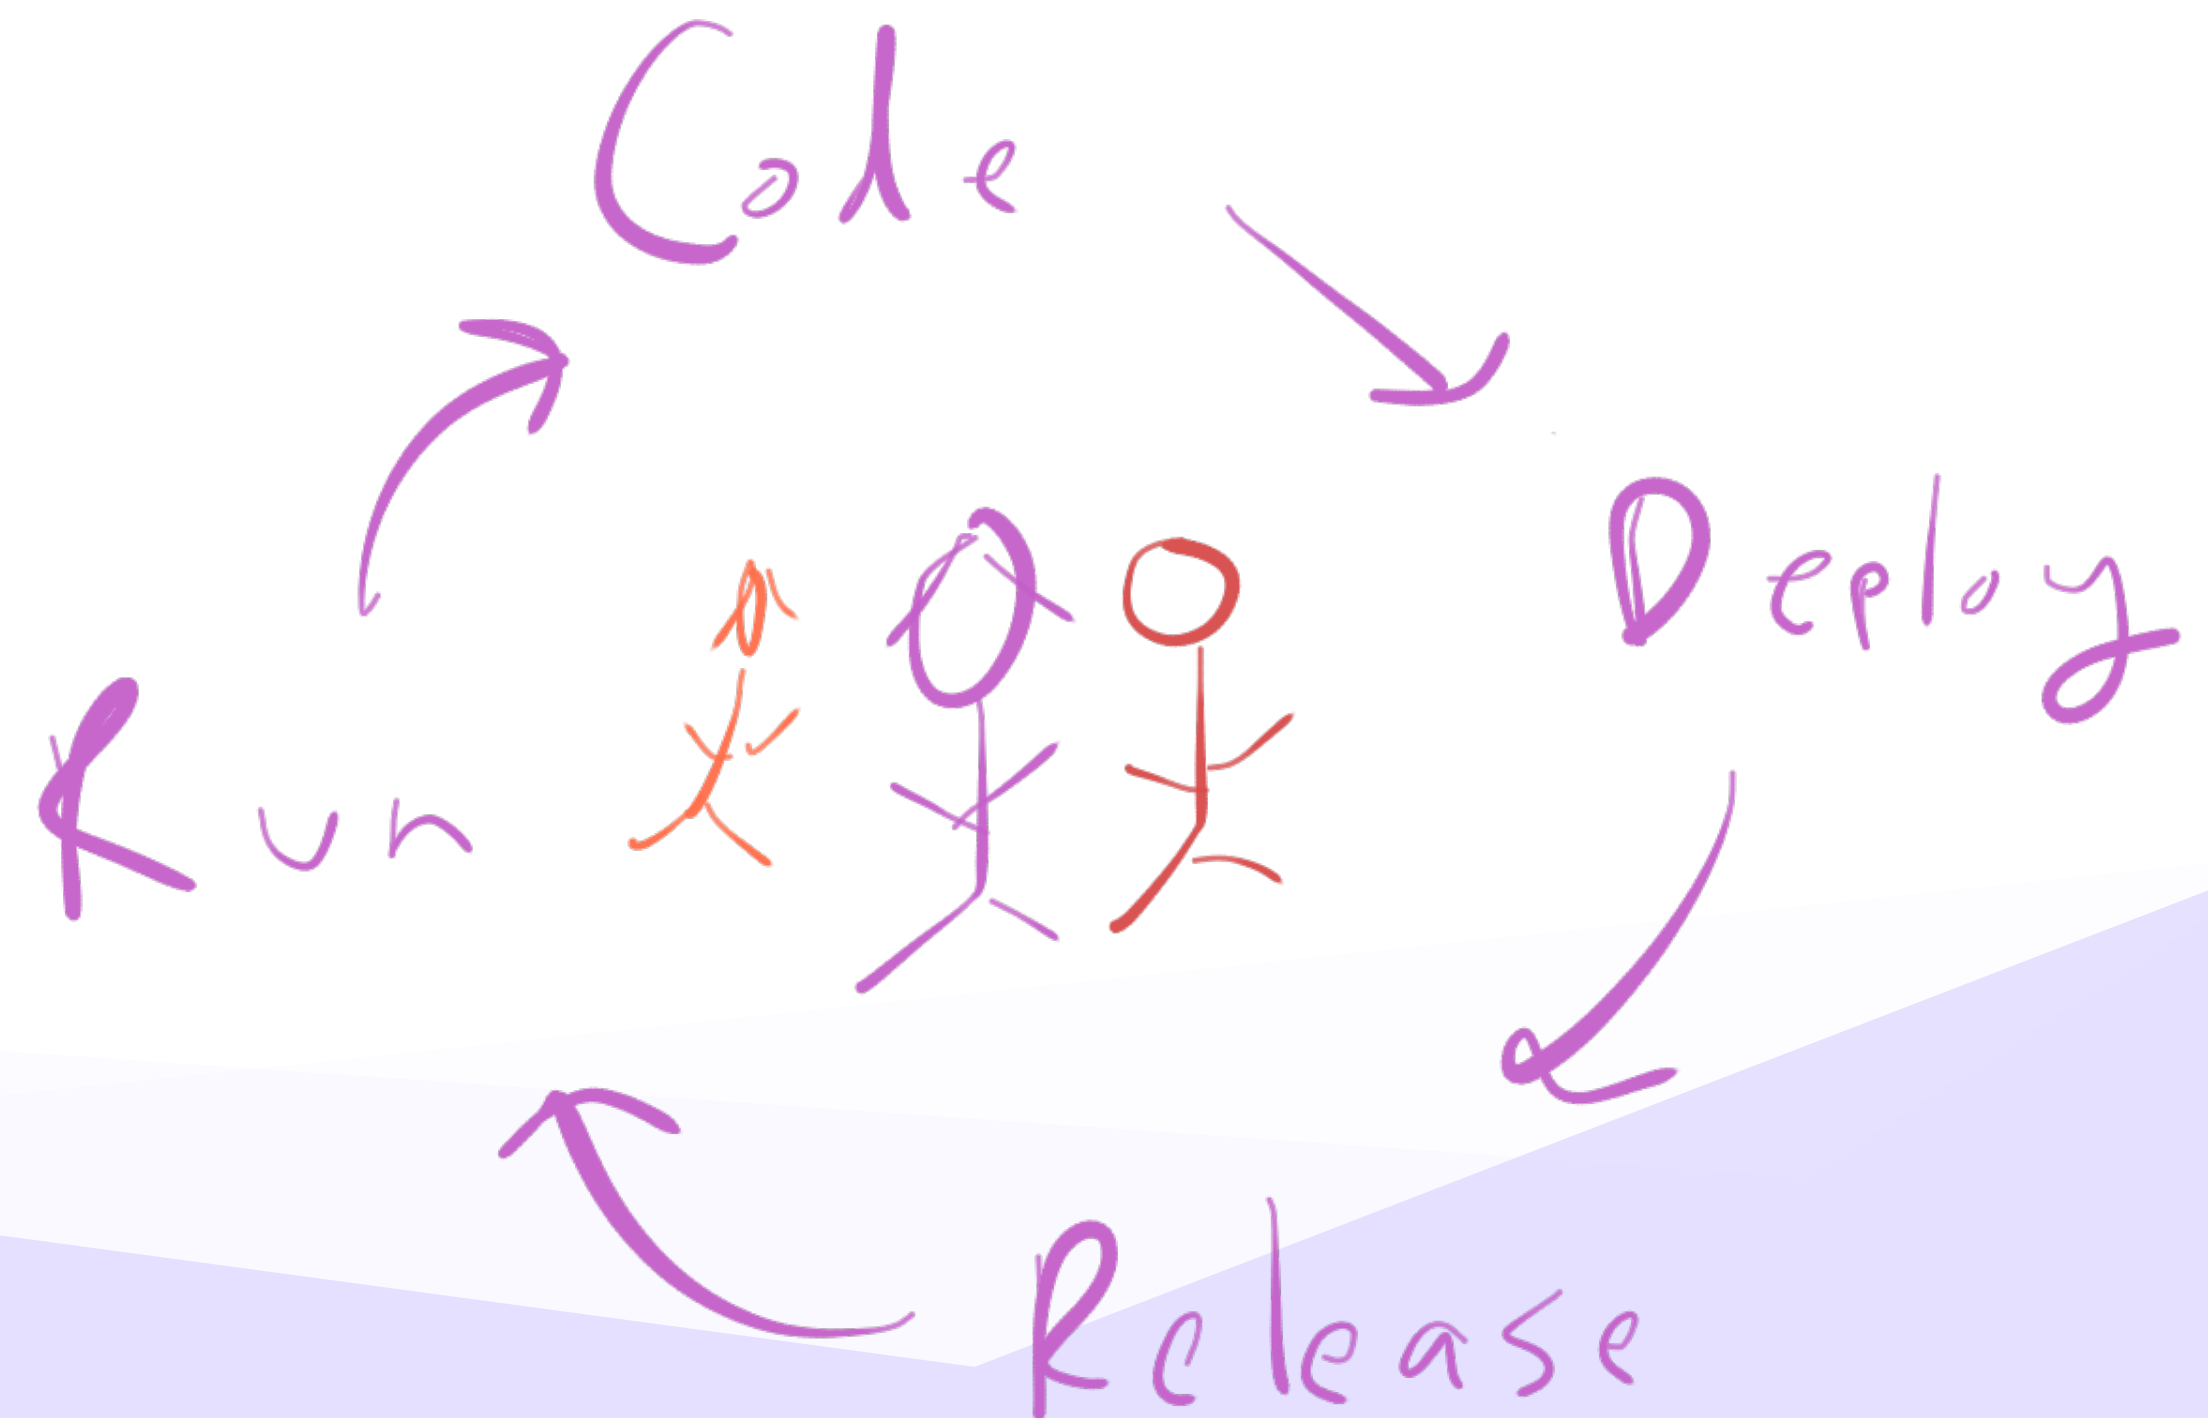


**Microservices lead to an even bigger change.**



**Developers are on call.**

# Microservices: Full Cycle Development



- App teams have full responsibility (and authority) for **delivering a service**
- Increases agility by accelerating the feedback loop.
- <https://netflixtechblog.com/full-cycle-developers-at-netflix-a08c31f83249>

A 3D arrow pointing to the right, carved into a cracked, textured surface. The arrow is made of a light-colored material, possibly wood or stone, and is set against a background of a light-colored, heavily cracked and textured material, likely concrete or stone. The arrow is positioned on the left side of the image, pointing towards the right. The text "This is a change in workflow." is overlaid on the right side of the image.

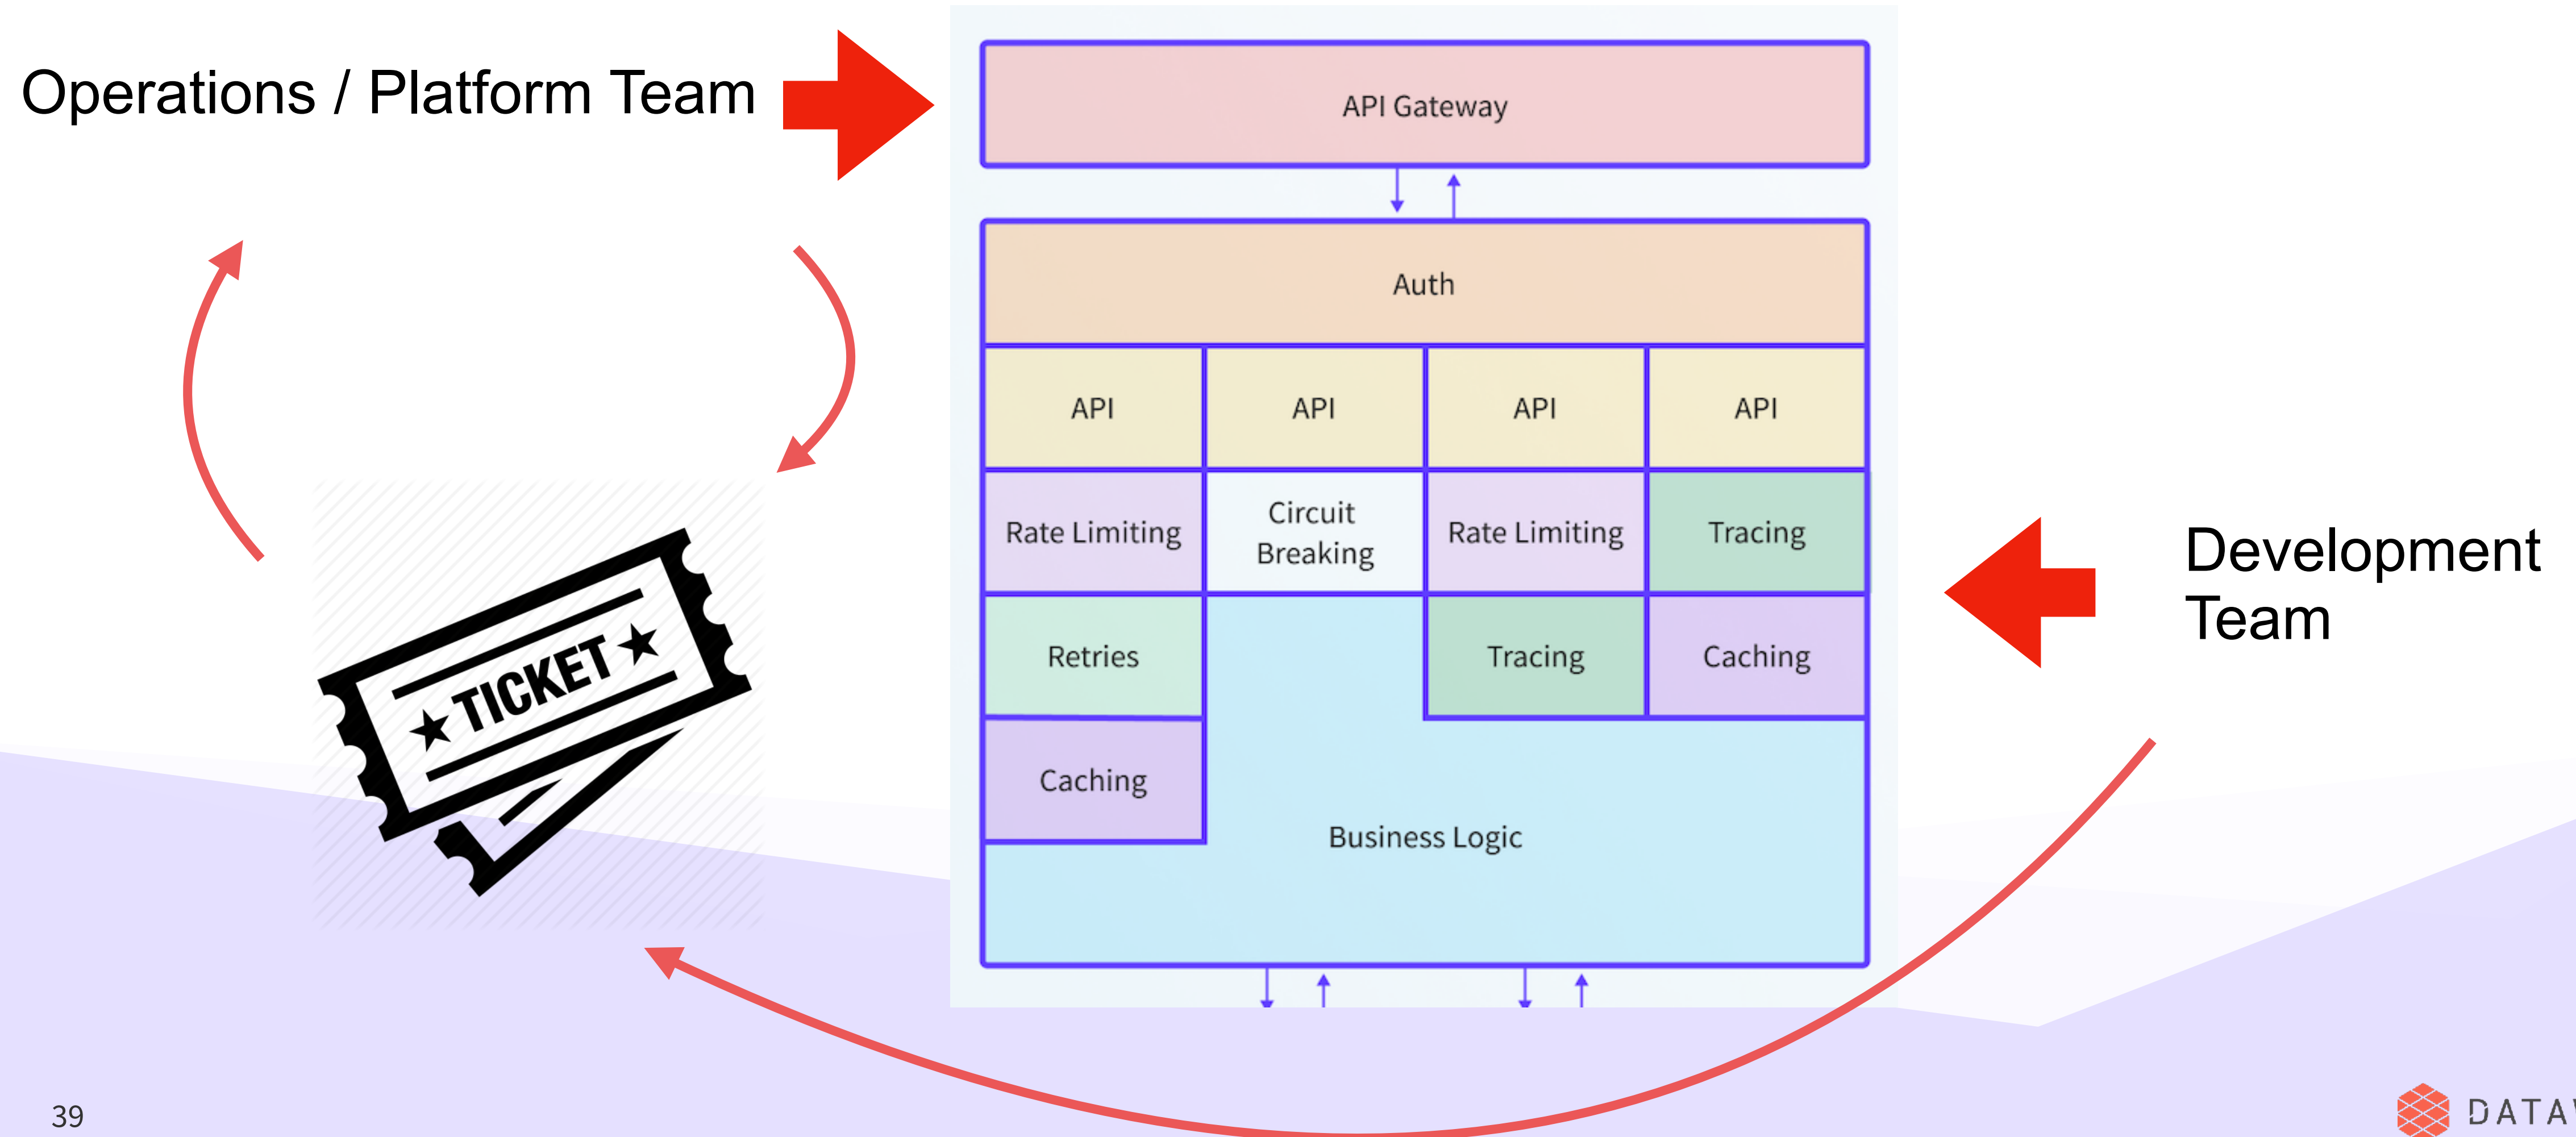
**This is a change in *workflow*.**

Thesis: The evolution of the edge ~~has been~~ **will be** driven by application architecture **and the application development workflow.**

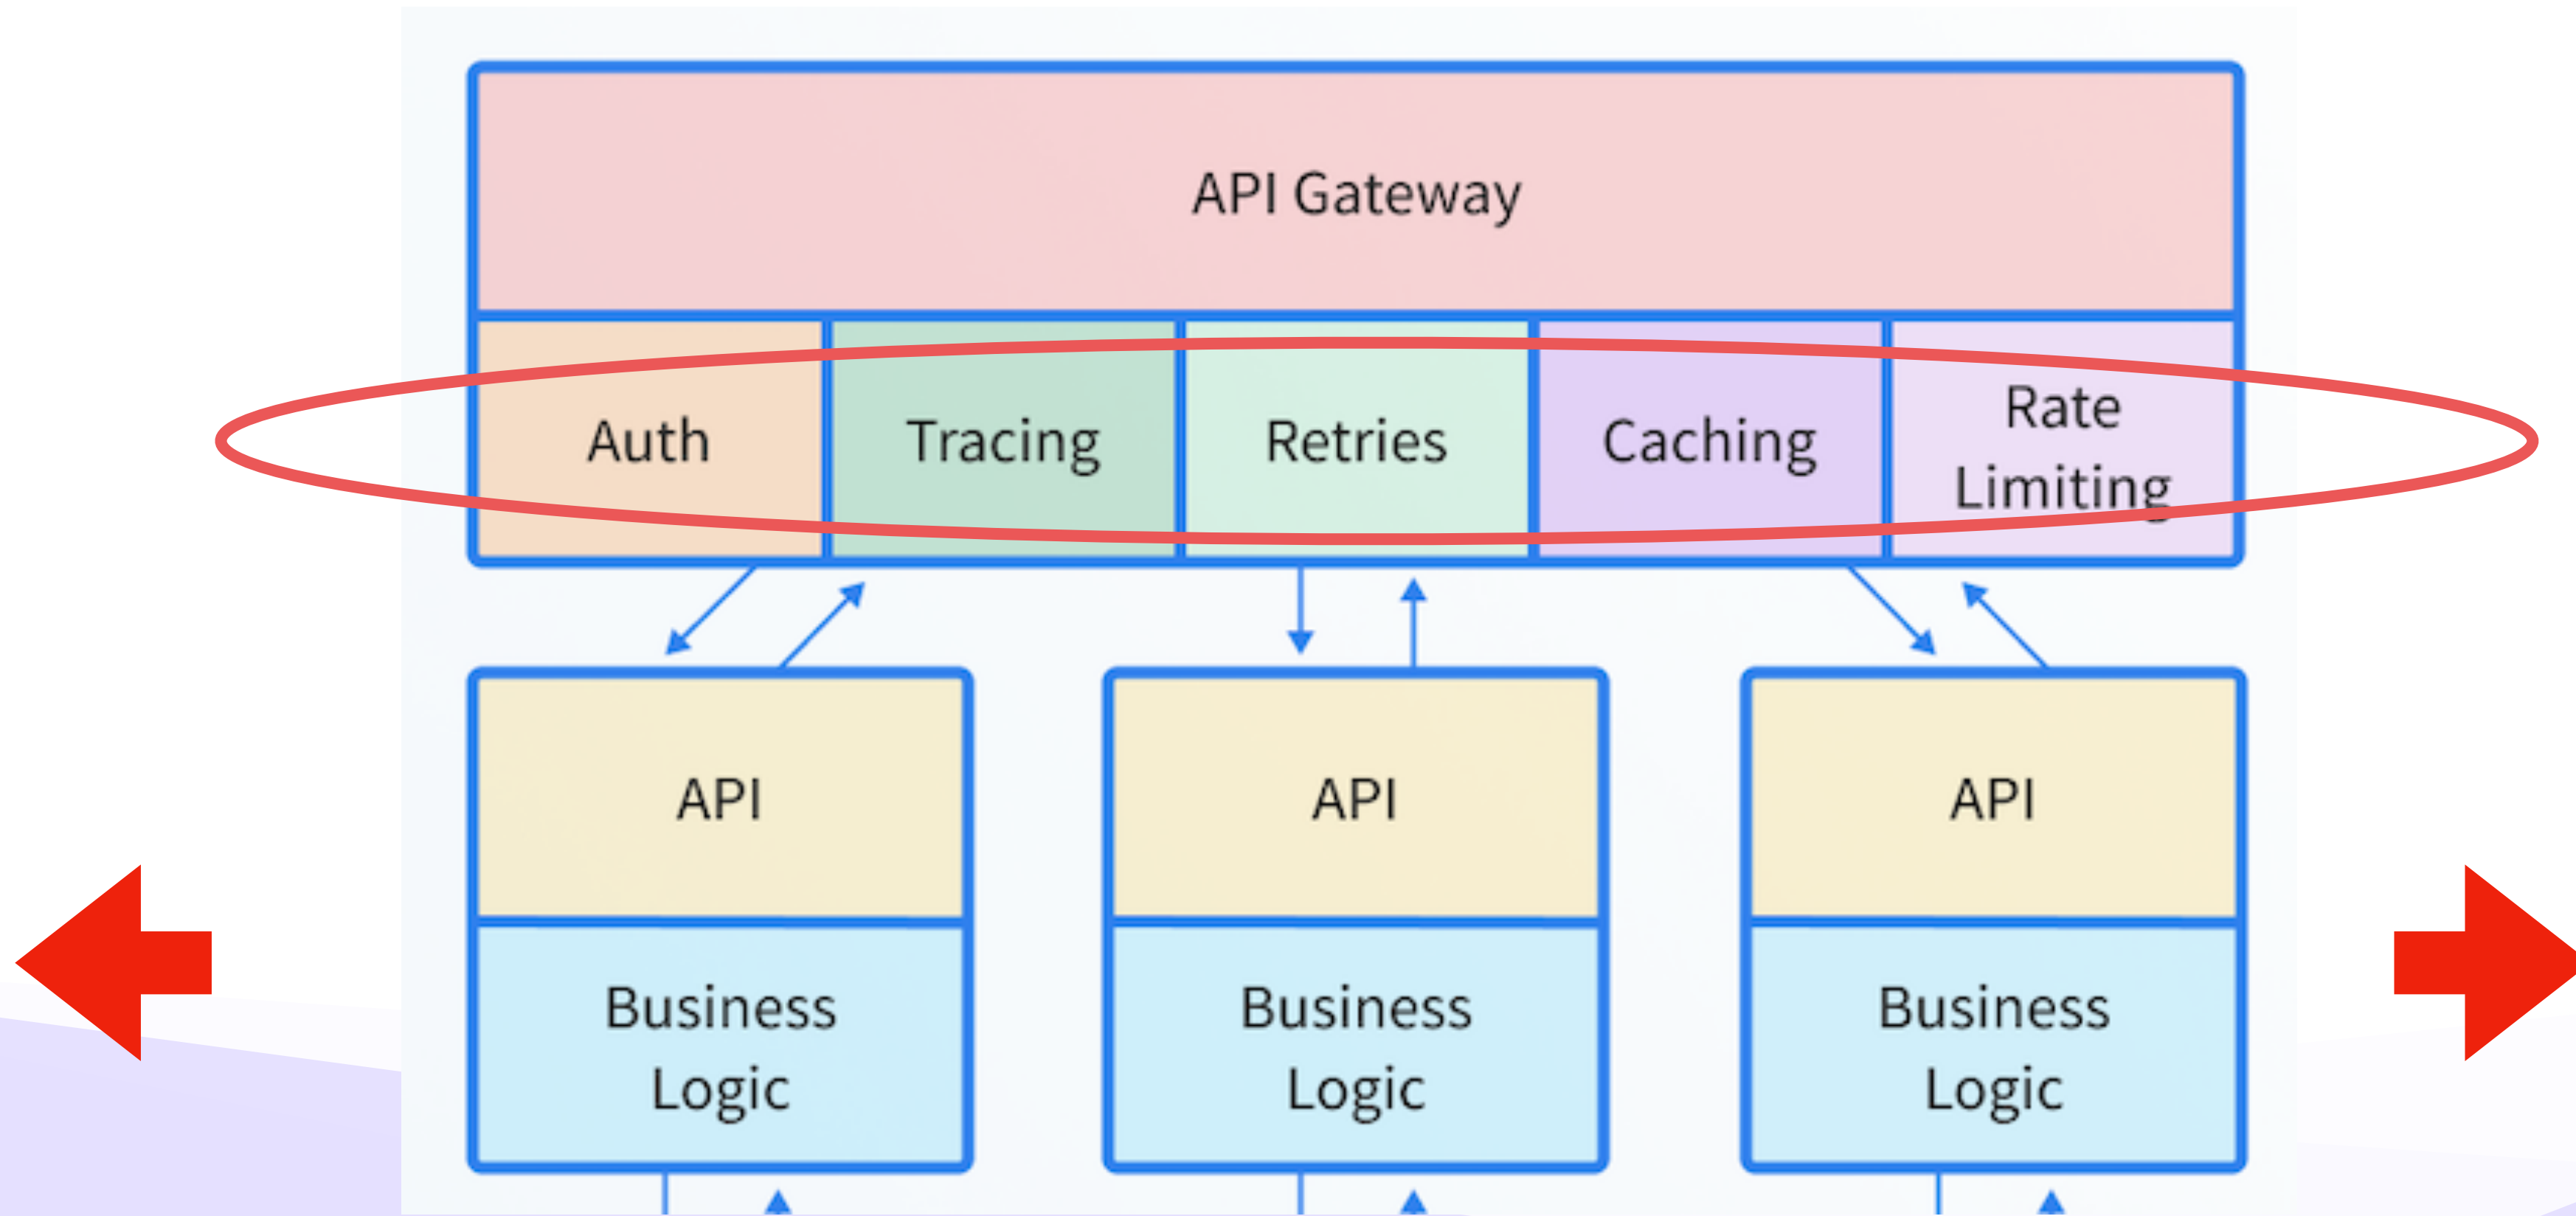


# Two Biggest Challenges

# Challenge #1: Scaling Edge Management

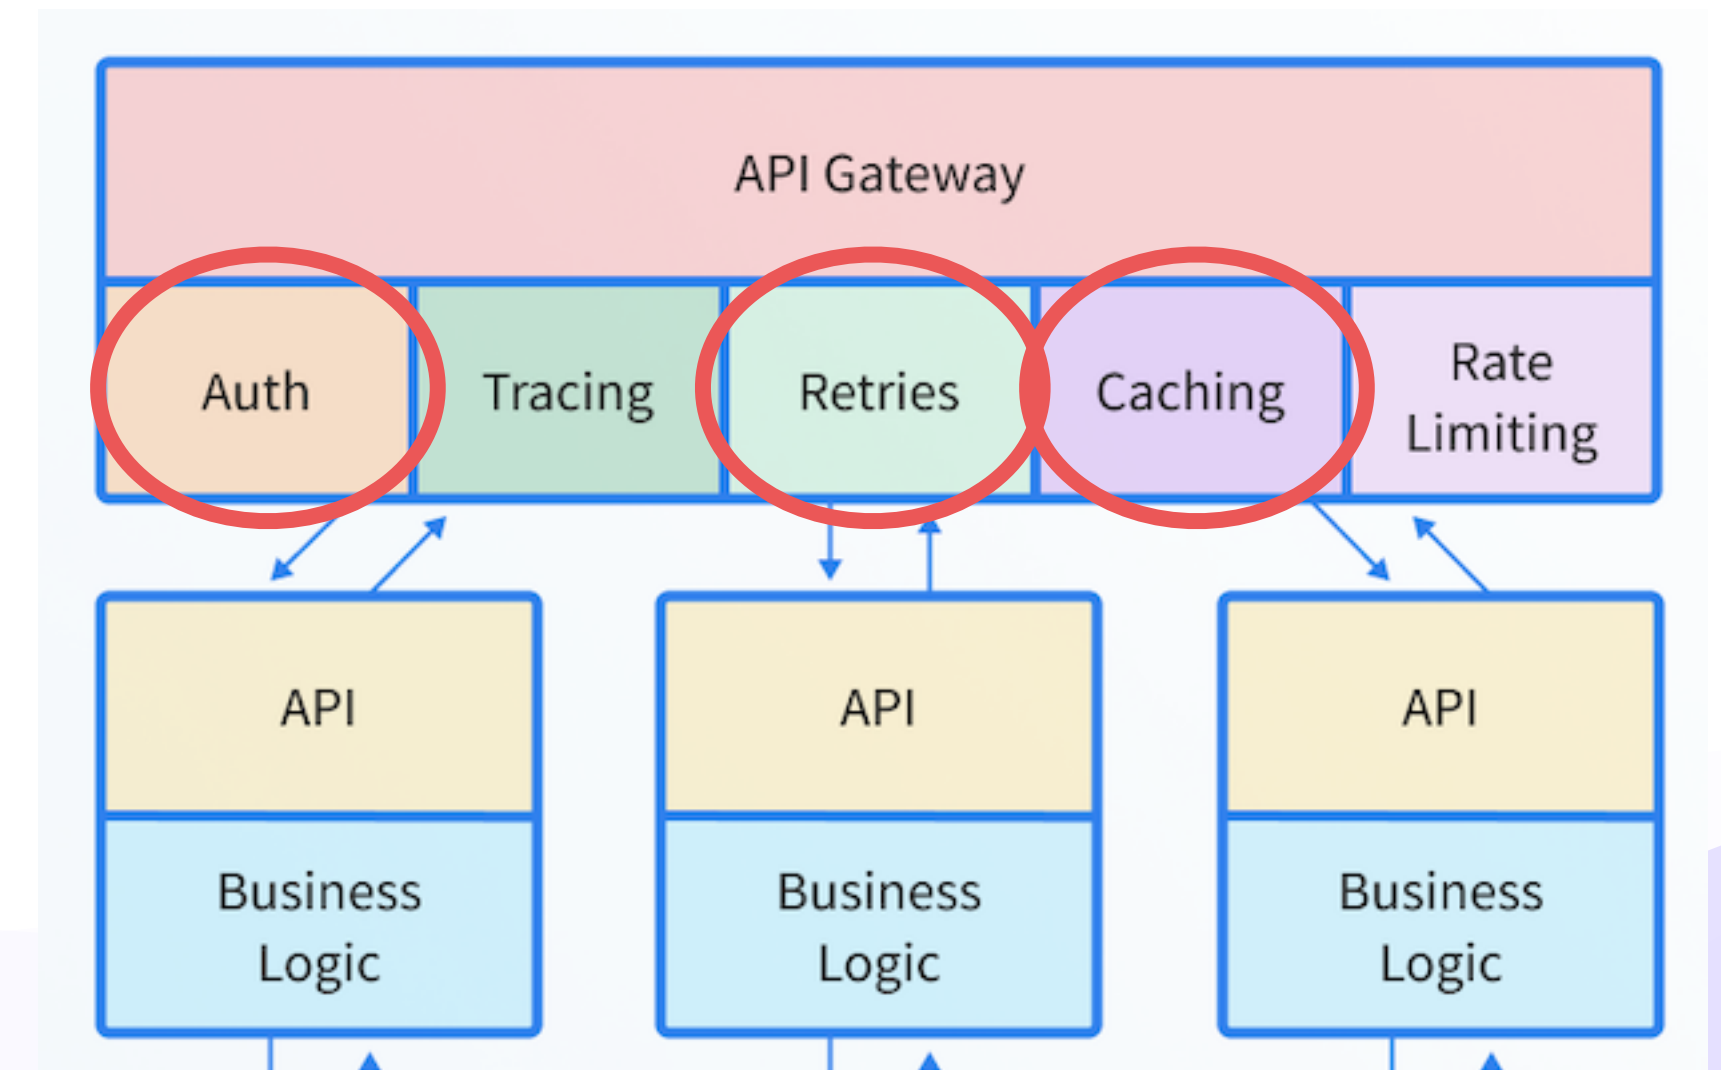
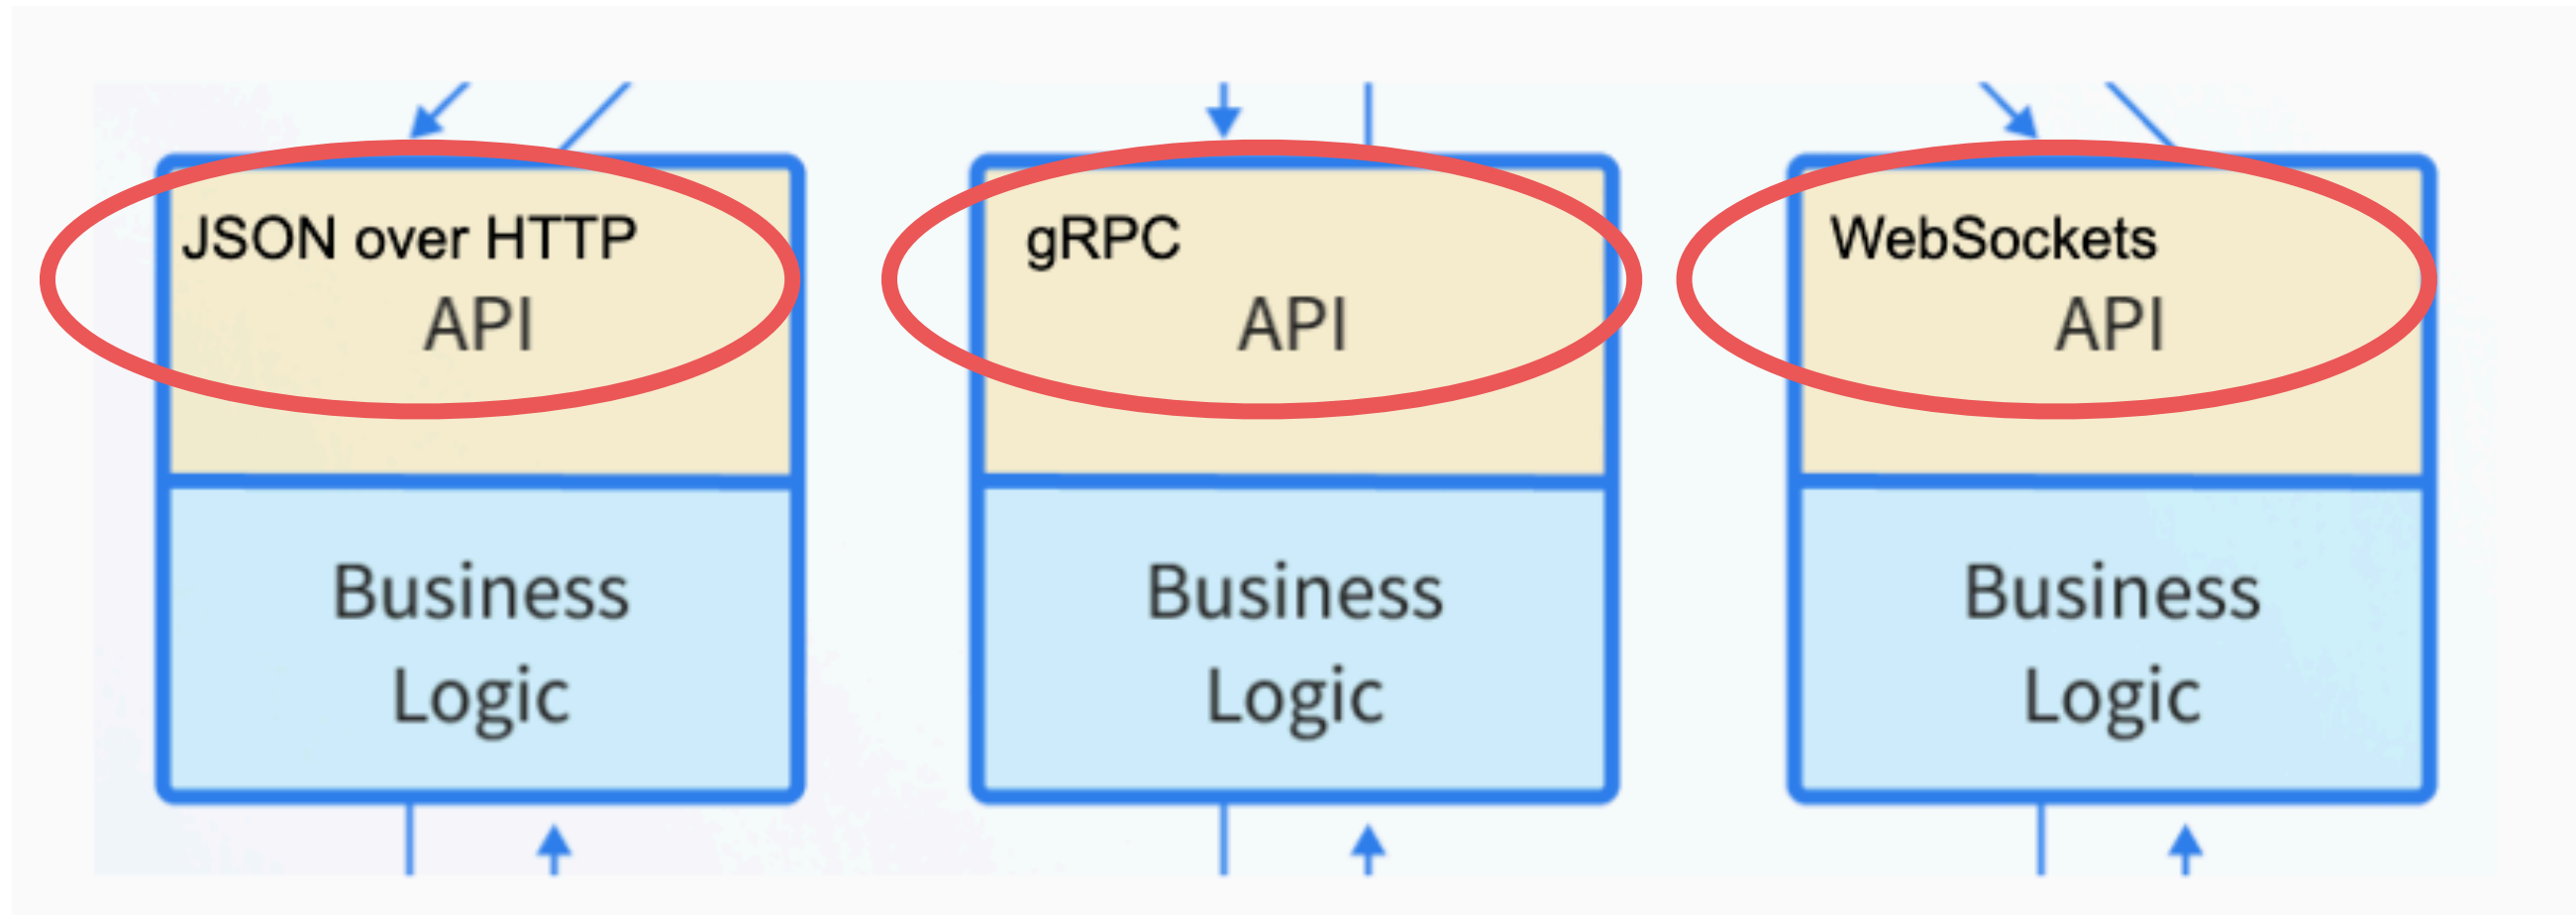


# Challenge #1: Scaling Edge Management





# Challenge #2: Supporting Diverse Edge Requirements





# Three Strategies

# Three Strategies for the Edge with Kubernetes

#1: Deploy an Additional Kubernetes API Gateway

#2: Extend Existing API Gateway

#3: Deploy an in-Cluster Edge Stack

**Three Strategies for Managing APIs and the Edge with Kubernetes**

Refactoring applications into a microservice-style architecture package within containers and deployed into Kubernetes brings several [new challenges](#) for the edge. In particular, as an increasing number of microservices are exposed to end users, the edge must support managing a multitude of configurations for a wide range of microservices. For more on these challenges, see the article "[The Two Most Important Challenges with an API Gateway when Adopting Kubernetes.](#)"

This article explores three strategies that engineering teams can apply in order to effectively manage the edge when migrating to microservices and Kubernetes: deploying an additional Kubernetes API gateway; extending an existing API gateway; and deploying a comprehensive self-service edge stack.

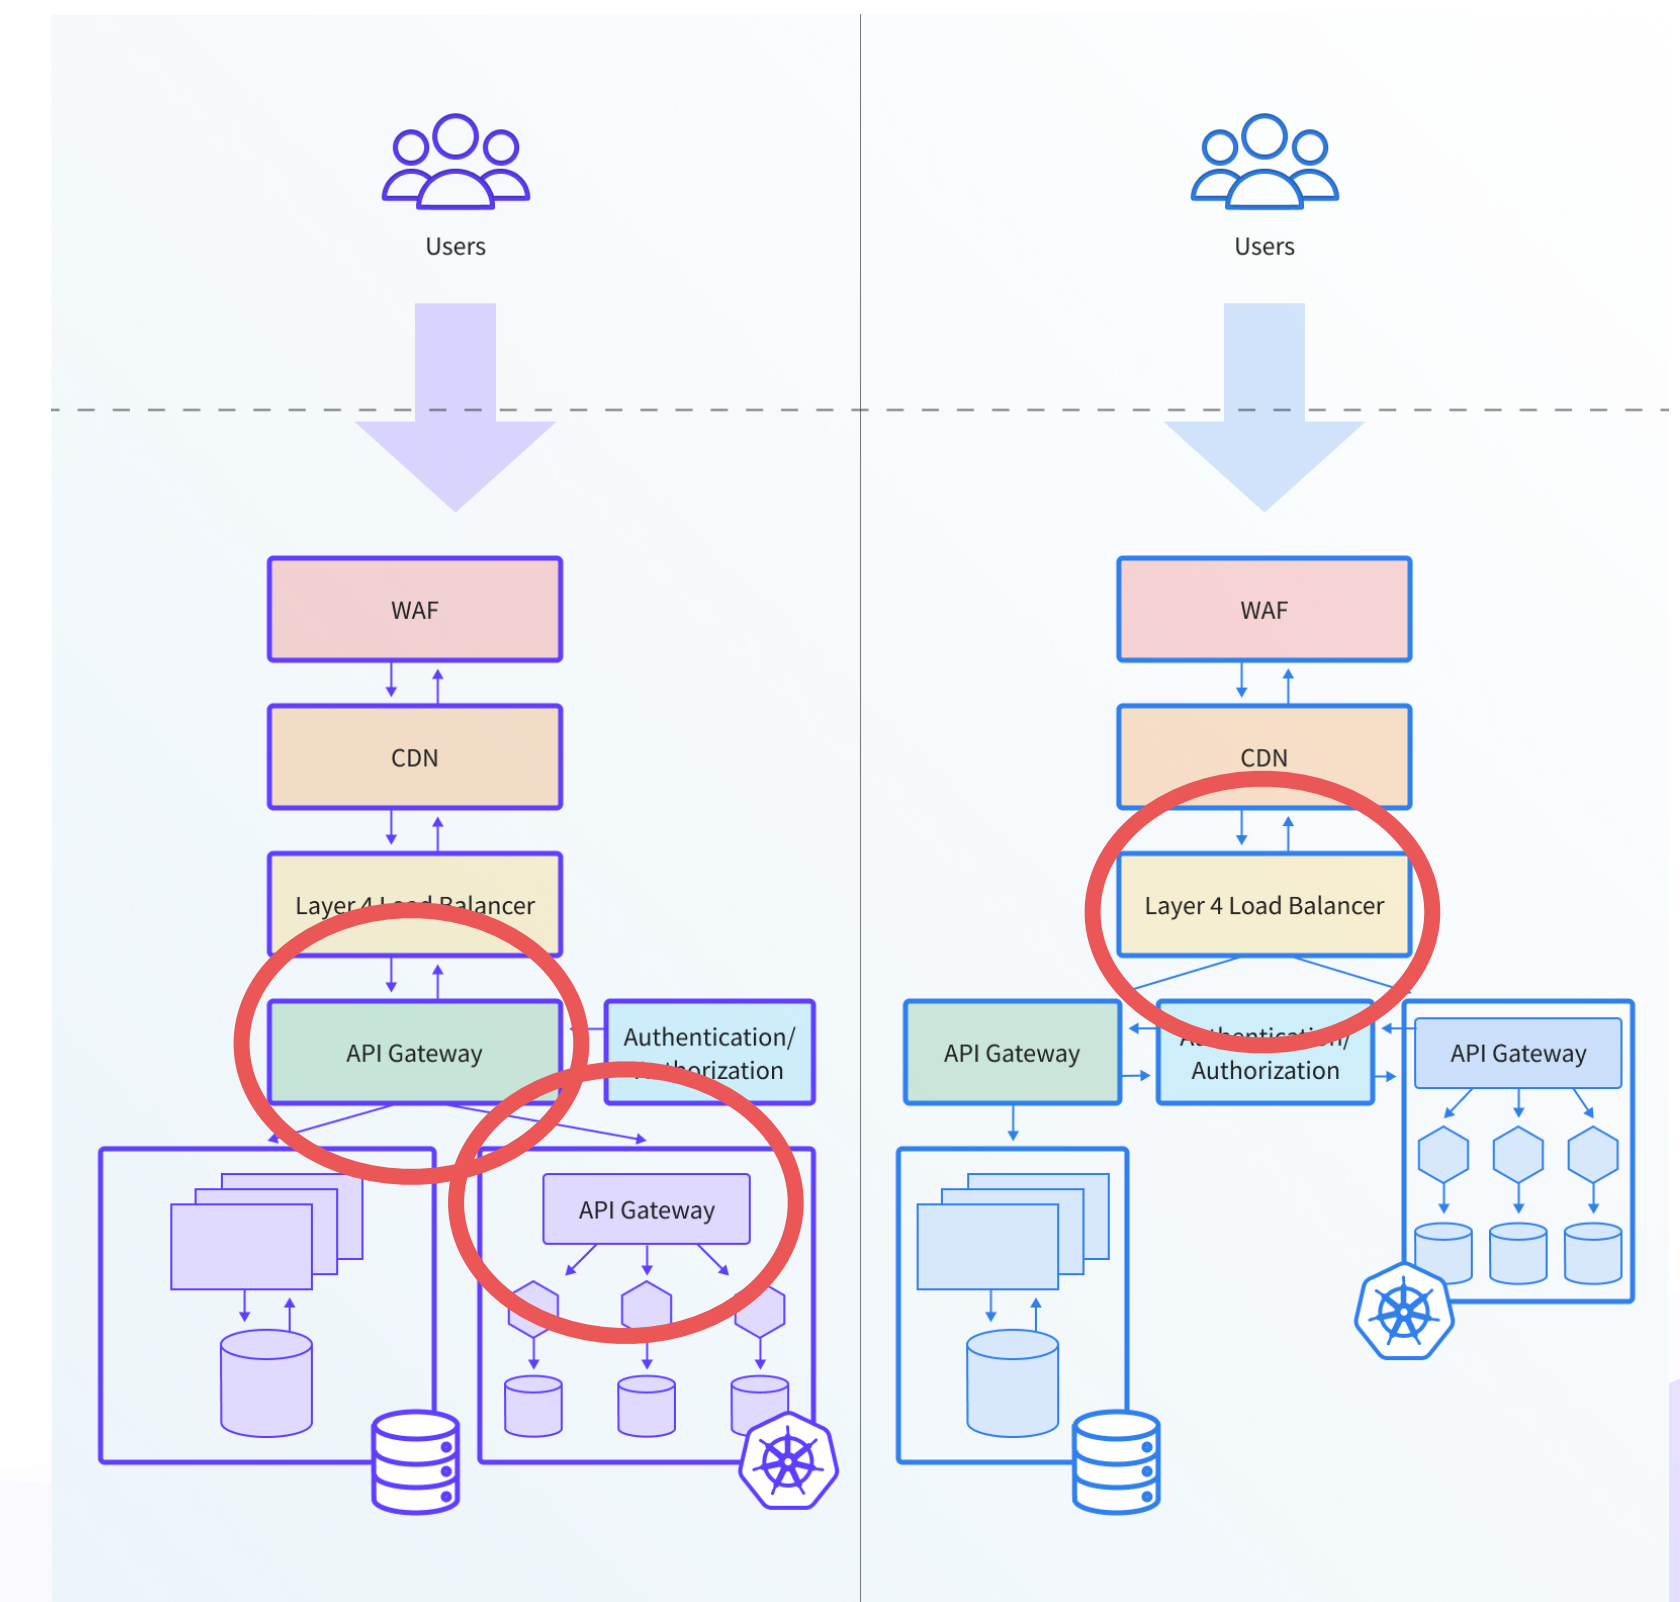
### 3 API AND EDGE MANAGEMENT STRATEGIES FOR KUBERNETES

- Deploy an Additional Kubernetes API Gateway**  
With the "deploy an additional Kubernetes API gateway" strategy, a platform team will simply deploy a completely separate gateway within the new Kubernetes clusters.
- Extend Existing API Gateway**  
The key goal here is to enable synchronization between a user's API endpoints and the location of the services deployed within the new Kubernetes clusters.
- Deploy a Comprehensive Self-Service Edge Stack**  
The edge stack is installed in each of the new Kubernetes clusters and replaces the majority of existing edge and gateway functionality that previously ran outside.

<https://www.getambassador.io/resources/strategies-managing-apis-edge-kubernetes/>

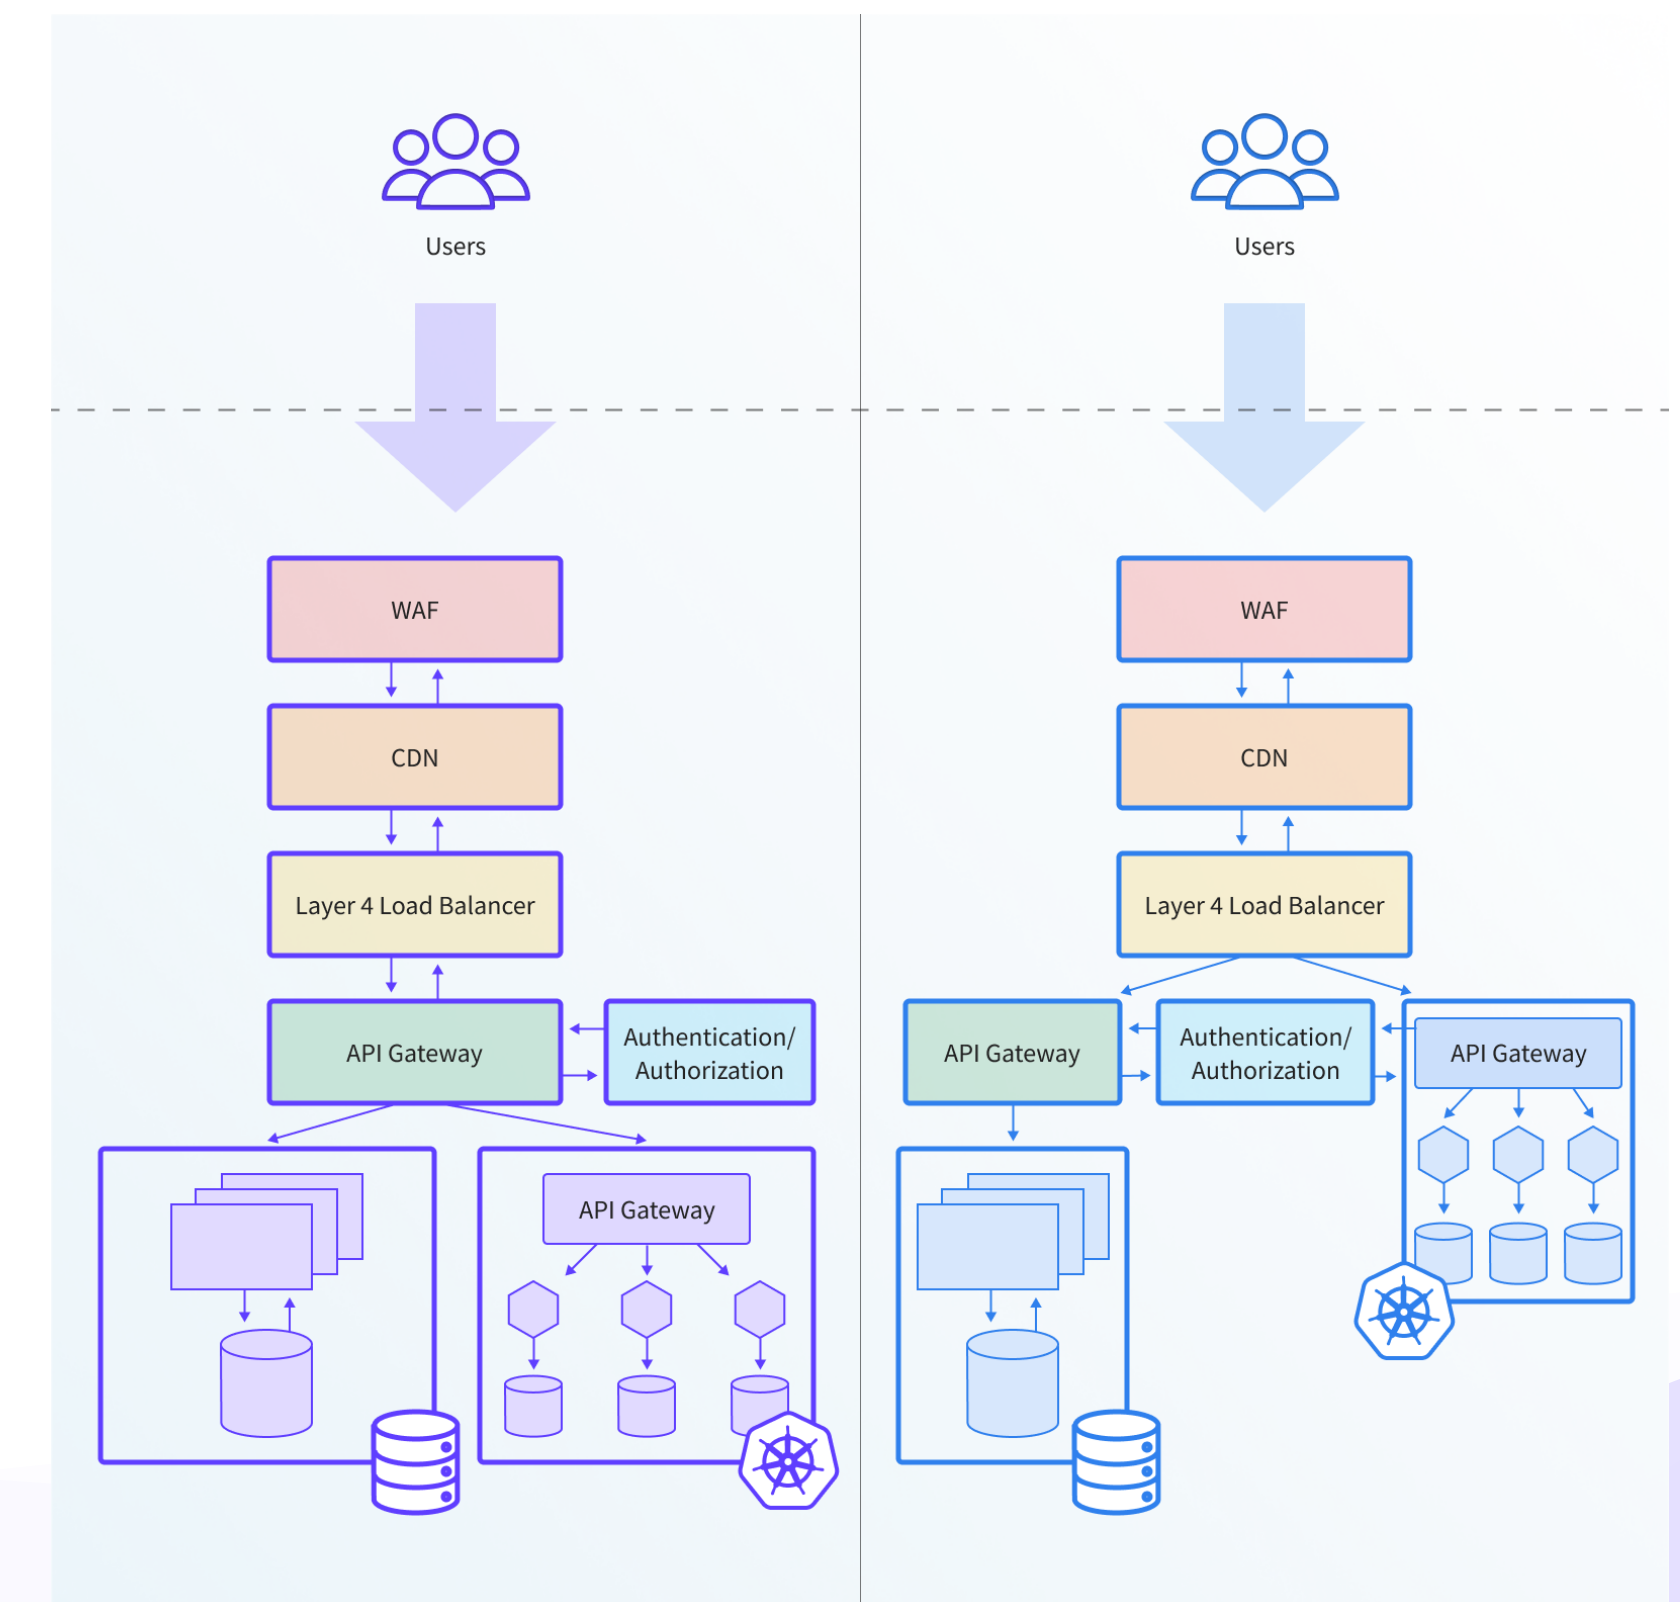
# #1 Deploy an Additional Kubernetes API Gateway

- Simply deploy an additional “in-cluster” gateway
  - Below the existing gateway
  - Below the load balancer
- Management
  - Development teams responsible
  - OR existing ops team manages this



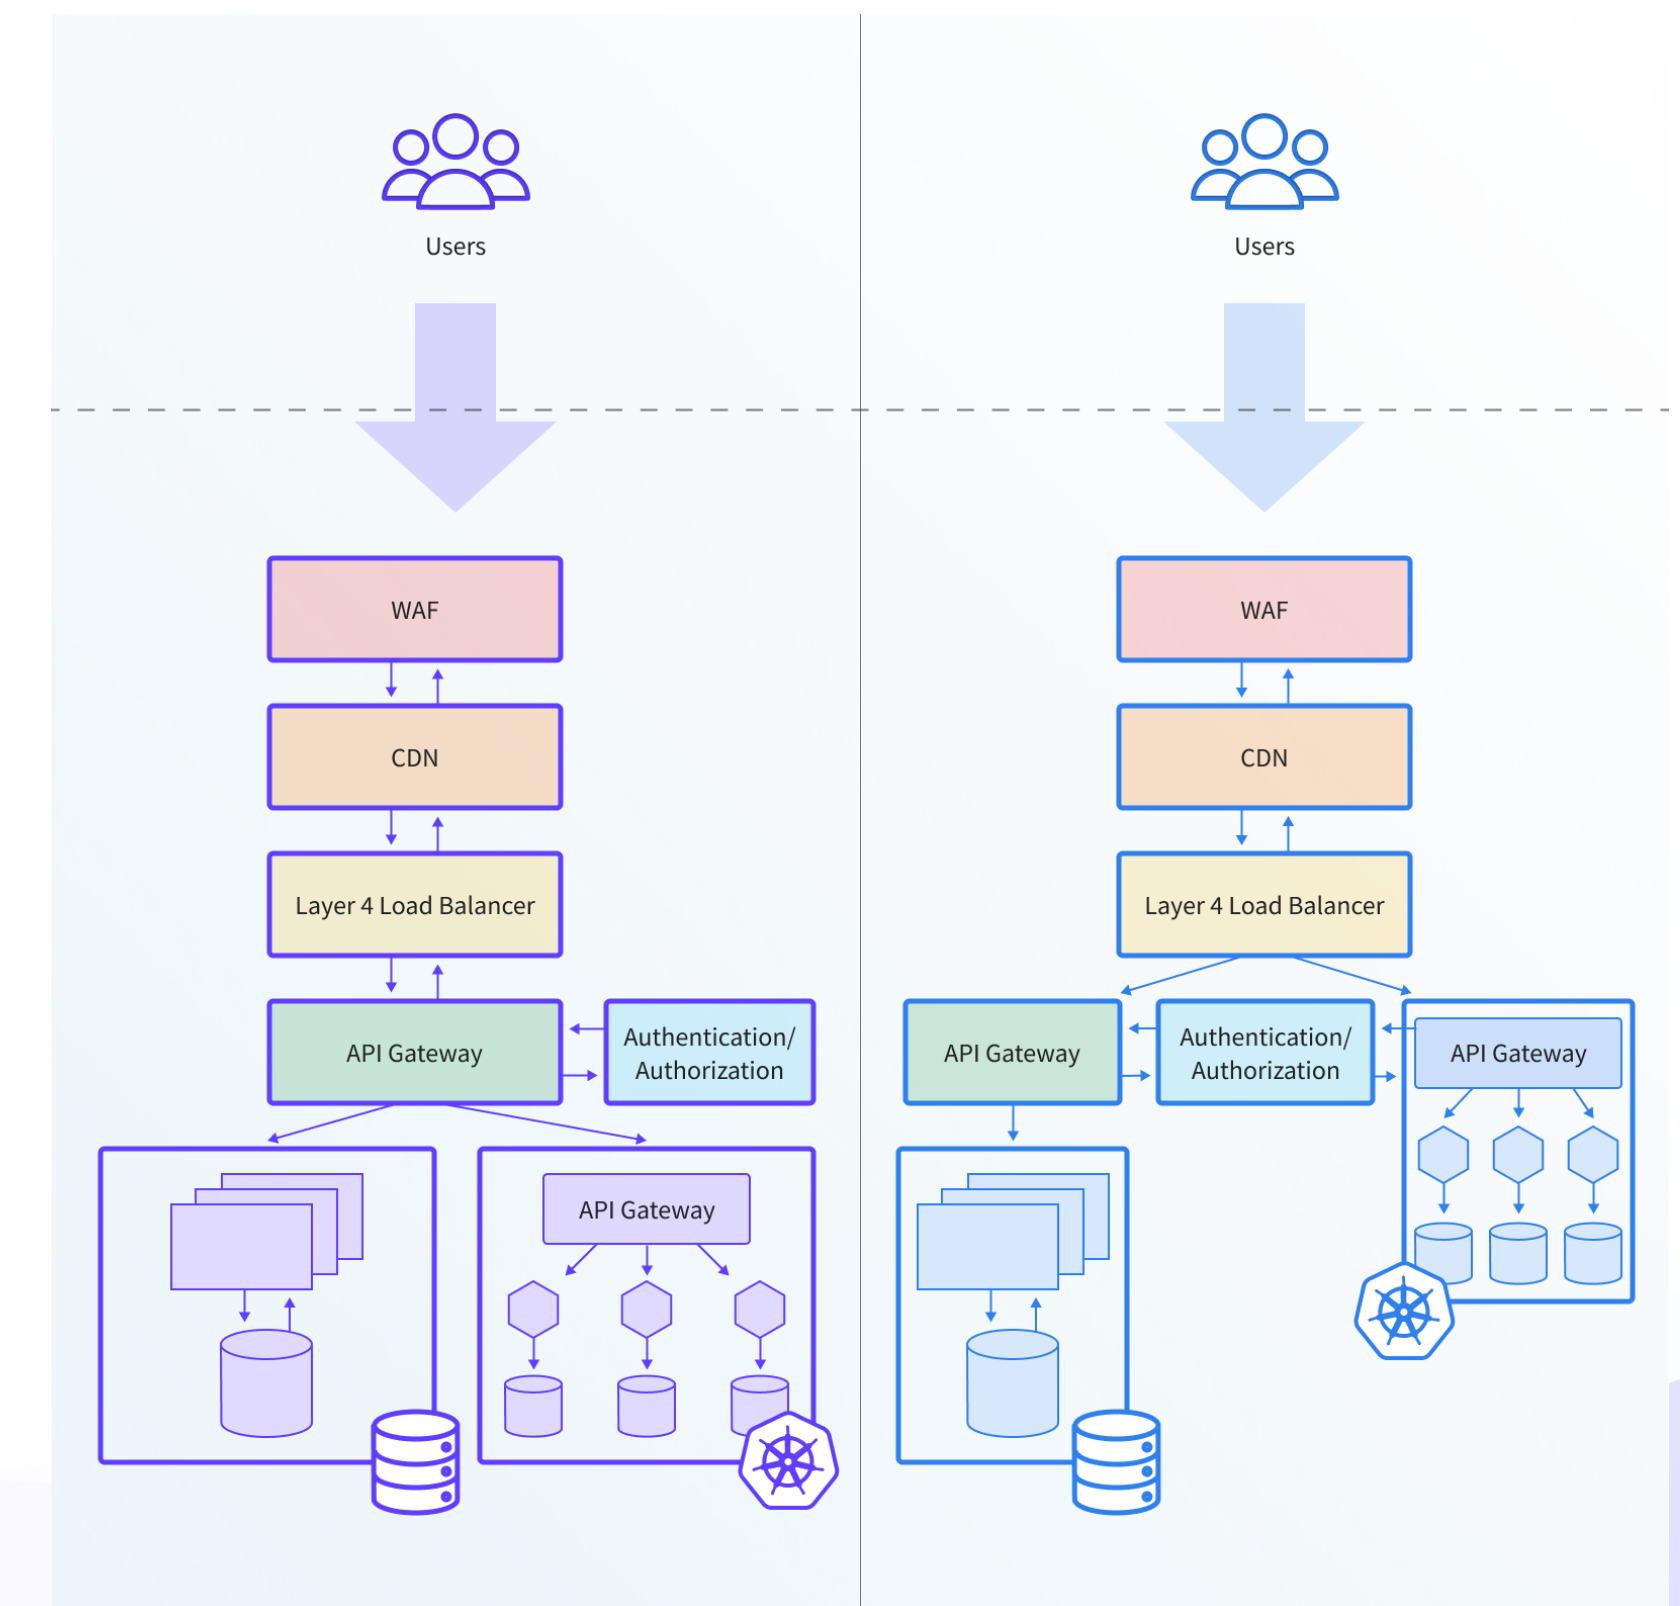
# #1 Deploy an Additional Kubernetes API Gateway

- Pros
  - There is minimal change to the core edge infrastructure.
  - Incremental migration easily
- Cons
  - Increased management overhead of working with different components
  - Challenging to expose the functionality to each independent microservice teams




# #1 Deploy an Additional Kubernetes API Gateway

- As much edge functionality as possible should be pushed into the Kubernetes API Gateway, and directly exposed to application developers
- For edge functionality that needs to remain centralized, the operations team should create a workflow for application developers, and support this with SLAs
- Application development teams should use these SLAs in their release planning to minimize release delays



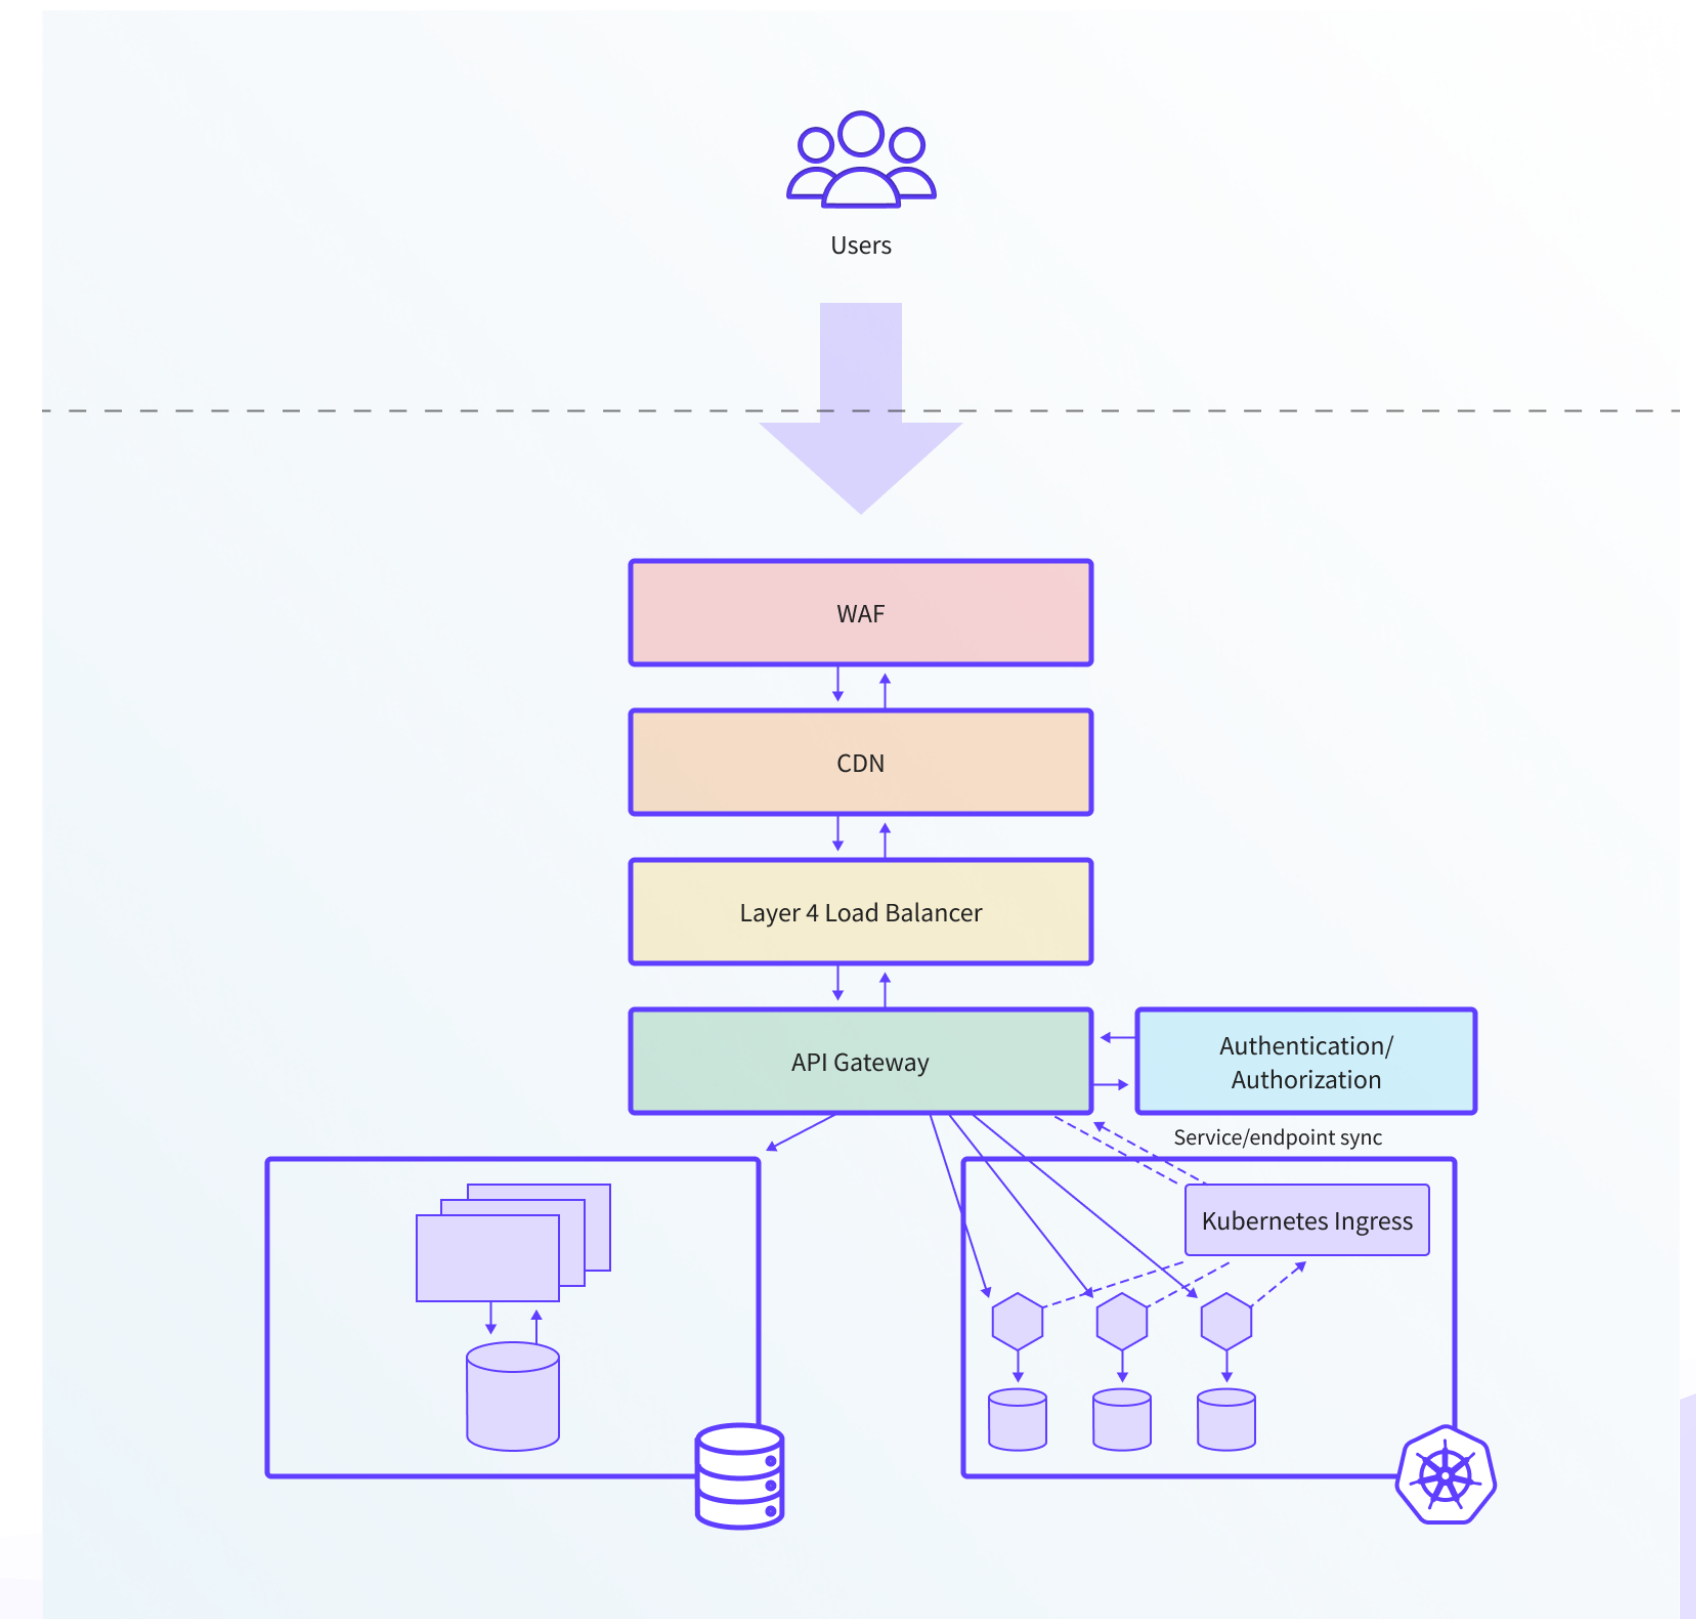
# #2 Extend Existing API Gateway

- Implemented by modifying or augmenting the existing API gateway solution
- Enable synchronization between the API endpoints and location of k8s services
- Custom ingress controller for the existing API Gateway or load balancer



# #2 Extend Existing API Gateway

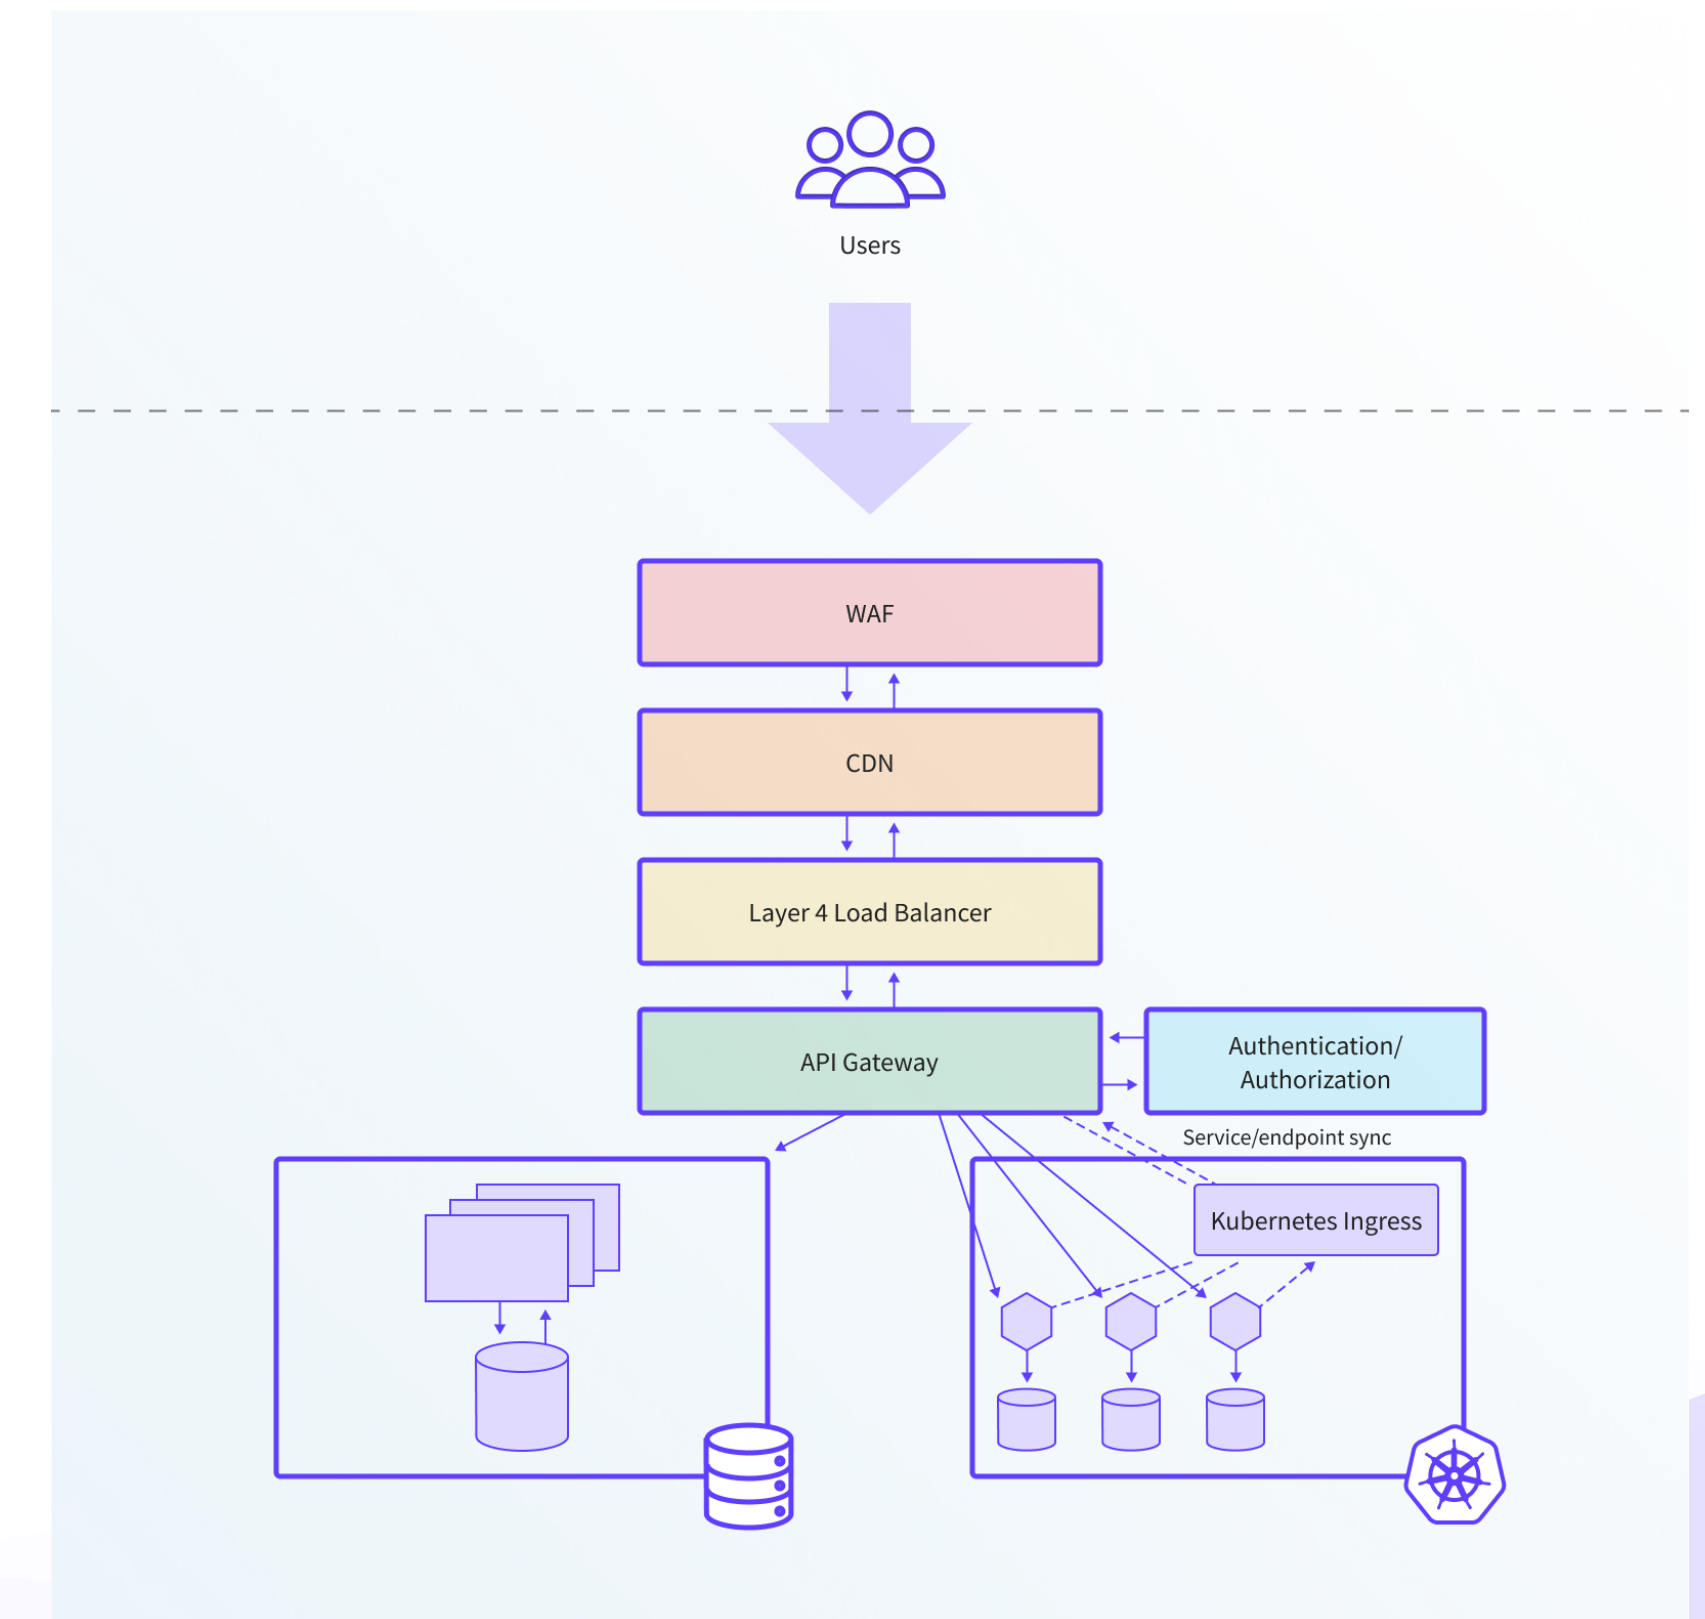
- Pros
  - Reuse the existing tried and trusted API gateway
  - Leverage existing integrations with on-premises infrastructure and services
- Cons
  - Workflows must change to preserve a single source of truth for the API gateway configuration.
  - Limited amount of configuration parameters via Kubernetes annotations





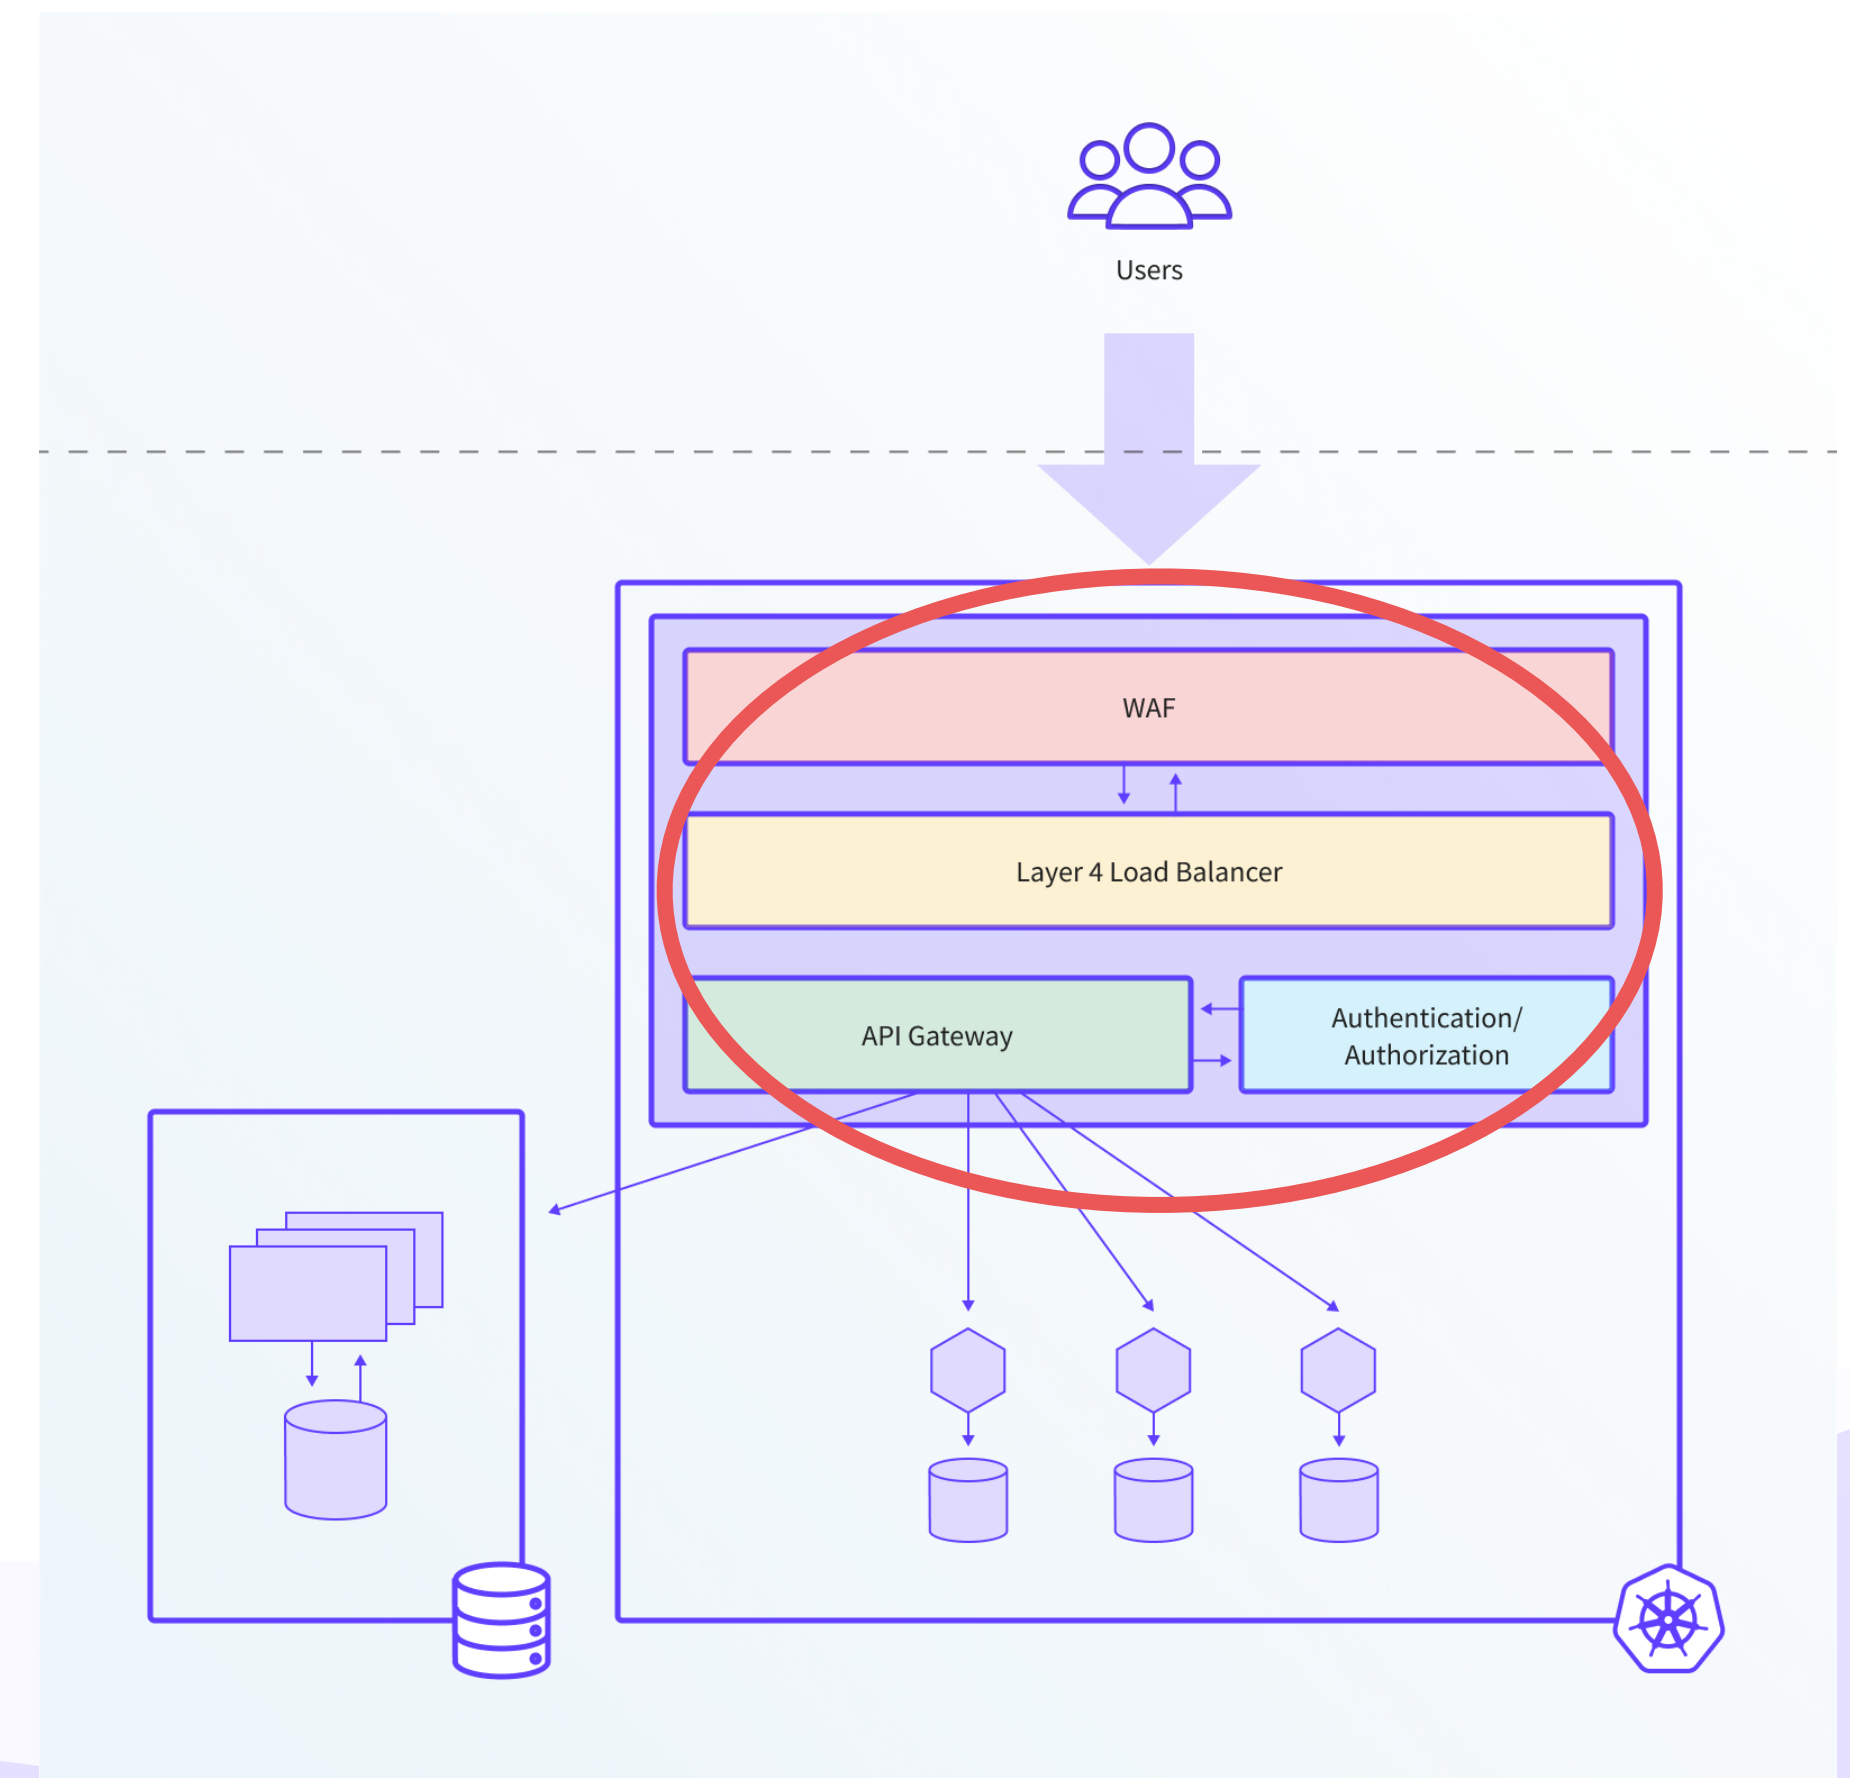
# #2 Extend Existing API Gateway

- Recommended to shift away from the traditional API/UI-driven configuration model of their existing gateway
- A standardized set of scripts should be used so any modification of routes to services running outside the Kubernetes cluster does not conflict with the services running inside the new cluster
- Before adopting the strategy, an architectural roadmap review of current and anticipated edge requirements for microservices is essential



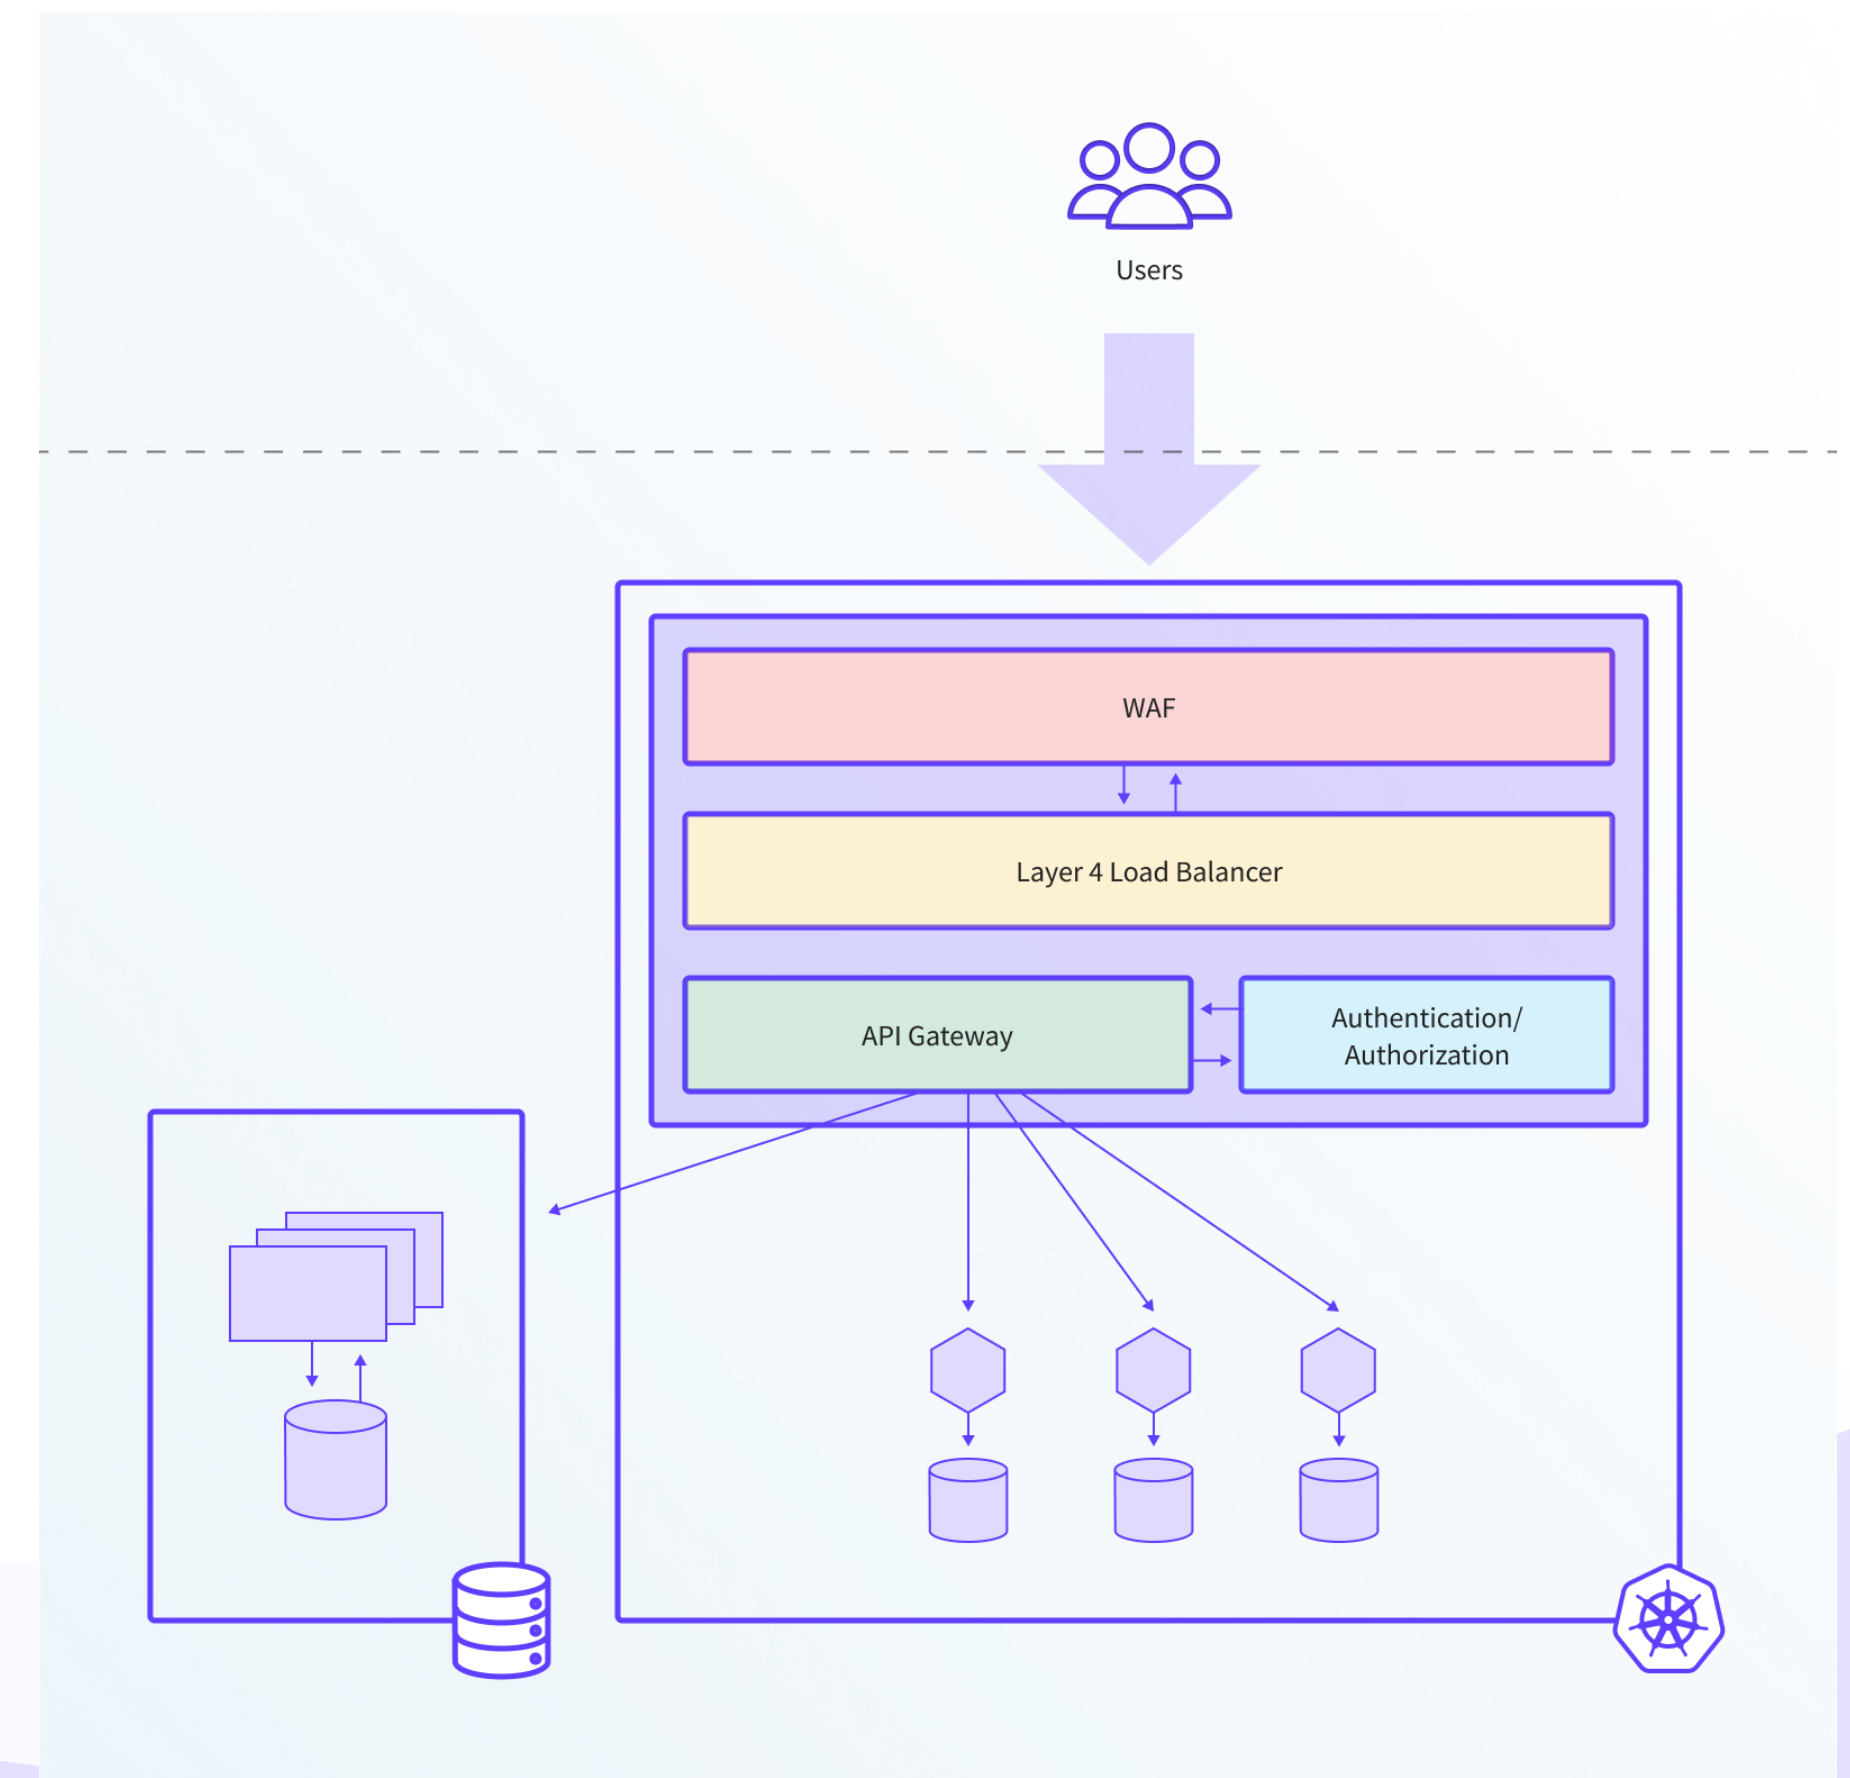
# #3 Deploy an In-Cluster Edge Stack

- Deploy Kubernetes-native API gateway with integrated supporting edge components
- Installed in each of the new Kubernetes clusters, replacing existing edge
- Ops team own, and provide sane defaults
- Dev teams responsible for configuring the edge stack as part of their normal workflow




# #3 Deploy an In-Cluster Edge Stack

- Pros
  - Edge management is simplified into a single stack
  - Supports cloud native best practices: “single source of truth”, GitOps etc
- Cons
  - Potentially a large architectural shift.
  - Platform team must learn about new proxy technologies and edge components



# #3 Deploy an In-Cluster Edge Stack

- Each microservice team is empowered to maintain the edge configuration specific to each of their microservices.
- The edge stack aggregates the distributed configuration into a single consistent configuration for the edge.
- To support the diversity of the edge services, adopt an edge stack that has been built on a modern L7 proxy with a strong community such as the Cloud Native Computing Foundation's Envoy Proxy.





# Wrapping Up

# In Conclusion

- Edge/API gateways have undergone a series of evolutions, driven by architecture
  - Hardware -> software
  - Networking Layer 4 -> Layer 7
  - Centralized management -> decentralised
- Adopting microservices/Kubernetes changes workflow
  - Scale edge management
  - Support multi-protocol and cross-functional requirements
- Chose your solution intentionally

# Many thanks!

- Learn more:
  - <https://www.getambassador.io/learn/building-kubernetes-platform/>
  - <https://www.getambassador.io/podcasts/>
  - <https://blog.getambassador.io/>
- Find me in:
  - Datawire OSS Slack: <http://d6e.co/slack>
  - Twitter [@danielbryantuk](https://twitter.com/danielbryantuk)

The collage consists of three overlapping screenshots. The top screenshot is the Ambassador API Gateway website, featuring a blue header with the logo and navigation links like 'RELEASES', 'CASE STUDIES', and 'ARCHIVE'. Below the header is a grid of article cards with titles such as 'Explore the Ambassador API Gateway with MicroK8s' and 'The Ambassador Edge Stack & Ambassador API Gateway 1.6 Now Available'. The middle screenshot is the SoundCloud profile for 'Ambassador Labs', showing a profile picture of a penguin and a banner for 'AMBASSADOR EDGE STACK'. Below the banner is a list of podcast episodes with titles like 'LOTE #14: Katie Gamanji on Kubernetes Tooling DX, GitOps, and the Cluster API'. The bottom screenshot is the Datawire logo, which is a red hexagonal pattern with the word 'DATAWIRE' in grey text to its right.