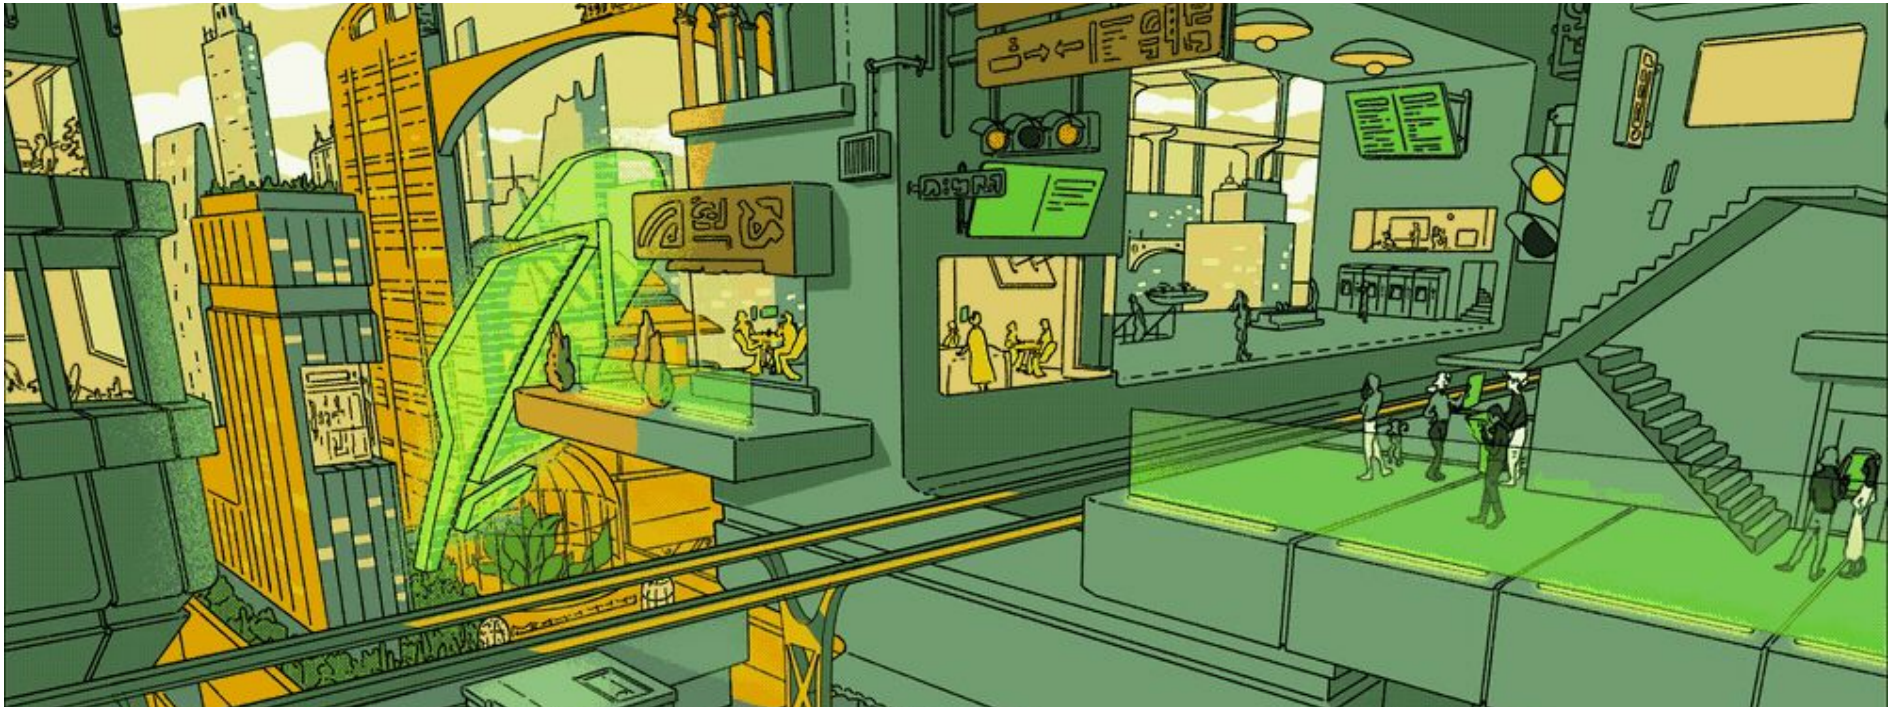
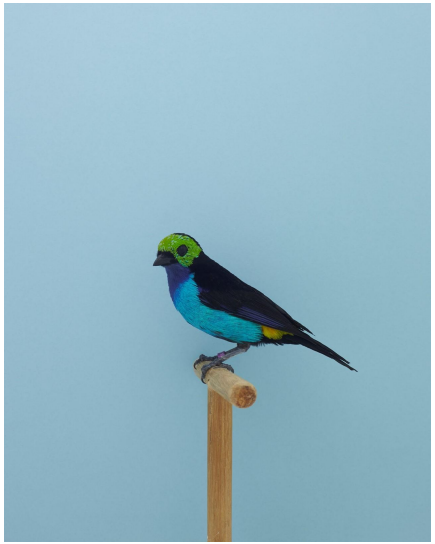


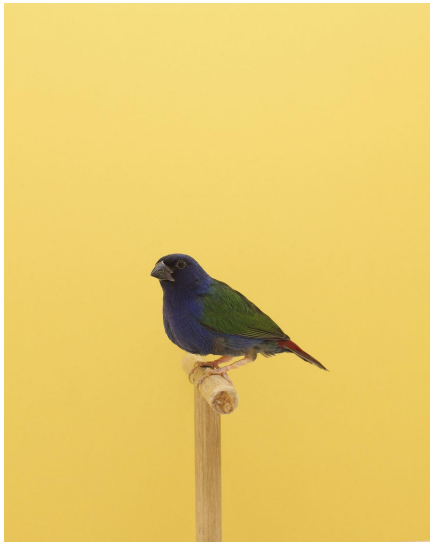
Robinhood's Use of WarpStream for Logging

How we safely migrated from Kafka to WarpStream

Ethan Chen
Renán Rueda



Ethan Chen
Software Engineer



Renán Rueda
Senior Software Engineer

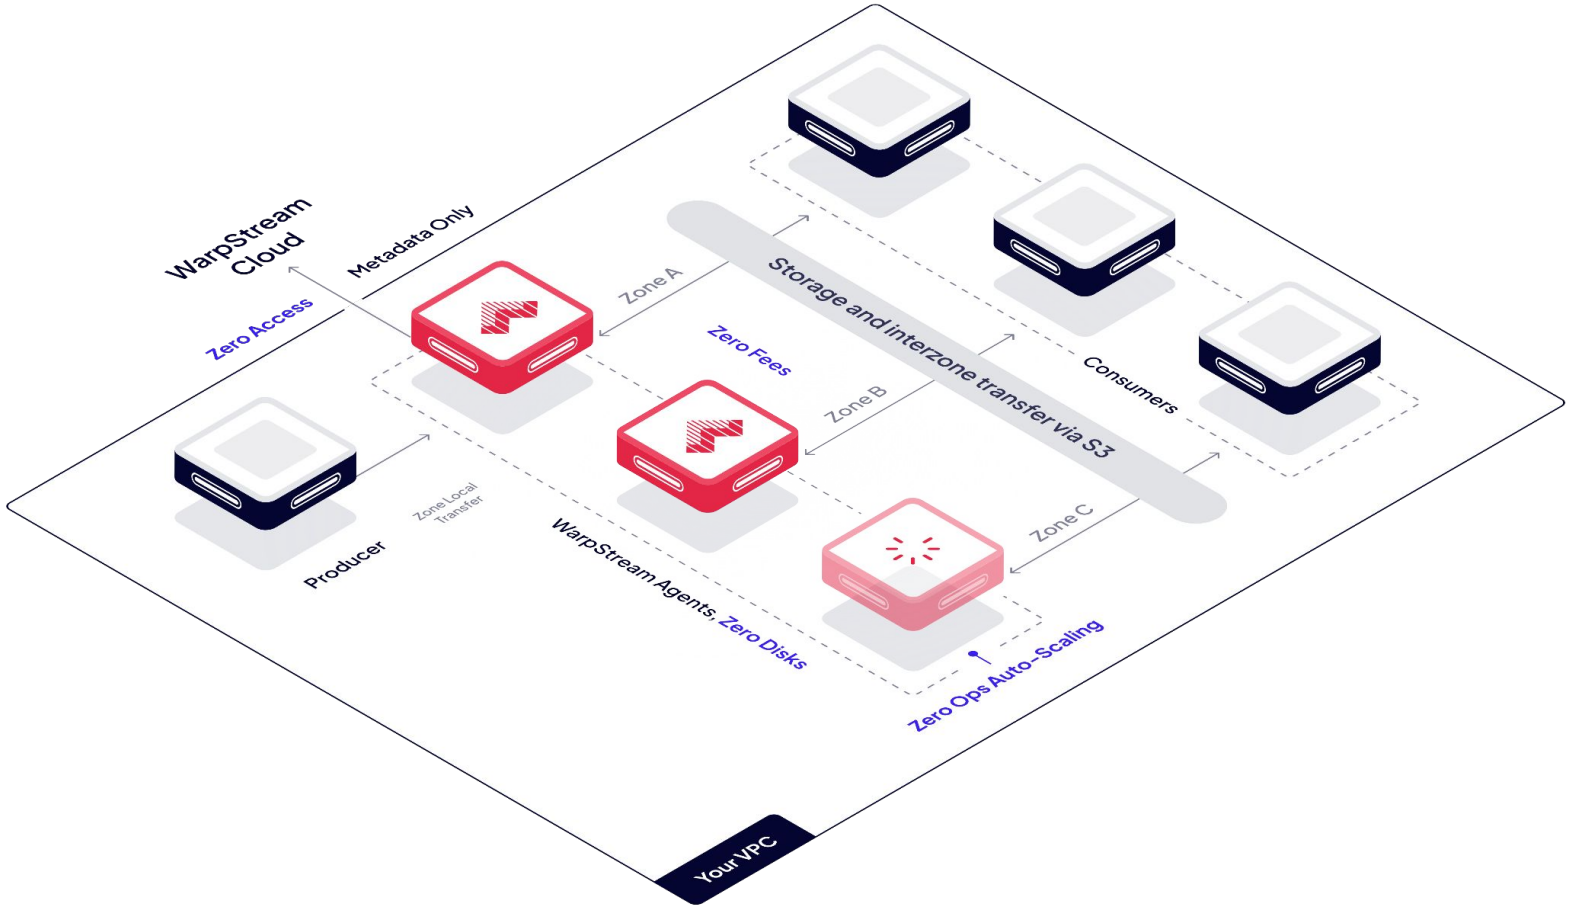
01
What is WarpStream?

What is WarpStream?

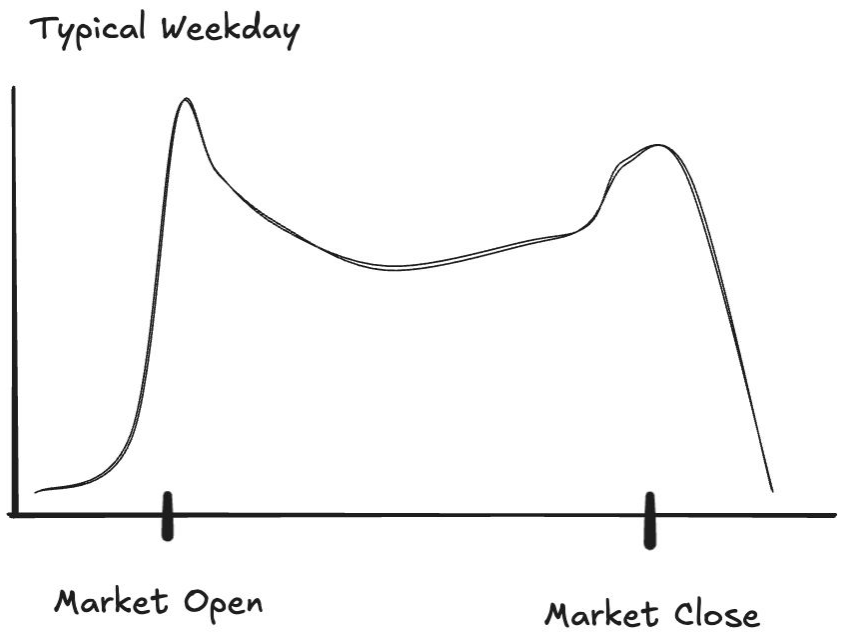


A Diskless, Kafka Compatible Data Streaming Platform

- Compatible with Kafka API.
  - If a project uses Kafka, it can probably use WarpStream.
- Clients interact with Agents.
  - Agents are managed by company using WarpStream.
- Zero Disk Architecture.
  - Can use S3 (or other object store) to persist data.
  - Potentially lower storage costs.
  - Contributes to lowering inter AZ costs.
  - Data never leaves your cloud.
- Control Plane exists in the cloud.
  - Managed by Confluent.
  - Only holds metadata.


WarpStream Cloud
Metadata Only
Zero Access
Zone A
Zero Fees
Storage and interzone transfer via S3
Consumers
Zone B
Zone C
Producer
Zone Local Transfer
WarpStream Agents, Zero Disks
Zero Ops Auto-Scaling
Your VPC

Traffic Patterns



What does logging traffic look like for us?

- Our cluster workload follows Market Hours.
  - Market hours result in higher workloads.
  - Pretty consistent pattern.
- However, traffic can be impacted by external events.
  - Load can vary throughout the day.
  - Sudden spikes are not unusual.
- Typically low traffic on nights and weekends.
- Logs follow the same traffic patterns.
  - We don’t always need the same logging capacity.
  - We should scale based upon our platform’s needs.

Reasons for Migrating



What are the pros and cons?

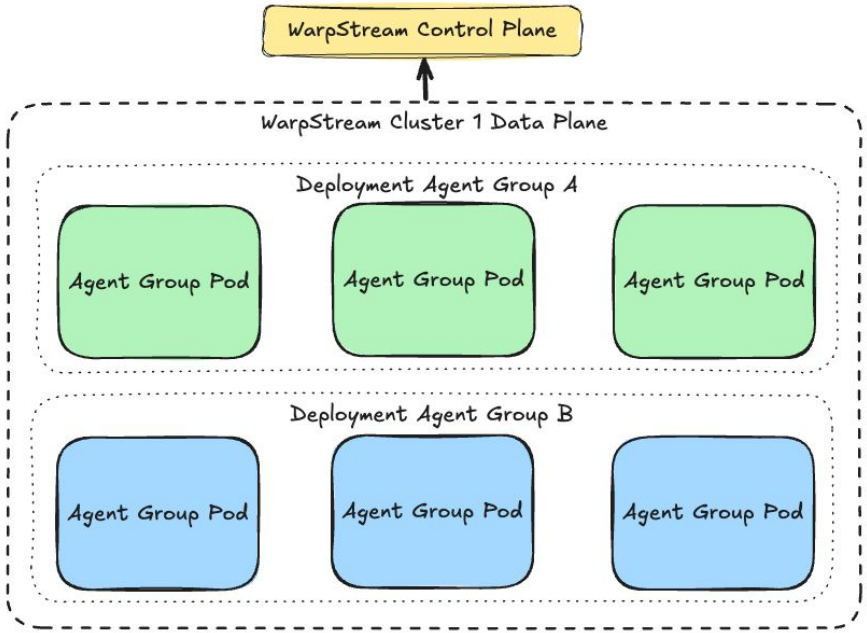
Pros:

- Cost savings.
  - Zero (or greatly reduced) inter AZ network costs.
  - Elasticity allows us to scale down during non-market hours.
- Benefits of using S3.
  - Data is replicated across AZs.
  - Costs can be lower than EBS.
  - No limit on storage capacity.
- Control Plane is managed.
  - Reduced operational complexity.
  - We only maintain Agent groups.

Cons:

- Increased latency.
  - Pick workloads that can tolerate the increase.
- Control Plane is managed.
  - “Is it us or is it the control plane?” Scenarios.
  - Support required to address issues and perform upgrades.

What's in a WarpStream Cluster



A look at the components that make up a virtual WarpStream cluster at Robinhood.

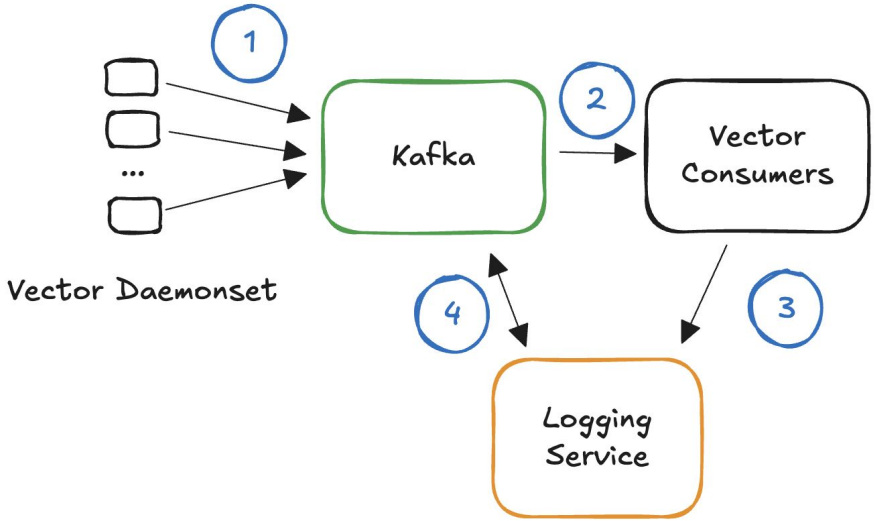
- Agents
  - Modeled as a K8s Pod.
  - Service responsible for interacting with Kafka clients.
  - One Pod per Node.

- Agent Groups
  - Modeled as a K8s Deployment.
  - Can be composed of multiple K8s deployments.
  - Clients can be assigned to a particular group.

- Common examples of agent groups:
  - SSL/TLS
  - Plaintext
  - Dedicated for noisy or important client.

02
Migrating Logging to WarpStream

Current Logging Setup (Kafka)



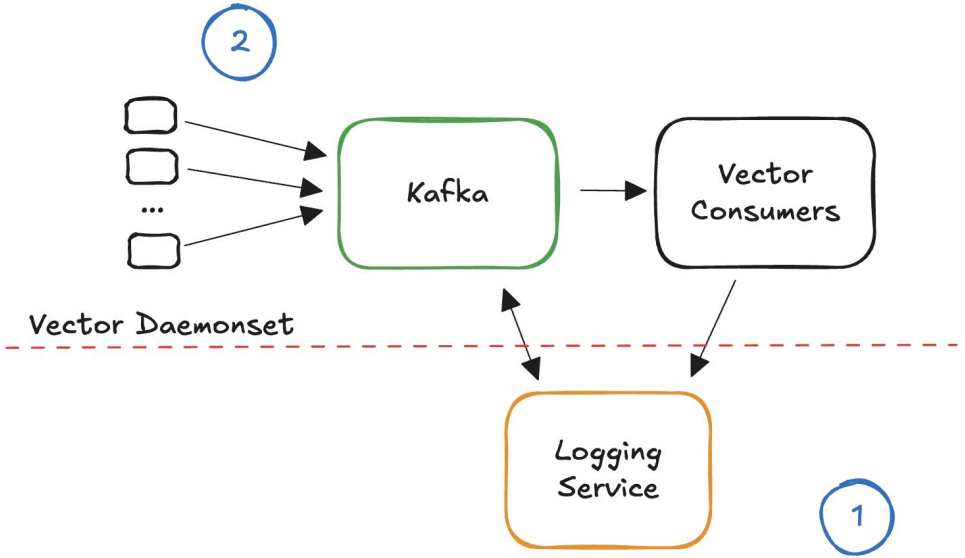
Background:

- Vector is a tool for building observability pipelines

Steps

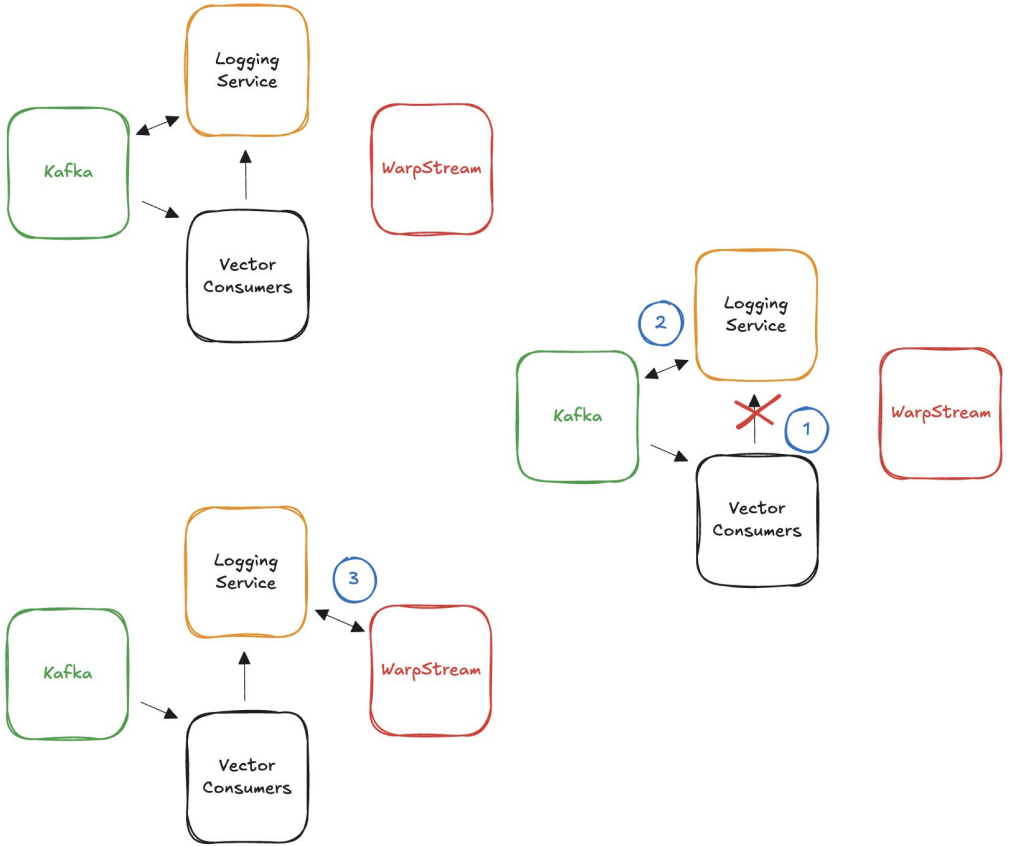
1. Produce logs to Kafka
2. Consume logs from kafka
3. Ship logs to Logging Service
4. Logging application uses Kafka as backend

How can we migrate?



Break it up

1. **Logging Service**
2. **Vector Log Shipping**



Log Application Migration

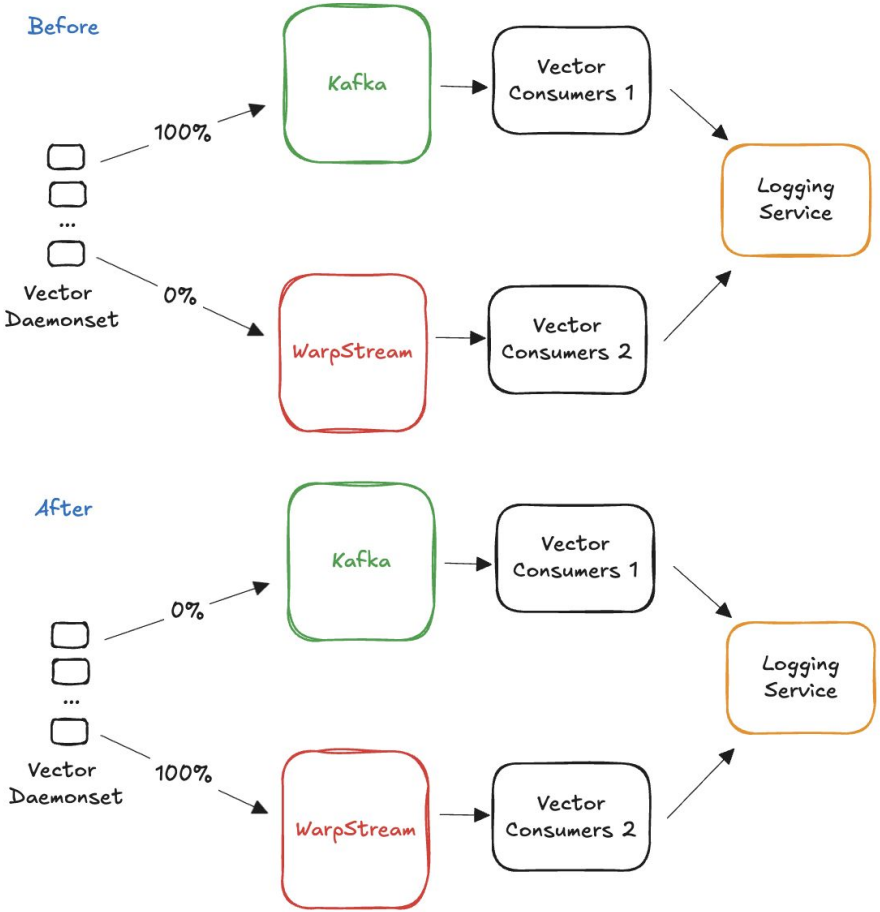
Background:

- Our logging service's kafka setup is "all or nothing"

Goal: Swap Kafka out for WarpStream with as little impact as possible

Strategy:

1. Shut off Vector ingestion and buffer logs in Kafka.
2. Wait until our logging application has finished processing.
3. Switchover to WarpStream and unblock log ingestion



Vector Migration

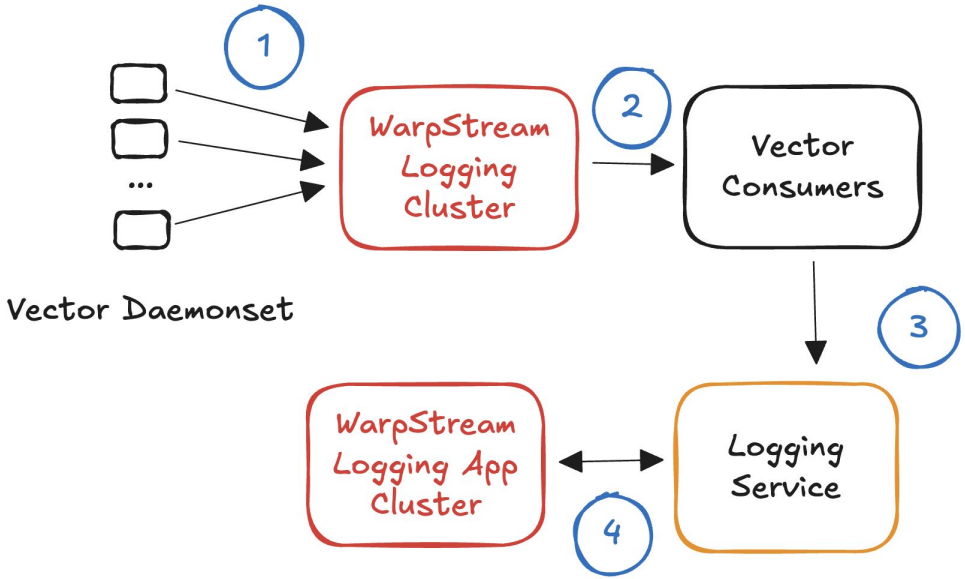
Goals:

- No duplicate logs.
- No impact on log viewing experience.
- Minimized disruption.
- Gradual rollout and rollback.

Strategy:

1. Temporarily duplicate our vector consumers
2. Gradually point the log producers to WarpStream.

Current Logging Setup (WarpStream)



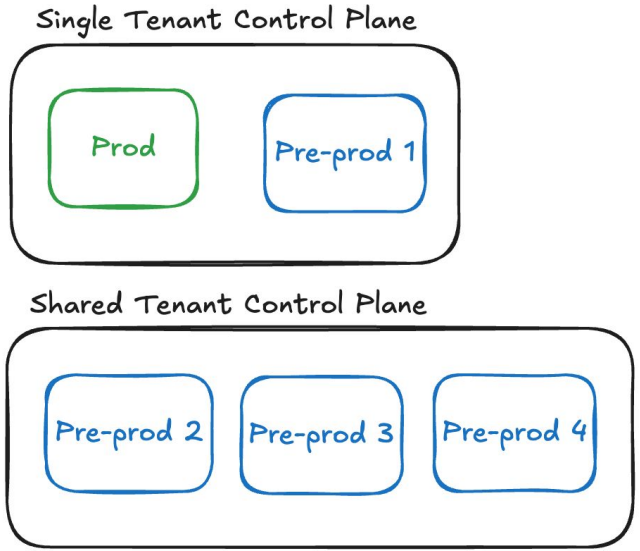
New

- Monolith cluster split into separate clusters

Steps

1. Produce logs to WarpStream
2. Consume logs from WarpStream
3. Ship logs to Logging Application
4. Logging application uses Kafka as backend

Control Plane



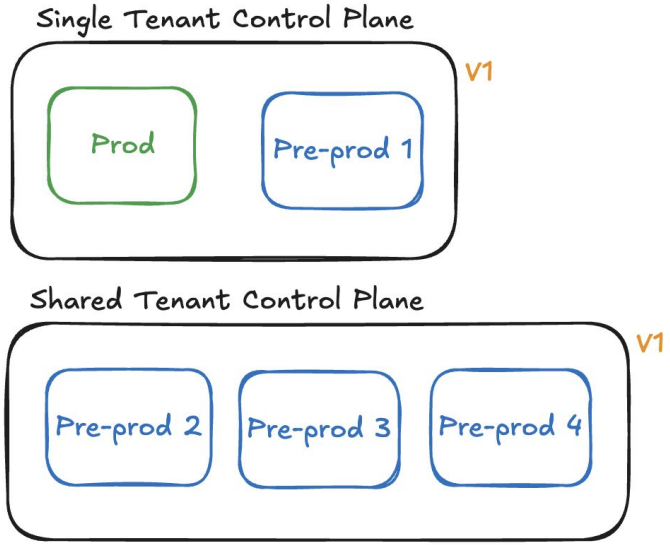
WarpStream Control Plane Setup

- Manages metadata for WarpStream
- Shared Tenant Control Plane
  - Regional, multi-tenant service

Change Management

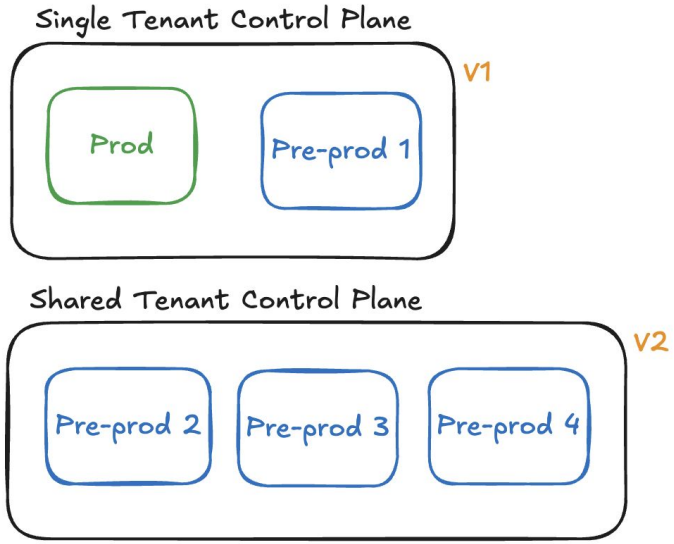
- Shared Control Plane Changes
  - All changes go out at once.
- Dedicated Control Plane Changes
  - Lagging version updates.
  - Maintenance windows.

Control Plane Upgrade Example



Base State

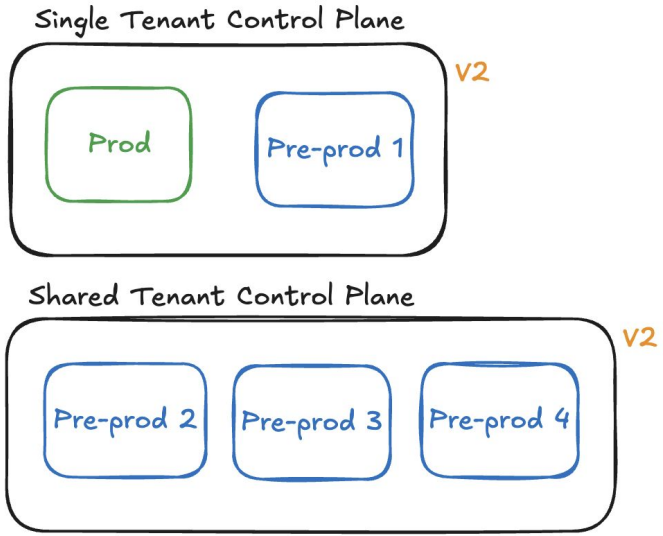
Control Plane Upgrade Example



V2 Hits Shared Tenant Control Plane

- Let this state bake for X amount of time

Control Plane Upgrade Example



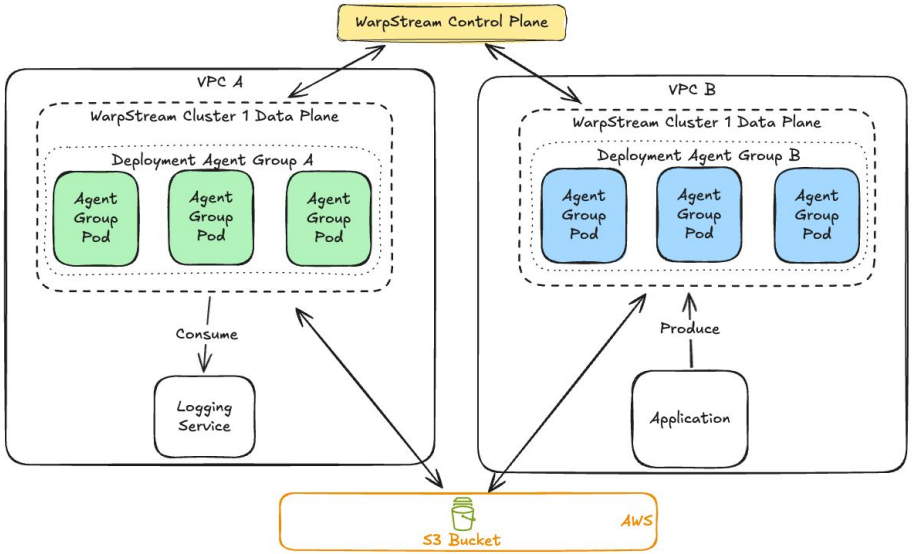
Single Tenant Upgrade

- After successful baketime, rollout the upgrade to production.

Why have prod and pre-prod in the single tenant?

- Mirror environment to prod

Cross AWS Account Logging



One WarpStream Cluster, multiple accounts

- A WarpStream cluster can have Agent groups in different AWS accounts.

- Storage can be centralized.
  - Allows us to provide logging infrastructure easily across AWS Account/VPCs.
  - No need to set up Private Link.

03
Challenges

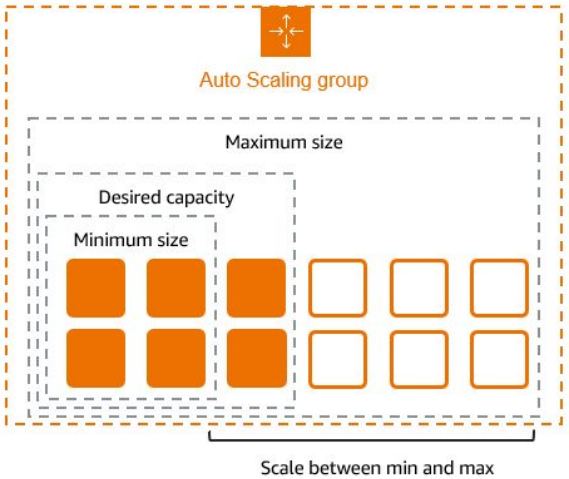
Finding the right components



And fine-tuning our compute

- One Pod per Node.
- Choosing an instance type.
- Bigger Instances vs Smaller Instances.
- Fewer Agents vs More Agents.
- Tuning HPA Policies.

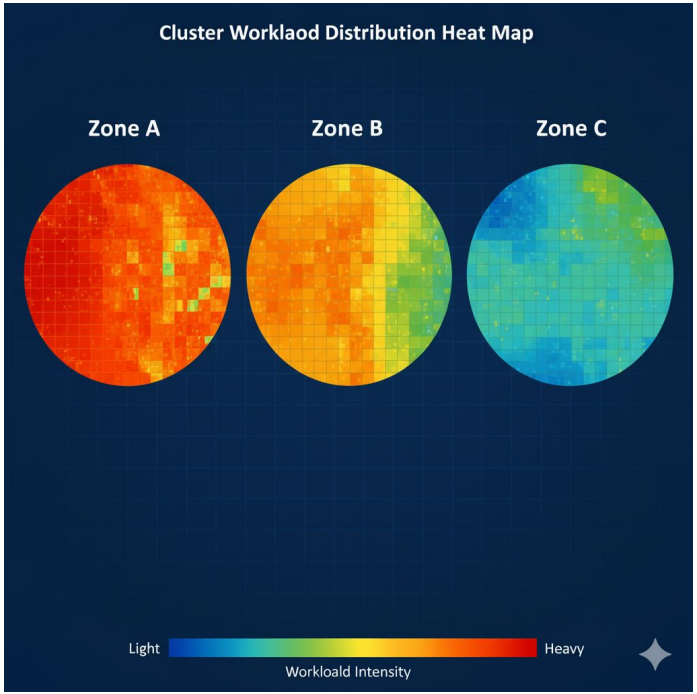
Auto Scaling



Using just the right amount of compute

- Scaling based on CPU.
- Fast scale ups.
- Slow, graceful scale downs.
  - Fast scale downs result in latency spikes.
  - Graceful shutdowns allow clients to move away.

Auto Scaling per AZ

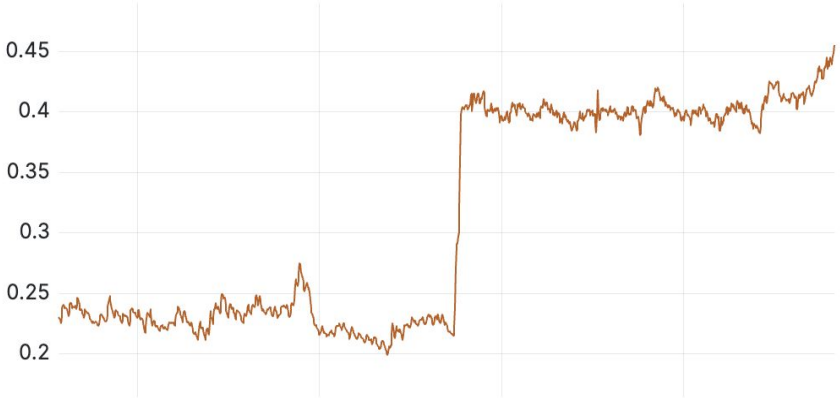


Matching capacity where the workloads need it

- WarpStream supports AZ awareness.
  - `client.id=foo.ws_az=zone-a`
  - Different AZs might have different workloads.

- WarpStream Agents should match the AZ footprint.
  - One k8s deployment per AZ.
  - One HPA per deployment to scale independently.

Latency Tradeoffs



Avg produce latency before and after batch tuning (seconds)

Tweaking batch settings

- Larger batch sizes –> Less S3 requests.
- Less S3 requests –> S3 savings $$$.

Tradeoff

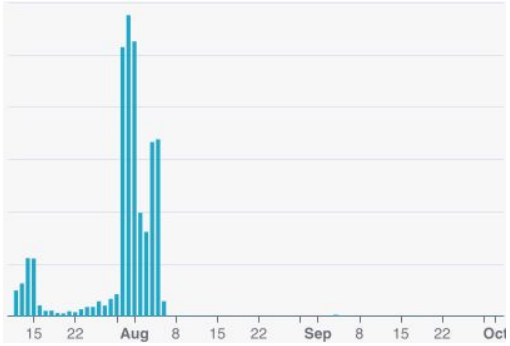
- Larger batch sizes –> Increased latency waiting for batches to fill up.

Other Options

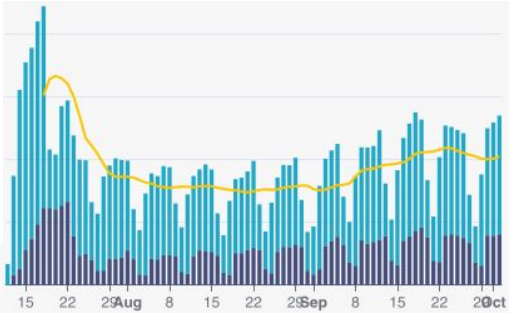
- If lower latency is required, WarpStream also provides compatibility with S3 Express.
- Disclaimer: We have not used this ourselves.

03
Outcomes

Cost Efficiency Over Time



Networking (Inter AZ)



Storage



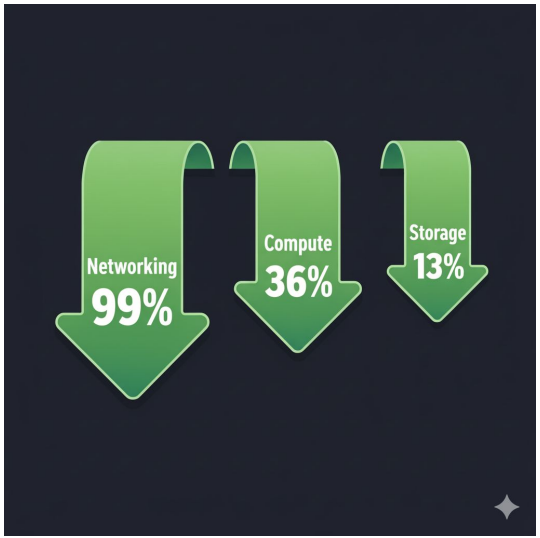
Compute

How did we do?

- We took the peak 7 Day Average after we started for Compute, Storage, and Networking costs.

- In our most recent month vs. the starting month:
  - Networking (Inter-AZ)
    - Down by 99%.
  - Compute
    - Down by 47%.
  - Storage
    - Down 34%.

Cost efficiency vs. Kafka



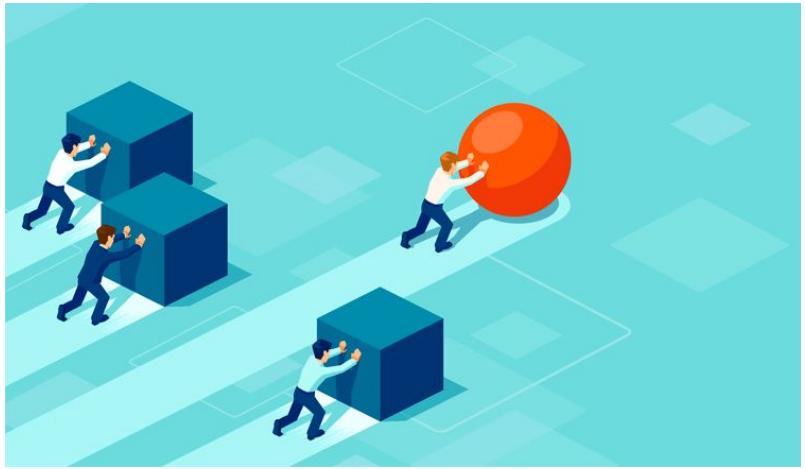
How did we do compared Kafka?

- Peak 7 Day Moving Average cost vs Kafka:
  - Networking (Inter-AZ)
    - 99% cheaper.
  - Compute
    - 36% cheaper.
  - Storage
    - 13% cheaper.

- Weekend savings are a big factor.
  - Kafka cluster is static.
  - WarpStream is elastic.

- These are cost comparisons.
  - Actual savings are realized by scaling down Kafka clusters.

Operational efficiency



How did we do?

- Added stability with ability to scale.
  - No more increasing EBS volumes.
- Reduced time spent maintaining Kafka brokers.
  - No more having to deal with URPs.
- Automation saves us time.
  - Create new clusters in hours.
- Reduced on-call burden.
  - We have to do less to keep our services in a healthy state.

We didn’t do this alone



Thanks to:

- Mun Yong Jang
- Nitesh Dhanpal
- Vi Tran
- Tony Chen
- Vikrant Rajput
- Kai Zhu
- Charlie Pisuraj
- Sreeram Ramji
- Sujith Katakam
- As well as the WarpStream folks from Confluent!

Thank you