

How **Asurion** Scaled and Simplified Airflow for Workload Orchestration

Rajesh Gundugollu
Karthik Raparathi



 **Airflow Summit**

Let's flow together

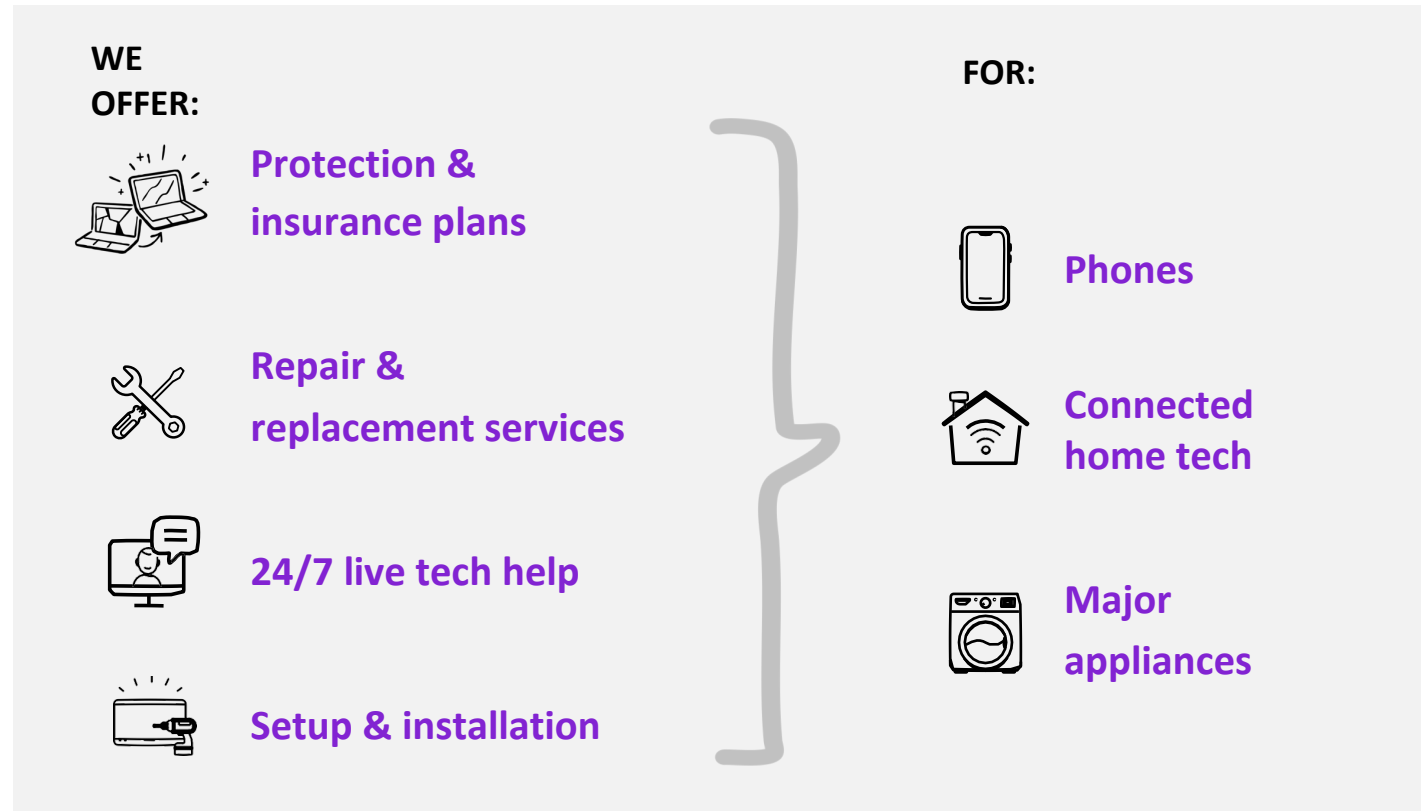
September 19-21, 2023,
Toronto, Canada

asurion

is a **tech care** company

We eliminate the fears and frustrations associated with technology, to ensure our 286 million customers get the most out of their devices, appliances and connections.

asurion



Our purpose

We care for
people and their
tech – whenever,
wherever, however.

PEOPLE

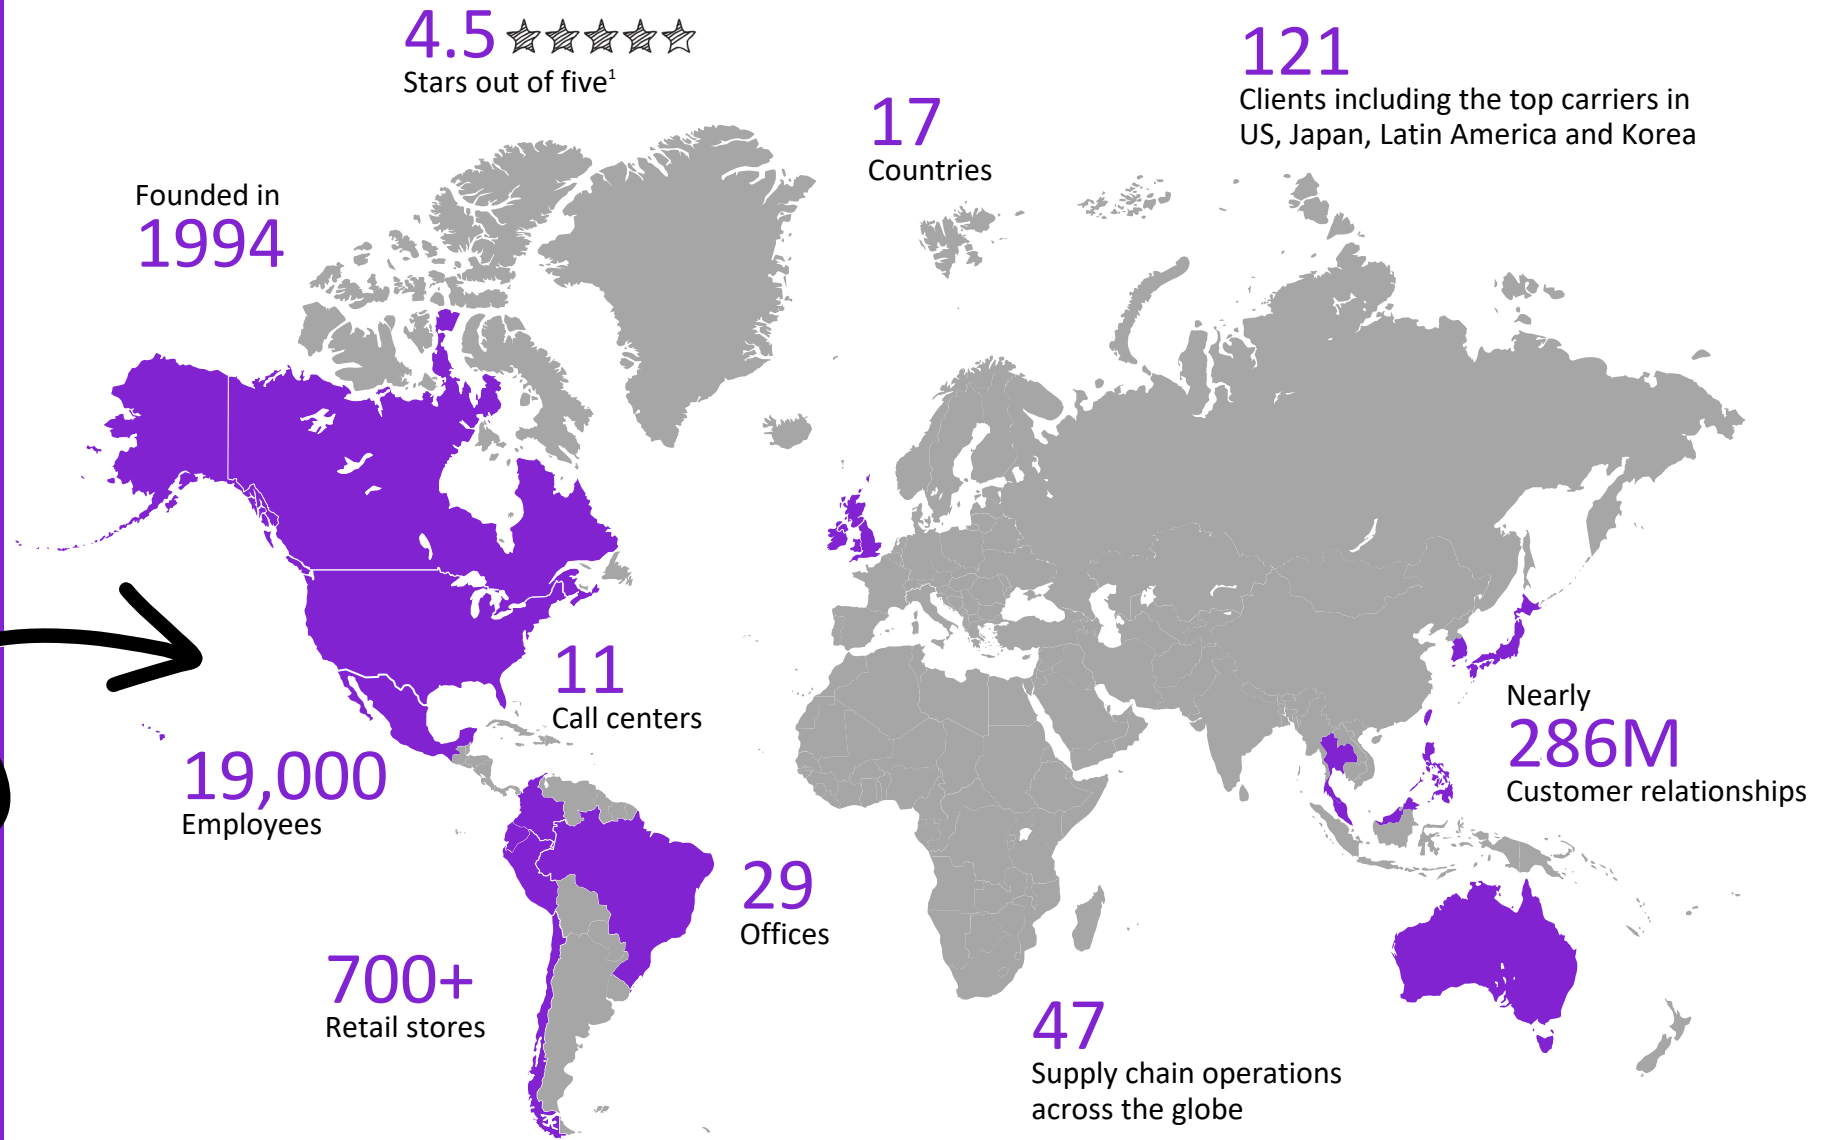
We support close to 300 million people all over the world

TECH

We care for thousands of tech devices, including your favorite phones, laptops, tablets, TVs and gaming consoles

asurion

We operate
on a global
scale



¹ based on 17.8M Asurion customer reviews

How we help our customers:

10,000+
experts



We've created the largest tech care network:

- ✔ **Virtual**
Call or chat with us 24/7
- ✔ **In-Store**
700+ stores nationwide
- ✔ **On the road**
We make house calls for 88% of US Population

asurion



We have built long-term partnerships

Mobile | Telco | Cable



asurion

Retail | OEM



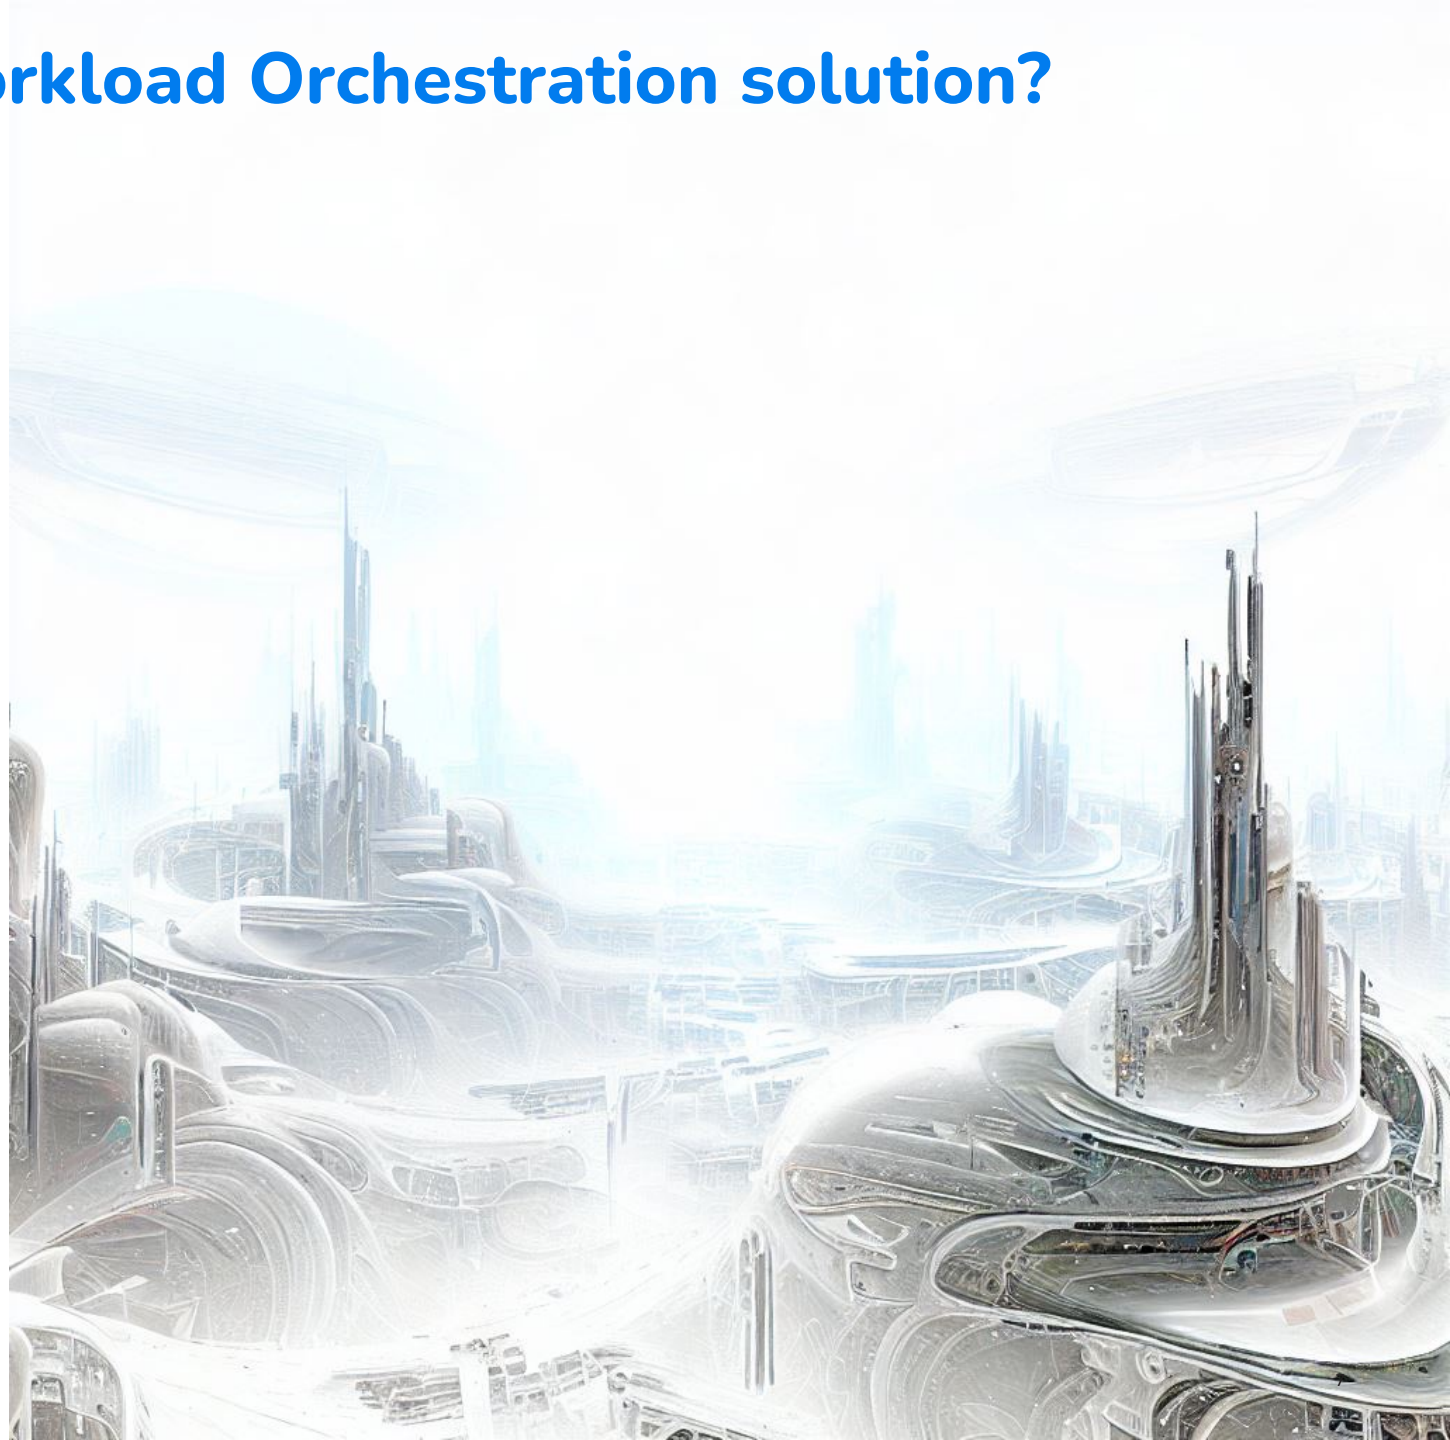
Why do we need a robust Workload Orchestration solution?

A Data Driven Future

Workload Orchestration is at the heart of a Data Platform Founded on Data.

Gravity of the Problem

1000s of Job runs that process **100s of Billions** of rows per day with **1,000s** of complex dependencies.



What's our Solution?

Orca

Its Asurion's flavor of Apache Airflow that's at the heart of our Data Platform.

With unparalleled reliability, accuracy, stability, and resiliency, **Orca** is the linchpin of our data-driven future—a future that embodies our vision.

DAGs

1000s

Job Runs

Millions

Availability

99.5%



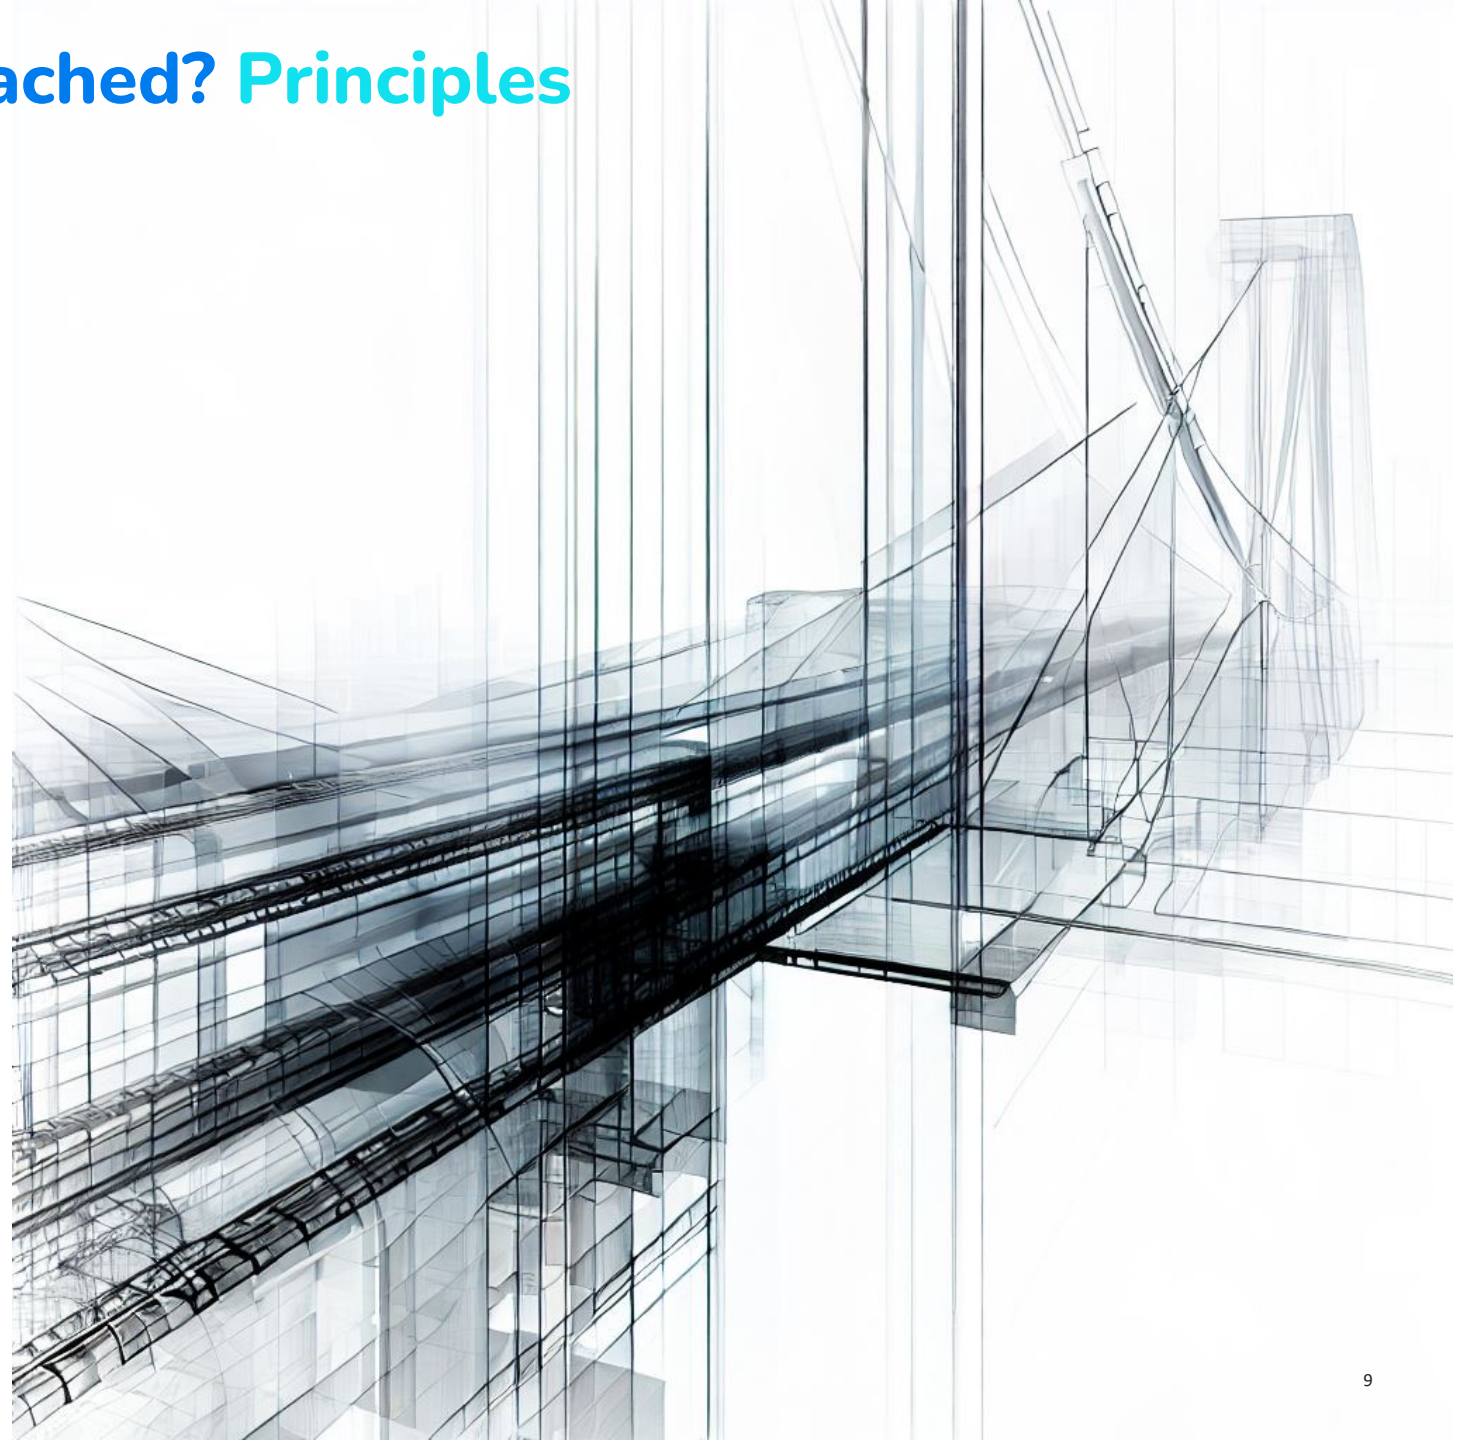
How was the solution approached? Principles

Founding Principles

Separation of Concerns
Metadata First

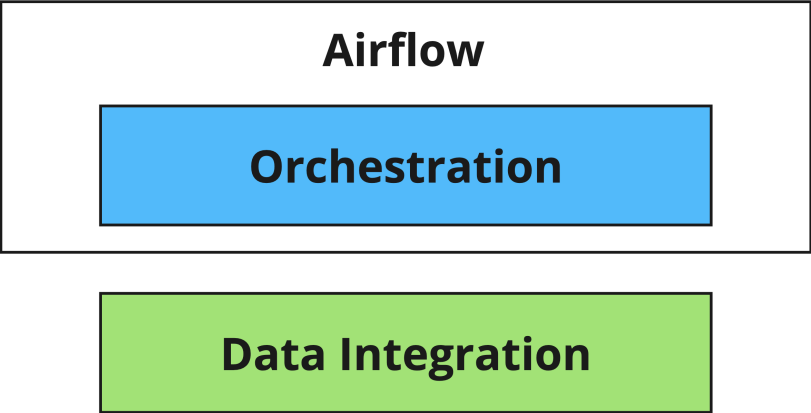
Evolved Principles

Universal Operator
Shared nothing Infrastructure
Metadata Rich
Automation

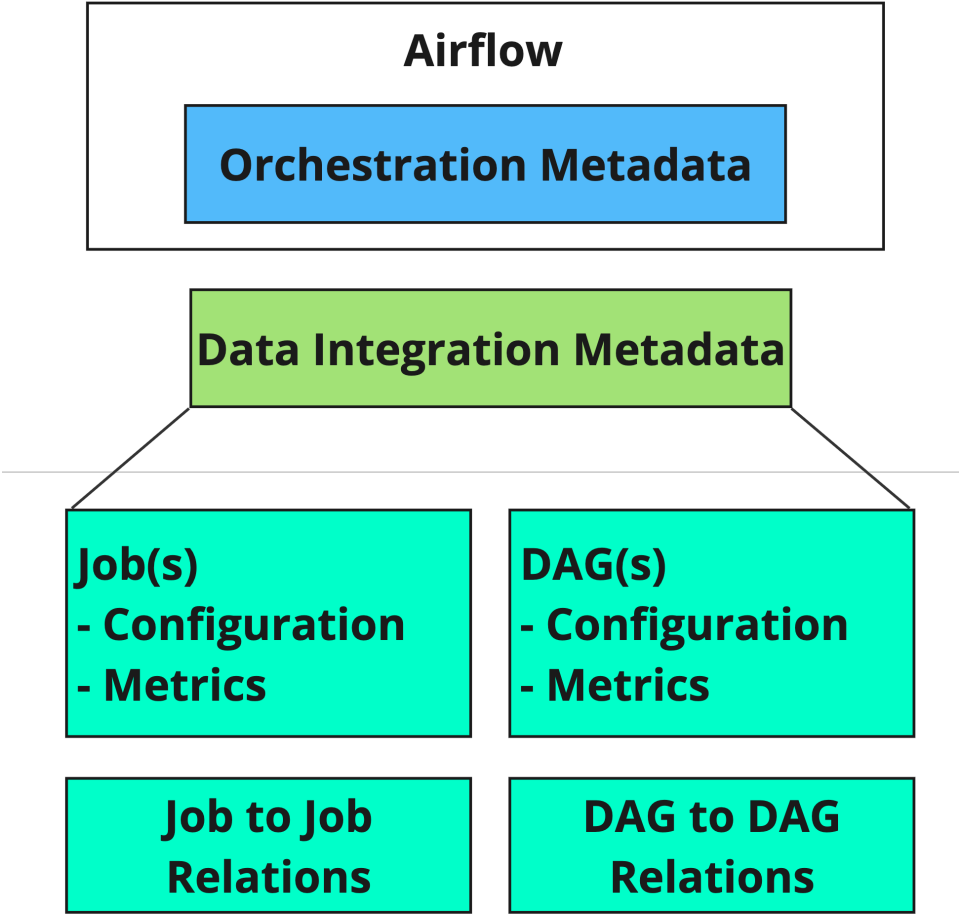


How was the solution approached? Founding Principles

Separation of Concerns



Metadata First



How was the solution approached? Experiment

Stress Testing over 250 Days

DAGs
1000

Peak Concurrency
700

Job Runs
5.2M

Job Runs per Day
22K

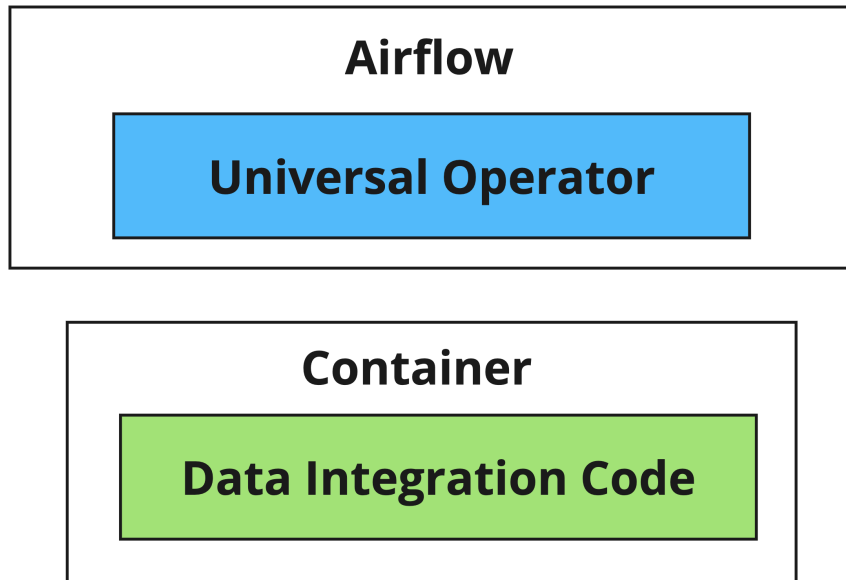
Results

Sizing of Worker nodes/concurrency
Auto scaling metrics

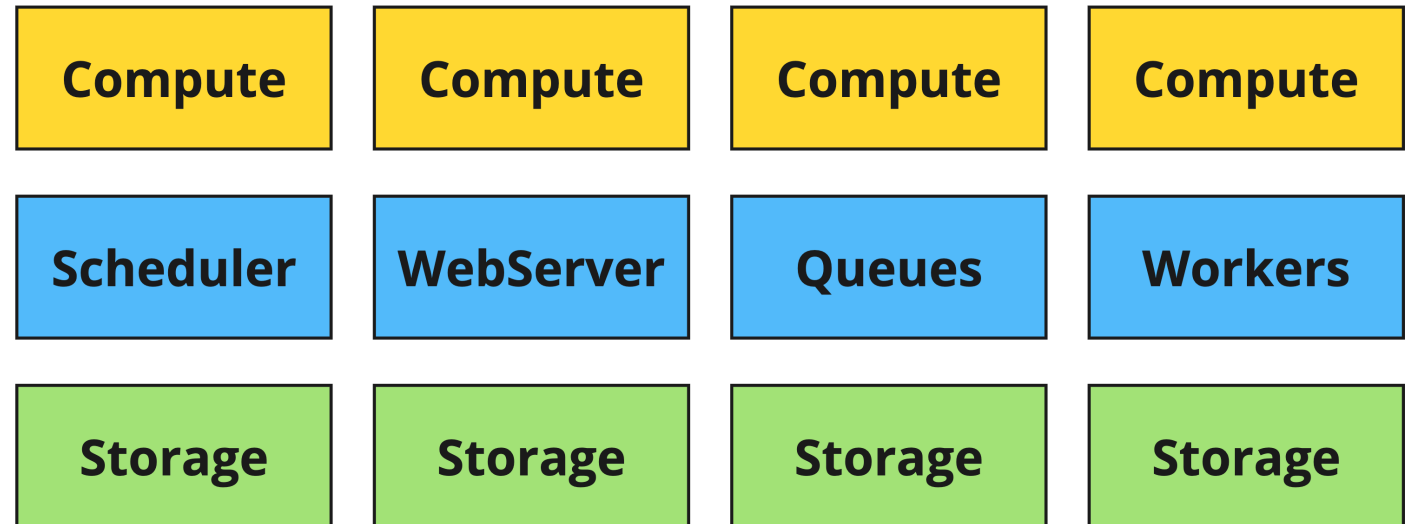


How was the solution approached? Evolved Principles

Single Operator

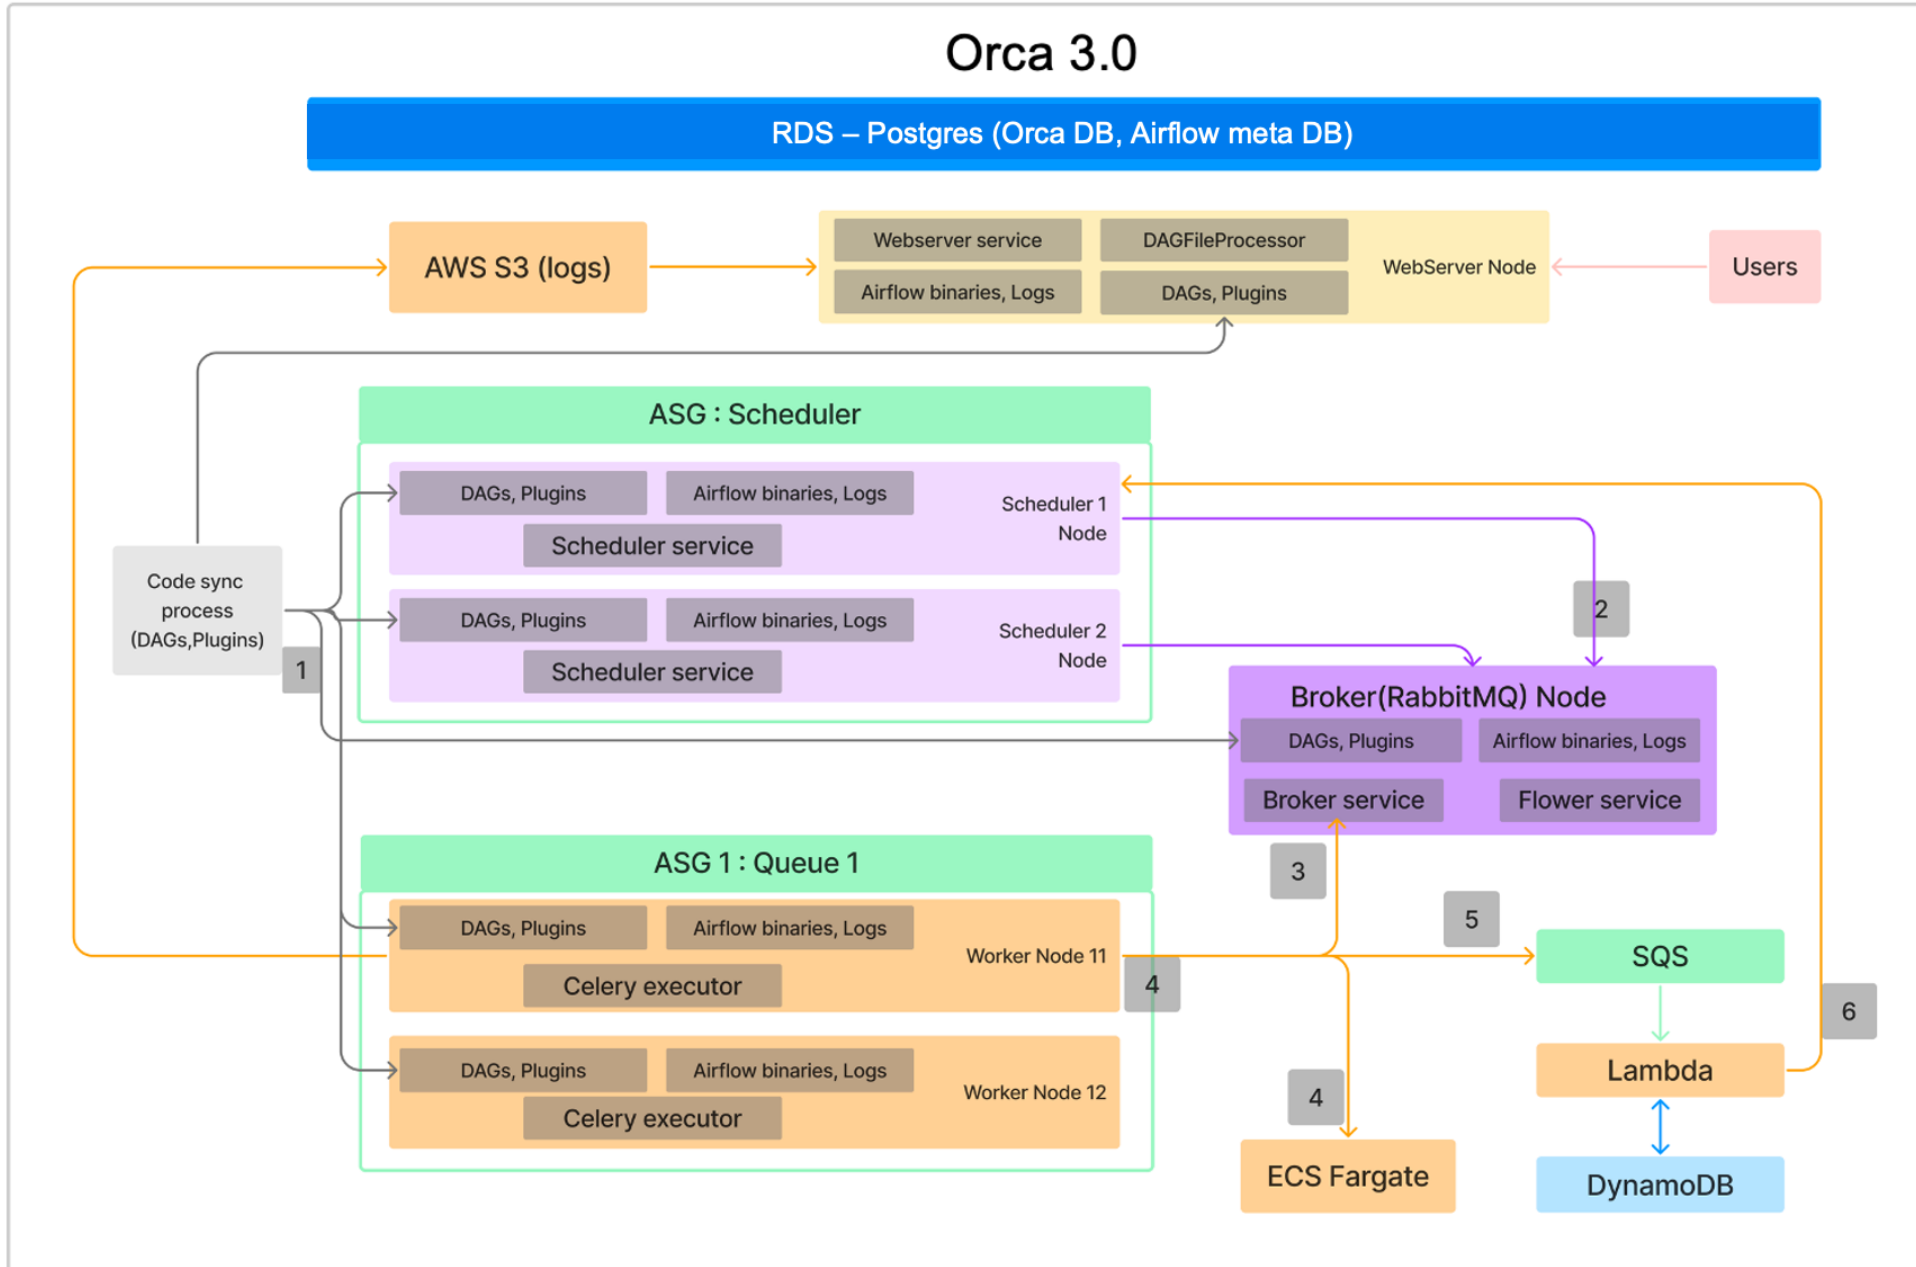


Shared Nothing Infrastructure



“Simple is Powerful”

How was the solution approached? Architecture



Resiliency?

Metadata Rich

Comprehensive Metadata

Real-time Metadata Events

Automation

Code Generation

Deployments

Monitoring and Recovery

Incident Management



User Accelerators

- 1- Metadata Validations
- 2- Metadata Registration

- 3- Intelligent Orchestration code embeds
- 4- DAG File creation
- 5- Event Generation

The screenshot displays the Airflow web interface for DAG deployment. At the top, the navigation bar includes the Airflow logo, 'DAGs', 'Security', 'Browse', 'Admin', 'Docs', and 'Orca Menu'. The current time is 00:06 EDT (-04:00) and the user is 'RG'. The main heading is 'DAG Deploy' with a search bar on the right. Below this, there is a search input field containing 'Search for Dag...', a dropdown menu set to 'dag', a date-time selector for 'Tue, 11 Apr 2023 18:27:05 UTC', and buttons for 'Go!', 'Download', and 'Clear'. Further down, there are input fields for 'Dag:', 'Type:', and 'Version:' (set to '2'). There is a 'Choose File' button with the text 'No file chosen', a 'Publish' button with a checkbox, and a blue 'Visualise' button. The main area shows a DAG visualization with nodes such as 'customer_case', 'interaction_line', 'service_request', 'activity', and 'customer_case' connected by arrows.

User Accelerators

[DAGs](#)[Security](#)[Browse](#)[Admin](#)[Docs](#)[Orca Menu](#)

00:10 EDT (-04:00)

RG

DAG Edit

Range of Utilities, Command Center Operations to use while operating Airflow or Orca Menu.

Edit Load Pointers

Metadata Profile

Disable Tasks

Disable Pods

Edit Metadata

Disable Dag Notification

Update Parameters

Download Dag Image

Delete Dag

Sync Dag Metadata

Create Dag file

User Accelerators



DAGs

Security ▾

Browse ▾

Admin ▾

Docs ▾

Orca Menu ▾

00:11 EDT (-04:00) ▾

RG ▾

Dag Overview

Dag Version

Latest Dag version published date

*author
published_by*

Dag Data Pointer

Data loaded up till

Dag Flow

Upstream

upstream dags

Downstream

downstream dags

User Accelerators

[DAGs](#)[Security](#)[Browse](#)[Admin](#)[Docs](#)[Orca Menu](#)

00:13 EDT (-04:00)





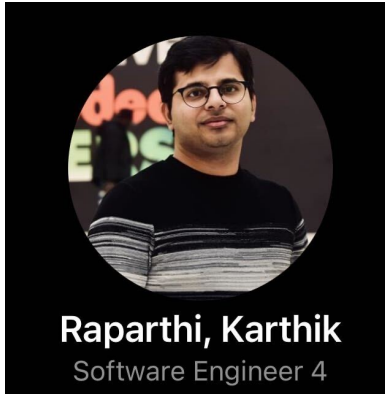
RG

Queue Information

updated: 2023-04-18 04:11:50 UTC

Cluster	Running	Waiting
Queue1	4	0
Queue2	0	0
Queue3	0	0
Queue4	0	0
Queue5	0	0
Queue7	0	0
Queue8	1	0
Queue9	0	0

The Team



Summary : Principles and Lessons Learned

Separation of Concerns


Metadata First

Universal Operator

Shared Nothing Infrastructure

Metadata Rich

Automation



Questions?

Rajesh.Gundugollu@Asurion.com
Karthik.Raparthi@Asurion.com