

Scaling Airflow for Data Productivity at Instacart

Anant Agarwal

Software Engineer, Instacart



About Me

- Been with Instacart ~4 years
- Worked across different platform teams building systems, frameworks and libraries along the way.
- Live in SF Bay Area
- <https://www.linkedin.com/in/anantag/>



Agenda

- **Part 1 – Instacart’s Airflow Journey**
- **Part 2 – How we use Airflow at Instacart**
- **Part 3 – Key Takeaways**



Instacart's Airflow Journey



State of Airflow (earlier)

- A combination of legacy Airflow clusters, used by multiple teams.
- No central team to manage them!
- Older versions
- Limited visibility and monitoring
- Scalability issues.



State of Airflow (now)

1. One Central Airflow Cluster
 - a. Managed by the Data Infra team
 - b. Airflow 2.7.3 with upgrades every 6 months
 - c. Infra as Code!
 - d. Support for multiple use-cases and remote executors.
 - e. High auditability with compliance controls.
2. All new workloads and use-cases onboard to the Central cluster.
3. Migration of workloads on legacy cluster - one set of workloads every quarter.



State of Airflow (future)

1. One Central Airflow cluster across the company.
2. New clusters only on a case by case basis.
3. Legacy clusters removed and all workloads migrated.



Use Cases powered by Central Airflow at Instacart



dbt pipelines



Non-dbt
pipelines



Snowflake <>
Delta
pipelines



GSheets, Jira,
Incidents etc.
data
movement

How we use Airflow at Instacart




Central Airflow Usage by the numbers

- ~3 years in production
- ~2200 DAGs
- ~16M tasks/month
- <3 critical incidents over the past 2 years
- Snowflake, Databricks, dbt, Custom ECS Operator
- Airflow version upgrades every 6 months.
- 99.5% task success rate
- Separate Production and Development environments



Airflow Worker Deployment at Instacart



Abstract away DAG Configurations

```
1  owner: anant.agarwal
2  start_date: 2022-01-04
3  schedule_interval: 0 * * * *
4  catchup: false
5  description: Example job that fail due to failed pretest condition
6  totem_system: data-productivity-data-pipelines
7  notify_on_failure:
8    slack:
9      team_channel: bot-data-infra-productivity-non-critical
10     description: Testing example that will fail please pause after testing.
11     pause_dag_on_error: true
12  deploy_env: local, dev, prod
13  models:
14    - name: my_dbt_fail_pretest_model
15      sensors:
16        - sensor_always_pass
17      custom_dependencies:
18        - name: my_dbt_fail_pretest_model
19          depends_on:
20            - test_always_fail
21
```

Abstract away DAG Configurations

```
! test_gsheet.pipeline.yml ×  
1  owner: anant.agarwal  
2  start_date: 2023-06-20  
3  schedule_interval: "30 * * * *"  
4  description: A pipeline for testing the gsheet to sf dag  
5  deploy_env: local, prod  
6  notify_on_failure:  
7    slack:  
8      team_channel: bot-data-infra-productivity-critical  
9  sheets:  
10 - Sheet1  
11 table_prefix: gs  
12 gsheet_url: "https://docs.google.com/spreadsheets/d/1jmdyVN5XxERAhTLfyk3I0-r6ipqLsRpAzGmHgTlOMXA/edit#gid=0"  
13
```


Abstract away DAG Configurations

```
! ingest_jira_to_snowflake.pipeline.yml ×
1  owner: anant.agarwal|
2  start_date: 2023-10-06
3  schedule_interval: "20 10 * * *"
4  description: Pull data from Jira into Snowflake
5  deploy_env: prod
6  notify_on_failure:
7    opsgenie:
8      responders: IC-SF-DataEng-Platform-Orchestration
9      priority: P3
10     description: Please check runbook https://instacart.atlassian.net/wiki/spaces/DATA/pages/3932356704/care-to-sf+Runbook
11  tasks:
12    - airflow_task_name: care-to-sf
13      isc_task_name: care-to-sf.data-eng
14      command: python /app/bin/ingest_jira_to_snowflake.py
15      task_properties:
16        timeout: 90m
17
```

Abstract away DAG Configurations

```
! example_delta_to_snowflake_sync.pipeline.yml ×
1  owner: anant.agarwal
2  totem_system: team-3-2396 # Compute Frameworks
3  start_date: 2024-03-26
4  schedule_interval: null
5  description: Example pipeline to demonstrate how the delta to snowflake sync is configured
6  deploy_env: local, dev, prod
7  notify_on_failure:
8    slack:
9      team_channel: bot-data-infra-productivity-non-critical
10 tags: examples
11 tasks:
12   - name: delta_to_snowflake_sync
13     type: spark
14     operator:
15       databricks_connection_id: DATABRICKS_CONNECTION_ID_DELTA_TO_SNOWFLAKE_SYNC
16       operator_params:
17         type: delta_to_snowflake_sync
18         delta_to_snowflake_sync_new_cluster_defaults: example_cluster
19         delta_to_snowflake_sync_timeout_seconds: 3600
20         delta_path: "{{ file_env_var('instacart_ads_dags/scheduling/examples/example_env_params.yml', 'delta_path') }}"
21         snowflake_table_fq: "{{ file_env_var('instacart_ads_dags/scheduling/examples/example_env_params.yml', 'snowflake_table_fq') }}"
22         snowflake_connection_id: snowflake_connection_id_delta_to_snowflake_sync
23         key_columns: [ identity_key ]
24         stage: instadata.public.delta_to_snowflake_stage_dev
25
```

Abstract away DAG Configurations

```
from datetime import datetime, timedelta


from airflow import DAG
from airflow.operators.bash import BashOperator
from utils.config import AirflowConfig
from utils.notifications.opsgenie import create_opsgenie_failure_callback

# Dag that should always pass. Alerts should be triggered when this dag fails.
with DAG(
    "current_build",
    description="Print the current build loaded to the worker",
    schedule_interval="*/5 * * * *",
    start_date=datetime(2022, 6, 10),
    default_args={
        "owner": "Data Platform Orchestration",
        "on_failure_callback": create_opsgenie_failure_callback(
            responder_team="IC-SF-DataEng-Platform-Orchestration",
            priority="P1",
        ),
    },
    catchup=True,
    is_paused_upon_creation=False,
) as dag:

    t1 = BashOperator(
        task_id="print_current_build",
        bash_command="cat /app/current-build",
        retries=AirflowConfig.task_retry,
        retry_delay=timedelta(seconds=AirflowConfig.task_retry_delay),
    )
```

Terraform-ized Airflow

- Infrastructure as code!
- Easy to scale and configure resources
- Easy to spin up new clusters
- Consistency across DEV and PROD environments.
- Auditability and Compliance controls



HashiCorp

Terraform



```
module "airflow2-test" {  
  count           = var.environment == "development" ? 1 : 0  
  source         = "github.com/instacart/terraform-airflow?ref=v1.20.3"  
  environment    = var.environment  
  domain        = var.domain  
  tags           = module.tags  
  cluster_name  = "infra"  
  has_autoscaling = false  
  name          = "airflow2-test"  
  app_name      = "airflow2-playground"  
  create_db_parameter_group = true  
  pgbouncer_security_group = module.pgbouncer-defaults.security_groups  
  rds_database_name = "airflow2test"  
  target_group_deregistration_delay = var.target_group_deregistration_delay  
  webserver_dashed_name = true  
  proxy_webserver = true  
  loadbalancer_security_group_ids = { "LoadBalancer Security Group" : data.terraform_remote_state.ecs-clusters.outputs.data_eng_ecs  
  }  
  ecs_cluster_names = [  
    "infra-dev"  
  ]  
}
```

```
rds_primary_instance_class      = var.rds_primary_instance_class
rds_primary_is_multi_az        = var.rds_primary_is_multi_az
rds_primary_skip_final_snapshot = var.rds_primary_skip_final_snapshot
engine_version                  = var.engine_version

redis_node_type                 = var.redis_node_type
redis_number_cache_clusters    = var.redis_number_cache_clusters

# replica config
create_replica_rds_instance    = var.create_replica_rds_instance
rds_replica_instance_class     = var.rds_replica_instance_class
create_replica_pgbouncer      = var.create_replica_pgbouncer

webserver_cores = var.webserver_cores
webserver_rammb = var.webserver_rammb
webserver_nodes = var.webserver_nodes

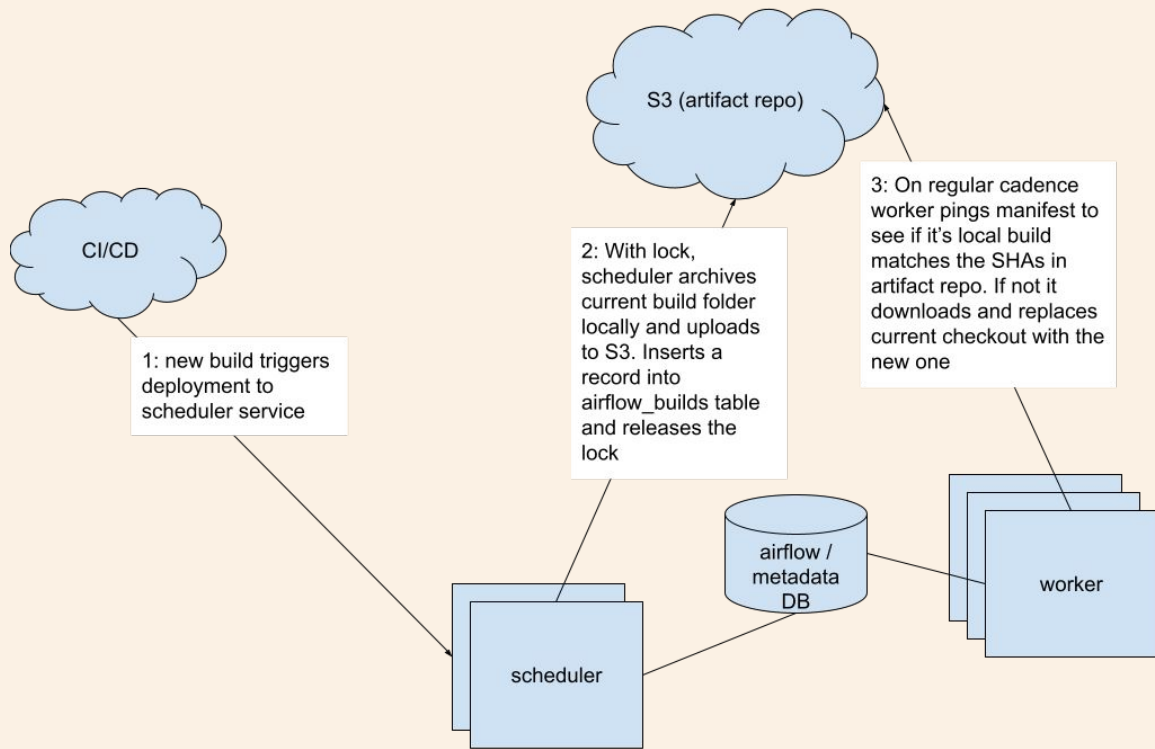
worker_cores      = var.worker_cores
worker_rammb     = var.worker_rammb
worker_nodes     = var.worker_nodes
worker_has_autoscaling = false
worker_min_capacity = 2
worker_max_capacity = 2

scheduler_cores = var.scheduler_cores
scheduler_rammb = var.scheduler_rammb
scheduler_nodes = var.scheduler_nodes

datadog_monitor_notification_list = ["@something"]
```

```
}
```


Development Friendly CI/CD



Monitoring and Alerting

- Monitoring as a first class citizen!
- Airflow Platform Monitoring
- DAG Monitoring
- Providing abstractions to create monitoring resources easily



Part 2 – How we use Airflow at Instacart

Core High Level Metrics

Number of D...

2182

Task Instance ...

99.3%

Executor Slots

DAG Run Task Duration

of DAGs over time

Running vs queued tasks

Slots available

Task Instances Successes & Failu...

Queued tasks

DAG Running Count

DAG_ID	SUCCESS	FAILURE
connected_store_2...	2	2
flyers_test_run_fly...	2	2
ma_fp_optimizatio...	6	2

Task Instance Success Rate Over ...

Running tasks trend

Average tasks run time hour

1.87	display_tao_campai...mpaign_cpms_daily
1.87	display_tao_campai...aign_summary_daily
1.86	model.instacart.sear...t_queries_flattened
1.86	model.instacart.use...user_coupon_abuse
1.72	event_ranked_item...event_ranked_items
1.66	fact_image_coverag...ge_retailer_insights

Useful Links:

- > Worker
6 widgets
- > Webserver
5 widgets
- > Scheduler

Container Count

Scheduler CPU Utilization (%)

Memory Utilization

Heartbeat

ECS OOMs
- > RDS
6 widgets
- > Redis
1 widget
- > Triggerer
6 widgets
- > Slack and Opsgenie Integration Health
5 widgets
- > dbt Pipeline Level Metrics
2 widgets

```
resource "datadog_monitor" "airflow_scheduler_heartbeat" {
  name      = "[Data Platform Airflow] Scheduler No Heartbeat"
  type      = "query alert"
  query     = "sum(last_10m):avg:airflow.scheduler_heartbeat{environment:production,ecs_cluster:data-eng}.as_count() <= 0"
  message   = <<-EOT
    service
      https://isc.fernet.io/services?page=1&per\_page=50&query%5Benvironment%5D=production&query%5Bservice%5D=data-platform-airflow

    UI
      https://data-platform-airflow.icprivate.com/home

    Runbook @opsgenie-IC-SF-DataEng-Platform-Orchestration-P1
  EOT

  tags = ["team:data-eng-platform"]

  monitor_thresholds {
    critical          = 0
    critical_recovery = 1
  }

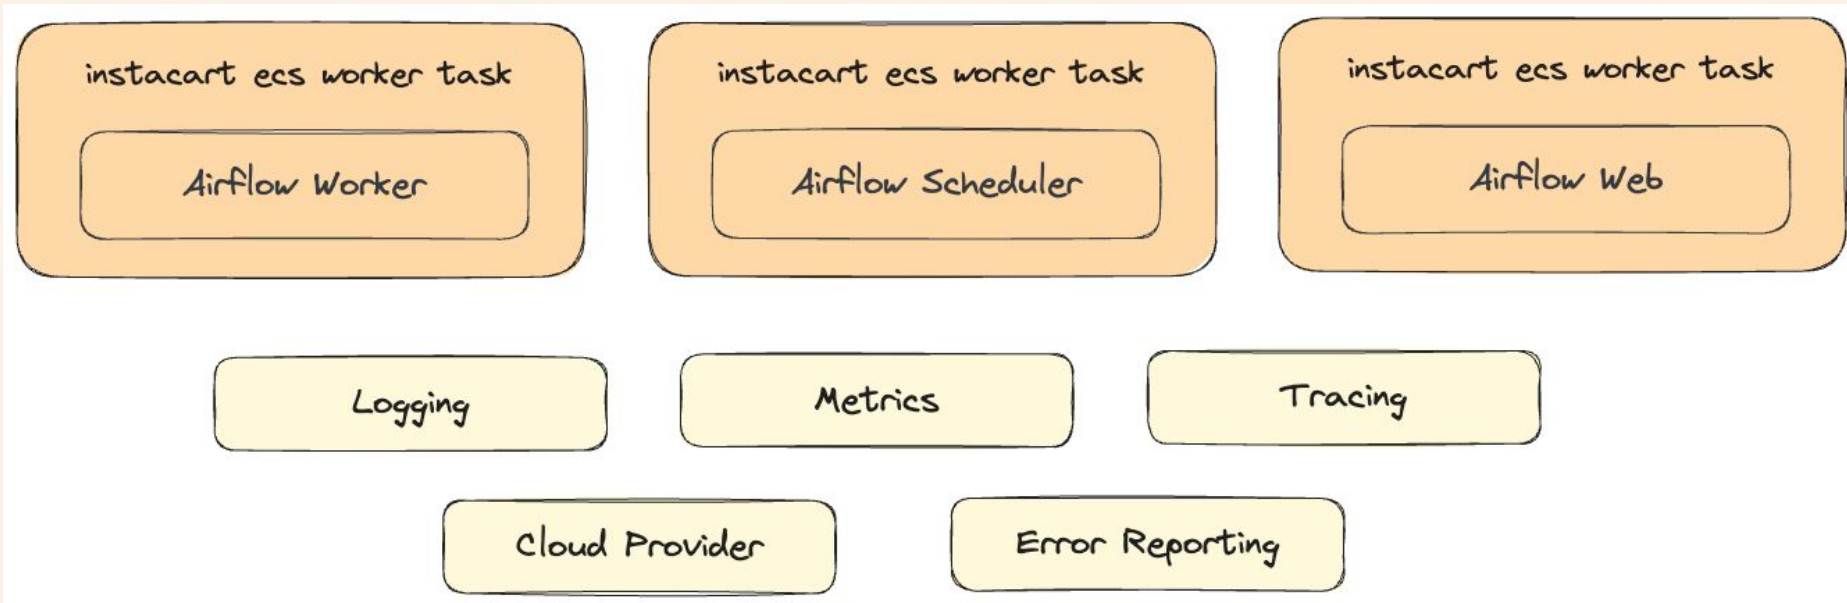
  notify_audit          = false
  require_full_window  = false
  notify_no_data       = true
  renotify_interval    = 0
  include_tags         = true
  no_data_timeframe    = 10
  priority              = 1
}
```

Key Takeaways

Our Recommendations!



Key Takeaway 1: Building Abstractions



Key Takeaway 2: Terraforming Resources




Key Takeaway 3: YAML DAGs & Remote Executors!

```
! ingest_jira_to_snowflake.pipeline.yml X
```

```
1  owner: anant.agarwal|
2  start_date: 2023-10-06
3  schedule_interval: "20 10 * * *"
4  description: Pull data from Jira into Snowflake
5  deploy_env: prod
6  notify_on_failure:
7    opsgenie:
8      responders: IC-SF-DataEng-Platform-Orchestration
9      priority: P3
10     description: Please check runbook https://instacart.atlassian.net/wiki/spaces/DATA/pages/3932356704/care-to-sf+Runbook
11 tasks:
12   - airflow_task_name: care-to-sf
13     isc_task_name: care-to-sf.data-eng
14     command: python /app/bin/ingest_jira_to_snowflake.py
15     task_properties:
16       timeout: 90m
17
```

Key Takeaway 3: DAG abstractions & Remote Workers!



Questions?

- Find me on Airflow's community slack
- <https://www.linkedin.com/in/anantag/>

