



Optimising Airflow Performance

Tips & strategies to enhance metadata database performance



Pankaj Singh Software Engineer @ Astronomer & Airflow Committer

Pankaj Koti Software Engineer @ Astronomer & Airflow Committer

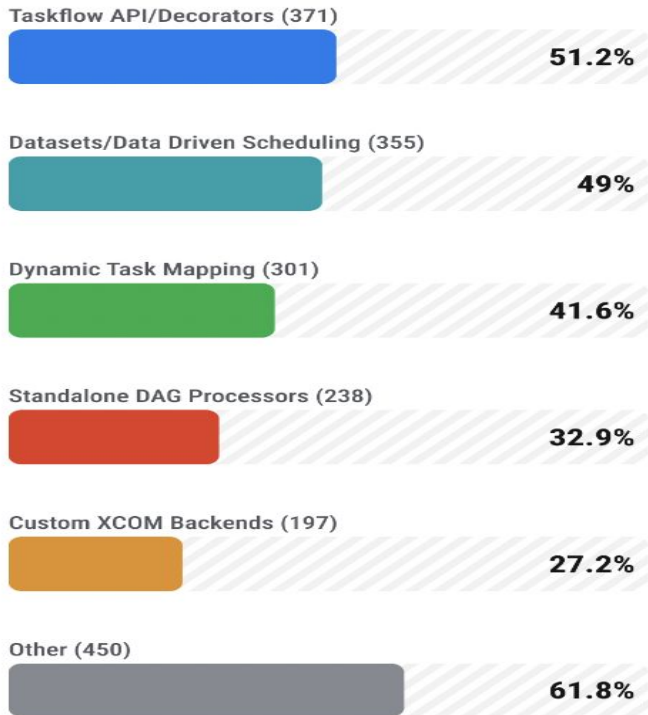


What We'll Cover

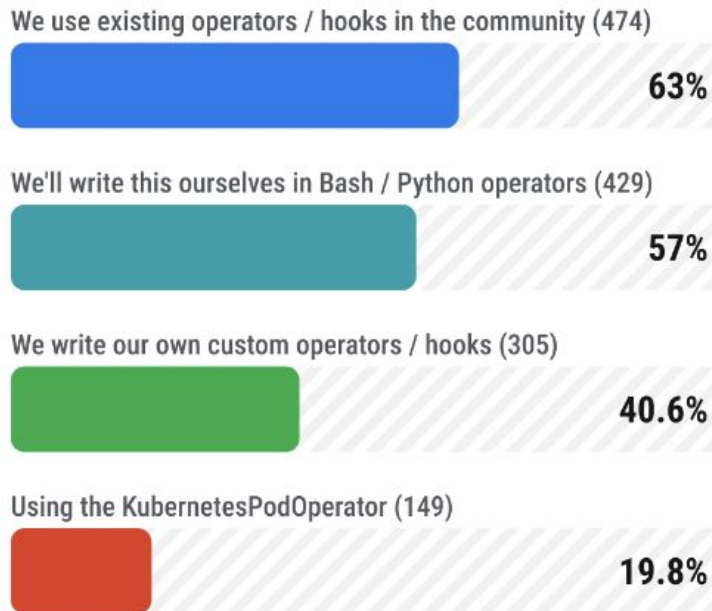
1. DAG Authoring Best Practices
2. Database Optimization
 - a. Unused Indexes
 - b. Missing Indexes
 - c. Table and Index Bloat

2023 Airflow Survey Result

Which features do you use?



How do you usually interface with other services from your Airflow DAGs?





DAG Authoring

Expensive Call

Avoid making network calls or performing heavy computations at the top level of the code.

Heavy Library Import

Avoid top-level imports for large libraries

Jinja Template

Use Jinja templates to access Airflow resources, as they are resolved at runtime

Detect Top-level Code

A simple test can be conducted by running `python <dag_file>.py`.



DAG Authoring

Example: Top-level expensive call

```
1 import pendulum
2 from airflow import DAG
3 from airflow.decorators import task
4 from time import sleep
5
6 def expensive_api_call():
7     print("Hello from Airflow!")
8     sleep(5) # Simulate an expensive call
9
10 # This call will be executed every time the DAG file is parsed
11 my_expensive_response = expensive_api_call()
12
13 with DAG(
14     dag_id="example_bad_practice",
15     schedule=None,
16     start_date=pendulum.datetime(2024, 1, 1, tz="UTC"),
17 ) as dag:
18
19     @task
20     def print_expensive_api_call():
21         print(my_expensive_response)
22
23     print_expensive_api_call()
```

```
real 0m7.136s
user 0m2.030s
sys 0m0.104s
```



```
1 import pendulum
2 from airflow import DAG
3 from airflow.decorators import task
4 from time import sleep
5
6 def expensive_api_call():
7     sleep(5) # Simulate an expensive call
8     return "Hello from Airflow!"
9
10 with DAG(
11     dag_id="example_good_practice",
12     schedule=None,
13     start_date=pendulum.datetime(2024, 1, 1, tz="UTC"),
14 ) as dag:
15
16     @task
17     def print_expensive_api_call():
18         my_expensive_response = expensive_api_call()
19         print(my_expensive_response)
20
21     print_expensive_api_call()
```

```
real 0m2.283s
user 0m2.041s
sys 0m0.228s
```



DAG Authoring

Example: Top-level import

```
1 import pendulum
2 from airflow import DAG
3 from airflow.decorators import task
4
5 # Expensive imports at the top level
6 import pandas as pd
7 import torch
8
9 with DAG(
10     dag_id="example_bad_imports",
11     schedule=None,
12     start_date=pendulum.datetime(2024, 1, 1, tz="UTC"),
13 ) as dag:
14
15     @task
16     def process_data():
17         # Use the imported libraries
18         df = pd.DataFrame({'a': [1, 2, 3]})
19         print(df)
20         print(torch.__version__)
21
22     process_data()
```

```
real    0m3.656s
user    0m4.213s
sys     0m0.213s
```



```
1 import pendulum
2 from airflow import DAG
3 from airflow.decorators import task
4
5 with DAG(
6     dag_id="example_good_imports",
7     schedule=None,
8     start_date=pendulum.datetime(2024, 1, 1, tz="UTC"),
9 ) as dag:
10
11     @task
12     def process_data():
13         # Import the expensive libraries inside the task
14         import pandas as pd
15         import torch
16
17         # Use the imported libraries
18         df = pd.DataFrame({'a': [1, 2, 3]})
19         print(df)
20         print(torch.__version__)
21
22     process_data()
23
```

```
real    0m2.267s
user    0m2.033s
sys     0m0.221s
```



DAG Authoring

Example: Jinja template

```
1 from airflow import DAG
2 from airflow.decorators import task
3 from airflow.models import Variable
4 import pendulum
5
6 # Top-level variable fetch
7 foo = Variable.get('foo')
8
9 with DAG(
10     dag_id="example_top_level_var_fetch",
11     schedule=None,
12     start_date=pendulum.datetime(2024, 1, 1, tz="UTC"),
13 ) as dag:
14
15     @task
16     def print_var():
17         print(foo)
18
19     print_var()
```

```
real    0m2.425s
user    0m2.120s
sys     0m0.295s
```



```
1 from airflow import DAG
2 from airflow.decorators import task
3 import pendulum
4
5
6 with DAG(
7     dag_id="example_jinja_template",
8     schedule=None,
9     start_date=pendulum.datetime(2024, 1, 1, tz="UTC"),
10 ) as dag:
11
12     @task
13     def print_var(foo):
14         print(foo)
15
16     print_var("{{ var.value.get('foo') }}")
17
```

```
real    0m2.168s
user    0m2.027s
sys     0m0.140s
```



DAG Authoring

Example: Detect top-level code

```
> python <my_dag.py>
```

```
> time python <my_dag.py>
```




Optimising Database Performance



Metadata Performance Degradation

High Disk Consumption

Ex - The TI table uses about 4.7GB, but its indexes add another 20GB

Slow Query

Larger index sizes can significantly slow down queries by increasing disk I/O, lock contention, and resource usage

Scheduler Liveness Failure

Scheduler fails to respond, often due to metadata poor performance



Unused Indexes



Identifying Unused Indexes

Identifying unused indexes can be challenging due to various factors

Index Scan Metrics Over Time

How frequently the indexes are used in query execution over a period can change

Dynamic Usage Patterns

The indexes used vary depending on the specific use cases and feature requirements

Version-Related Usage

Indexes may be added or removed depending on the version of Airflow



Unused Indexes

SELECT

```
    schemaname || '.' || relname AS table,  
    indexrelname AS index,  
    pg_size_pretty(pg_relation_size(i.indexrelid)) AS index_size,  
    idx_scan as index_scans
```

FROM

```
    pg_stat_user_indexes ui
```

JOIN

```
    pg_index i ON ui.indexrelid = i.indexrelid
```

WHERE

```
    NOT indisunique AND idx_scan = 0
```



Unused Indexes: Time Is Also a Factor

Index	12/04/2024	15/04/2024	17/04/2024
idx_last_scheduling_decision (size MB)	5047	5125	5183
idx_last_scheduling_decision (scan #)	0	0	0
idx_log_dag (size MB)	2291	2353	2399
idx_log_dag (scan #)	6	6	6



Deleting Unused Indexes

Deleted index

- Table: dag_run
- Index:
idx_last_scheduling_decision

Airflow 2.9.2

<https://github.com/apache/airflow/pull/39275>

Potential candidate

- Table: log
- Index: idx_log_dag



Impact of Index Deletion

More disk space

Deleting the `idx_last_scheduling_decision` indexes freed up 5GB of disk space

Fast query

Improving the performance of write queries (such as insert and update)

Scheduler liveness failure

Gain acceptable performance to prevent frequent scheduler failures



Future: Exporting Stats for Index Usage

Reliable Index Stats

- To identify unused indexes in PostgreSQL, query the `pg_stat_user_indexes` view
- To export `pg_stat_user_indexes` via Prometheus, use the PostgreSQL exporter tool

```
custom_queries:  
- query: |  
    SELECT  
        schemaname || '.' || relname AS table,  
        indexrelname AS index,  
        pg_size_pretty(pg_relation_size(i.indexrelid)) AS index_size,  
        idx_scan as index_scans  
    FROM  
        pg_stat_user_indexes ui  
    JOIN  
        pg_index i ON ui.indexrelid = i.indexrelid  
    WHERE  
        NOT indisunique  
metrics:  
- table:  
    usage: "LABEL"  
    description: "Name of the table"  
- index:  
    usage: "LABEL"  
    description: "Name of the index"  
- index_size:  
    usage: "LABEL"  
    description: "Size of the index"  
- index_scans:  
    usage: "COUNTER"  
    description: "Number of index scans"
```



Missing Indexes



Slow Queries

Slow DAG List Page

Time take to loading dag list page was proportional to size of metadata.

Slow Stale Metadata Deletion

Stale metadata deletion of 1 week's data took 7 mins and Astronomer Support needed to delete data of 1 year for a customer which could take around 6hrs.



Slow Query Side Effect

Turnaround Time to Fetch Results

Time taken to get the results increases.

High CPU Utilization

Sequential scanning of the metadata database raises CPU utilization



Identify Slow Query

```
# Add Extension
CREATE EXTENSION pg_stat_statements;

# Restart DB

# Get Stats
SELECT query,
       calls,
       min_time,
       max_time,
       mean_time,
       total_time
FROM   pg_stat_statements
ORDER BY mean_time DESC;
```



Identify Slow Query - Result

_query	calls	min_time	mean_time	total_time
SELECT substring(query, \$1, \$2) AS trimmed_qu	+ 2	1.274807	1.654795	2.929602
SELECT substring(query, \$1, \$2) AS _query,cal	+ 1	1.157411	1.157411	1.157411
SELECT substring(query, \$1, \$2) AS _query,cal	+ 1	1.02273	1.02273	1.02273
SELECT dag_code.fileloc_hash, dag_code.fileloc, da	871	0.001832	0.721667	20.819917999999987
SELECT ab_permission.name, ab_view_menu.name AS na	43	0.14725	0.49942	7.827014999999985
SELECT trigger.id FROM trigger JOIN task_instance	+ 2422	0.002959	0.4748050000000003	21.481291999999999
SELECT ab_user_1.id AS ab_user_1_id, ab_role.id AS	48	0.154585	0.452504	10.257408000000002
SELECT dag.dag_display_name, dag.dag_id, dag.root_	790	0.006876	0.301583	41.520620000000006
SELECT dag_priority_parsing_request.id, dag_priori	7658	0.000708000000000001	0.284584	13.009313000000125
SELECT ab_role.id AS ab_role_id, ab_role.name AS a	165	0.002417	0.26421	13.192050000000004
SELECT a.attnname, pg_catalog.format_	+ 6	0.088876	0.225668	0.8616269999999999
SELECT dag.dag_display_name, dag.dag_id, dag.root_	81	0.034542	0.195295	6.510171
SELECT c.relname FROM pg_class c JOIN pg_namespace	27	0.033126	0.17929299999999998	2.3051829999999995
SELECT ab_role.id, ab_role.name, ab_permission_vie	2	0.120209	0.16796	0.288169
SELECT anon_1.dag_display_name, anon_1.dag_id, ano	11	0.078249	0.163544	1.136466
SELECT dag.dag_display_name, dag.dag_id, dag.root_	2769	0.008291000000000001	0.155531	54.92970500000011
SELECT dag.dag_display_name, dag.dag_id, dag.root_	81	0.07234800000000001	0.143972	7.516767999999999
SELECT \$1 AS anon_1 FROM serialized_dag WHERE se	+ 1357	0.001584	0.140863	14.921677999999982
SELECT t.oid, typparray FROM pg_type t JOIN pg_name	+ 3722	0.004667	0.13902799999999998	40.376448000000025



Adding Missing Indexes

Table	Index
dag_tag	idx_dag_tag_dag_id
dag_warning	idx_dag_warning_dag_id
dag_schedule_dataset_reference	idx_dag_schedule_dataset_reference_dag_id
dag_schedule_dataset_reference	idx_dataset_dag_run_queue_target_dag_id
dag_schedule_dataset_reference	idx_task_outlet_dataset_reference_dag_id

Airflow 2.10

<https://github.com/apache/airflow/pull/39638>



Impact of Index Addition

Not all metadata index is require by everyone

Slow Stale metadata deletion

Improved by the addition of an index which reduced the time of deletion of 1 year data to 36sec from ~6 hours.

Future Works of a slow load of DAG list page

Split the page into multiple components that can execute parallel queries instead of the serial execution.



Table & Index Bloats

Causes, Detection, &
Mitigation Strategies



Introduction

What is a Database bloat?

Table Bloat

Excessive unused space in tables due to deleted or outdated data that hasn't been reclaimed

Common in systems with frequent updates & deletes

Index Bloat

Due to deleted or outdated index entries

Can significantly degrade performance as indexes grow larger than necessary



Causes of Table Bloats

Frequent Updates/Deletes

Lack of autovacuum or incorrect autovacuum settings

Lack of Proper Maintenance

Failure to regularly vacuum and analyze the database

Inefficient Storage

Over-allocated space during table creation or after significant changes in data volume



Credits:

<https://hakibenita.com/postgresql-unused-index-size#index-and-table-bloat>



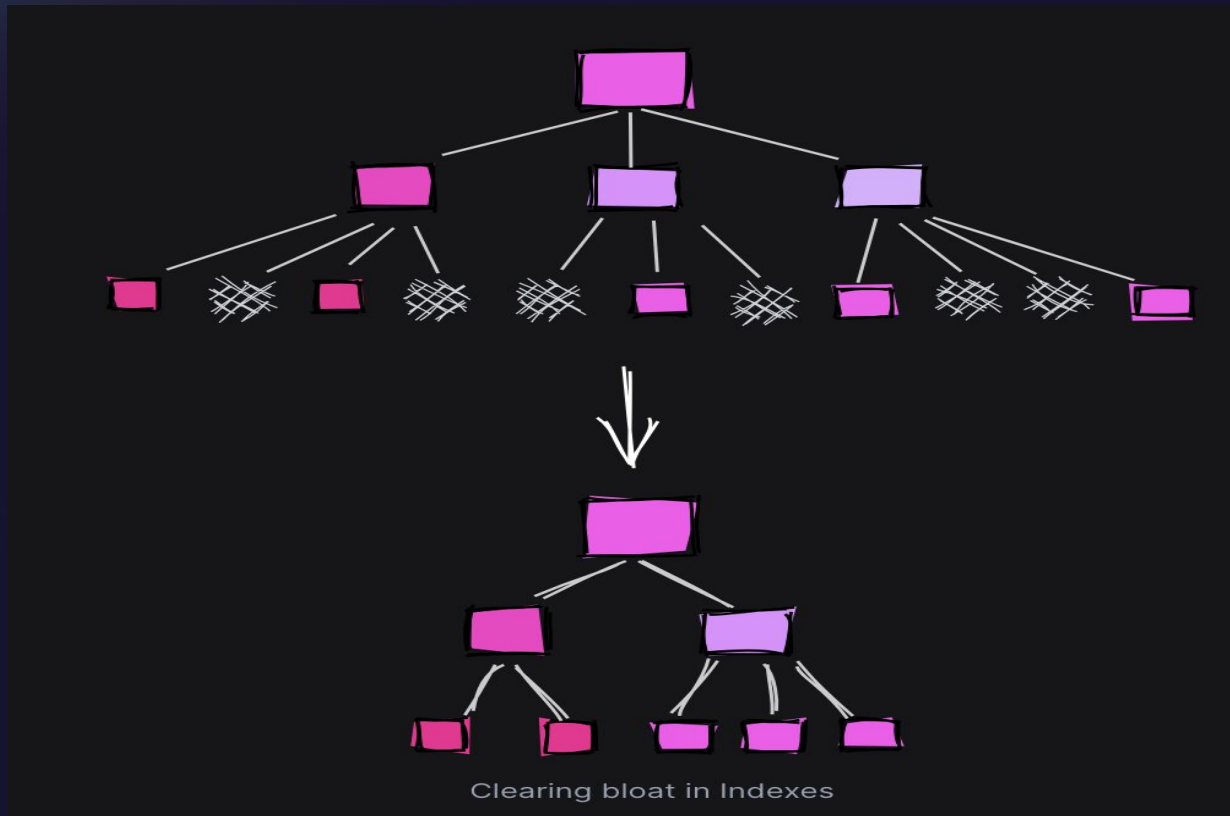
Causes of Index Bloats

Frequent Updates/Deletes on Indexed Columns

Indexes don't shrink automatically after deletions

Poor Index Management

Over-indexing and lack of regular index maintenance



Credits:

<https://hakibenita.com/postgresql-unused-index-size#index-and-table-bloat>



Impact of Table & Index Bloats

Performance Degradation

Slower query execution times

Increased I/O operations & memory usage

Increased Storage Costs

Larger than necessary database files

Maintenance Overhead

Longer backup & restore files



Detecting Bloats

Tools for Detection

PostgreSQL: `pg_stat_all_tables`,
`pgstattuple`, `pg_repack`

MySQL: `OPTIMIZE TABLE`,
`ANALYZE TABLE`

Key indicators

- Difference between table/index size and the actual data size
- Increasing table/index size without proportional data growth



Mitigating Table Bloats

Re-Create the Table

Often requires a lot of development, especially if the table is actively used as it's being rebuilt

Vacuum the Table

Query:

```
VACUUM FULL table_name
```

Will lock the page briefly

Using pg_repack

```
create EXTENSION pg_repack;
```

```
$ pg_repack -k --table  
table_name db_name
```

Mitigating Index Bloats

- Look for queries to detect index bloats based on your database
E.g. For PostgreSQL, below is a helpful query

https://github.com/ioguix/pgsql-bloat-estimation/blob/master/btree_btree_bloat.sql

- Reindex indexes with bloats

```
REINDEX INDEX index_name
```




Schema Name	Table Name	Index Name	Real Size (MB)	Extra Size (MB)	Extra (%)	Fill Factor	Bloat Size (MB)	Bloat (%)
airflow	dag_run	idx_last_scheduling_decision	8822	8789	99.62	90	8785.96	99.58
airflow	task_instance	task_instance_pkey	9928	7647	77.02	90	7374.58	74.27
airflow	xcom	idx_xcom_task_instance	9718	5190	53.41	90	4649.91	47.85
airflow	job	job_type_heart	5416	4725	87.24	90	4647.12	85.80
airflow	task_instance	ti_state_lkp	7275	4830	66.40	90	4559.35	62.67
airflow	task_instance	ti_job_id	2872	2292	79.81	90	2226.77	77.52
airflow	dag_run	idx_dag_run_dag_id	2245	2185	97.33	90	2178.83	97.03
airflow	job	idx_job_dag_id	2944	2253	76.53	90	2174.82	73.87
airflow	task_instance	ti_dag_run	3456	1856	53.72	90	1673.41	48.42
airflow	task_instance	ti_trigger_id	2159	1570	72.76	90	1505.84	69.75
airflow	xcom	xcom_pkey	4403	1550	35.20	90	1230.00	27.93
airflow	task_instance	ti_state_incl_start_date	2921	1157	39.62	90	954.65	32.67
airflow	task_instance	ti_dag_state	1686	761	45.14	90	655.26	38.86
airflow	task_instance	ti_pool	1637	712	43.50	90	606.38	37.04
airflow	log	idx_log_dag	4285	930	21.71	90	546.12	12.74
airflow	job	job_pkey	871	488	55.98	90	445.04	51.05
airflow	task_instance	idx_laminar_ti_end_date	976	388	39.77	90	323.21	33.10
airflow	xcom	idx_xcom_key	1554	387	24.95	90	258.84	16.65
airflow	task_instance	ti_state	674	254	37.69	90	207.27	30.71
airflow	job	idx_job_state_heartbeat	798	262	32.81	90	202.72	25.38
airflow	dag_run	dag_run_dag_id_execution_date_key	259	185	71.76	90	177.61	68.53
airflow	dag_run	dag_run_dag_id_run_id_key	309	182	59.08	90	168.19	54.39
airflow	log	idx_log_event	2531	398	15.75	90	162.70	6.43
airflow	dag_run	dag_run_pkey	120	87	72.48	90	83.91	69.40
airflow	dag_run	idx_laminar_dagrun_end_date	104	57	55.41	90	52.66	50.47
airflow	dag_run	dag_id_state	117	44	37.92	90	36.33	30.82
airflow	log	idx_log_dttm	1700	175	10.31	90	4.43	0.26
airflow	log	log_pkey	1700	175	10.31	90	4.42	0.26



Caveat: Do It at Your Own Risk

- Airflow metadata database expects the DDLs to be unaltered and only be modified via the migrations
- But if you're cautious and can take care of any potential conflicts then you're good to apply your findings and solutions



Thank you!

Questions?
#airflow-performance