snowflake SUMMIT 2023

Optimizing Snowflake Spend by Prioritizing User Experience

Teresa Kovich & Laura McKinley
DAS42 Principal Consultants

June 27, 2023

Who We Are

Teresa Kovich

Principal Consultant
DAS42

Laura McKinley

Principal Consultant
DAS42

DAS42

Testimonials

"The DAS42 team really helped us think through how to leverage Snowflake to design and architect an enterprise solution."

Tracy Smith
Sr. Director Data & Analytics,
Be the Match

BE THE MATCH®

"DAS42 is a key partner in data technology and data pipelines. They bring a wealth of experience which allows them to tailor their solutions to each client. DAS42 is unique, agile, and helps customers find the true value and best practices with the Data Cloud."

Katie Ecklund
Sr. Director, Amers Partner Sales,
Snowflake

snowflake®

DAS42

"Our strategic partnership with DAS42 to bring Snowflake to the table and bring the modern platform to us helped us accelerate our roadmap… we are able to gain business insights at a pace that matches the rapid growth rate of our business. We have a strategic design, an enterprise-grade solution, and a partner that will take us far into the future."

Jon Moss
AVP of Edge Software Engineering,
Zayo

Zayo®

Agenda

- Benefits of optimization
- The basics
- The user-first approach
- Case studies
- Wrap-up
- Q&A

www.das42.com/snowflakesummit2023

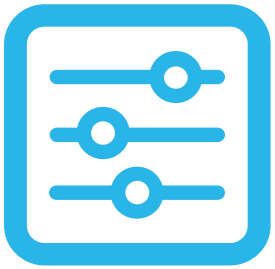

Benefits of Optimization

Why Optimize?

- Focus on effectively using Snowflake
- Can result in cost savings
- Free up credit usage for new use cases—do more with less

- Faster queries = happier users
- People are more likely to use data in decision making if it's easy

The Basics

Don’t Forget Your Basics

Can you confirm that:

- The code has performant SQL, including DDL
- Table structures favor long-skinny vs. short-wide
- Ingest follows least possible uptime
- Auto-suspend and auto-resume are enabled
- Timeouts are set appropriately for workloads
- Account statement timeouts are set
- Warehouses deviating from the seven-day average are identified
- Warehouses approaching the cloud service billing threshold are monitored
- Unused tables are dropped
- Dormant users are purged
- Resource monitors are applied to warehouses

Don’t Forget Your Basics

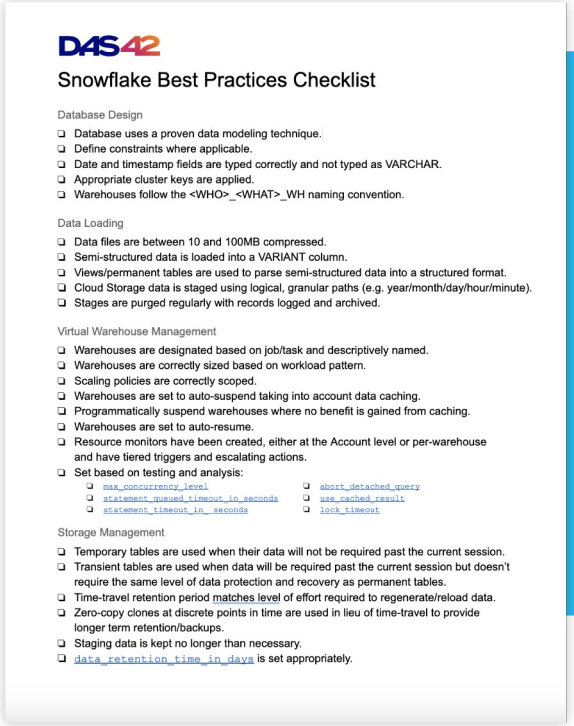

DAS42

Snowflake Best Practices Checklist

Database Design

- ❑ Database uses a proven data modeling technique.
- ❑ Define constraints where applicable.
- ❑ Date and timestamp fields are typed correctly and not typed as VARCHAR.
- ❑ Appropriate cluster keys are applied.
- ❑ Warehouses follow the <WHO>_<WHAT>_WH naming convention.

Data Loading

- ❑ Data files are between 10 and 100MB compressed.
- ❑ Semi-structured data is loaded into a VARIANT column.
- ❑ Views/permanent tables are used to parse semi-structured data into a structured format.
- ❑ Cloud Storage data is staged using logical, granular paths (e.g. year/month/day/hour/minute).
- ❑ Stages are purged regularly with records logged and archived.

Virtual Warehouse Management

- ❑ Warehouses are designated based on job/task and descriptively named.
- ❑ Warehouses are correctly sized based on workload pattern.
- ❑ Scaling policies are correctly scoped.
- ❑ Warehouses are set to auto-suspend taking into account data caching.
- ❑ Programmatically suspend warehouses where no benefit is gained from caching.
- ❑ Warehouses are set to auto-resume.
- ❑ Resource monitors have been created, either at the Account level or per-warehouse and have tiered triggers and escalating actions.
- ❑ Set based on testing and analysis:
  - ❑ max_concurrency_level
  - ❑ statement_queued_timeout_in_seconds
  - ❑ statement_timeout_in_ seconds
  - ❑ abort_detached_query
  - ❑ use_cached_result
  - ❑ lock_timeout

Storage Management

- ❑ Temporary tables are used when their data will not be required past the current session.
- ❑ Transient tables are used when data will be required past the current session but doesn’t require the same level of data protection and recovery as permanent tables.
- ❑ Time-travel retention period matches level of effort required to regenerate/reload data.
- ❑ Zero-copy clones at discrete points in time are used in lieu of time-travel to provide longer term retention/backups.
- ❑ Staging data is kept no longer than necessary.
- ❑ data_retention_time_in_days is set appropriately.

When to check your spend and adhere to best practices:

- New data source is added
- New aggregate modeling needed
- With some regularity—
  use Snowflake’s Data Cloud monitoring dashboard

How Snowflake Makes it Easier

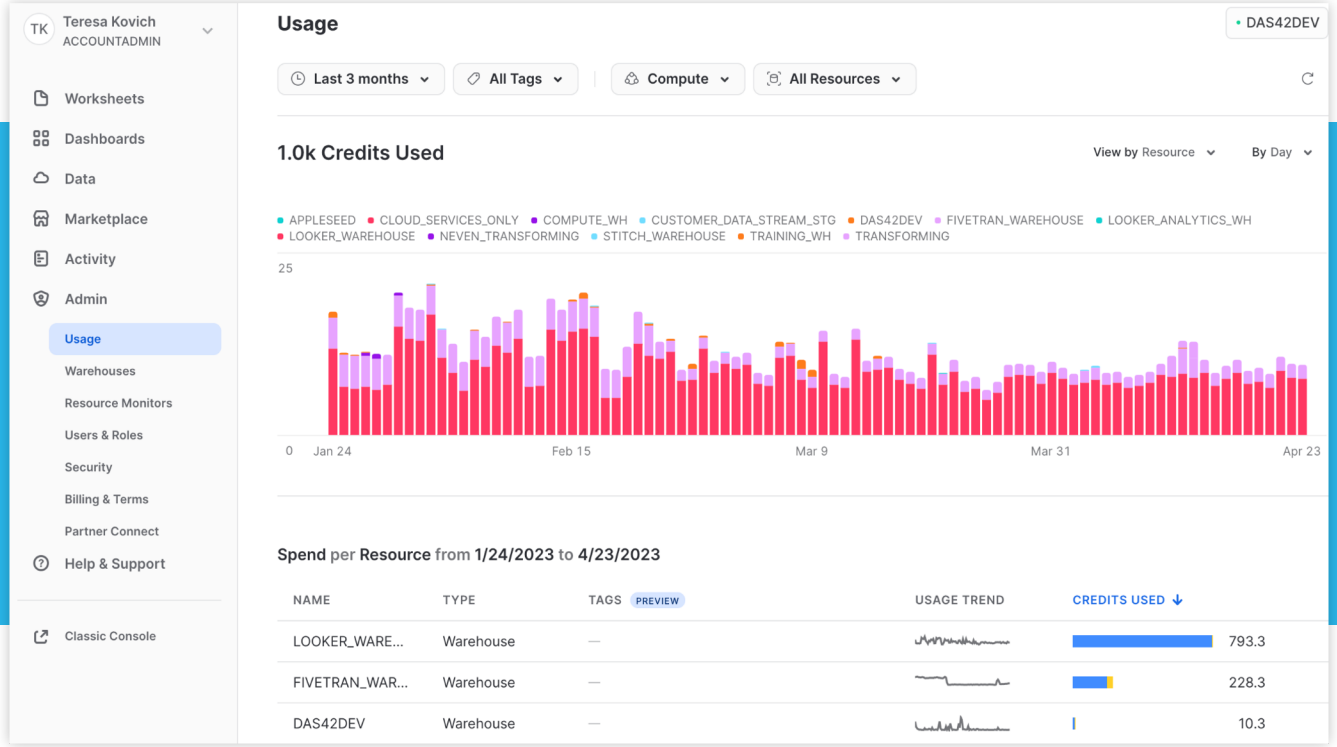

Best Practices from DAS42

- Plan for success (benchmark and set goals)
- Use smart build strategies
- Balance / right-sizing mindset for workloads
- Take advantage of time travel
- Be mindful of redundancies

And, of course…

- **Prioritize the user experience**

The User-First Approach

Why the User-First Approach?

- As your data grows, so does the complexity of your instance and organization's needs
- Understanding how and why the data is used by stakeholders leads to better decisions for the data engineering team
- Helps prioritize optimization projects
- More efficient queries = happier users = increased likelihood to use data in decisions
- Make sure that you're spending on things that can help your business grow and improve
- How can you make it easier for end users to get what they want? (*You need to ask the users!*)

How to Get Started

- Differentiate your warehouses and tag your queries to enable smart tracking
- Set up monitoring mechanisms and practices
- Design a spend strategy driven by business goals and growth areas
- Talk to end users in Snowflake and downstream tools (e.g. business intelligence/data science)
  - Identify pain points
  - Common usage patterns
  - Hours of activity
  - Confirm their understanding of SLAs

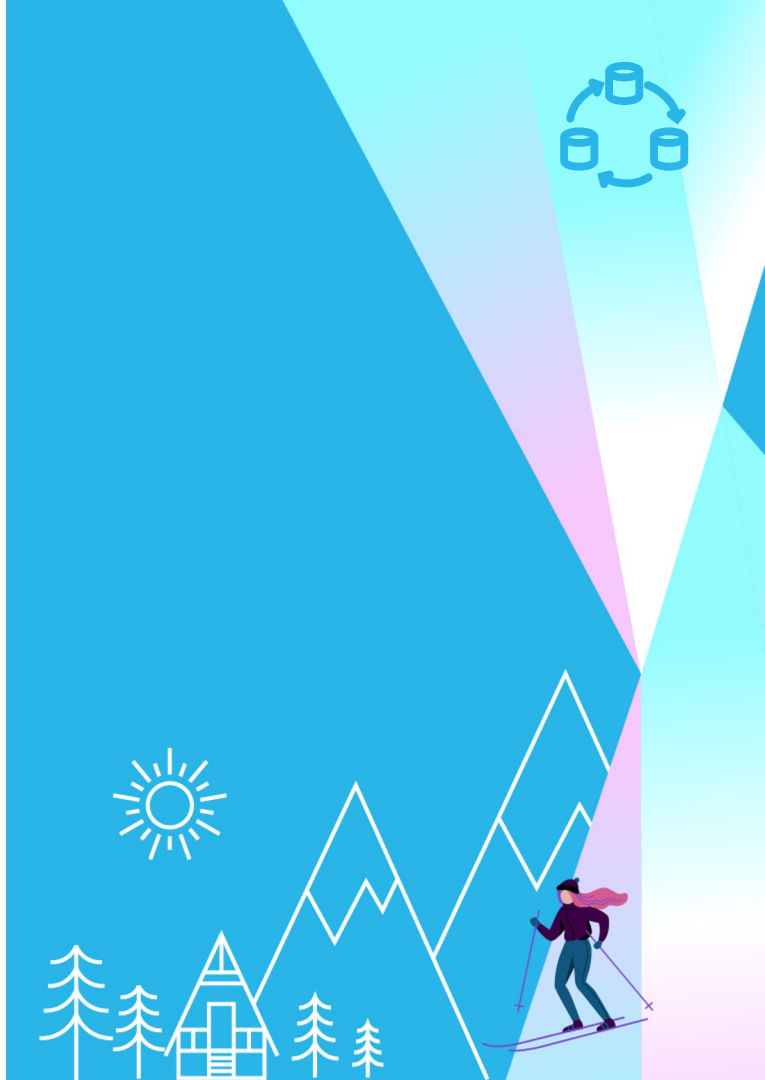

Use Your Data

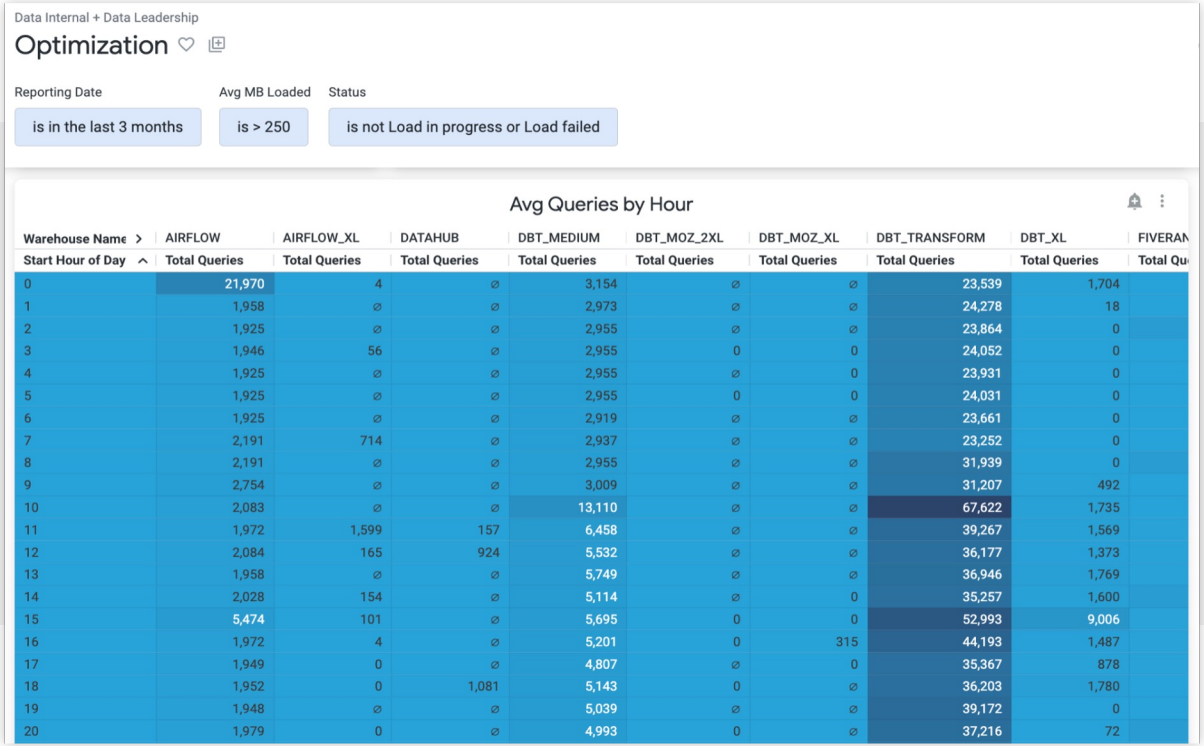

Data Internal + Data Leadership

Optimization

Reporting Date: is in the last 3 months

Avg MB Loaded: is > 250

Status: is not Load in progress or Load failed

Avg Queries by Hour

Warehouse Name	AIRFLOW	AIRFLOW_XL	DATAHUB	DBT_MEDIUM	DBT_MOZ_2XL	DBT_MOZ_XL	DBT_TRANSFORM	DBT_XL	FIVERAN
Start Hour of Day	Total Queries	Total Queries	Total Queries	Total Queries	Total Queries	Total Queries	Total Queries	Total Queries	Total Qu
0	21,970	4	∅	3,154	∅	∅	23,539	1,704	
1	1,958	∅	∅	2,973	∅	∅	24,278	18	
2	1,925	∅	∅	2,955	∅	∅	23,864	0	
3	1,946	56	∅	2,955	0	0	24,052	0	
4	1,925	∅	∅	2,955	∅	0	23,931	0	
5	1,925	∅	∅	2,955	0	0	24,031	0	
6	1,925	∅	∅	2,919	∅	∅	23,661	0	
7	2,191	714	∅	2,937	∅	∅	23,252	0	
8	2,191	∅	∅	2,955	∅	∅	31,939	0	
9	2,754	∅	∅	3,009	∅	∅	31,207	492	
10	2,083	∅	∅	13,110	∅	∅	67,622	1,735	
11	1,972	1,599	157	6,458	∅	∅	39,267	1,569	
12	2,084	165	924	5,532	∅	∅	36,177	1,373	
13	1,958	∅	∅	5,749	∅	∅	36,946	1,769	
14	2,028	154	∅	5,114	∅	0	35,257	1,600	
15	5,474	101	∅	5,695	0	0	52,993	9,006	
16	1,972	4	∅	5,201	0	315	44,193	1,487	
17	1,949	0	∅	4,807	∅	0	35,367	878	
18	1,952	0	1,081	5,143	0	∅	36,203	1,780	
19	1,948	∅	∅	5,039	∅	∅	39,172	0	
20	1,979	0	∅	4,993	0	∅	37,216	72	

Case Studies

Optimization Within Media & Entertainment

BUSINESS INSIDER

- Cost reporting, budgeting, baselining, and regular check-ins
- Fine-tuning warehouse sizing
- Uptime and ingest frequency—hours of usage and monitoring

Results

- Reduced Snowflake spend by **23% month-over-month**

OTHER OPTIMIZATIONS:

- Identified over- and under-clustering
- Consolidation within BI tool
- Granularity control and aggregate awareness in business intelligence applications

Optimization Within FinTech

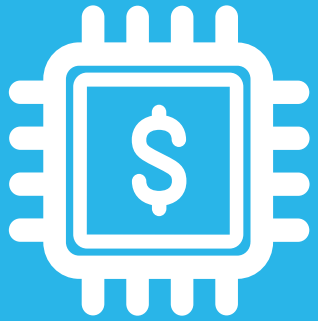

- Focus of project was to optimize user experience of business intelligence (BI) platform
- Conducted stakeholder interviews across business to discover needs and help form priorities—focus on personas and usage patterns
- Business users wanted faster queries, easier navigation for analytics, and cleaner content
- Resulted in **40% drop month-over-month** in BI warehouse credit usage (~$30k reduction in monthly spend)

Optimization Within FinTech

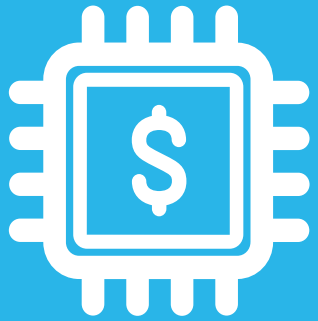

- Provisioned warehouses for personas:
  - Increased analytics usage for leadership
  - More granular scaling of appropriate warehouse sizes, based on personas and use cases
  - Transparent credit usage across use cases and user groups
- Configured ingest schedules to align with necessary timeframes, reduced unnecessary compute spend and warehouse uptime
- Improved compiled SQL (e.g.: joins on dates)
- Redesigned data model to hub-and-spoke:
  - Reduction of the number of data warehouse models, and inclusions in each spoke to improve usability and increase performance, **reducing unnecessary queries by 24-26%**
  - **Reduced** queue times and overall warehouse up-time for high-usage tables, reducing queue times **by 37% (average of 8s)**

DAS42

www.das42.com/snowflakesummit2023

snowflake®

SUMMIT 2023

THANK YOU!