



Migrating and Modernizing Victoria's Secret & Co. Analytics and Reporting Workloads from On-prem to Azure and Snowflake Stack

Driving Performance, Scalability, and Efficiency

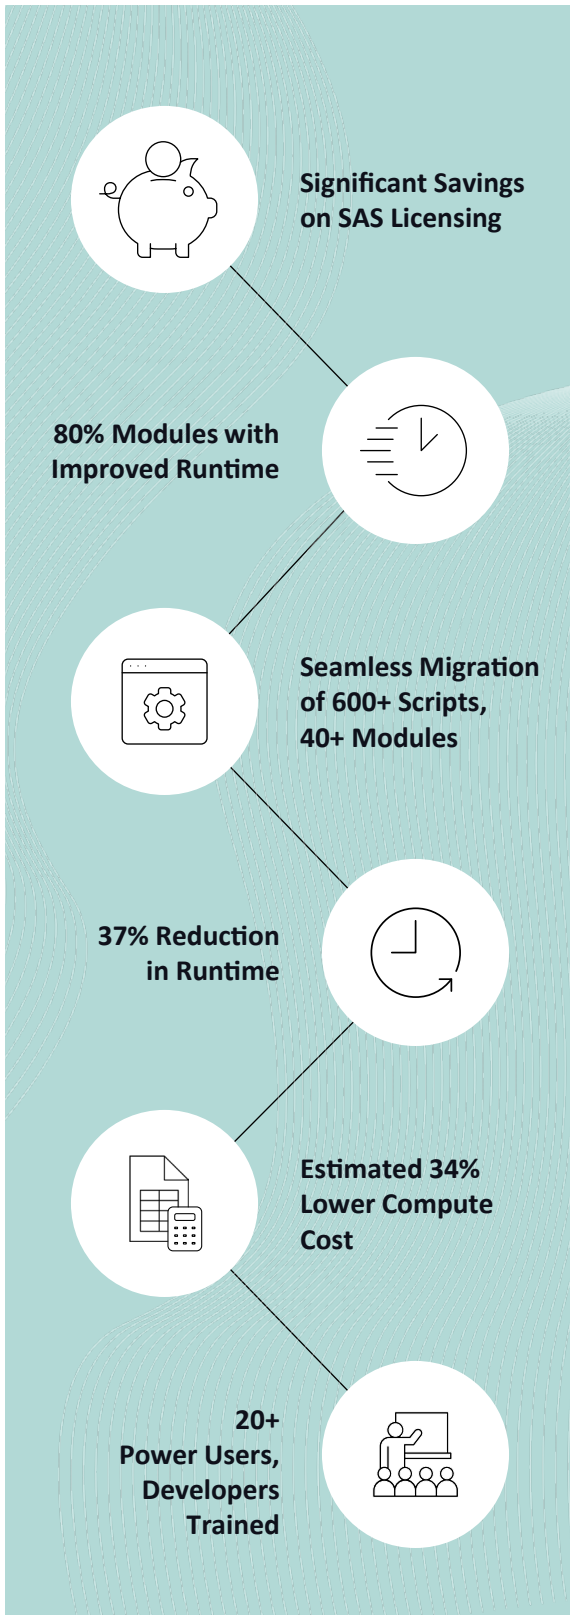


Contents

01	Executive Summary	3
<hr/>		
02	Key Drivers for Migration and Modernizing Analytics Workloads	4
<hr/>		
03	Migration Strategy	5
<hr/>		
04	Snowflake and Azure for Migration and Modernization	12
<hr/>		
05	Performance Optimization and Monitoring	15
<hr/>		
06	Future Opportunities and Conclusion	18



01 Executive Summary



The evolving landscape of data analytics necessitates that organizations adopt more flexible, scalable, and cost-effective solutions. Traditional on-premises systems, while robust and reliable, often face limitations in handling the increasing volume and velocity of data. To thrive, organizations must modernize their critical applications, replacing legacy systems with cloud-native solutions that offer better scalability, fault tolerance, and security.

The whitepaper will discuss how Victoria’s Secret & Co and Tiger Analytics executed a digital transformation of VS&Co’s analytics ecosystem and outline the challenges, pitfalls and strategies to execute a seamless migration from the on-premises model to Azure and Snowflake.

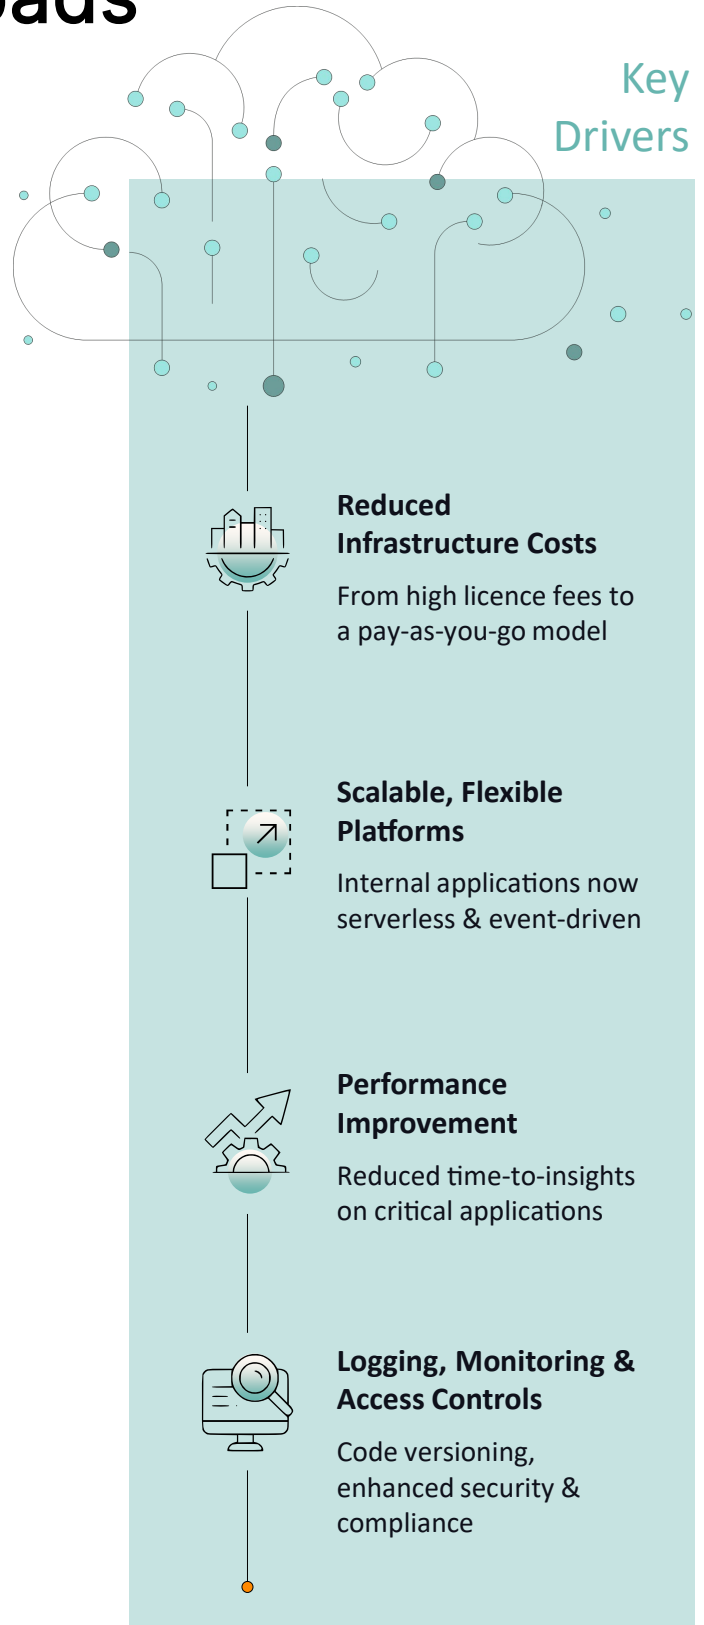
This whitepaper can be used as a standard playbook for any organization attempting to migrate from legacy data analytics platforms and adopt more sophisticated technologies. The framework to formulate and execute migration strategy has proven to be successful for VS&Co and can be tweaked for each organization based on their priorities, processes and current ecosystem.

Over a period of 10 months, VS&Co was able to transition out of legacy on-prem systems with 100% success rate on all workloads and save significant costs along the way.

02 Key Drivers for Migration and Modernizing Analytics Workloads

In today's fast-paced business environment, the ability to quickly adapt to consumer needs is paramount. Migrating analytics and reporting workloads from on-prem to cloud empowers data science teams to be more agile and address complex use cases with greater efficiency. Transitioning to the cloud provides a high-performance, scalable platform that supports growth, reduces infrastructure cost and fosters a culture of innovation. Businesses need to future-proof their operations, ensuring they remain competitive and responsive to evolving market demands.

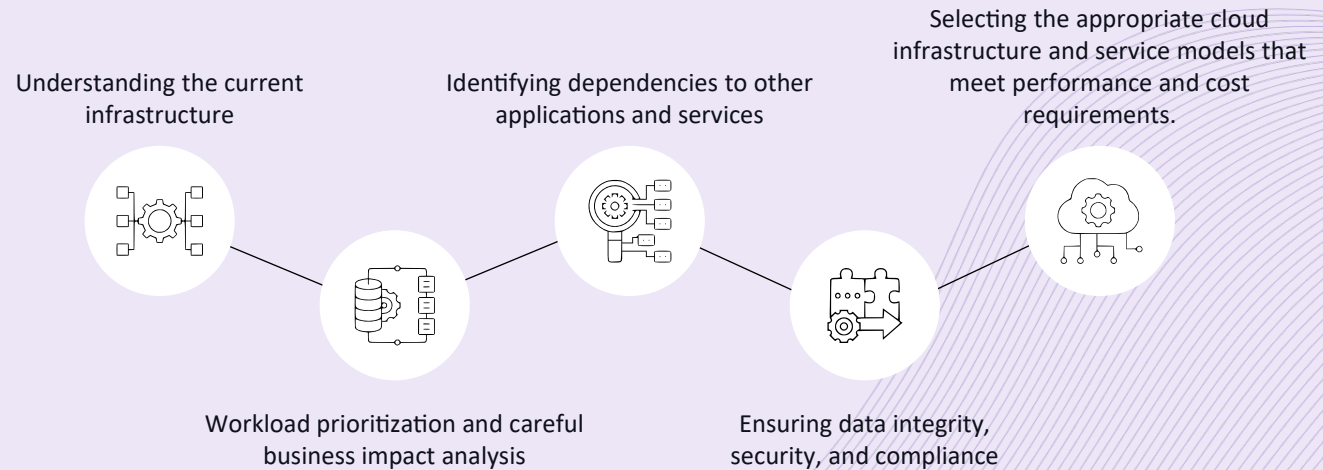
VS&Co is on a mission to build a modern ecosystem that will provide the right tools to Data Science, App Engineering and IT organizations to improve customer experience and enhance productivity. VS&Co adopted Microsoft Azure in 2022 and migrated their enterprise data from Teradata to Snowflake in 2023. With the right foundation in place, VS&Co seized the opportunity to build a modern data platform and shed the legacy technology debt.



03 Migration Strategy

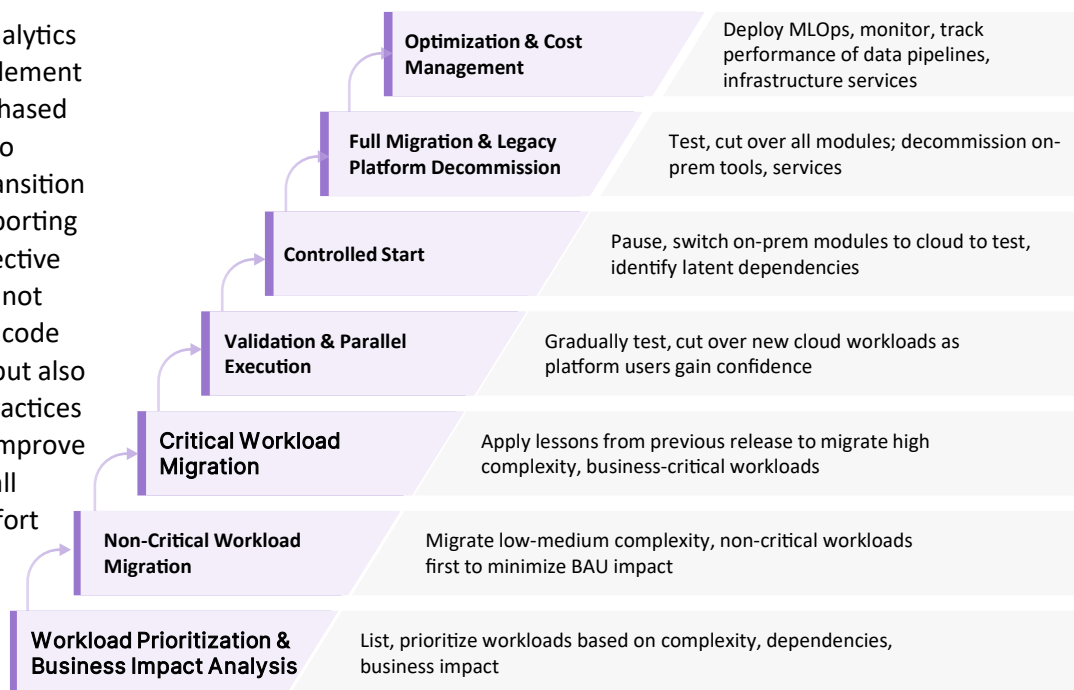
Adopting a migration strategy for moving workloads to the cloud involves thorough assessment and careful planning to align with business objectives. This ensures that the migration not only supports current business needs but also positions the organization for future growth and innovation.

Key Considerations



Effective change management, upskilling, and stakeholder engagement are also essential for a smooth transition, while continuous performance monitoring, cost management, and optimization help maintain efficiency and keeps costs in check in the new cloud environment.

VS&Co and Tiger Analytics collaborated to implement a well-structured, phased migration strategy to ensure a smooth transition of analytical and reporting workloads. The objective of this exercise was not only to refactor the code to a new platform, but also to introduce best practices and efficiencies to improve current processes, all while minimizing effort and costs.



Workload Prioritization and Business Impact Analysis

On-prem R, Python and SAS platforms had been vital in supporting VS&Co with day-to-day operations, as well as with larger and more strategic initiatives.

Data prep, filtering and feature engineering to create Business KPIs

Summary Reports and Performance Tracking



Calculate expected value of the customer spend in near future

Demand Forecasting



Product Affinity, Email affinity, Market Basket

Affinity Analysis



Help the Legal team manage customer info

GDPR Compliance



Customer Segmentation

Split customers in multiple segments based on their transactions, loyalty and demographics



Promo and Campaign Analysis

Promo performance, Targeting, Marketing Mix Modelling



Web Behavior and Funnel Analysis

Digital transactions, page/category views, funnel metrics

The first step in the migration strategy was a comprehensive workload prioritization and business impact analysis. This process involved collaboration with business stakeholders, IT and the Data Science (DS) team to thoroughly understand the current landscape.

The following factors were considered during this phase:



Number of Scripts per Module

Each module was analyzed to determine the volume of scripts involved, which provided insight into the effort required for migration.



Need for Automation

The necessity for automation within each module was assessed to streamline processes and reduce manual intervention during and after the migration.



Module Dependencies

Understanding the interdependencies between different modules was crucial to ensure that migrating one module did not adversely affect the functionality of others.



Impact on BAU


The overall impact of each module on daily business operations was analyzed to prioritize modules that could be migrated with minimal disruption to BAU activities.



Complexity of Implementation

Modules were evaluated for complexity based on factors such as the intricacy of logic, data transformations, and integration points with other on-prem systems.


This detailed analysis enabled the team to prioritize workloads effectively, ensuring a strategic and phased approach to migration.



PHASE 1 MIGRATION

Non-critical Workloads

Phase 1 of the migration involved the refactoring and rebuilding of applications from on-prem/SAS to cloud-native technologies. Non-critical and low-medium complexity workloads were picked as they offer the following **advantages**:



Mitigates Risk to BAU

Low-impact workloads reduce risk of business disruptions and allow the team to be comfortable with the new systems, understand potential pitfalls and refine the migration process early on



Process Refinement and Validation

Identifying challenges and opportunities for optimization on performance, reliability, and security of the new system early on ensures the migration of high impact workloads goes smoothly



Building Momentum and Confidence

Successfully migrating low-impact workloads helps build confidence within the team and other stakeholders and provides evidence to the efficacy of new system



This phase focused on the following activities:

Code Conversion and Refactoring

Each module's SAS scripts were meticulously converted to Snowpark, ensuring that the logic and functionality were preserved. Tiger Analytics utilized bespoke accelerators to speed up code refactoring and conversion on most commonly used SAS procedures and functionalities.

PROC FREQ

PROC COMPARE

PROC MEANS

SAS vs Python output validation

Send automated emails to a distribution list

Building Pipelines and Automation

Identifying modules that require automation and building data and ML pipelines, orchestrating workflows, and setting up event-based and scheduled runs.

Cutover Preparation

Once a module was thoroughly tested and validated, it was prepared for cutover. This involved finalizing the conversion, ensuring all dependencies are resolved, and making integrations are made to transition to the production environment.

Testing and Validation

The converted scripts underwent rigorous testing to validate their accuracy and performance. Data scientists were actively involved in this process to ensure that the outputs aligned with the expected results.

Performance Optimization

Supporting each application with right compute and memory requirements that provides optimum performance at a reasonable cost.

PHASE 2 MIGRATION

Critical Workloads

Phase 2 followed a similar approach to Phase 1 but focused on medium-high impact modules and SAS Stored Processes. This phase ensured that all analytical workloads were converted and validated before the final cutover. Key steps included:



Continued Code Conversion and Refactoring

The remaining SAS scripts were converted to Snowpark, maintaining the integrity and functionality of each module.



Ongoing Testing and Validation

Continuous testing and validation were conducted for each converted module, with data scientists ensuring the accuracy and reliability of outputs.



App Hosting

SAS Stored Procedures were redesigned to Streamlit apps that run on Azure App Services, hosted on VS's custom domain, and backed by stringent access controls and security protocols.



Final Cutover

Modules that passed validation were transitioned to the production environment. This phase ensured that all workloads were running smoothly on Snowpark.



POST-MIGRATION ACTIVITIES

Once all modules were successfully migrated and replicated in the new environment, it was essential to ensure that the old platform and services could be sunset without affecting business operations.

VS&Co and Tiger Analytics undertook the following steps to decommission SAS and on-prem servers:



Control Start

After replicating all applications, switch off select workloads on SAS/ on-prem and switch on the applications setup in Azure and Snowflake to identify potential points of failure, hidden dependencies and ensure integration with new services is successful.



Decommissioning of legacy platform

Once the control start is successful, the legacy platform and its applications are ready to be decommissioned. VS&Co was able to decommission SAS with no impact on operational and strategic workloads.

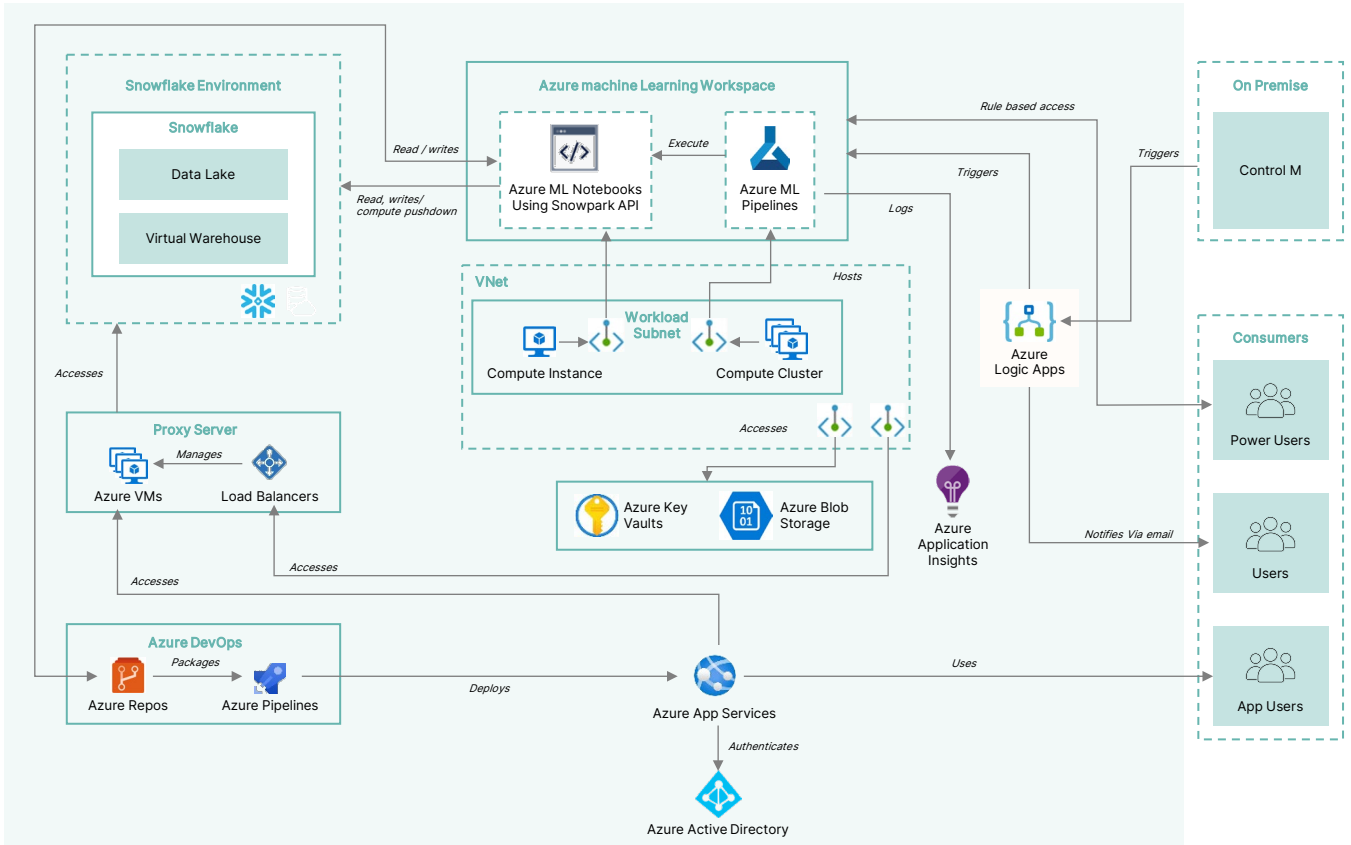


Monitoring and Optimization

With the new managed infrastructure in place, continuous monitoring and tracking were important for business continuity and identifying further opportunities for optimization. Tiger Analytics helped VS&Co set up a Platform Monitoring Dashboard to monitor compute and cost consumption across each application and track usage on both Azure and Snowflake.

This phased migration strategy ensured a systematic and controlled transition from on-prem to cloud, minimizing risks and disruptions while maximizing the benefits of the new analytics environment.

04 Snowflake and Azure for Migration and Modernization



New Architecture Overview

The migration from on-prem to cloud introduced a streamlined and efficient architecture that leveraged Azure services for scalable and secure setup while maximizing the use of Snowflake's powerful compute resources.

The key components of the new architecture are as follows:

Azure VM Instances

Setup and Connectivity: VM instances on Azure were configured to connect to Snowflake over a public network. The IP addresses of these VMs were whitelisted on Snowflake to ensure secure and seamless connectivity.

Workload Execution: These VMs were tasked with running analytics and reporting workloads. They effectively managed data processing tasks, leveraging Snowflake's compute power.

Snowflake Compute (Virtual Warehouses)

The architecture made extensive use of Snowflake's Virtual Warehouses (VWs), which are scalable compute resources that can be dynamically adjusted based on workload requirements.

Application Hosting

SAS Stored Processes to Streamlit Apps: The SAS stored processes were converted to Streamlit applications, enhancing usability and integration with modern web technologies.

Hosting on Azure App Services: These Streamlit apps were containerized using Docker and hosted on Azure App Services, ensuring high availability and scalability. They were hosted on a custom domain to enable SSO and accessed through a private endpoint for added security.

Virtual Warehouses Assignment and Performance Evaluation

The assignment of Virtual Warehouses (VWs) to each module was a critical aspect of the migration strategy. Here's how the process was executed:

Performance Evaluation

Runtime Comparison: The team conducted comprehensive performance evaluations, comparing module runtimes on Snowpark across different VW sizes (Small, Medium, Large, X-Large) with the runtimes on SAS.

Evaluation Criteria: The evaluation considered multiple factors, including the criticality of each module, Service Level Agreements (SLAs), and the cost-performance balance.

Assignment Based on Criticality and SLAs

Critical Modules: Modules with stringent SLAs and high business impact were assigned larger VWs to ensure rapid processing times and meet performance expectations.

Cost vs. Performance: The decision-making process involved balancing costs incurred against the performance gains, ensuring efficient use of resources and optimizing both performance and cost.



Security Implementation

Ensuring robust security was a critical aspect of the new architecture, especially as these workloads connected to sensitive PII data.

Single Sign-On (SSO) with Azure AD

The applications were integrated with Azure Active Directory (Azure AD) for Single Sign-On (SSO). This integration provided a seamless and secure login experience, enhancing user convenience and security.

Secure Credentials Management

To secure connections between Azure VMs and Snowflake, user credentials were stored in Azure Key Vault. This ensured sensitive information was protected and managed according to best practices.

Allowlisted IPs

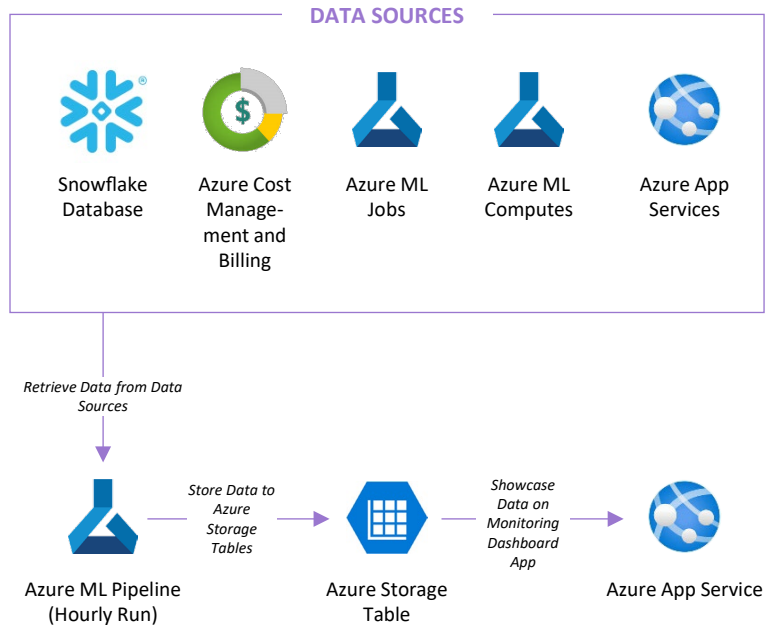
The IP addresses of the VMs were whitelisted on Snowflake, so that only authorized machines could connect to the Snowflake environment. This measure added a layer of network security to protect against unauthorized access.

Performance Optimization and Monitoring

Optimization and monitoring are not one-time activities but ongoing processes that drive continuous improvement and innovation. By regularly analyzing application performance and compute usage patterns, organizations can identify areas for enhancement and implement best practices to optimize their cloud environment. This approach enables businesses to stay agile and respond to changing demands.

VS&Co recognized this need and worked with Tiger Analytics to build a comprehensive dashboard. This dashboard delivers near-real-time updates on cost and compute consumption, runtime performance benchmarks, and failure events alerts and maintaining an ongoing log of executed pipelines.

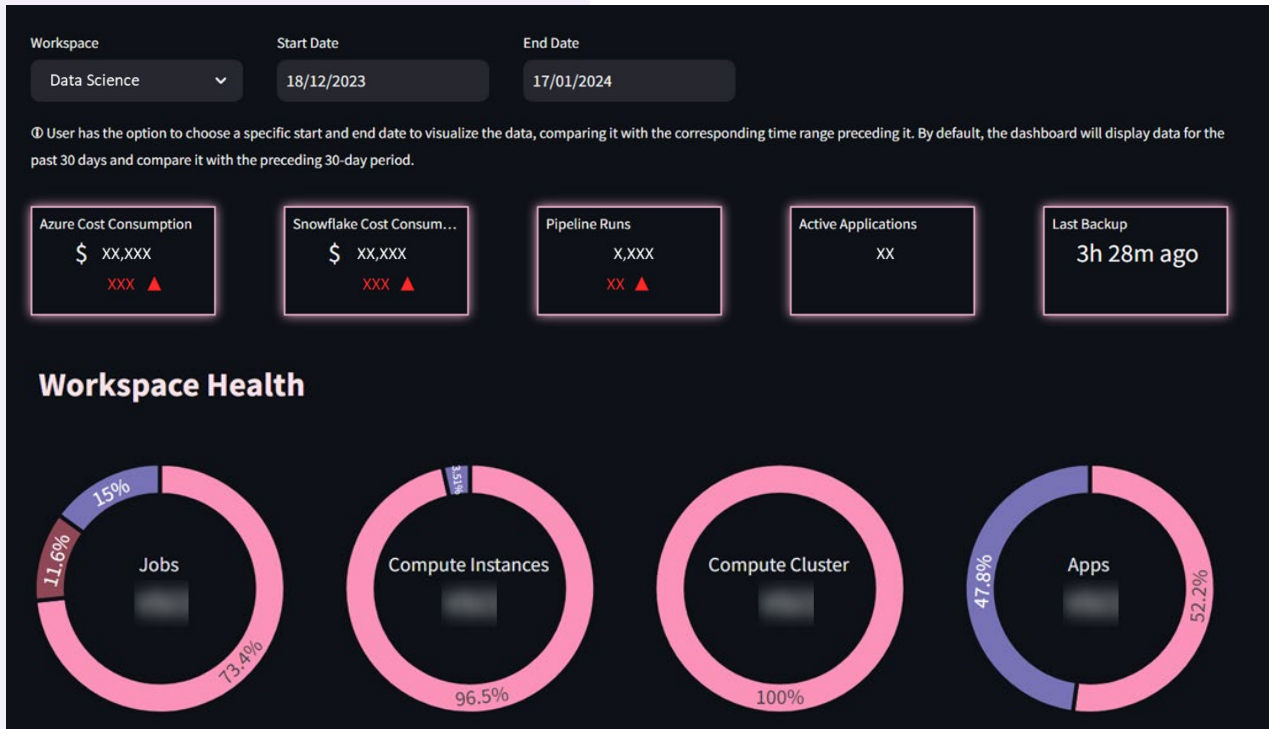
The custom dashboard was built using Streamlit and hosted on Azure App Services. The app connected to necessary APIs and data sources to collect logs and events.



Here are the key components delivered by the Platform Monitoring Dashboard:

Summarized view of Cost consumption, pipeline runs and active applications

This view enables senior stakeholders to monitor and track resource consumption metrics and flag any upticks.



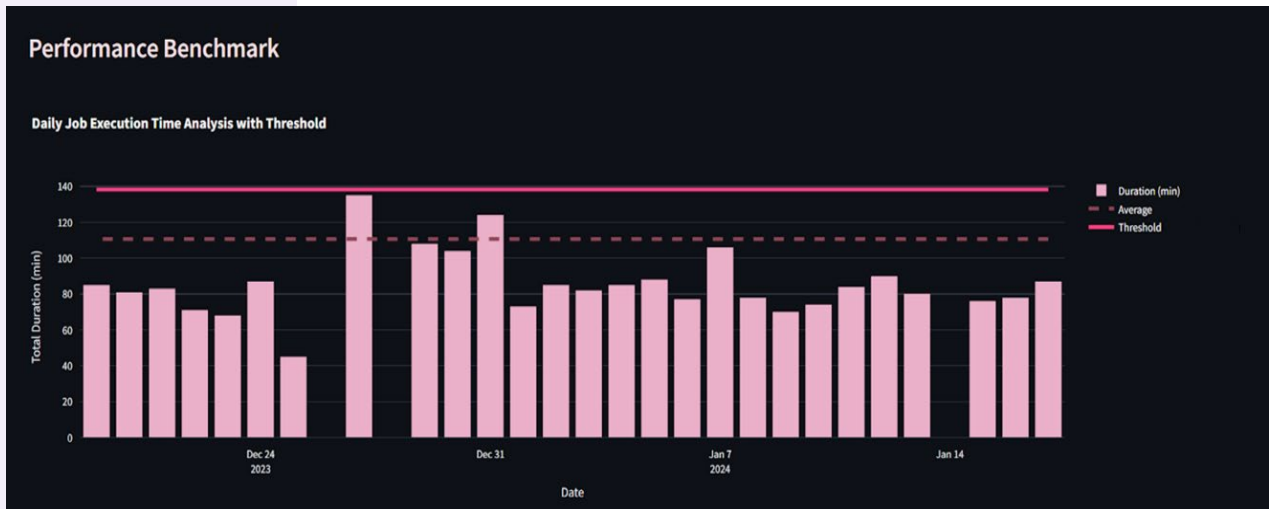
The table below provides information on all pipeline runs as well as compute cluster uptime and utilization. This helps IT Ops understand the status of the job executed and helps identify idle compute resources that result in unnecessary costs.

Resource Utilization and Status

Resource		Status					
Pipeline		All					
Experiment	Name	Status	Created On (EST)	Created By	Duration (min)	Job Type	
0		Completed	2024-01-16 11:41:42-05:00		472	Pipeline	
1		Completed	2024-01-16 08:42:08-05:00		130	Pipeline	
2		Completed	2024-01-16 08:42:07-05:00		40	Pipeline	
3		Completed	2024-01-16 07:26:50-05:00		75	Pipeline	
4		Completed	2024-01-16 07:12:11-05:00		14	Pipeline	
5		Completed	2024-01-16 06:01:00-05:00		78	Pipeline	
6		Completed	2024-01-16 03:16:32-05:00		6	Pipeline	
7		Failed	2024-01-16 03:08:34-05:00		0	Pipeline	
8		Failed	2024-01-16 03:01:40-05:00		1	Pipeline	
9		Completed	2024-01-15 09:39:48-05:00		84	Pipeline	

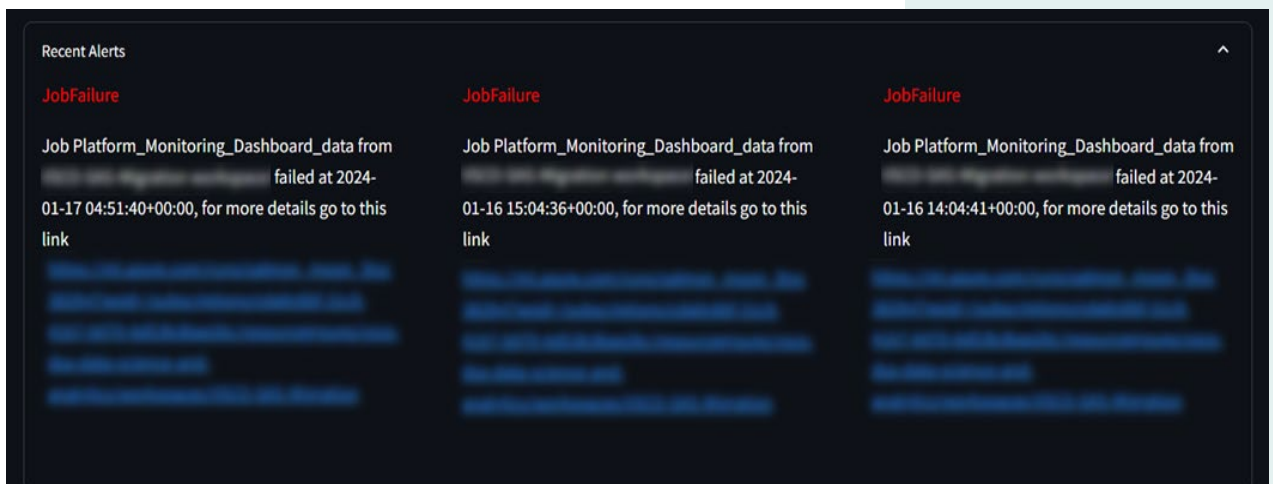
Runtime performance benchmark

Each pipeline's runtime is measured and evaluated against its previous runs. This helps understand if the runtime is exceeding previous runtimes or breaching a predefined threshold. This could either be due to a feature/functionality change or the infra services coming under a strain. By performing trend analysis, the IT Operations team can pinpoint anomalies and investigate further to ensure optimal performance and stability.



Providing alerts on critical workloads helps in timely mitigation of issues. For automated pipelines, this becomes vital as there might be a domino effect on other services getting affected by such events. Similar alerts are also published in an email to sys admin users.

Critical Alerts



06 Future

Opportunities and Conclusion

The migration has set the stage for exciting opportunities to further enhance and optimize this setup. So far, significant benefits have been realized in the 6 months since migration, including a 34% reduction in compute costs, the establishment of a secure and scalable platform, and high adoption within the Data Science team to train, test and deploy critical workloads on Azure and Snowflake stack.

Building upon these successes and embracing cutting-edge technologies will enable a more sophisticated data analytics and machine learning ecosystem, which in turn will drive better and faster decision making an enhanced customer experience. Listed aside are some of the key opportunities and use cases that VS&Co is exploring since migration.



Automated Model Deployment

Implement CI/CD for all analytics, ML workloads



Snowpark ML

Train, deploy ML models in Snowflake to reduce data movement, latency



Real-Time Analytics

Use dynamic tables, views to provide immediate insights for decision-making



Feature Store

Ensure feature consistency, reusability across models, teams



Vector Store

Build in vector store to support Generative AI use cases

In conclusion, organizations embark on cloud migration to enhance business agility through automation, provide richer customer experiences, and introduce new products and services. Pivoting towards a modern, scalable, and reliable ecosystem is crucial, as it allows organizations to quickly identify and capitalize on business opportunities and generate value for internal and external stakeholders. VS&Co has embarked on this journey and are already reaping benefits of the modernized analytics and machine learning platform by saving significant costs, increasing productivity and establishing industry best practices as the norm.

About the Authors




Venkat Gundu

Manager
Apps and Data Services
Victoria's Secret & Co.

Venkat Gundu is an Applications Manager at Victoria's Secret & Co, where he leads Data Analytics, Enterprise Data Integration and Machine Learning platforms. He has expertise in steering strategic initiatives that transform data into actionable insights, streamline operations, and foster innovation.

Venkat is always learning new technologies that lead to desired business outcome for customers – e.g. AI/ML, GenAI and Data Analytics. He also enjoys traveling to new places and playing tennis whenever he can find time.



Sahil Khanna

Senior Manager Analytics
Consulting
Tiger Analytics

Sahil Khanna is a Senior Manager, Analytics Consulting at Tiger Analytics where he helps customers achieve digital transformation, adopt emerging technologies and solve complex business problems through the use of Gen AI and machine learning. He's an avid Formula 1 enthusiast and enjoys watching movies in his spare time.

About Tiger Analytics

Tiger Analytics is a global leader in AI and analytics, helping Fortune 1000 companies solve their toughest challenges. We offer full-stack AI and analytics services & solutions to help businesses achieve real outcomes and value at scale. We are on a mission to push the boundaries of what AI and analytics can do to help enterprises navigate uncertainty and move forward decisively. Our purpose is to **provide certainty to shape a better tomorrow.**

Being a recipient of multiple industry awards and recognitions, we have 4000+ technologists and consultants, working from multiple cities in 5 continents.

www.tigeranalytics.com