



**EMERGERE TECHNOLOGIES**  
TECHNOLOGY ONWARDS AND UPWARDS



## **WHITEPAPER**

---

# **METL MIGRATION TO DATA PRODUCTIVITY CLOUD (DPC): A HYBRID AZURE STRATEGY FOR RETAIL CUSTOMER**



# METL MIGRATION TO DATA PRODUCTIVITY CLOUD (DPC): A HYBRID AZURE STRATEGY FOR RETAIL CUSTOMER

## EXECUTIVE SUMMARY

A leading retail organization successfully migrated its monolithic Matillion ETL workloads from a constrained Azure Virtual Machine (VM) deployment to the elastic, modern Matillion Data Productivity Cloud (DPC). This strategic move, executed in partnership with Emergere Technologies, leveraged a hybrid architecture utilizing Azure Kubernetes Service (AKS) for containerized execution.

The migration achieved immediate cost reductions by eliminating infrastructure management overhead and delivered critical operational agility through Kubernetes-driven elasticity and native optimization within Snowflake.



*By integrating Matillion's generative AI layer, Maia, the client accelerated pipeline creation and ensured superior resource governance, establishing a future-proof data supply chain that maximizes return on their cloud investments.*



## INTRODUCTION

This whitepaper outlines the strategic need and successful execution of a leading retail customer's transition of existing Matillion ETL (METL) workloads to the Matillion Data Productivity Cloud (DPC). This migration established a hybrid SaaS model utilizing Azure Kubernetes Service (AKS) for execution, with a clear roadmap already defined for the final transition to Azure Container Apps.

This move represented a fundamental shift toward a cloud-native, elastic, and AI-powered data operating model that was fully implemented to achieve complete customer understanding through Snowflake and maximize operational agility.

*"The greatest danger in times of turbulence is not the turbulence; it is to act with yesterday's logic."*

*- Peter Drucker, Father of modern management.*





## EMERGERE TECHNOLOGIES: THE EXECUTION PARTNER

Emergere Technologies served as the dedicated implementation partner, specializing in data transformation, cloud enablement, and modern data stack architecture. The firm's role in this project was central, covering strategic consulting, detailed execution planning, and the technical refactoring required for the successful shift from the legacy Azure VM architecture to the cloud-native DPC environment.



The expertise of our certified team was critical in

### STRATEGIC BLUEPRINTING

All existing Matillion Projects and custom scripts that required rewriting were documented.

### AZURE CONTAINER PROVISIONING

Successfully provisioning and integrating the Customer-Hosted Agents (CHAs) within the retail customer's existing Azure Kubernetes Service (AKS) environment.

### MAIA INTEGRATION AND OPTIMIZATION

Leveraging DPC features, including the Maia AI engine, to refactor legacy ETL jobs into highly optimized, cost-efficient data pipelines in Snowflake.

This partnership ensured that the migration was not just a lift and shift, but a fundamental redesign that maximized the benefits of Azure's container services and Matillion's Data Productivity Cloud.



## THE STRATEGIC IMPERATIVE: DECOUPLING, ELASTICITY, AND AZURE ALIGNMENT



The prior METL architecture, an Azure Virtual Machine (VM) deployment (using an image from the Azure Marketplace or a custom VHD), was proven to clash with the organization's modern Azure strategy:

**Fixed Compute Footprint:** Infrastructure scaling was tied to a single VM, which led to high wasted costs during slow times and resource crashes during peak retail events. This lack of elasticity often halted urgent analysis.

**Infrastructure Headaches:** The data team was burdened with basic maintenance work including patching the VM, managing the OS. This time was diverted away from high-value data pipelines.

**Non-Native Azure Architecture:** The METL setup did not natively utilize Azure's modern container services (AKS/Container Apps), resulting in missed benefits related to superior resource orchestration and automated scaling.

The migration to DPC successfully resolved these problems by establishing a true SaaS Control Center that separates the management layer from the execution layer (Agents). These Agents were deployed on Azure's platform, delivering a significant, measurable jump in Data Productivity.



## KEY BENEFITS ACHIEVED USING SNOWFLAKE AND AZURE

| Feature              | Matillion ETL (METL) – Prior State | Matillion Data Productivity Cloud (DPC) on Azure – Current State     | Outcome  |
|----------------------|------------------------------------|--|--|
| Execution Platform   | Fixed VM/AMI                       | Customer-Hosted Agents (CHA) on AKS (Path to Container Apps defined) | Fully aligned with Azure investments; future-proof container strategy        |
| Data Warehouse       | Standard Snowflake connection      | Native, Optimized Snowflake Pushdown                                 | Maximized performance; minimized Snowflake warehouse usage and costs         |
| Scaling & Resilience | Manual, VM-limited                 | Kubernetes-driven Elasticity   | Instant, automatic scaling; absorbed peak sales loads with zero contention   |
| AI Integration       | Limited custom code                | Maia Integration (AI Data Assistant)                                 | Faster pipeline creation; smart error-fixing; optimized Snowflake SQL output |



## THE MAIA FACTOR: SUPERCHARGING DATA PIPELINES

The deep use of Maia, Matillion's generative AI layer, was a critical differentiator for DPC. For the retail customer, Maia was integrated across the entire data process and successfully drove down the **Time-to-Insight (TTI)** within the Snowflake environment.

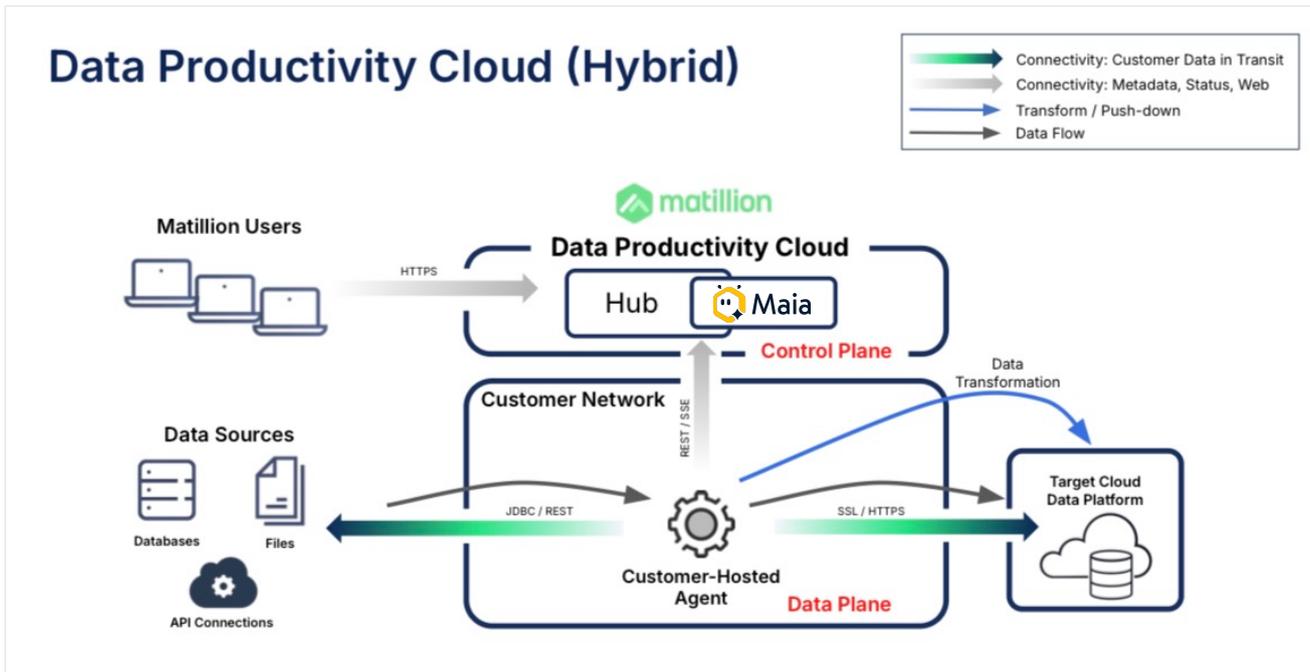


Image Credit : matillion

### PIPELINE AUTHORIZING AND TRANSFORMATION:

Maia successfully translated natural language descriptions of desired outcomes (e.g., finding high-value customers) into efficient Matillion Transformation Job blueprints, resulting in highly optimized Snowflake SQL for maximum performance.

### EXECUTION AND AZURE RESOURCE MANAGEMENT:

Maia monitored pipeline execution profiles and recommended the optimal resource settings for the AKS Pods/Container Apps running the Agents, ensuring efficient utilization of the Azure compute layer.

### PROACTIVE DATA QUALITY:

Maia automatically suggested and integrated Check components into pipelines to verify data quality rules before data landed in Snowflake's consumption layer, guaranteeing data accuracy.

**Autonomous Debugging:** When a job failed, Maia provided root-cause analysis based on logs from both the pipeline and the underlying AKS Agent, successfully suggesting the fix.

### CONTAINER APP READINESS:

Maia flagged pipeline patterns that were highly suitable for the future move from AKS to the simpler, serverless Azure Container Apps, smoothing the transition path.



## MIGRATION PLAYBOOK:

# A PHASED AZURE STRATEGY (EXECUTION SUMMARY)

The move from METL to DPC utilized a structured approach that successfully ensured business continuity and data integrity. A Trickle Migration strategy was implemented, running critical METL jobs in parallel with their DPC/AKS counterparts for a rigorous validation period.

### PHASE I: PREPARATION AND AZURE AGENT PROVISIONING

#### INVENTORY AUDIT

All existing Matillion Projects and custom scripts that required rewriting were documented.

#### AZURE SETUP (HYBRID DPC)

The DPC Tenant was provisioned, and the Customer-Hosted Agents (CHAs) were successfully deployed onto the existing AKS cluster.

#### SNOWFLAKE CONNECTIVITY

Optimized connections to Snowflake were reinitiated, and appropriate role-based access control (RBAC) was configured for the new DPC Agents.

### PHASE II: PHASED WORKLOAD MIGRATION AND AKS VALIDATION

#### PILOT SELECTION (LOW-RISK)

Non-critical reporting pipelines were used to establish the migration pattern and validate Agent performance within AKS.

#### SELF-SERVE MIGRATION

The native Matillion migration tool was used for lift-and-shift of existing logic.

#### REFACTORING FOR CONTAINERS

Slow METL jobs were identified and refactored to use DPC's concurrent, containerized execution, successfully maximizing performance within the AKS cluster.

### PHASE III: CUTOVER AND FUTURE-STATE ROADMAP

#### PARALLEL RUN AND DATA RECONCILIATION

DPC pipelines ran alongside METL jobs, and a rigorous Data Reconciliation process was executed.

#### BUSINESS VALIDATION

Sign-off was secured from key retail stakeholders confirming downstream reports were unaffected.

#### FUTUREPROOFING (CONTAINER APPS)

A small pilot was begun to test deploying a DPC Agent to Azure Container Apps, to validate the path for future decommissioning of the AKS cluster.

#### DECOMMISSIONING

Once stability was confirmed, the old METL VM will be terminated, realizing immediate Azure infrastructure cost savings.



## RISK MITIGATION AND CONSIDERATIONS (AZURE/SNOWFLAKE FOCUS):

### NETWORKING AND SECURITY (AZURE AKS):

| RISK  | MITIGATION   |
|---|--|
| DPC Agents lacked necessary network access to source systems or Snowflake due to restrictive enterprise Virtual Network (VNet) policies during initial setup. | A dedicated Azure VNet/Subnet review was conducted. The AKS cluster was deployed with Private Link connectivity to Snowflake (where required) and Network Security Groups (NSGs) were correctly configured to allow egress traffic only to required endpoints, strictly adhering to zero-trust principles. |

### SNOWFLAKE RESOURCE GOVERNANCE AND COST MANAGEMENT:

| RISK  | MITIGATION   |
|---|--|
| Inefficiently refactored pipelines threatened to trigger massive or sustained use of larger Snowflake Virtual Warehouses (VW), leading to unexpected consumption costs. | Strict DPC scheduling was implemented, and Maia-recommended job optimization techniques were used to manage and successfully reduce Snowflake VW usage based on workload priority and defined cost limits. |

### SECURITY CREDENTIAL EXPOSURE (TECHNICAL STACK):

| RISK  | MITIGATION   |
|---|--|
| Sensitive credentials (for Snowflake, source systems, etc.) were initially stored on the legacy VM and posed a security risk during migration and execution in the new environment. | All secrets were migrated to a centralized Azure Key Vault. DPC Agents running on AKS were configured with Managed Identity access to the Key Vault, eliminating hardcoded credentials and securing the environment. |



## CONCLUSION

By successfully adopting the Matillion Data Productivity Cloud on the Azure platform and utilizing the power of Snowflake, this retail customer built a future-proof data supply chain that is now flexible, intelligent, and primed for the next generation of AI-driven retail experiences. The deep integration of Maia transformed the operational cost of data engineering into a genuine strategic asset.

**Architect your data  
infrastructure with us!**

 +1 469 401 7117

 [hello@emergertech.com](mailto:hello@emergertech.com)

 [www.emergertech.com](http://www.emergertech.com)

 **US (Corporate Office)**

8105 Rasor Blvd, Suite 60, Plano, Texas 75024