ENGINEERING THE AI FLYWHEEL: INTEGRATING AI OUTCOMES INTO YOUR DATA LAKEHOUSE FOR SCALABLE INTELLIGENCE

Steven Mellare

Head of Data, AI & Architecture

Resimac Group

ABOUT RESIMAC

Resimac is one of Australia and New Zealand's premier non-bank lenders

Resimac Group Ltd (ASX:RMC) is one of Australia's longest standing non-bank lenders, recognised for our consistency, disciplined credit approach and strong performance across market cycles.

We originate through a national broker distribution network, giving us access to a wide and varied flow of customers. Our focus is on responsible credit decisions and high-quality underwriting, supported by long established processes and a specialised credit team.

Our portfolio mix provides scale, diversification and resilience.

- prime borrowers with strong credit profiles
- near prime customers who do not always meet traditional lender criteria
- commercial borrowers through our asset finance and business lending division

A SUCCESSFUL AI PILOT IS JUST THE BEGINNING

Scaling innovation for **sustained** ROI requires a networked AI architecture

ARCHITECTING FOR MOMENTUM

Isolated Models Can Lack Impact

Single AI pilots often solve narrow problems without broader integration, limiting sustained business value.

Siloed Outputs Inhibit Momentum

When AI outputs remain siloed, we prevent insights from influencing other processes, potentially leading to diminishing returns over time.

Networked AI Systems Offer Momentum

Organizations can create interconnected AI systems where models share outputs for cumulative learning and scaling.

FOUNDATION: UNIFIED DATA PLATFORM

ENABLING THE MOMENTUM NEEDED FOR SUSTAINED ROI

Unified Data Platform

A unified data platform enables seamless sharing of model outputs, telemetry, and insights across AI systems.

Creating network effects and exponential gains.

Data as First-Class Products

Model outputs and telemetry treated as versioned, governed, and discoverable data products that enable reuse and reliability. I.e. First Class Data Products

Interconnected AI Models

Models improve through feedback loops and accelerate other models' performance via shared signals.

Governance and Compliance

Embedding governance, lineage, and observability ensures compliance while enabling innovation.

AI Flywheel

Integrating AI models transforms decisions into reusable data products, fueling continuous innovation.

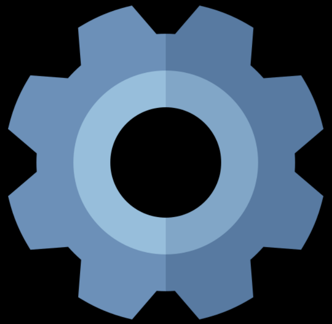

WHAT IS A FLY WHEEL?

A flywheel is a heavy rotating disk used in engines and machines to store rotational (kinetic) energy.

It uses the conservation of angular momentum

Over time, its own weight keeps it spinning with less effort.

At rest, it's hard to move, but once you start pushing consistently, it gains momentum.

FROM SINGLE-MODEL IMPROVEMENT TO A NETWORK EFFECT

AI Flywheel Concept

The AI flywheel creates a network effect where multiple AI models interact through a unified data platform for continuous improvement.

i.e. We minimise isolated pipeline architectures in favor of a **networked architecture**

It does this in 2 ways

- **Individual Feedback loop:** leverage single model telemetry to improve the AI feature store and improve a model's performance
- **Cross-Model Integration:** Outputs from one model serve as inputs or signals for others, enabling personalized and efficient business workflows.

The result: Sustained Momentum for Innovation

The flywheel effect enriches the data ecosystem, fostering smarter decisions and faster innovation across domains.

CONCEPTUAL COMPONENTS

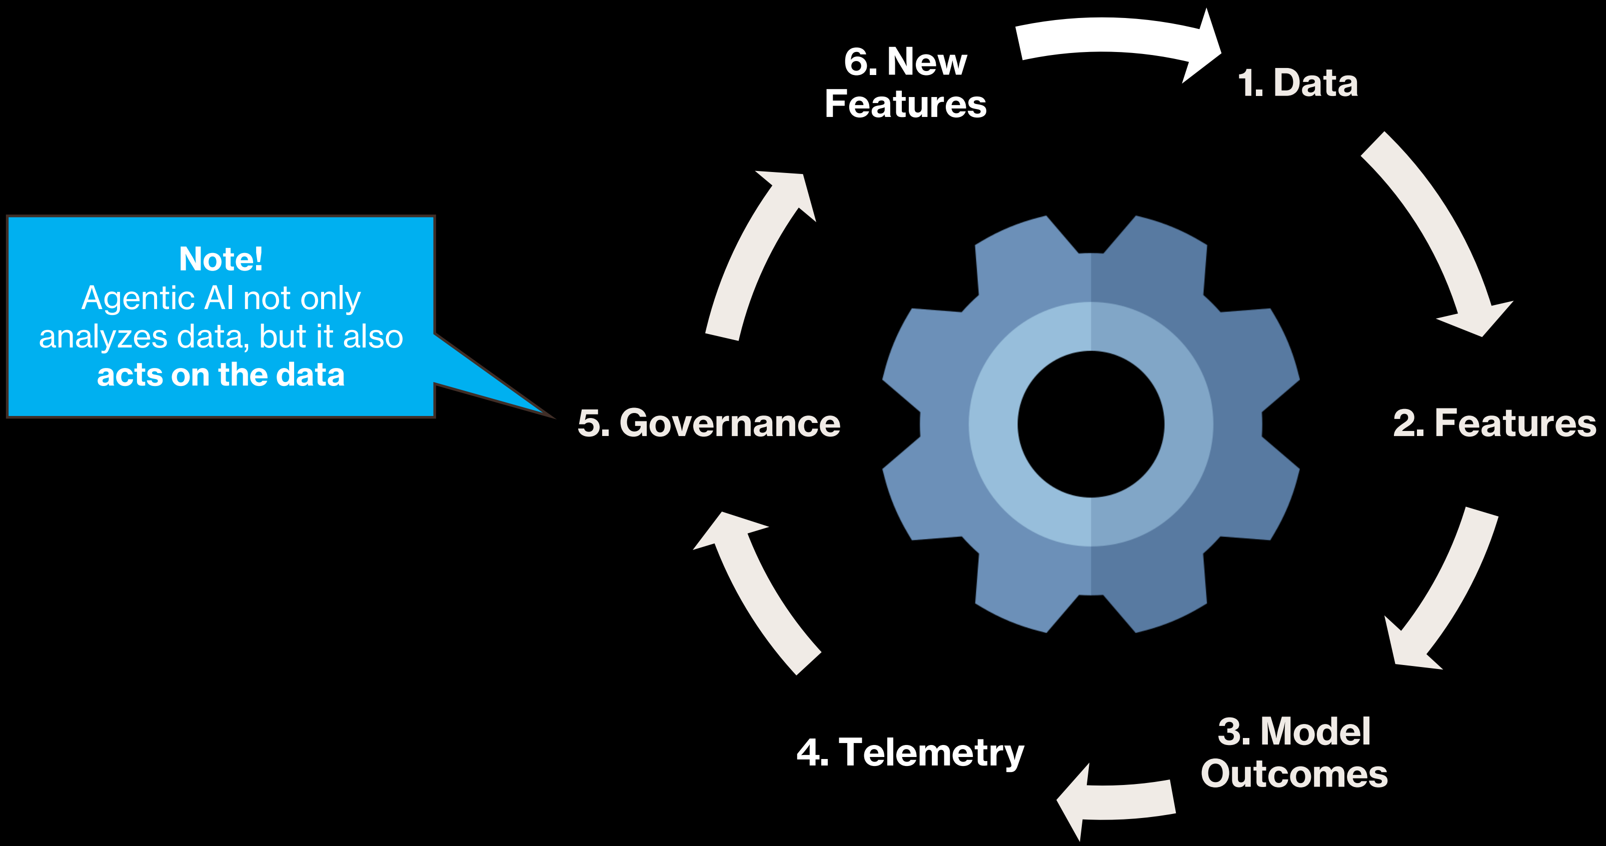

2 WAYS TO GENERATE MOMENTUM

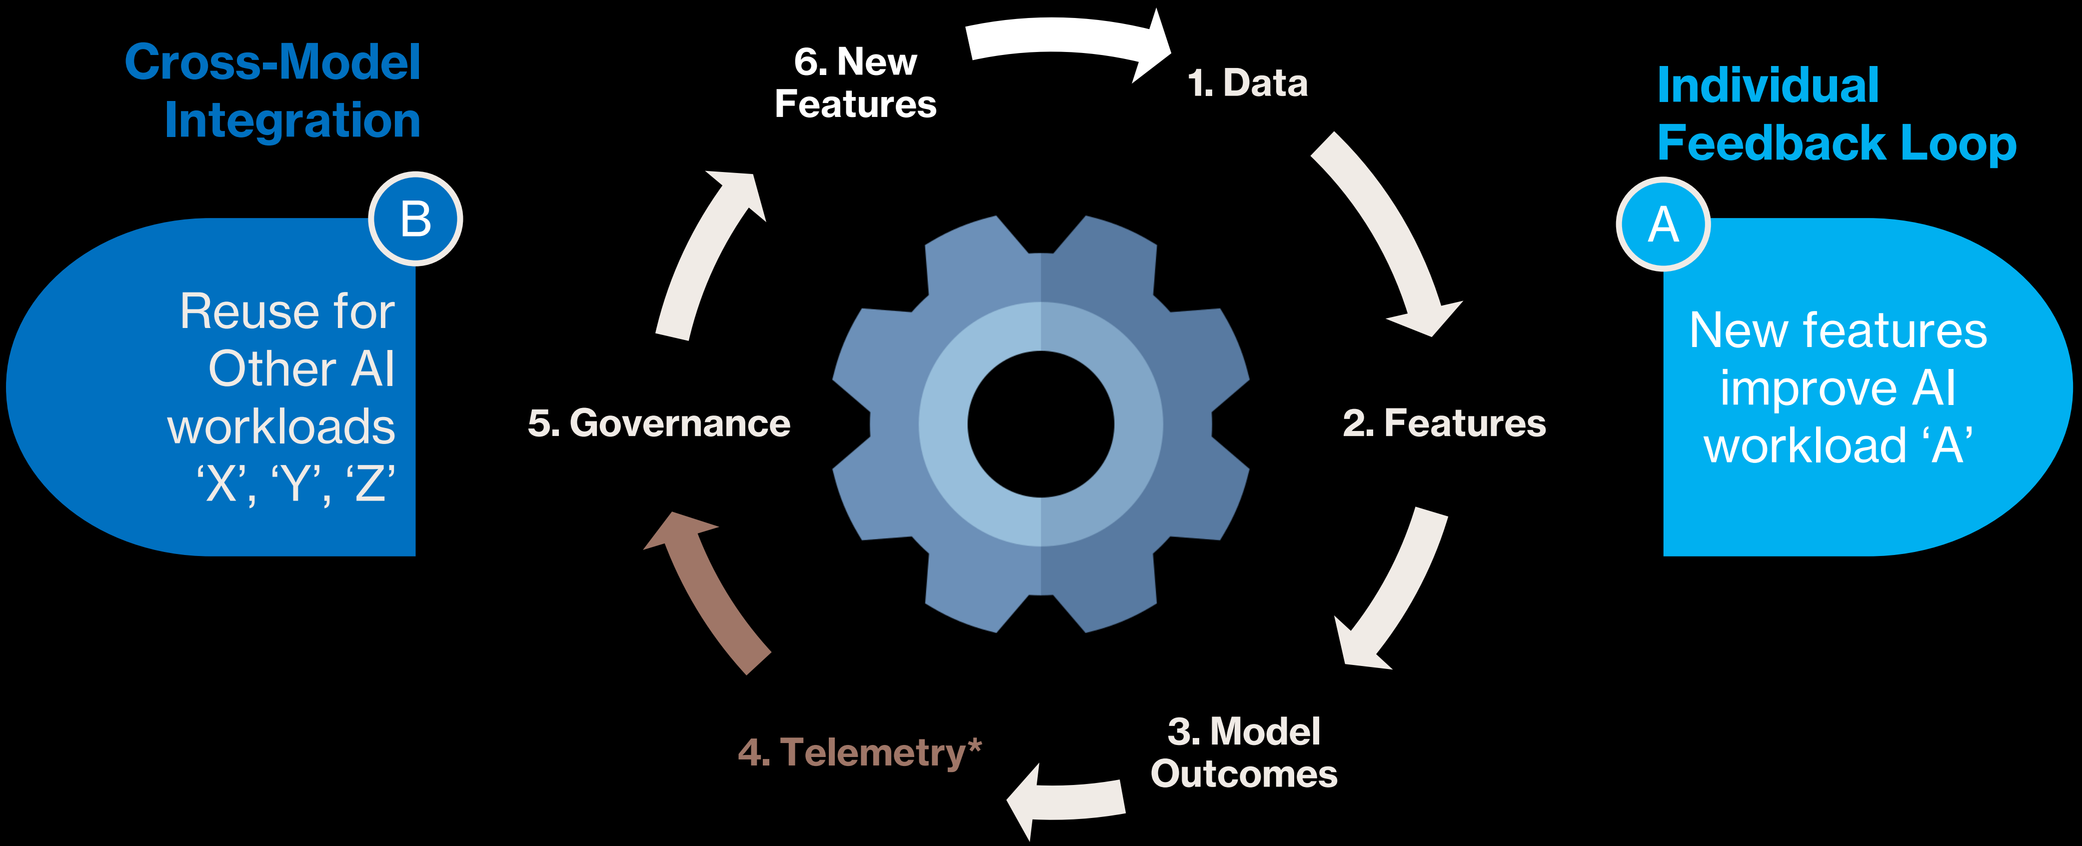

E.g. Ignore model 'A' predictions where a successful outcome was the result of a human override

CROSS-MODEL INTEGRATION: OPTIMISING MOMENTUM B

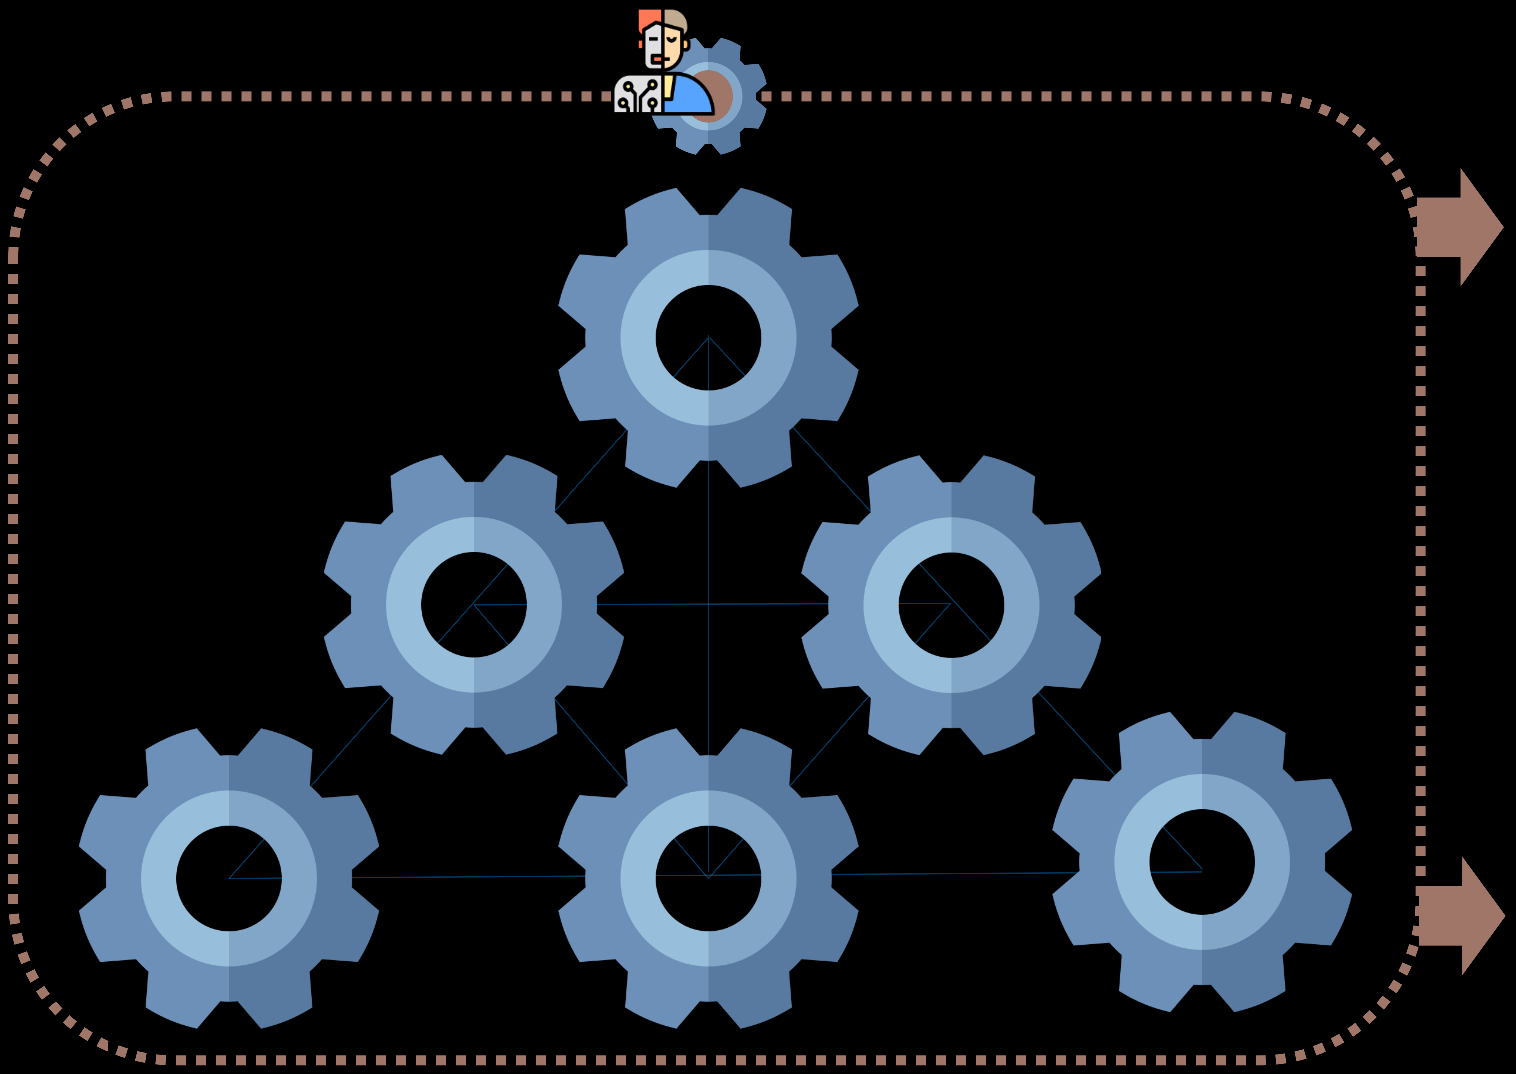

Agentic AI running the ‘engine room’

- Recommendations for the ‘cross pollination’ of telemetry generated features
- Optimising the order in which recommendations are implemented to maximise ROI as a function of the model’s network potential
- Managing fault tolerance in what’s become a very complex interdependent system

KPI Publication

- Publish key performance indicators with lineage visibility in BI dashboards to demonstrate business impact.

COMPONENT EXAMPLES

PROACTIVE CUSTOMER RETENTION → CX BOT → HARDSHIP RISK MODEL → COLLECTIONS → OUTCOMES

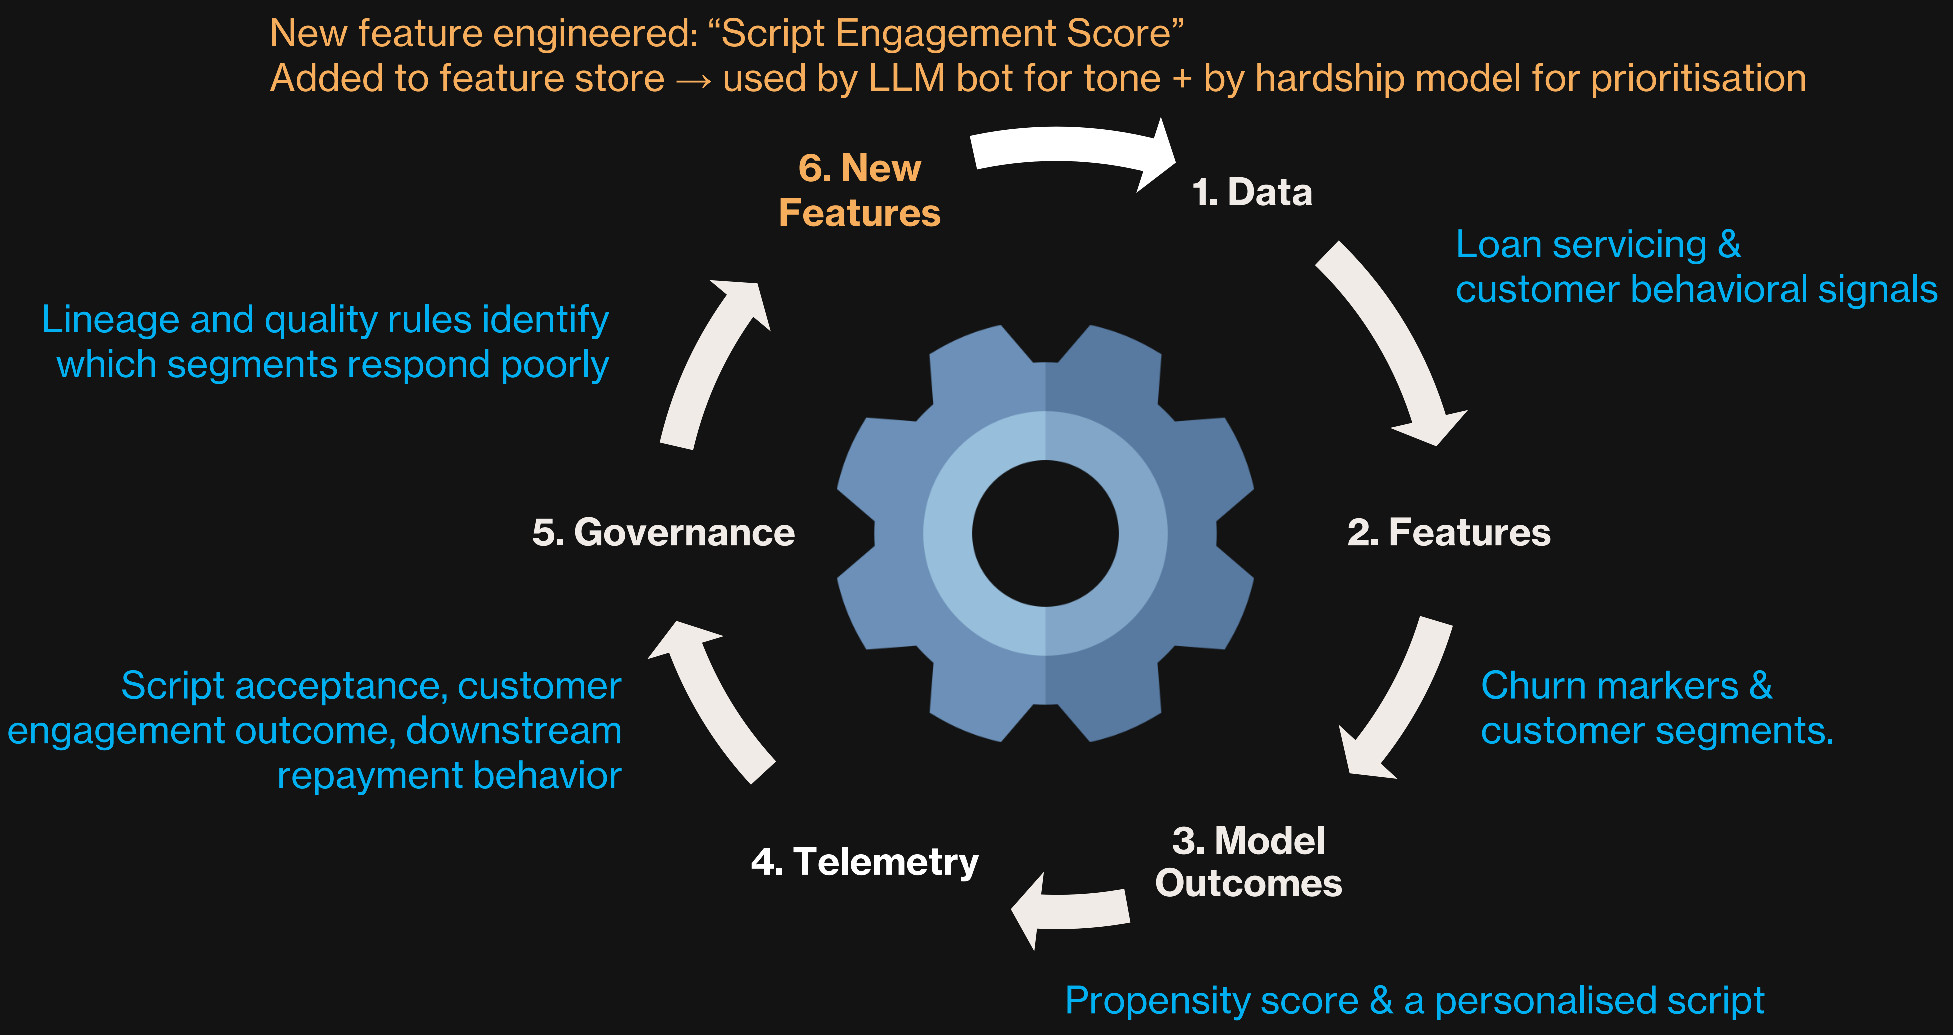

ROLE OF TELEMETRY AND MODEL OUTPUTS

Telemetry as “First Class Data Assets”

Telemetry captures decisions and real-world outcomes as structured, versioned data assets critical for model improvement.

Feedback Loop for Model Training

Logged telemetry events create a feedback loop that supports retraining and development of models.

Use Cases of Telemetry

Customer interactions and risk scores feed into predictive models enhancing marketing and retention strategies.

Continuous Learning Ecosystem

Telemetry transforms AI initiatives into continuously learning systems with compounding insights.

FIRST CLASS
DATA PRODUCTS

"First-class data products" is a synthesis of two well-established ideas:

1. **Data as a product** (from Data Mesh) – treat data/outputs like products with owners, contracts, and SLAs; and
2. **First-class citizen** (from programming language theory) – treat something with full rights (versioned, referenced, passed around, reused) rather than treated as an afterthought.

Together: Model outputs and telemetry are not ephemeral logs; they're governed, discoverable, versioned, and reusable data products in the lake-house – so other models and automations can consume them.

ARCHITECTING THE UNIFIED DATA PLATFORM

LAKEHOUSE ARCHITECTURE FOR CROSS-MODEL INTEGRATION

Unified Lakehouse Platform

The Lakehouse integrates the scalability of data lakes with the reliability of DV2 data warehouses for unified data management.

Data Governance and Pipeline Automation

Microsoft Purview manages data lineage and metadata, while Azure DevOps automates CI/CD for pipelines and deployments.

Feature Store and Reusable Data Products

Snowflake and Snowpark enable reusable, versioned features for batch and real-time inference, fostering cross-model data use.

Data Flow and Integration

Data flows through a 'data storefront' standardizing and transforming data before analytics and feature engineering in Snowflake.

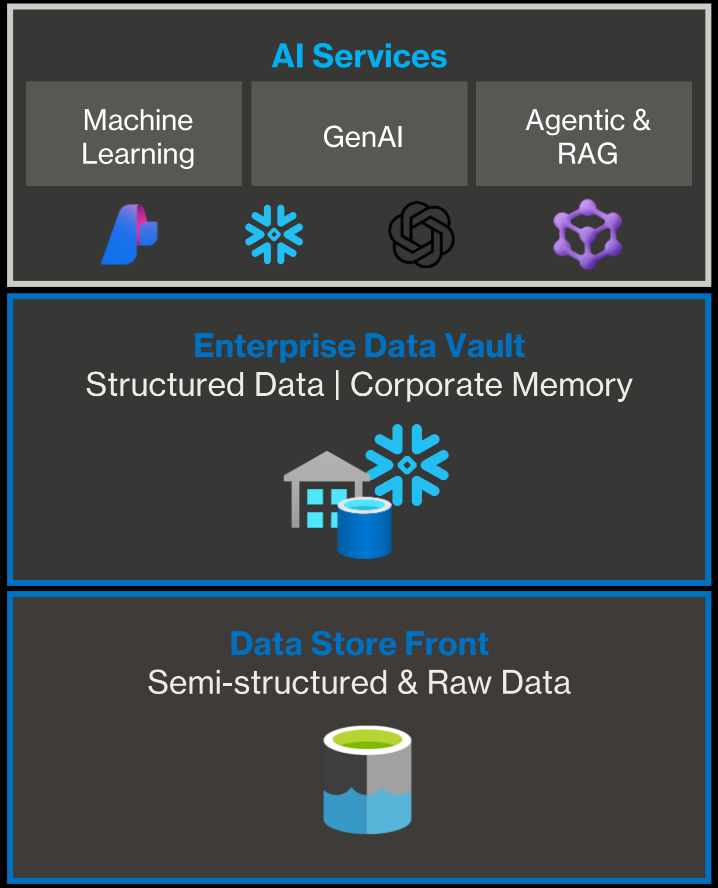

Resimac Data Hub | Unified Data Platform

DATA MODEL ENGINEERING TO ENABLE THE NETWORK EFFECT DATA VAULT.

Data Vault 2.0

Key Components:

1. **Hub & Links:** Maintain entity relationships and keys.
2. **Satellites**: Capture descriptive attributes and changes over time.
3. **Raw Vault**: Stores immutable, historical business keys and relationships. ***Includes Telemetry.***
4. **Business Vault**: Adds derived, calculated, and enriched data >> engineered features.

Feature Store as Feature Vault

1. **Features** stored in Business Vault Satellites
   - Each feature can be stored as a satellite linked to a hub (e.g., Customer, Loan, Vehicle).
   - Time-stamped changes allow point-in-time feature retrieval for training.
2. **Feature Groups** with Logical Links
   - Group related features under a link structure for specific ML models.
3. **Model Metadata** in Separate Hubs/Satellites
   - Feature definitions, owners, and versions stored in a metadata hub.

Feature Engineering:
Performed in Snowflake using Snowpark / Python.

Feature Store Layer:
Implemented as a dedicated schema or set of tables in Snowflake with metadata for:

- Feature definitions
- Versioning
- Time-aware snapshots

Access Patterns:

- Batch: For model training pipelines (via ADF or Azure ML).
- Real-Time: For scoring via Snowflake's external functions or Snowpark APIs

Key Practices

- Use time travel for historical feature snapshots.
- Implement versioning via metadata tables.
- Policy as code access control for governance.
- Integrate ADF pipelines for feature refresh and batch scoring.

Integration

- Training: Pull batch features from Snowflake into Azure ML.
- Inference: The Originations application will initiate asynchronous REST API calls to the UDP Azure ML instance.
- Monitoring: Track feature drift and model performance using Snowflake queries + PowerBI dashboards.

DATA MESH CONSIDERATIONS

- Data mesh **already demands strong metadata foundations**, especially because data products are created and owned across multiple domains.
- Without reliable cataloguing, lineage, and discoverability, the mesh quickly fragments
- The challenge for generating the network - ensuring domain data products are model-ready, integrating telemetry across domains, and enabling model outputs to flow back into the mesh as first-class data assets

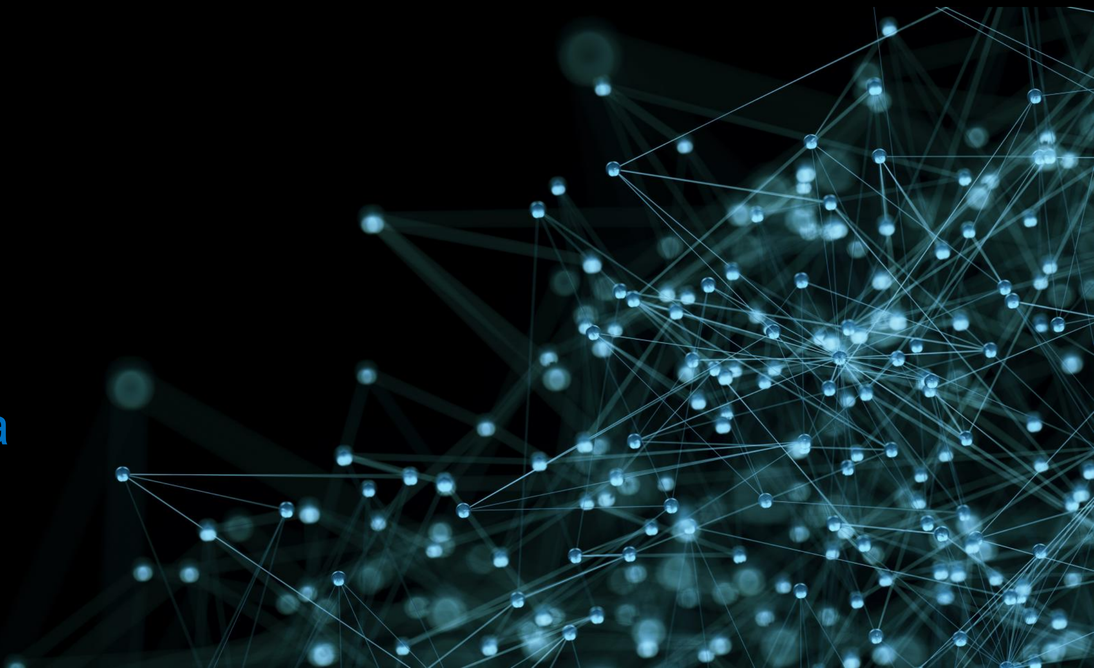

ARCHITECTURE ROADMAP

ROAD MAPPING THE FLYWHEEL

Architectural imperative: order workloads to amplify downstream model performance

1. Model Interdependency Map

Component	What they are	Why they matter	Resimac Context
Producers	AI workloads (or upstream services) that **emit reusable signals** – predictions, classifications, risk tiers, scripts, intents, embeddings, or enriched attributes.	Their outputs have **high reuse** potential across other models and processes; they seed the network effect.	Document Intelligence Agent (DIA) producing structured features from PDFs/bureau data used by scorecards and recommendation bot.
Consumers	Downstream **models, automations, or decision services** that **ingest producer outputs** as features or policies and generate new outcomes/telemetry.	They **compound** the flywheel by turning producer signals into decisions that are logged and reused again.	**LLM customer service bot** using segmentation/retention/scorecard signals to personalise tone and guidance.
Interfaces	The **contracted pathways** that make producer outputs **discoverable, reliable, and consumable** by other workloads.	They create the **operational fabric** of the flywheel – without robust interfaces, signals don't travel.	▪ Feature Store tables ▪ Telemetry/event tables ▪ APIs / orchestration ▪ Governance & discovery

ROAD MAPPING THE FLYWHEEL

Architectural imperative: order workloads to amplify downstream model performance

---

2. Calculate Networking (Flywheel) Score

- Reusability (R): number/variety of consuming models
- Consumer Count (C): confirmed downstream consumers
- Data Readiness (D): quality/coverage
- Business Impact (B): KPI lift potential
- Governance Complexity (G↓): lower complexity → earlier

Prioritise high (R×C×D×B)/G workloads

ROAD MAPPING THE FLYWHEEL

Architectural imperative: order workloads to amplify downstream model performance

3. Sequence in waves

#	AI Workload	Primary Outputs	Key Consumers	Networking Score (why)	Wave
1	DIA / Doc Intelligence	Structured features from documents	Scorecards, Collections, LLM Assist	High D; medium R; feeds many models	0
2	MOSAIC Segmentation	Segment, Persona	Marketing, **LLM Bot**, Offer Optimisation	**Very high R×C** across CX	1
3	Credit Scorecard (Vehicle)	Risk Tier, PD, Override	Collections, Pricing, Credit Bot	High impact on risk & ops	1
4	Elula Customer Retention	Churn risk, scripts, outcomes	Agentic CX flows, marketing suppression	High reuse in CX orchestration	1
5	LLM Customer Service Bot	Interaction outcomes, intents	Hardship prediction, CX analytics	Consumer that produces valuable telemetry	2
6	Collections Prioritisation	Priority, scripts, outcomes	CFO risk views, Service ops	Consumes scorecards & DIA; produces outcomes	2
7	Offer/Campaign Optimisation	Offer decisions, uplift	Marketing, CRM	Consumes segmentation & retention	2
8	Hardship Risk Prediction	Hardship likelihood	Service ops, ASIC analytics	Consumes bot/collections outcomes	3
9	NRT Scoring in Apps	Real-time decisions	Originations/Servicing	Reinforces loop; high ops value	3

FINANCIAL SERVICES
USE CASE CONTEXT

CREDIT SCORING MEETS CUSTOMER ENGAGEMENT

Credit Scoring and Risk Tiers

- Credit scorecard models generate risk tiers and override decisions logged as telemetry in the data lake-house.

Marketing Segmentation Using Risk Data

- Segmentation models use risk tiers to refine marketing campaigns targeting specific customer groups effectively.

AI Chatbot for Personalized Engagement

- LLM-powered chatbots leverage segmentation and credit risk signals to personalize communication and product recommendations.

Collections Workflow Prioritization

- Collections teams use credit risk signals to prioritize outreach for high-risk customers, improving efficiency.

SYMBOLIC V MATHEMATICAL PROCESSING

Proactive Customer Retention

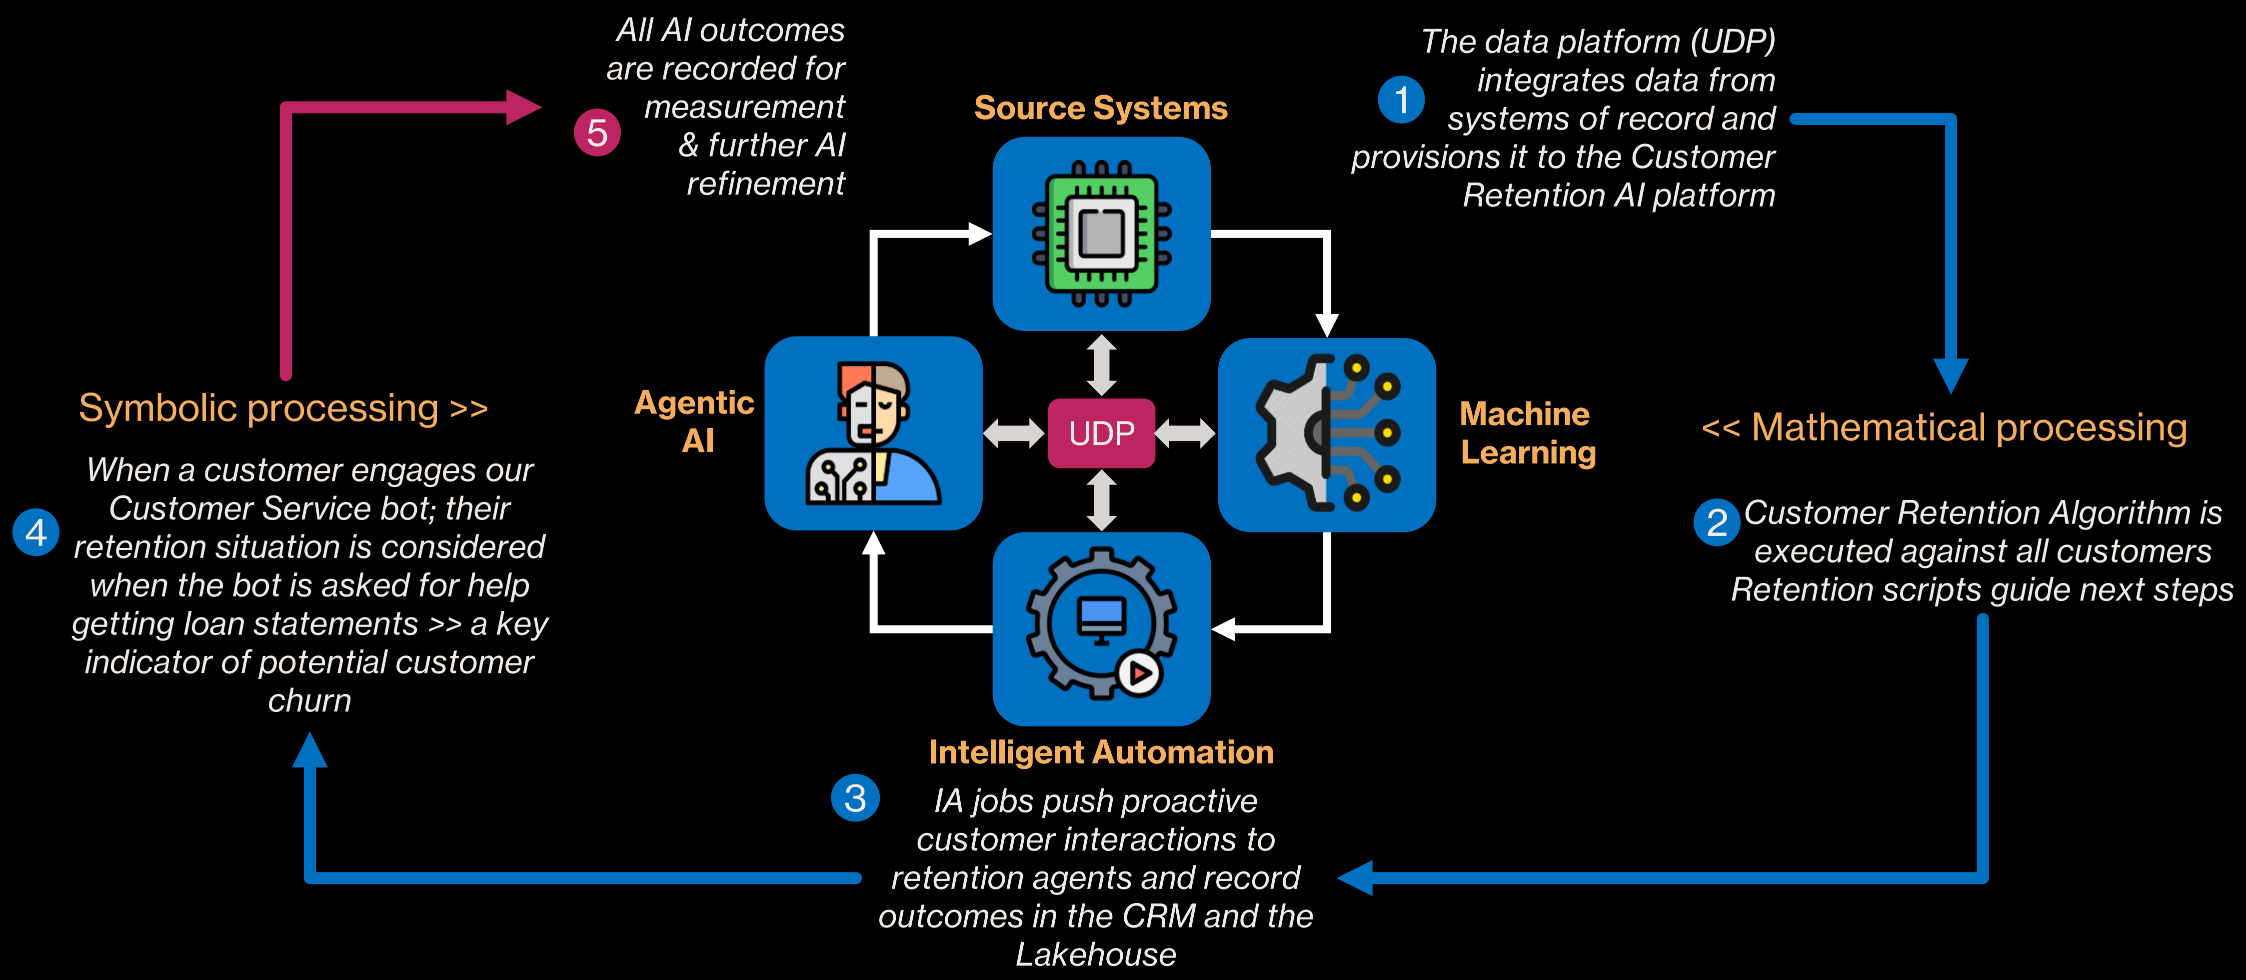

SYMBOLIC V MATHEMATICAL PROCESSING

Credit Scorecard

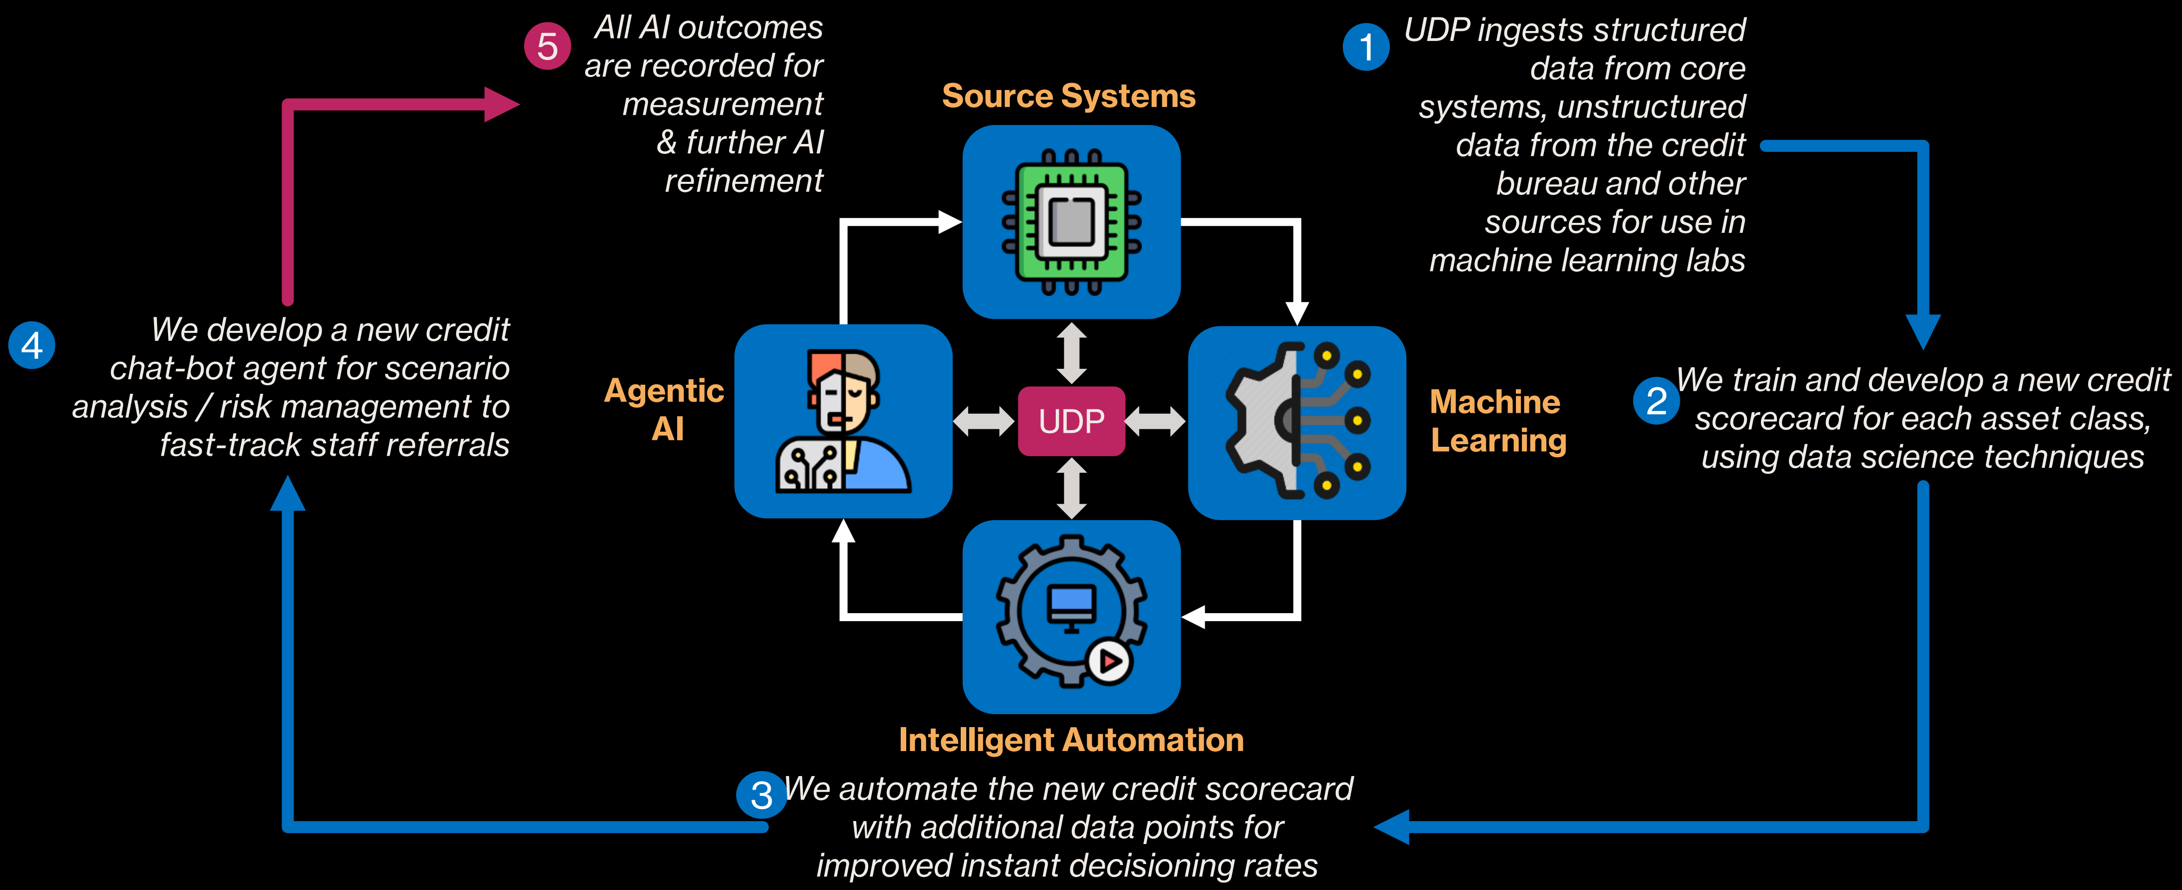

SYMBOLIC V MATHEMATICAL PROCESSING

CUSTOMER SEGMENTATION

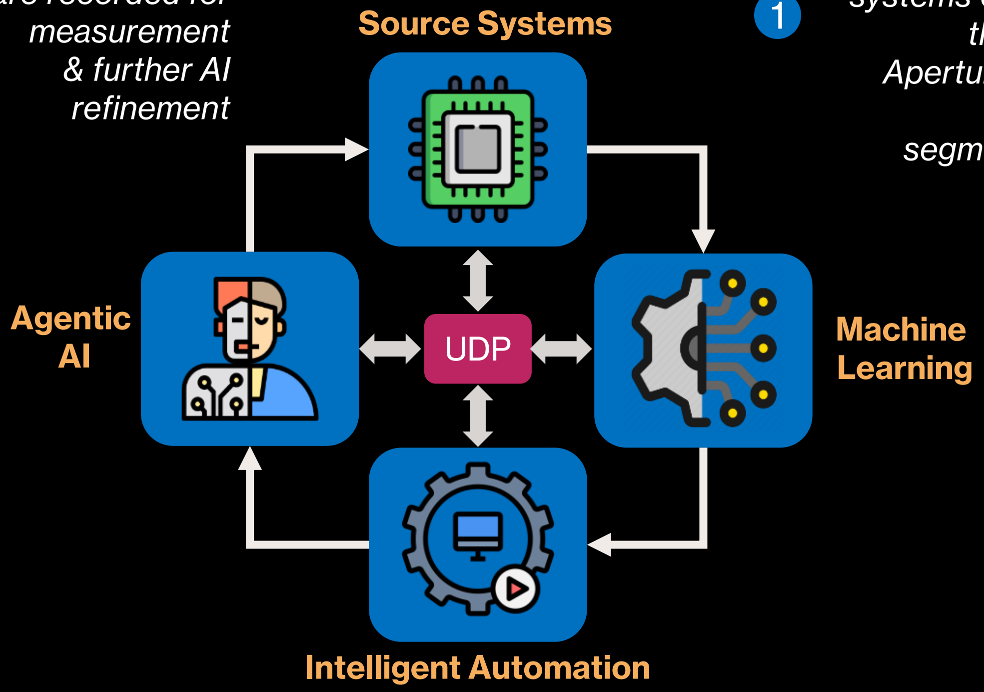

1 *UDP integrates data from systems of record and runs it through our Experian Aperture engine to append MOSAIC customer segmentation data points*

2 *Experian cluster market research and other data assets to create the Mosaic customer Segmentation Model & Persona descriptions*

3 *We Ingest MOSAIC persona descriptions into LLM training for customer preferences*

4 *Customer engages our Customer Service Agentic AI Bot: Their personal preferences are considered when the bot is asked for help*

5 *All AI outcomes are recorded for measurement & further AI refinement*

IN SUMMARY: OPTIMIZING THE ECOSYSTEM

AI Ecosystem Paradigm Shift

- **The AI flywheel shifts focus from single models to an interconnected ecosystem creating compounding insights**

Unified Data Platform

- **A unified data platform enables storing and sharing model outputs, accelerating continuous improvement.**

Call to Action

- **Identify a production model and show how it improves two other processes to prove the flywheel effect.**

FIVE-STEP IMPLEMENTATION PLAYBOOK

Data Asset Instrumentation

Treat decisions, outcomes, and model outputs as valuable data assets for better tracking and analysis.

Feature Store Contracts

Define clear ownership and SLAs for feature stores and output registries to ensure data governance.

Automated Retraining Pipelines

Use orchestration tools to automate model retraining and reuse pipelines for efficient AI lifecycle management.

Explainability and Override Workflows

Implement explainability and override workflows to improve customer-facing AI application transparency.

KPI Publication with Lineage

Publish key performance indicators with lineage visibility in BI dashboards to demonstrate business impact.

THANK YOU