IQZ SYSTEMS

Claims Accelerator

Real-Time Claims, End-to-End Efficiency

Prepared by IQZ Systems | June 27, 2025

Agenda

- Objective, Value & Key Outcomes
- Implementation Schedule
- Architecture and Lifecycle
- Observability, FinOps, and Governance
- Q & A

Objectives, Value & Key Outcomes

Claims Processing – Current State Challenges

Below is a set of industry challenges relating to claims processing that will be resolved by the IQZ Claims Accelerator.

Challenge	Description	Example
Fragmented Technology Landscape	Reliance on legacy systems, disconnected data sources, and paper-based processes impedes efficiency and scalability	Data is stored in multiple systems such as SQL databases, sharepoint sites, Excel files, and others. Unifying data storage and processing in a single platform reduces overhead and drives informed, timely business decisions
Manual, Labor-Intensive Operations	Significant time and resources are spent on repetitive data entry and non-automated workflows, increasing the risk of errors	Significant effort is expended entering claims data, combining customer data from various sources, validating claims, and assessing potential for fraud. Automation of these steps reduces hours per claim
High Operational Costs and Long Cycle Times	Elevated cost per claim and prolonged processing timelines negatively impact profitability and customer satisfaction	Claims must be entered, validated, and in some cases investigated, before a settlement can be determined. Automated classification of claims by risk reduces time to process low-risk claims and brings focus to high-risk claims for quick investigation
Limited Automation and Workflow Optimization	Lack of straight-through processing and intelligent automation slows down claim resolution and limits throughput	All incoming claims must be evaluated for potential risk prior to determining potential settlements. Low-risk claims automatically receive straight-through approval to increase throughput
Poor Integration with Modern Platforms	Inability to connect seamlessly with cloud-native applications and digital tools restricts innovation and adaptability	Fabric will integrate with structured, semi-structured, and unstructured data from a wide variety of real-time and batch sources

What is the IQZ Claims Accelerator?

Automatically flag suspicious claims in real-time using an ML detection system

Modernized cloud platform with unified digital processes and data storage

10%-30% reduction in claims processing time overall

Low risk claims automatically approved within 24 hours or less without human intervention

5%-10% reduction in fraud losses

Automated notifications of claims status

Who Cares?

Persona	Accelerator Value
Customer	▪ 10%-15% reduction in calls to check claims status ▪ Same-day approval of low-risk claims ▪ Transparent notifications of claim status in real-time
Claims Adjuster	▪ 40%-50% reduction in manual data entry ▪ 10%-25% reduced hours per claim ▪ Automated identification of high-risk claims with ML support ▪ Centralized data for investigations ▪ Automation of customer messaging
Claims Specialist	▪ 30%-50% reduction in administrative overhead resulting from documentation and multiple information systems ▪ Unified storage of policy details, claims, communications, etc. In a single system ▪ Automated identification of high-risk claims with ML support ▪ 10%-15% reduction in calls to check claims status
CISO	▪ 100% visibility and coverage of critical data assets ▪ 50%-80% decrease in manual audit time ▪ Centralized and modern security solutions under a single product ecosystem ▪ Centralized governance of PII, PHI, and financial data

Implementation

Implementation Schedule

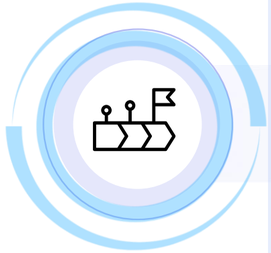

Timeline

- 4-6 weeks

The Team

- Project Manager
- Solution Architect (4)
- Data Scientist

Deliverables

- End-to-end claims processing solution in MS Fabric
- Automated detection of high-risk claims
- Automated approval of low-risk claims
- Automated quarantine of malformed claims

IQZ SYSTEMS

Architecture & Lifecycle

High Level Flow Diagram

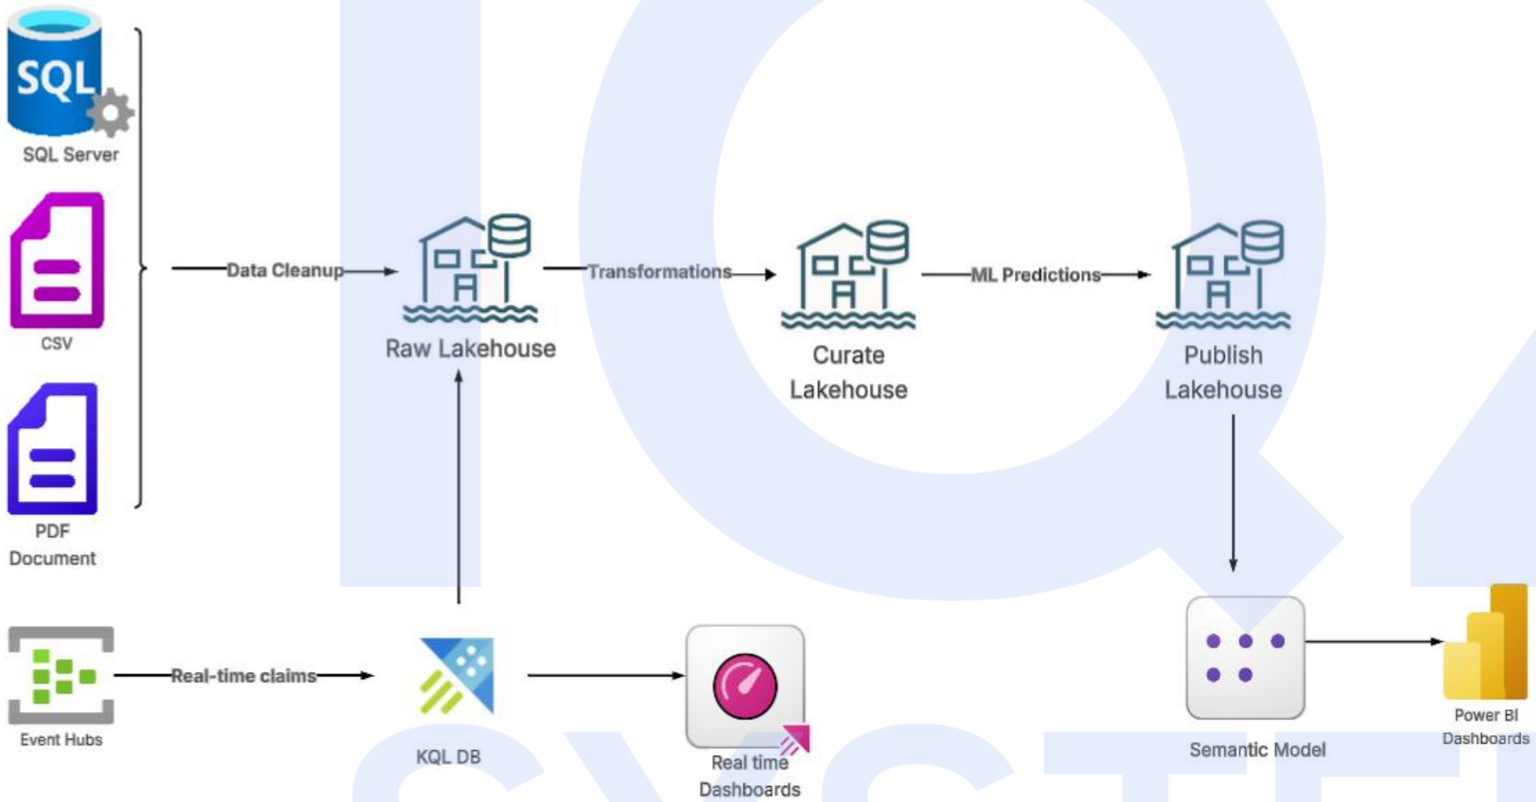

Ingests structured, semi-structured, and unstructured data into the Raw Lakehouse.

Performs data transformations and feature engineering in the Curate Lakehouse.

Applies ML models for fraud detection and stores results in the Publish Lakehouse.

Feeds a Semantic Model to drive insights through Power BI dashboards.

Handles real-time claims via Event Hub and KQL DB for instant dashboard updates.

Solution Architecture

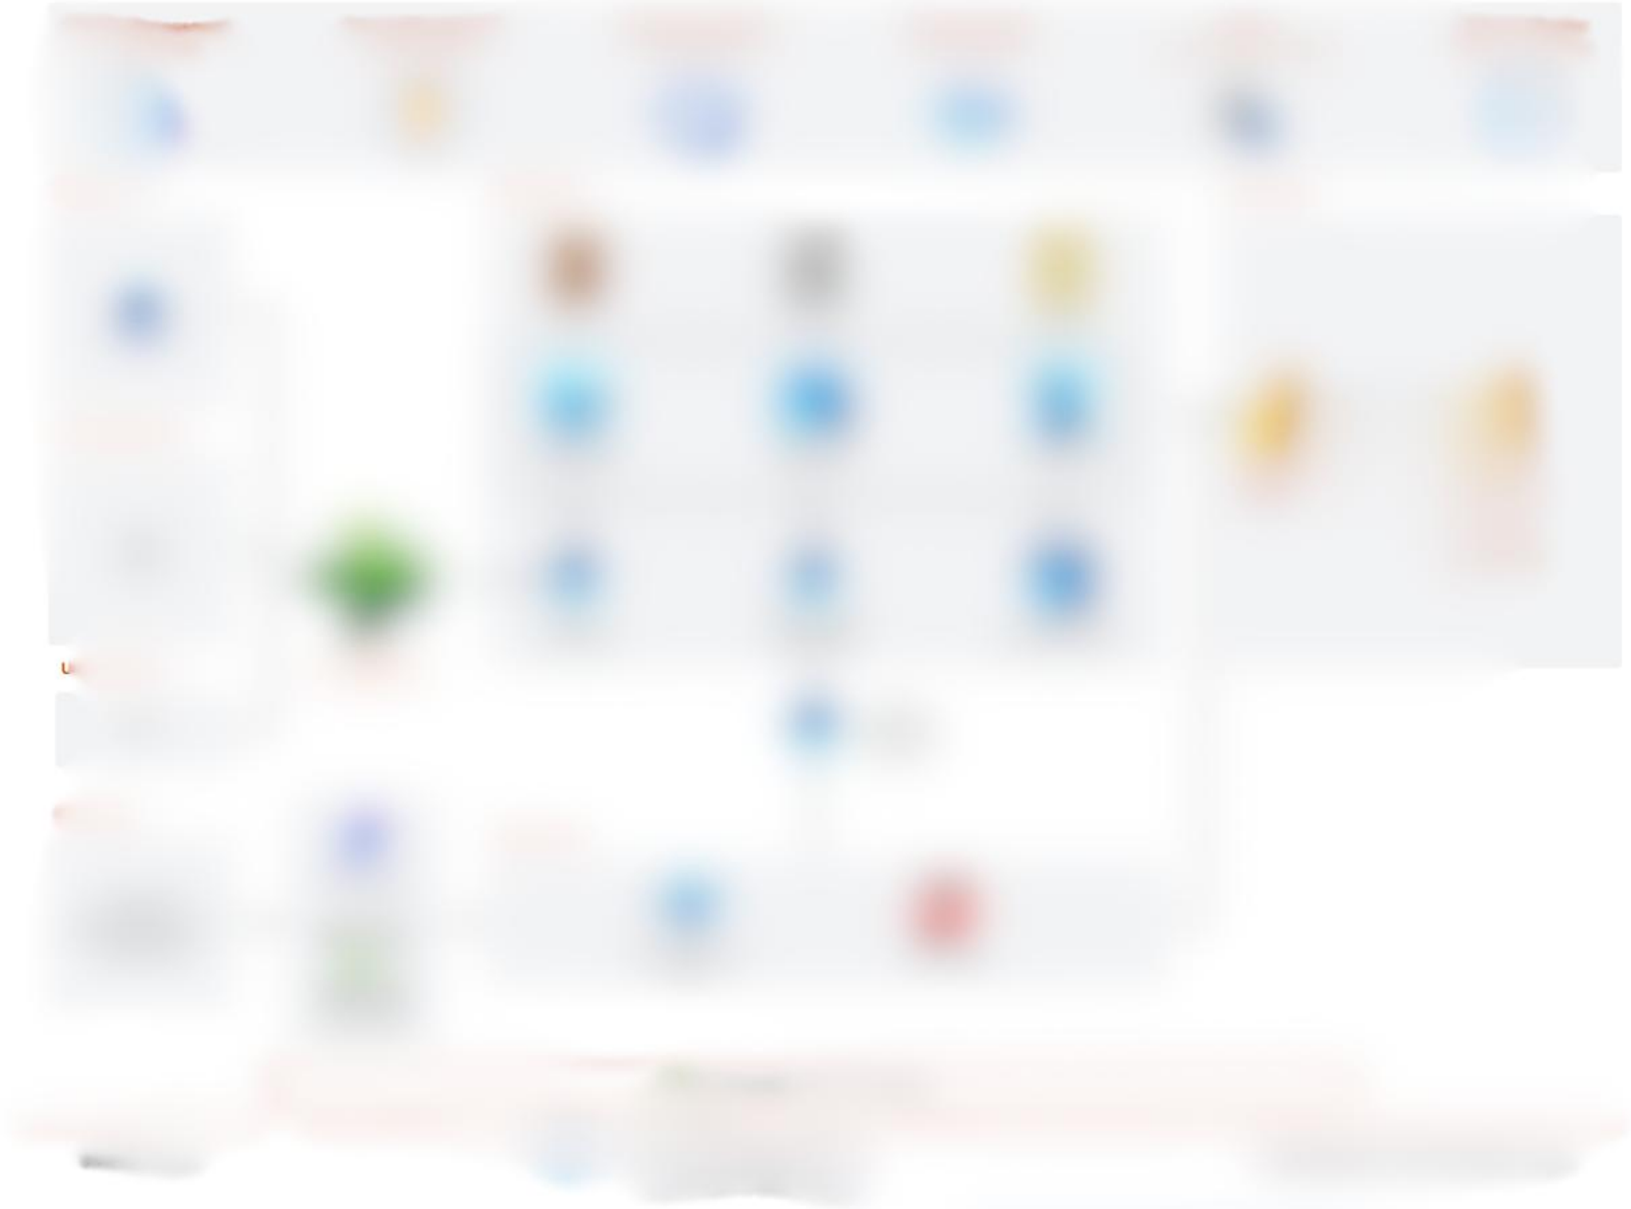

- Modular & Loosely Coupled Design
- Unified Under the Microsoft Fabric Ecosystem
- Real-Time Fraud Detection with Tailored ML Models
- Ingestion of Structured, Semi-Structured, and Unstructured Data
- Low Latency, Stream-Based Claims Processing
- Scalable Lakehouse Storage with Governance
- Automated Orchestration and Action
- Analytics-Ready Outputs via Data Warehouse and Power BI

Claims Lifecycle

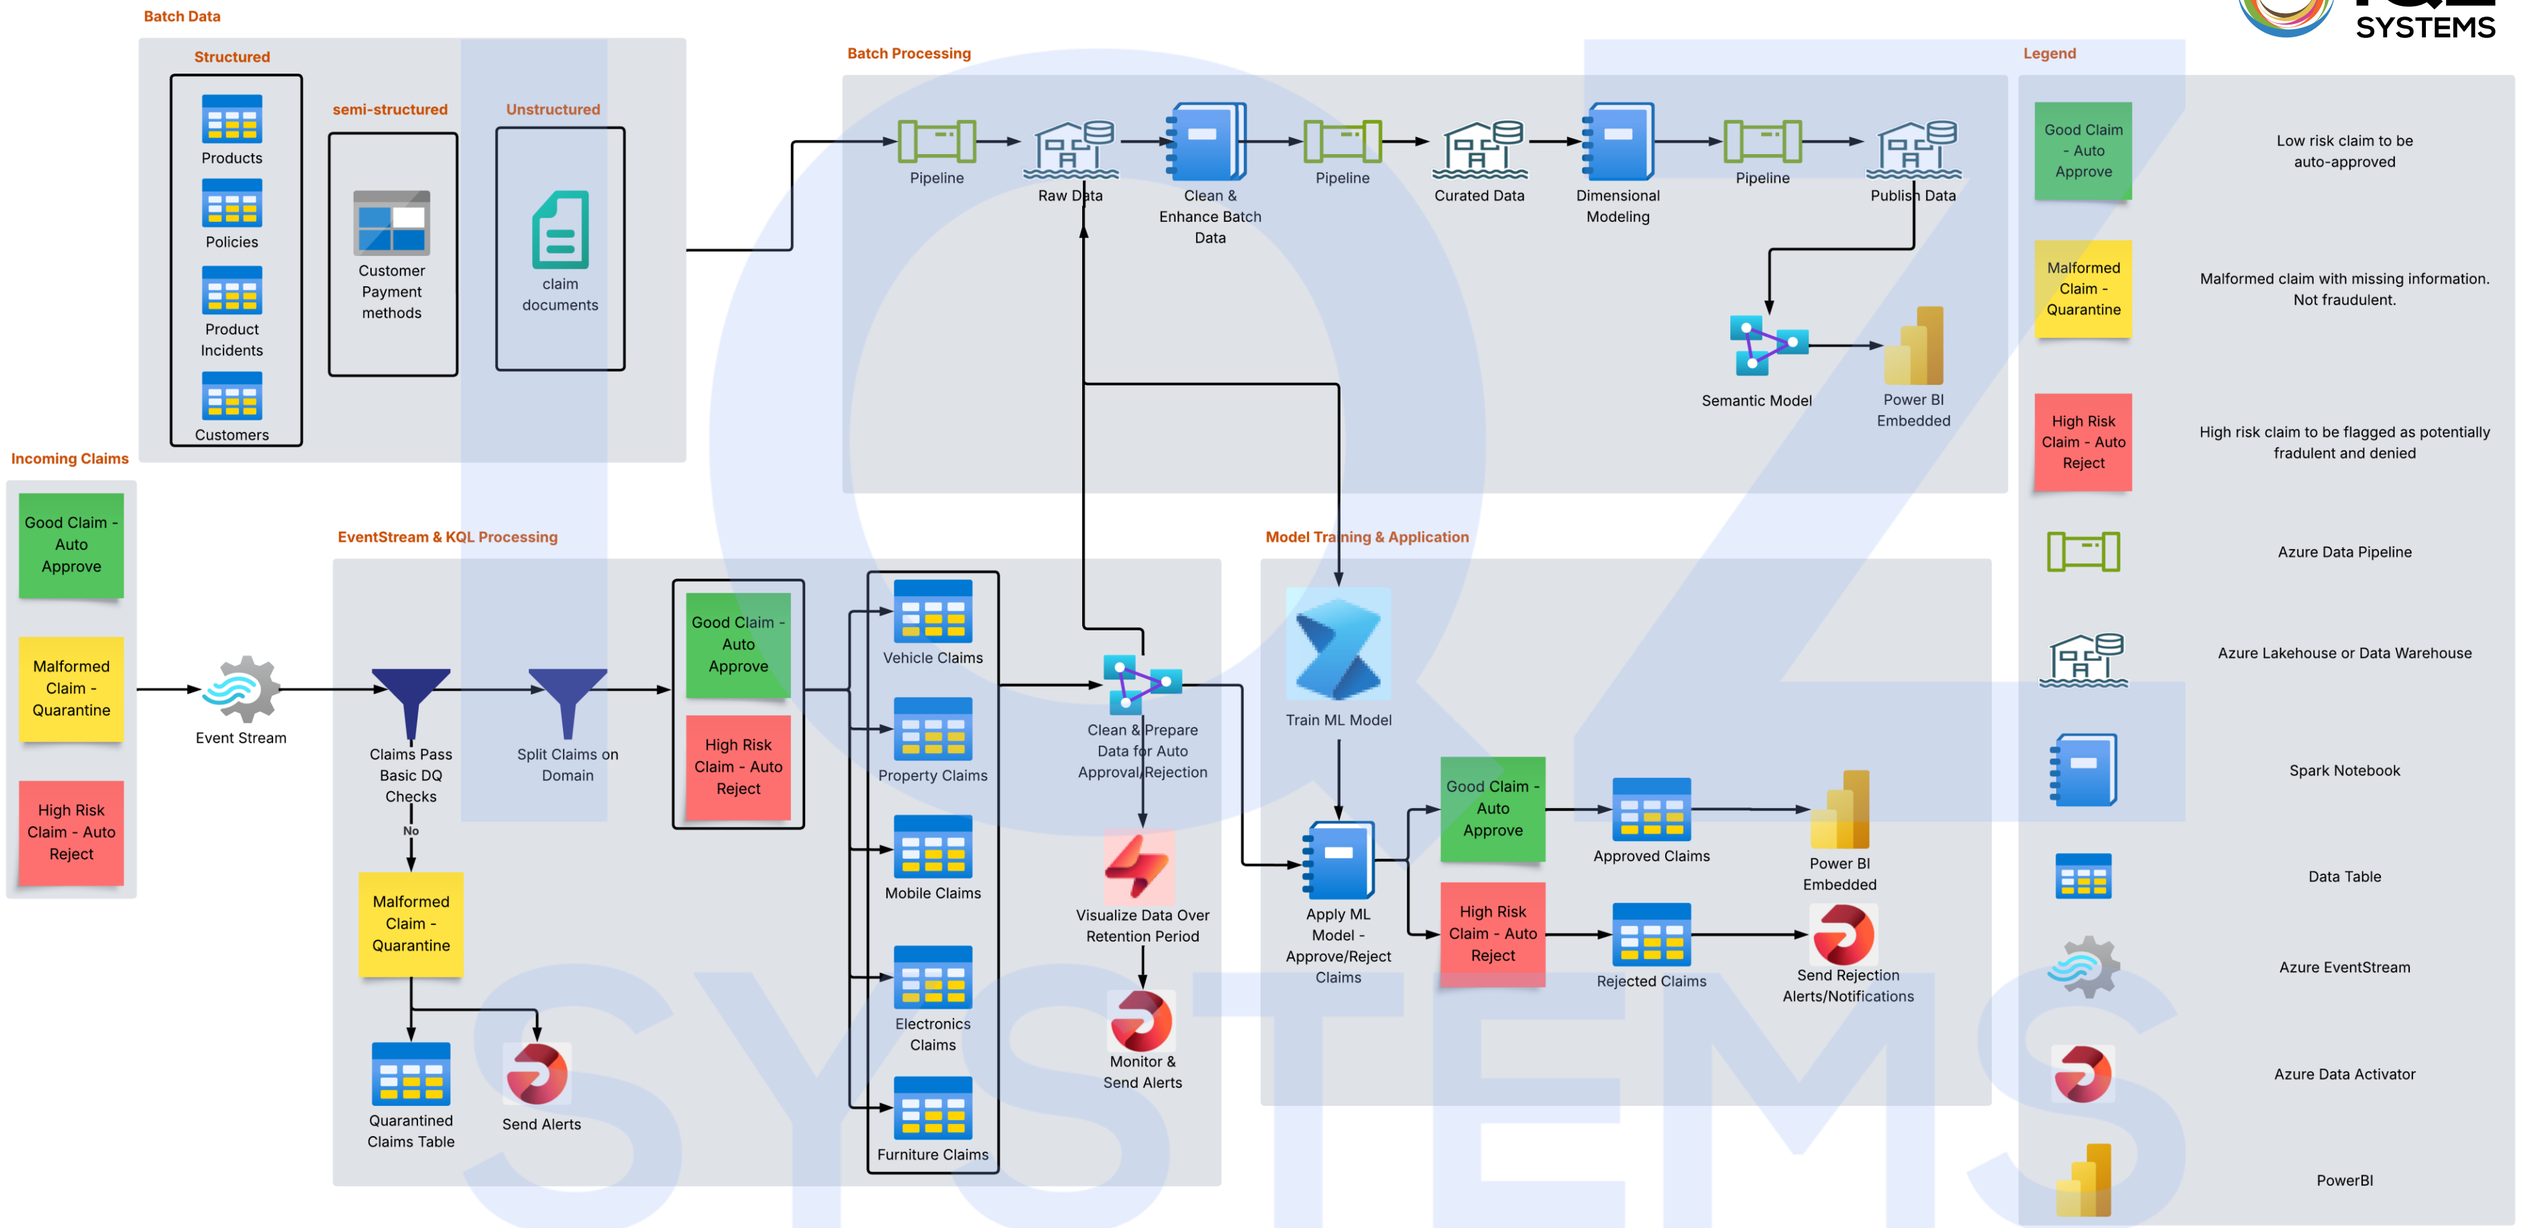

Fabric Services & Data Lifecycle

The IQZ Claims Accelerator is a unified platform within the MS Fabric ecosystem. It utilizes the services below to achieve our stated objectives within the data lifecycle.

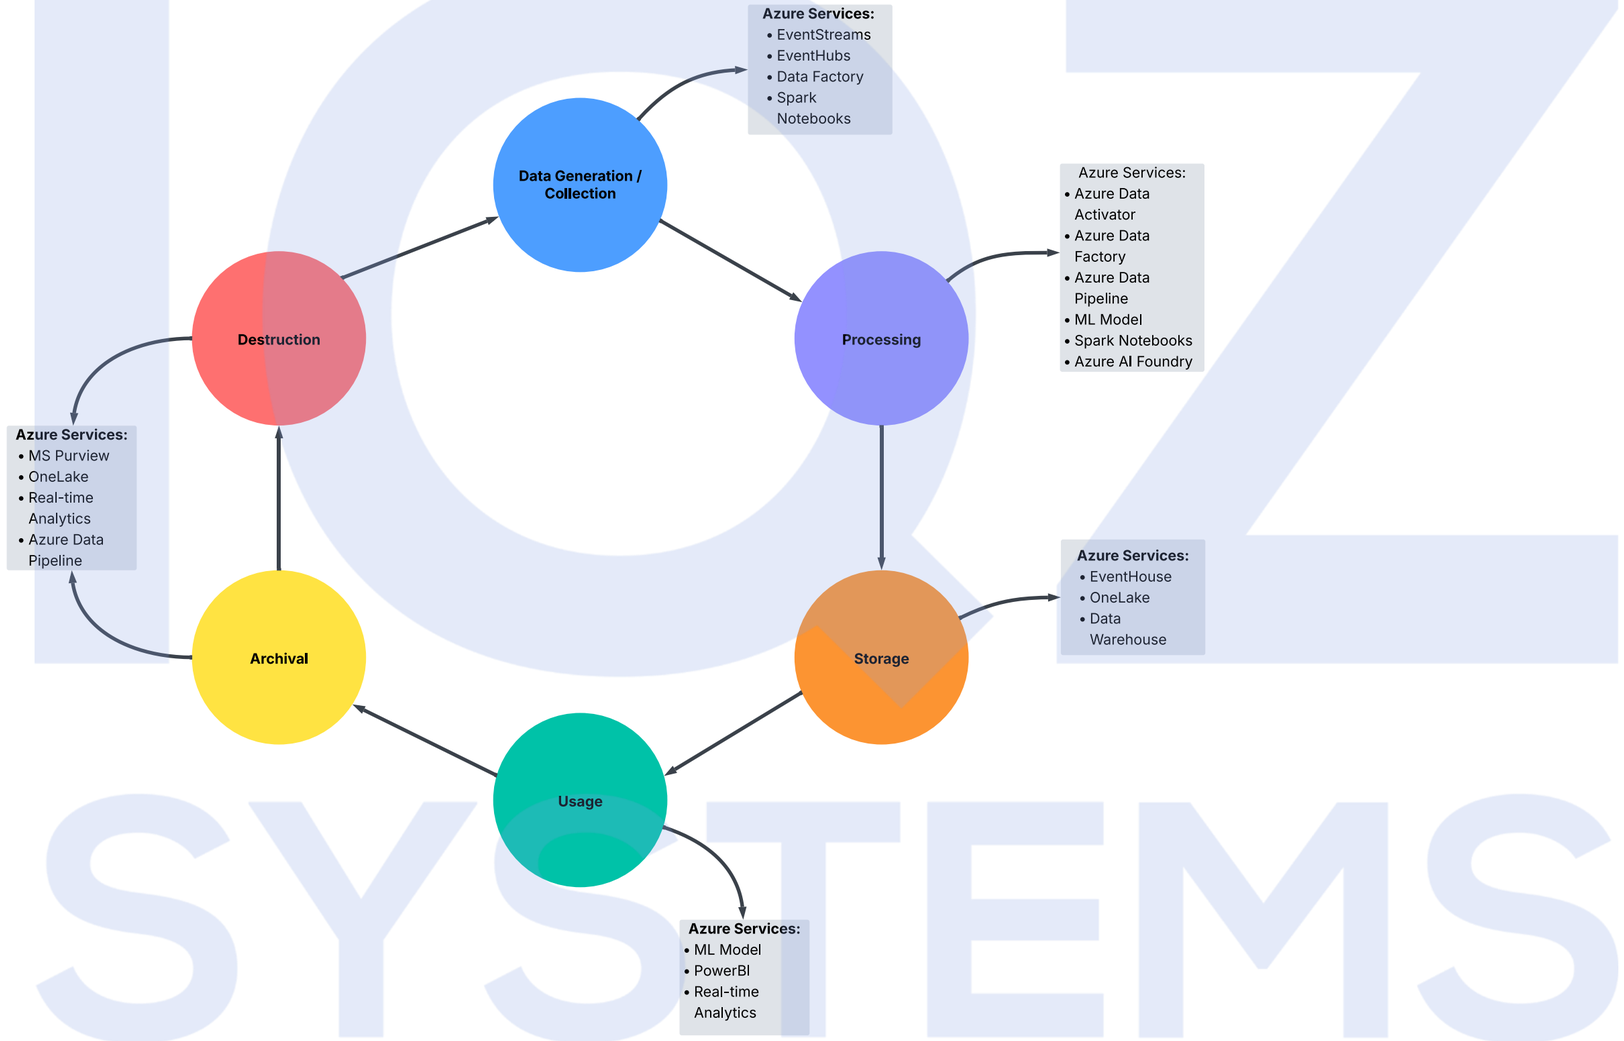

Why Azure & Fabric?

In light of the wide variety of data solutions available on the market today it is important to evaluate the benefits of using the Microsoft Azure & Fabric ecosystems to craft our solution. The benefits of using these services and offerings is explored below.

Capability	Value
Unified, end-to-end data platform	Streamlines data operations, reduces tooling complexity, and accelerates time to insight
Real-time data ingestion and processing	Enables faster decision-making and improved responsiveness in critical workflows like claims
Built-in governance and compliance (via MS Purview)	Reduces regulatory risk and simplifies audits with centralized data oversight and lineage tracking
Seamless integration across Microsoft ecosystem	Enhances productivity and collaboration by connecting data workflows with familiar tools (e.g., Power BI, Excel, Teams)
Low-code/no-code automation and orchestration	Accelerates solution delivery, reduces IT backlog, and empowers business users to act on data
Scalable storage with OneLake and Lakehouse architecture	Provides a single source of truth for structured and unstructured data, supporting growth and agility
AI and ML capabilities natively integrated	Delivers actionable insights and automation at scale, such as real-time fraud detection or risk scoring
Optimized analytics with Power BI and Data Warehouse	Improves operational visibility and strategic reporting with interactive, enterprise-grade dashboards

Azure AI Foundry

Azure AI Foundry is used to intelligently enhance data ingestion techniques, providing powerful capability to read and standardize a wide variety of structured, semi-structured, and unstructured data into existing business ready formats.

Capability	Usage
Quickly evaluate and deploy AI models	Identified ideal models for optical character recognition. Settled on PHI-4 Multi-modal model to process text and images.
Read structured, unstructured, semi-structured data	Data from nearly any source, including hand-written sources, can be read and standardized in business ready formats like JSON for processing in existing systems.
Monitor cost and performance of models	The usage, cost, and performance of models is evaluated in real-time within AI Foundry
Data Security	AI Foundry offerings are hosted entirely by MS servers with no runtime connections to model providers, preventing sharing of secure data.

ML Model Details

After evaluating three machine learning models it was decided to use a logistic regression model for classifying claims as potentially fraudulent

Why Logistic Regression Was Chosen

Fast and Lightweight

Delivers rapid training and inference, even on large insurance datasets — enabling near real-time fraud scoring at scale

High and Consistent Accuracy

Demonstrated the most balanced and reliable performance across all key model evaluation metrics

Superior at Detecting Fraud

Outperformed Random Forest in identifying fraudulent claims, improving fraud capture rates and reducing false negatives

Advantages of Logistic Regression

Model Advantage	Business Value
Interpretability – Model coefficients clearly show how each data feature influences the fraud prediction	Instantly detect and validate new claim submissions to accelerate intake and reduce data quality issues at the source
Training Speed – Trains quickly, even on large-scale claim datasets	Reduces model development cycles and enables frequent retraining to adapt to emerging fraud patterns
Strong Baseline for Comparison – Serves as a reliable benchmark for evaluating future, more complex models	Helps prioritize enhancements while ensuring continuity of performance and business impact
Probabilistic Output – Provides fraud likelihood scores instead of just binary results	Enables threshold tuning based on business risk tolerance, supporting dynamic fraud investigation strategies

IQZ SYSTEMS

Quality, Observability & Governance

Data Observability & Quality

Implementation	Benefit & Value
Established monitoring tables to track data freshness, volume trends, schema changes, lineage, and quality metrics	Delivers real-time detection of data anomalies (e.g., missing fields, delayed pipelines, duplicates), ensuring high-quality input data. This leads to more reliable fraud predictions, reduced false outcomes, and protects against losses from misclassified claims
Developed Power BI dashboards to provide unified visibility into data health and quality across all pipeline layers	Provides quick, centralized visibility into data health, empowering teams to proactively address issues. This boosts model accuracy, strengthens stakeholder trust in analytics, and supports informed decision-making
Implemented automated alerts for pipeline failures, delays, or data quality breaches	Enables rapid identification and resolution of data issues, minimizing manual intervention and reducing operational downtime. This proactive approach improves efficiency and safeguards against disruptions that could impact fraud detection
Integrated with Microsoft Data Activator to trigger real-time alerts for critical data incidents	Accelerates issue detection and resolution by triggering alerts the moment critical data problems are detected. This real-time response reduces manual effort, prevents SLA breaches, and minimizes risk exposure in fraud investigations

Pillars and Critical Features of Data Observability

5 PILLARS OF DATA OBSERVABILITY

- **Freshness:** Are your tables updating at the right time?
- **Volume:** Do you have too many or too few rows?
- **Distribution:** Is the value within a normal range?
- **Schema:** Has the organization of the data changed?
- **Lineage:** How are data assets connected across your data stack upstream and downstream?

Critical Features of Data Observability

Source: Gartner

FinOps

FinOps enhances financial control, reduces waste, and improves resource planning. It empowers stakeholders with real-time cost insights and ensures cloud resources scale efficiently while staying within budget.

Cost Monitoring & Management

Implementation	Business Value
▪ Exported Azure cost data into a Microsoft Fabric Lakehouse ▪ Developed a custom Semantic Model for usage and cost analysis ▪ Built a dedicated Power BI dashboard for real-time FinOps reporting	▪ Improves cost transparency across data projects and services ▪ Enables self-service cost visibility for business and technical stakeholders ▪ Supports cost optimization by identifying unused resources and overprovisioned workloads ▪ Detects anomalies in cost and usage early, reducing budget surprises

Capacity Consumption Monitoring

Implementation	Business Value
▪ Integrated Fabric Capacity Metrics App with existing Semantic Model ▪ Built a performance and usage monitoring dashboard in Power BI	▪ Provides detailed insight into actual resource consumption across Fabric services ▪ Enhances forecasting and capacity planning by surfacing usage trends ▪ Prevents under- or over-provisioning, improving platform efficiency and reliability ▪ Supports early detection of performance bottlenecks or system strain

Data Governance, Part 1

Data Catalog and Discovery

Implemented Capability	Business Value
All data assets — including lakehouses, notebooks, KQL queries, and models — are centrally indexed in a searchable catalog	Improves discoverability and reduces time spent locating data, increasing user productivity
Enables users to quickly locate and understand data across raw, curated, and published layers	Enhances data accessibility and trust, enabling faster analysis and decision-making

Access and Security Governances

Implemented Capability	Business Value
Role-based access controls restrict sensitive and raw data to authorized users only	Protects confidential information and supports compliance with industry regulations
Only approved, business-ready datasets (e.g., published layers) are made available to reporting users	Ensures users access only high-quality, validated data, reducing the risk of misinformed decisions

Glossary and Business Context

Implemented Capability	Business Value
All data assets — including lakehouses, notebooks, KQL queries, and models — are centrally indexed in a searchable catalog	Improves discoverability and reduces time spent locating data, increasing user productivity
Enables users to quickly locate and understand data across raw, curated, and published layers	Enhances data accessibility and trust, enabling faster analysis and decision-making

Data Governance, Part 2

IQZ SYSTEMS

Data Quality and Observability

Implemented Capability	Business Value
Automated profiling and rule-based checks validate data for completeness, accuracy, and freshness across pipelines	Ensures high-quality data is used in analytics and decision-making, reducing rework and costly errors
Early detection of data pipeline issues before they impact downstream reports or machine learning models	Minimizes disruption to business operations by proactively addressing data quality issues

Insights and Usage Tracking

Implemented Capability	Business Value
Tracks scan coverage, glossary usage, and sensitivity labeling across the data environment	Gives visibility into governance effectiveness and identifies areas for improvement in data stewardship
Allows governance teams to monitor adoption of the catalog and usage of certified data products	Provides measurable ROI on data initiatives and informs decisions on where to focus governance efforts

End-to-End Lineage

Implemented Capability	Business Value
Provides clear, visual lineage from data sources (e.g., event streams or batch files) to final Power BI dashboards	Enables faster root cause analysis, reduces investigation time, and builds trust in analytics outputs
Supports end-to-end impact analysis and audit readiness with traceable data flows	Helps meet compliance requirements and simplifies documentation for audits and regulatory reviews

2025

thank you

IQZ SYSTEMS

CONNECTING PEOPLE TO POSSIBILITIES

www.iqzsystems.com