



INVISO

Data Intelligence Platform Migration Guide

A practical framework for evaluating platform modernization before committing to migration.

Contents

Purpose	3
When the platform becomes the problem	3
What this guide covers	3
What we consistently see in platform migrations	3
Inventory and planning	4
Know why you are moving and what needs to come along	4
Migration paths and tradeoffs	5
Platform fit and what happens next	6
Getting migration right, not just fast	6
Use cases	7
Architect your migration before you execute it	9

Purpose

This guide is for organizations actively considering changes to their data platform and looking for a practical way to evaluate their options before migrating.

When the platform becomes the problem

Legacy data warehouse and analytics platforms were built to support traditional business intelligence workloads, but they were not designed to handle modern operational, analytical, and AI use cases. As complexity and costs increase, trust in the data erodes, leaving teams under pressure to modernize without a clear, structured path forward.

Common warning signs your platform isn't working any longer:

- Analytics backlogs that delay insights and decisions
- High-cost platforms that are difficult to operate and maintain
- Ongoing data quality and trust challenges
- Lack of clarity on what to migrate, rebuild, or replace

What this guide covers

Most migrations stall because teams select a target platform before understanding what actually needs to change. This guide helps you think through:

- What to inventory before migrating
- Migration paths and modernization tradeoffs
- How to sequence change to reduce risk

This is not a step-by-step migration manual or a tool comparison. It is a framework for thinking through data platform decisions that can enable real business and competitive advantage.

What we consistently see in platform migrations

- Data dependencies discovered too late in the process
- Business logic tightly coupled to legacy systems
- Governance treated as a post-migration clean-up effort
- New platforms replicating old architecture patterns

Inviso can help with Microsoft investment opportunities for organizations migrating to Fabric.

[Contact us →](#)

Inventory and planning

Know why you are moving and what needs to come along

Platform migration is not simply a lift-and-shift to new technology. It is an opportunity to enable new use cases that drive organizational growth, while ensuring core business, operational, and analytical workloads remain stable throughout the transition.

These four areas determine whether a migration will accelerate the business — or introduce risk and rework:



Data

Understand what data you have, how it is used, and how it could be used.

- ✓ Source systems, volume, and growth trends
- ✓ Data freshness and latency requirements
- ✓ Sensitivity, compliance, and governance needs



Data code and transformations

Identify how business logic is implemented today.

- ✓ SQL, notebooks, stored procedures, and scripts
- ✓ Complexity of transformations and dependencies
- ✓ Areas of brittle or duplicated logic
- ✓ Operational support capabilities and constraints



Orchestration and scheduling

Map how data flows and what keeps it running.

- ✓ Pipelines, schedules, and dependencies
- ✓ Failure handling and recovery processes
- ✓ Operational ownership and SLAs



Consumption

Clarify how data is actually used.

- ✓ BI reports, dashboards, and ad hoc analysis and by which people
- ✓ Operational reporting and downstream applications that require data integration
- ✓ Machine learning, automation, and APIs

Why this matters

Migrations fail to deliver value when data, insights, and consumption patterns break or are discovered too late. A clear inventory across these four areas leads to more realistic timelines, better platform decisions, and fewer surprises during migration.

Migration paths and tradeoffs

Moving to a lakehouse or data intelligence platform

When organizations migrate data platforms, our recommendation is to move toward a data intelligence platform leveraging a modern lakehouse. A data intelligence platform provides the foundation required to support modern analytics, operational workloads, and AI use cases that drive value.

That said, not every workload should be migrated the same way. Successful migrations balance value, risk, and effort by applying different approaches to different workloads. Most migrations use a combination of three paths:

Rebuild

Rebuild workloads when existing logic no longer supports current or future business needs.

- Models designed only for historical reporting
- Complex or brittle logic that limits change
- Opportunities to redesign for operational analytics or AI

Rebuilding enables teams to fully take advantage of a lakehouse or data intelligence platform.

Refactor

Refactor workloads that still deliver value but need architectural improvement.

- Core analytical workloads with stable business logic
- Pipelines that can be adapted to Fabric or Databricks
- Code that benefits from improved performance or governance

Refactoring balances speed, risk, and modernization.

Retire

Retire workloads that no longer justify migration cost.

- Low-usage reports or datasets
- Redundant pipelines created over time
- Logic that no longer supports active business decisions

Retirement reduces complexity and accelerates progress.

Why this matters

A lakehouse or data intelligence platform delivers value only when workloads are intentionally rebuilt, refactored, or retired. Lift-and-shift alone rarely achieves the outcomes organizations expect.

Platform fit and what happens next

Turning decisions into action

Once workloads have been evaluated and migration paths are clear, the next step is selecting and sequencing the right platform approach. Our recommendation is to move toward a **lakehouse or data intelligence platform**, implemented using **Microsoft Fabric, Databricks, or a hybrid model**, based on workload needs and operating constraints.

The goal is not tool consolidation for its own sake, but a platform that supports analytics, operational use cases, and future AI initiatives in a sustainable way.

Platform fit at a high-level

- **Microsoft Fabric:**
Well suited for organizations that want tight integration across analytics, BI, governance, and operations, particularly in Power BI-centric environments.
- **Databricks:**
Well suited for data engineering, SQL analytics, large-scale processing, and ML-driven workloads that require flexibility and depth.
- **Hybrid:**
Common during transition periods or when different workloads have different requirements, especially with Power BI as the consumption layer.

Platform fit depends less on features and more on workload complexity, team maturity, and long-term operating model.

Getting migration right, not just fast

A successful migration isn't about moving fast—it's about making the right decisions early so the platform supports the business long after the migration is complete.





Use cases

Context

A retail organization relies heavily on transactional reporting. Several hundred users consume roughly 50 Power BI reports built on a combination of Azure Synapse, Oracle, and Power BI.

While the current platform supports historical reporting, it is costly to operate, difficult to evolve, and increasingly limiting the organization's ability to use transaction data beyond dashboards.

The organization is planning a move to **Microsoft Fabric** to consolidate analytics workloads, simplify the reporting stack, and modernize how transaction data is stored and processed. **Power BI being part of Fabric is a clear efficiency gain**, but the broader goal is to unlock new revenue-generating use cases from existing transaction data.

Why they're moving

Beyond simplifying reporting, the organization wants to use its transaction data more effectively to influence customer behavior and increase sales. Today, most insights are retrospective. The goal is to move toward more proactive, customer-facing use cases without building and maintaining separate systems.

USE CASE 1

Personalized product recommendations in the mobile app

Point of sale transaction data is used to train machine learning models that predict relevant product recommendations for customers using the retailer's mobile application. These recommendations are based on purchase history, basket patterns, and seasonality, and are refreshed regularly as new data arrives.

Value delivered

- More relevant product suggestions for customers
- Increased engagement within the mobile app
- Higher average basket size driven by personalized recommendations

Fabric provides a unified platform where transaction data, feature engineering, and model outputs can live together and be served back to the application with low latency.

USE CASE 2

Targeted coupons to drive incremental revenue

Instead of broad, static promotions, the retailer uses transaction data to identify customers who are likely to respond to targeted discounts or incentives. Machine learning models score customers and transactions to determine when a coupon is likely to drive incremental spend rather than subsidize an existing purchase.

Value delivered

- More effective promotions with measurable lift
- Reduced over-discounting
- Clear linkage between data, offers, and revenue impact

Fabric enables this by centralizing transaction data, model outputs, and governance, making it easier to operationalize insights across marketing and digital channels.

Why Fabric makes sense

By moving to Fabric, the organization consolidates reporting, analytics, and AI workloads into a single, governed platform. This reduces operational complexity while creating a clear path from transaction data to customer-facing experiences and revenue-driving use cases, all while continuing to use Power BI as the primary reporting layer.


Architect your migration before you execute it

Most migration delays stem from misaligned sequencing, hidden dependencies, or governance gaps discovered too late.

The [Fabric Readiness Workshop](#) provides a technical assessment of your platform across the four dimensions outlined in this guide — helping you identify:

- Where lift-and-shift creates risk
- Where refactoring accelerates value
- Where rebuilding unlocks new use cases
- What can be retired to reduce complexity

The output is a structured readiness scorecard and summary to support executive and architectural alignment.



Contact info@invisocorp.com or visit www.invisocorp.com for more information.



NOTICE

The information contained in this document represents the current view of Inviso on the issues discussed as of the date of publication and is subject to change at any time without notice. This document and its contents are provided AS IS without warranty of any kind and should not be interpreted as an offer or commitment on the part of Inviso and Inviso cannot guarantee the accuracy of any information presented.

INVISO MAKES NO WARRANTIES OR REPRESENTATIONS, EXPRESS OR IMPLIED,
IN THIS DOCUMENT.

All trademarks are the property of their respective companies. All rights reserved