


Enterprise BI Migration Playbook

From Excel chaos to an AI-ready data platform — in four disciplined phases.



01	Why Migrate Now	03
	The error rate under your decisions, the adoption gap, and what AI changes	
02	The 4-Phase Playbook	06
	Audit · Data model · Dashboards & adoption · Governance	
03	The Fabric Factor	13
	What the platform actually removes from a migration	
04	ROI & Next Steps	14
	Where the return comes from, and the five ways migrations fail	

The error rate under your decisions

94%

of spreadsheets used in business decisions contain errors, according to a 2024 systematic review covering more than thirty-five years of research.

The best-known example remains a bank's risk model, where a spreadsheet error contributed to a loss measured in billions. The mechanism was ordinary: a formula nobody re-checked.

2–5% of formula cells contain errors even when written by experienced users.


~50% of operational spreadsheet models in large companies carry a material defect.

86% of study participants made at least one error — while estimating their own rate at 10–18%.

30–40% of active Excel reports turn out to be duplicates once an audit is run — our own consistent finding.

Spending grows. Usage does not.

BI adoption has been stuck at roughly a quarter to a third of the workforce for over a decade, while the market kept expanding. That divergence is the whole argument for treating migration as a discipline rather than a licence purchase.



87% of organizations report rising analytics use

~29% of employees actually use the BI tools their employer pays for

The platform is not the problem. The migration is.

Everyone is about to become an analyst

90%

of today's consumers of analytics content are forecast to become creators of it by 2026, with AI doing the authoring. The question is no longer who can build a report — it is what they will be building on.

Your semantic model is the brain your AI agents will think with. Excel formulas cannot be governed. Measures can.

Business logic buried in VLOOKUP chains cannot be taught to an agent. Centralised measures in a governed semantic layer can. That is why Phase 2 of this playbook is, in practice, the construction of your company's analytical brain.

Copilot in Power BI	Generates DAX, builds visuals and summarises reports — available on all paid Fabric SKUs from F2 since April 2025.
Fabric data agents	Converse in natural language with data in OneLake; pluggable into Copilot Studio and M365 Copilot.
Fabric IQ	Answers questions against reports and semantic models in natural language; exposed to M365 Copilot and Agent 365 as a tool.

Four phases, in this order

The sequence matters more than the tooling. Every failed migration we have been asked to rescue skipped or compressed one of these phases.

2-4 weeks

4-8 weeks

6-10 weeks

ongoing

PHASE 01

Audit

Inventory every active report, map its sources, extract the business logic hidden in formulas and VBA, classify by tier.

Output: report inventory with migrate / merge / retire verdicts

PHASE 02

Data model

Star schema, slowly changing dimensions, incremental refresh — and the semantic layer your AI will later depend on.

Output: governed semantic model with centralised measures

PHASE 03

Dashboards & adoption

Tier 1 reports first, self-service by design, row-level security, and a deliberate period of parallel running.

Output: reports in daily use, power users trained

PHASE 04

Governance

Data stewards, report certification, refresh monitoring, a request process that prevents shadow BI, quarterly retirement.

Output: a platform that stays trustworthy

Find out what you actually have

Nobody knows their report estate at the start. The audit is not paperwork — it is the only way to avoid rebuilding work that should be retired.

-
- 01 **Inventory every active report** — including the ones maintained on someone's desktop


 - 02 **Map the sources** — where each number originates, and how often it moves

 - 03 **Extract the business logic** from formulas and VBA before it is lost with the person who wrote it

 - 04 **Interview the owners** — what the report is for, and who reads it on a Monday

 - 05 **Classify into tiers** and issue a verdict: migrate, merge or retire
-

Source: Deka delivery experience across BI migration audits



TIER 1 Decisions depend on it. Migrate first, with the owner in the room.

TIER 2 Useful, overlapping. Candidates for merging into fewer, better reports.

TIER 3 Opened rarely or never. Retire — and let the absence prove the point.

3-5×

more reports in existence than leadership expects, on every audit we have run.

The report inventory matrix

One row per report. The verdict column is the point of the exercise – an audit that ends without verdicts becomes a document nobody uses.

REPORT	OWNER	SOURCE	BUSINESS LOGIC	TIER	VERDICT
Weekly sales summary	Sales ops	ERP + manual CSV	Regional split, FX conversion	1	MIGRATE
Regional sales detail	Regional leads	Same ERP extract	Duplicate of the above, filtered	2	MERGE
Legacy KPI pack	unclaimed	Unknown / stale	Last opened 14 months ago	3	RETIRE


Expect the duplicate rate to sit between 30 and 40 percent. Every retired report is scope you do not have to rebuild, test, document or maintain.

Build the model, not the report

A star schema is not academic preference. It is what makes measures reusable, refreshes cheap and the model comprehensible to a person joining in two years — or to an agent answering a question next week.

SCD TYPE	WHEN TO USE IT
Type 1	Overwrite. History does not matter — corrected spellings, fixed typos.
Type 2	New row per change. The default when reporting must reflect the past as it was.
Type 6	Hybrid — current and historical views of the same attribute side by side.

Incremental refresh is configured here, not later. Full reloads are what turn a healthy model into a nightly incident.



The semantic layer is the AI foundation

Everything above this layer — a report, a Copilot answer, an agent's recommendation — inherits its correctness from the layer itself. Agents are only as good as the semantic model beneath them.

CONSUMERS

Reports & dashboards

governed, certified, one definition each

Copilot

DAX, visuals and summaries on request

Data agents · Fabric IQ

natural-language answers, in other tools

THE SEMANTIC MODEL

One definition of every measure

centralised measures · KPI definitions · relationships · row-level security · business logic in DAX, not in cells

STORAGE

OneLake — one copy of the data, every engine reading it

Read the diagram downward and the risk becomes obvious: skip the middle band, and every consumer above it is improvising a definition of "revenue" on its own.

Adoption is designed, not announced

Tier 1 first

Start where decisions are made. Credibility earned on the reports leadership reads carries the rest of the programme.

Self-service by design

Build for the question after the current one. If every follow-up needs IT, the estate will drift back to Excel.

Row-level security

Configured before rollout, not after the first complaint. It is also what lets one report serve every region.

PARALLEL RUNNING — 4 TO 6 WEEKS

WEEK 1

WEEK 2

WEEK 3

WEEK 4

WEEK 5

6

Excel reporting continues

New BI reports in production

OVERLAP — NUMBERS RECONCILED SIDE BY SIDE

The overlap is where trust is built. Users compare the two outputs themselves, find the discrepancies, and either accept the new number or expose a modelling error worth fixing. Switching off Excel before this has happened is what produces quiet, permanent shadow reporting.


Train power users, not audiences. Two or three capable people per function will answer more questions than any documentation you write.

Keeping it trustworthy

Governance is the phase most often cut for time, and the only one that determines whether the estate looks the same in three years. It is a loop, not a document.

The shadow-BI test

If a manager cannot get a new report request into a queue with a visible answer date, they will build it in Excel — and the migration will have been a detour. A request process is not bureaucracy; it is the alternative to reverting.



What the platform removes from the work

Platform choice does not rescue a badly sequenced migration. It does decide how much plumbing you build by hand — and lately that share has fallen sharply.

#1

Fastest-growing analytics product in Microsoft's history, by its own reporting.

25k

paid customers as of the FY25 annual report.

28k+

organizations by the end of 2025 — up from 11k in March 2024.

~70%

of the Fortune 500 reported using Fabric.

F2

the SKU floor for Copilot since April 2025 — no longer F64.

WHAT YOU NO LONGER BUILD BY HAND

- **OneLake** — one copy of the data, read by every engine
- **SharePoint / OneDrive shortcuts** — in preview; reach data in Excel files without copying it
- **Mirroring, generally available** — SQL Server 2016–2025, PostgreSQL, Cosmos DB, Snowflake Iceberg
- **Migration Assistant** — generally available for Synapse dedicated SQL pool to Fabric Data Warehouse
- **SAP mirroring** — in preview via Datasphere
- **OneLake security** — unified row-, column- and object-level security, on the way to GA

Organizations consolidating a multi-vendor stack onto Fabric report platform spend reductions of 30–45%. Treat that as industry-reported rather than a planning figure — your own baseline is the only number worth budgeting against.

Where the return comes from

Four components, in descending order of what we typically measure. Establish the baseline before go-live — a return nobody recorded is a return nobody believes.


Manual reporting hours




Fewer data errors



Faster decision cycles



Reporting-only licences



Relative contribution as typically observed in our engagements — not a benchmark.

WORKED EXAMPLE — CONSERVATIVE ASSUMPTIONS

Finance and operations team	25 people
Manual reporting, per person, per week	6 hours
Annual manual reporting effort	~7,800 hours
Recovered at 70% automation	~5,500 hours

8–14

months to pay for itself

The band we plan against. Long-running independent research puts returns on analytics investment at roughly \$6–13 per dollar spent, with payback often around nine to ten months.

Five ways these migrations fail

Ask your BI partner how they avoid each of these — including us. The answers are more informative than any capability deck.

☐ **The big-bang migration**

Everything switches over one weekend. There is no path back, and the first bad number costs the programme its credibility.

☐ **An IT-only project**

Built without the report owners, it reproduces the technical shape of the old reports and none of the reasoning behind them.

☐ **Go-live without a baseline**

Hours spent, error rates and cycle times unrecorded beforehand — so the improvement can never be demonstrated afterwards.

☐ **Self-service without governance**

Six versions of “revenue” within a quarter. The tool did what it was asked; nobody owned the definition.

☐ **Migrating everything**

The most expensive mistake of the five. Move the roughly 15% that is genuinely in use; retire the rest and let the silence confirm it.

Ready to leave Excel behind — and get AI-ready doing it?

Book a 30-minute BI migration assessment. We look at your report estate, name the phase you should start with, and tell you plainly what it will take. No sales pitch.

E info@dekatechs.com

T +90 216 606 69 57

W dekatechs.com

150+

engineers

ISO

27001 certified

5×

Deloitte Fast 50

We ran this playbook end to end for a tier-1 beverage producer, replacing an Excel reporting estate with a governed Power BI platform.

SOURCES

1. Poon et al., systematic review of spreadsheet errors, 2024
2. Panko, "What We Know About Spreadsheet Errors"; EuSpRIG research summaries
3. Powell, Baker & Lawson, material defects in operational spreadsheets, 2007
4. Gartner, BI & analytics adoption research, 2017-2019; BARC / Eckerson, 2022
5. Gartner, data & analytics predictions – consumers to creators by 2026
6. Microsoft FY2025 Annual Report – Fabric growth and paid customers
7. Microsoft Fabric blog – Ignite / FabCon 2025 updates
8. VentureBeat, Fortune 500 Fabric adoption, Dec 2025
9. Nucleus Research, analytics ROI series, 2011-2023
10. Deka Technology delivery experience – duplicate rate, payback band, worked example

