Prophecy

Move:
Alteryx → Agentic Data Prep

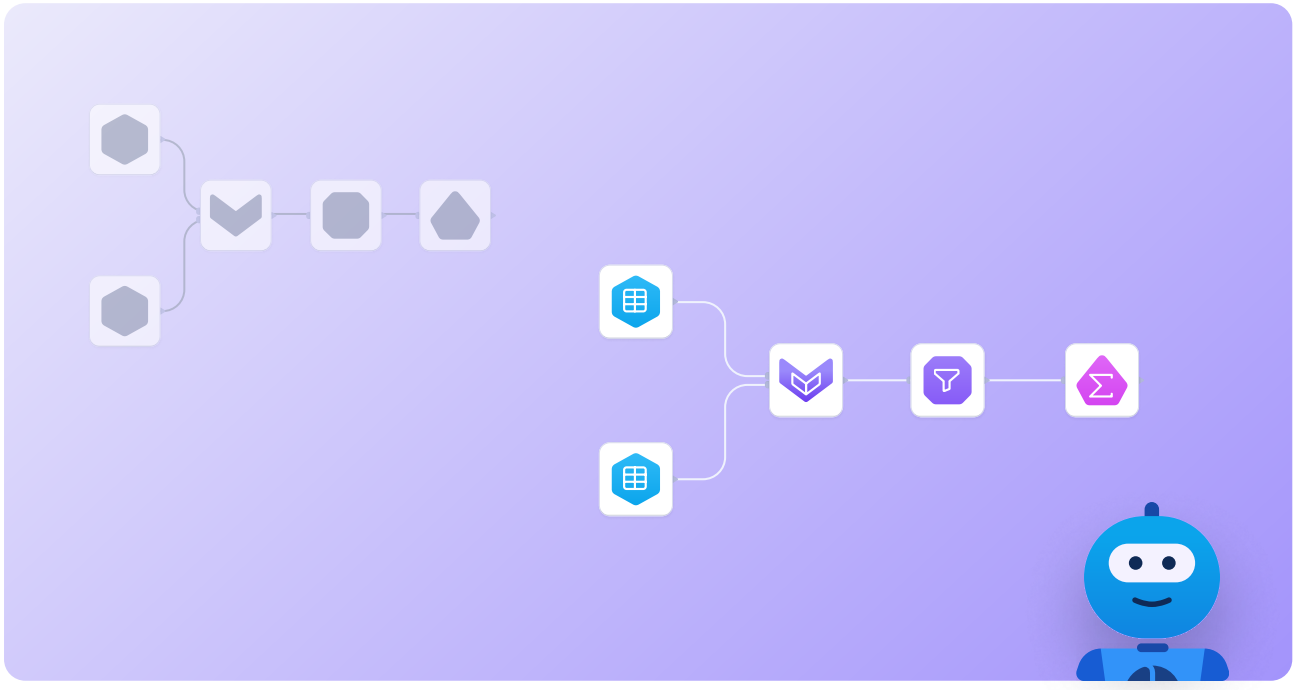

AI brings game-changing value to data prep. This paper shows what to expect from Agentic AI, and how it delivers results. A practical path to make the transition follows.

Introduction

AI transformed how software is written in 2025.
Data preparation is next: by the end of 2026, working without AI will feel unthinkable.

Why we're confident: data prep is visual programming applied to data, and we are seeing data analysts get results more impressive than manual coding.

We **distill what we've learned** helping large enterprises modernize from Alteryx to Prophecy.

What you'll learn

- The AI landscape for data teams
- AI impact on data prep and analysis
- How Prophecy's Agentic AI delivers results
- Can Alteryx match it?
- A pragmatic modernization path

Data prep needs specialized AI Agents

Data prep and analysis is a specialized form of programming. Agents for it must specialize and work hand-in-hand with the analysts.

Data context & semantics

Agents must understand universal data operations such as types, quality rules, relationships, and lineage.

Then, they must understand business-specific definitions for your business to build workflows that compute the right values.

Visual work canvas

Analysts primarily use visual workflows to develop business logic and agents must be able to understand & modify this.

In addition, spreadsheets, charts, SQL, and docs are common and the agents must understand each and the relations across them.

Specialized execution

Data workflow execution spans specialized systems & tasks.

Agents must know, for example, how to read from SharePoint and write to Tableau or Power BI.

Agents must also know how to execute at scale across Databricks, Snowflake, and BigQuery.

Enterprise governance

Every step must respect controls ensuring correct user access and the use of governed datasets for compliance and security.

The workflows deliver analyses that drive important decisions, needing workflows to have a single source of truth, be reproducible, and drive trust.

AI Landscape for Data Teams

It is daunting to pick a product in a market where numerous vendors are vying for your attention, with each promising the world with AI. Let's break down the AI products for data into a few categories.

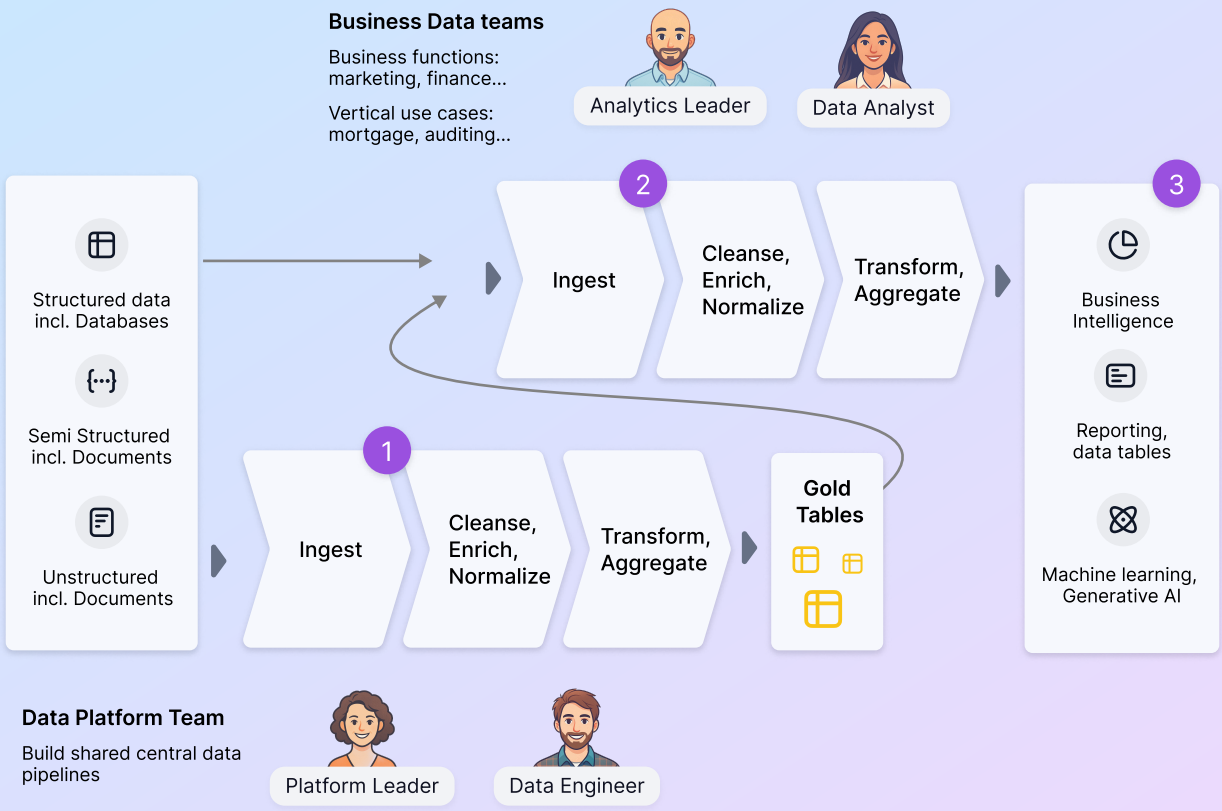

The distinct use cases

Working with data involves multiple teams, each trying to achieve different goals at different times in the data lifecycle:

1. AI for core data engineering
2. AI for business data prep & analysis
3. AI for BI querying & decisioning

AI for core data engineering (central datasets)

Purpose	Build and maintain shared, governed datasets
Primary Users	Central IT, data platform, and data engineering teams
AI interfaces	IDE coding agents (e.g., code copilots), notebook-native assistants
Notes	Not for day-to-day business analysis or ad-hoc prep

AI for business data prep & analysis

Purpose	Drive insights for business problems from many data sources
Primary Users	Data analysts, business data users
AI interfaces	Visual drag-and-drop pipelines, documents, chat, and SQL
Notes	Prophecy focuses here, replacing Alteryx, SAS EG, BI data prep

AI for BI Querying & Decisioning (last mile)

Purpose	Ask questions in natural language and get charts/tables from data prepared for a specific business problem
Primary Users	Business consumers, executives
AI interfaces	Text/chat over a semantic model; auto-generated visuals
Notes	Platform AI/BI offerings and AI-enhanced features within Tableau, Power BI, Sigma, and others

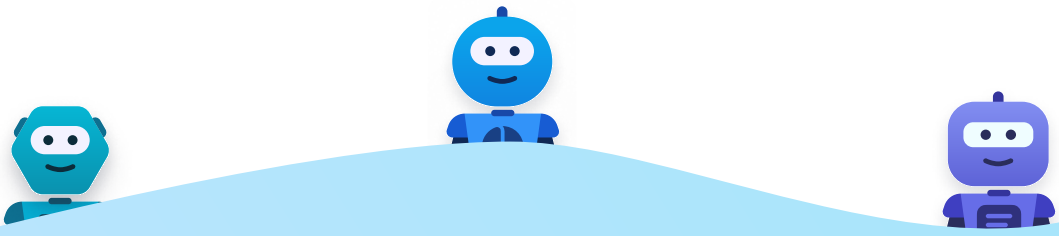

Impact of AI on data prep & analysis

Platform changes are a big lift, and you want a step-change in productivity. AI is the driver to modernize from desktop tools to a governed, cloud-native stack.

Faster time to insights

Agents draft transforms, tests, and documents; interactive loops shrink from days/weeks to minutes/hours.

Higher productivity

Agents handle mapping, joins, deduping, code generation, and documentation, unlocking ~2× analyst throughput.

Lower cost of ownership

Consolidate onto one platform; eliminate desktop infrastructure and rewrites, reduce license sprawl, and use elastic cloud compute.

Stronger governance, trust

Native policies, lineage, versioning, CI/CD, and auditability reduce errors and close security gaps common in desktop workflows.

Net effect: faster outcomes, fewer manual steps, lower spend, and production-grade practices by default. AI also provides an opportunity to fix long-standing architectural issues.

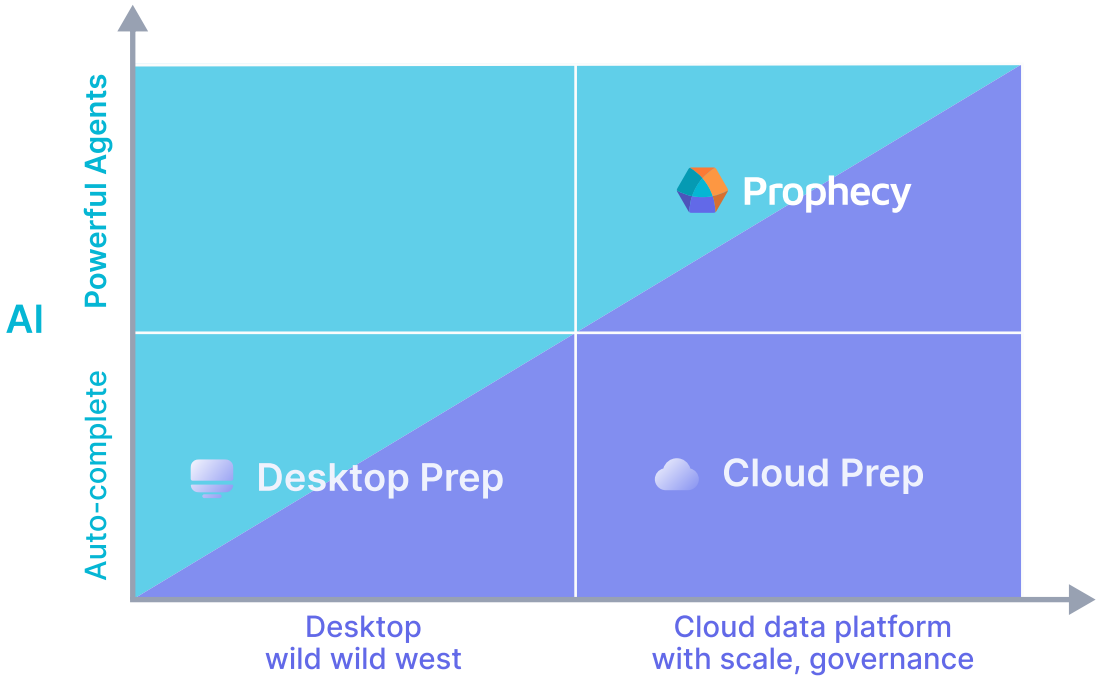

How Agentic AI delivers results

AI data prep and analysis uses visual drag-and-drop interfaces as the primary mechanism of developing business logic for data, where products such as Alteryx are currently used by business data users. AI does two things in data prep & analysis, similar to application programming:

- Brings a step-change in productivity on medium-to-larger tasks (like Cursor has done for programming)
- Lowers the technical bar for simpler tasks (like vibe-coding has done for programming)

Let's see the mechanics of how AI brings these changes.

AI-driven data lifecycle

AI moves a lot of user work from writing the initial business logic to inspecting and refining the results produced by AI. This leads to a very simple data lifecycle, but one that requires excellent technology for each phase.

Generate

Agents interpret the business prompt in the context of schemas, metadata, and prior work.

They produce results that are easy to review as visual workflows.

Refine

Users quickly inspect and edit in a visual interface, previewing data after each step and making changes to steps or adding new ones: iterating until the result is correct.

Deploy

Designs move from interactive runs to scheduled, governed production on the central data platform with git versioning, tests, approvals, CI/CD, monitoring, and alerting.

Practical usage of Generate → Refine → Deploy

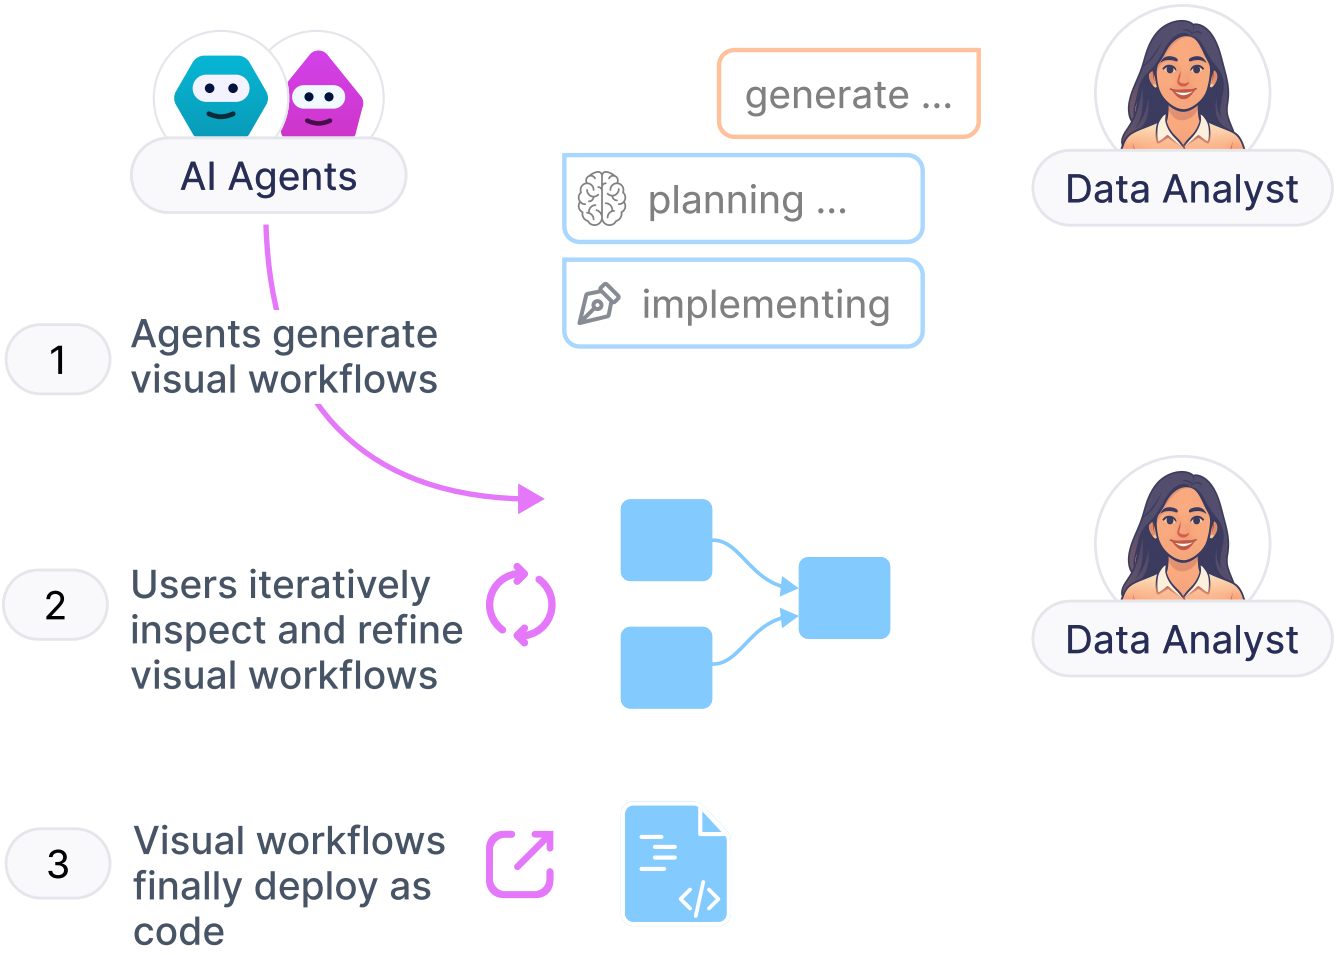

What does it take to build this lifecycle?

The agents have to be experts at the work being done by analysts, in the interfaces used by them, and must have the final robust deployment infrastructure.

Specialized Agents

Our agents understand data context (schemas, types, lineage, quality rules) and use that to generate accurate transforms, tests, and documentation.

Native to Visual

Native visual canvas integration for agents make it easy for the results to be reviewed with human-in-the-loop to ensure trust.

Execution, Ops

Designs move from interactive runs to production on the central data platform with governance, git versioning, tests, CI/CD, monitoring, and alerting.

Primary agents for data prep

Our agents are specialized in data prep and experts in working with interfaces that data users are familiar with. Most tasks will get accelerated, and some automated, with our specialized agents.

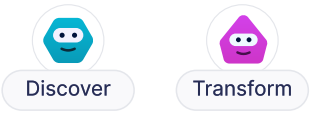

Discover, Transform Agents

Help users find relevant datasets, ingest them, and transform (cleanse, normalize, enrich, aggregate) them for insights.

These agents bring the full power of coding agents with data context to the visual layer, providing industry-leading results.

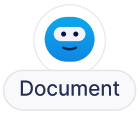

Document agent

Helps users document the data transformations applied to datasets for regulated industries such as pharma, accounting, or finance.

The agent understands the data operations, the target template, and completes the objective sections, and walks the user through subjective questions.

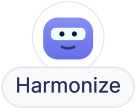

Harmonize Agent

Help data users onboard datasets from various sources, customers into a common data format so downstream transforms and processing can be done uniformly.

It uses definitions of the target common data model, and an example to produce high quality results.

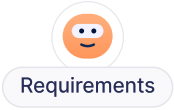

Requirements Agent

Helps generate data workflows from requirements documents and then keeps the requirement documents up to date as the workflows change.

This is essentially a Live document with the three layers - document, visual pipeline and the code always kept in sync.

Can Alteryx build a competitive product?

As we look at advances in technology, we're convinced that just like coding, in 12-24 months, no one will want to use a product that does not have advanced AI for data prep & analysis. Alteryx is a much-loved product that we also like, so we'll present technical requirements—not opinions. There are two main architectural options:

Apply AI at the visual layer

Agents can directly read and edit node-and-edge "boxes and wires", but it runs into:

- **Training data mismatches:** Foundation models are trained on text/code, not proprietary visual graphs; supervision for node/edge edits is sparse and inconsistent.
- **Determinism & testability:** Graph edits are hard to validate automatically; small changes often cascade in non-obvious ways.

Empirical note: In our own experiments (2024–2025), agents operating only on visual workflows produced brittle outcomes and significantly poorer generalization than code-aware approaches.

Apply AI at the code layer, use a compiler that unifies visual, code

Agents reason in code (SQL, Python, Spark, etc.). A bidirectional compiler guarantees lossless round-tripping between code and the visual canvas. Minimum requirements for a viable solution:

- **Data context:** Use best-in-class code agents grounded in schemas, types, quality rules, lineage, and business definitions.
- **Compiler & execution tools:** Agents must have a compiler that restructures code to match the visual components in the user's current context. Agents must use tools to check compile-time and run-time success and fix errors.
- **Visual canvas compiler:** A code-backed canvas requires compiler assistance: interactive feedback, iterative static analysis, parameter autocompletion, configuration validation, and inline errors/warnings.

Delivering this level of quality requires coordinated expertise across AI agents, compilers, UX, and large-scale data execution.

A vendor can provide a product with some AI built on the visual layer, with promises of better AI to hold their customers through an extra renewal cycle. Eventually, they must move to the Prophecy model to deliver quality results and this is a heavy lift, requiring a new product, built by a new team.

What's the path to AI modernization?

As you consider adopting Prophecy to bring agentic, AI-driven data prep & analysis aligned with your data platform, there are three primary questions to answer:

- **Architecture of Prophecy:** What is provided and how does it fit into your stack?
- **Change management:** What does the adoption path look like?
- **Financial case:** What are the dimensions along which value will be added?

Let's look at each of these in turn.

Architecture of Prophecy

Prophecy is designed to have familiar user interfaces, with specialized AI that run on cloud data platform providing governance and scale, giving you the perfect architecture you can ask for in a new product.

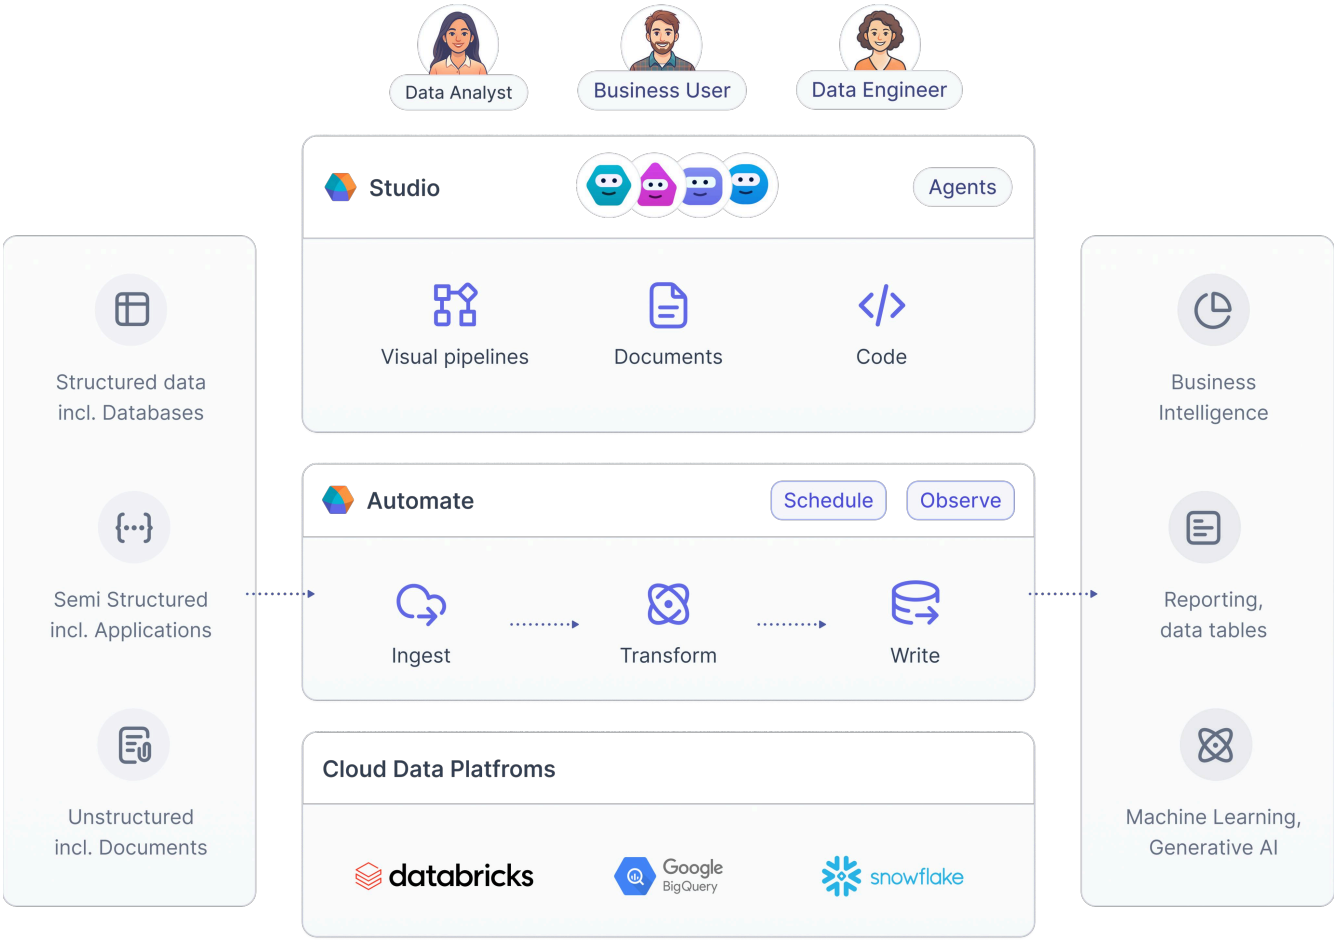

Prophecy Studio

Studio is the multimodal interface where a data workflow is represented as a visual workflow, as a document, and as code.

- **Rich & extensible visual interface:** A familiar canvas with visual operators (“gems”) for common operations; custom gems can be developed quickly as code packages.
- **AI agent integration:** Agents are integrated into the visual canvas with a prompt interface and can understand, edit, and write visual workflows (and documents/code).
- **Rich tools:** A sophisticated, bidirectional compiler keeps visual ⇄ code representations in sync and can express any code using built-in or user-developed gems.

Prophecy Automate

Automate provides the complete execution layer, where transformations are developed and executed natively on the underlying cloud data platform.

- **Connectors:** Built-in connectors for common analyst sources (e.g., read from SharePoint; write to BI tools such as Power BI/Tableau). Connectors are simpler for business users and AI-assisted to configure.
- **Transforms:** Data logic is represented as code native to Databricks, Snowflake, or BigQuery and executed there with platform governance and best practices.
- **Scheduling & observability:** Schedule end-to-end workflows (connectors plus platform transforms), monitor runs, and use the visual interface to diagnose and fix production issues.

Our agents brainstorming to get you the best results

Change Management

Change management is essential for enterprise adoption. We take a holistic approach that covers:

Moving Users

Current Alteryx users get a familiar visual interface, with very easy to use agents that they can chat with to make asks.

Moving Assets

Prophecy has an “Import” button that converts workflows with high (90% +) accuracy. The visual canvas can be used to make final changes.

Fwd Deployed Eng.

Prophecy’s expert team will ensure the first success with you in 90 days, and enable your users to be self-sufficient as to scale modernization.

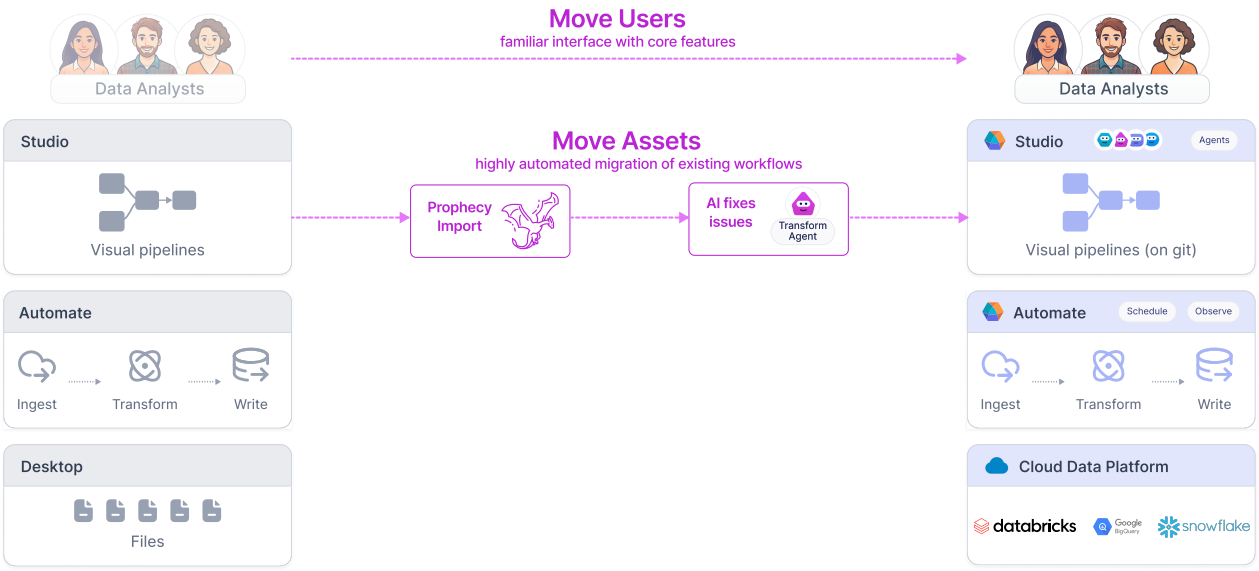

Accelerated Adoption - Enterprise Express

Designed for rapid results on real world use-cases, without long and expensive POVs

Fast Acquisition	Discounted offering at $4K/month with a dedicated SaaS deployment
Discounted Seats	20 user seats so your whole team can try the new experience
Fwd deployed eng.	Our engineers work hand-in-hand with your team for 90 days to deliver the first new or migrated workflows
Outcomes	
Fast win	Clear success on a priority use case in your environment
Change experience	Real, in-environment learning to plan a larger rollout with confidence

Financial Case

The financial case for moving from Alteryx desktop to Prophecy is straightforward and compelling. Here's a level deeper:

Faster time-to-insight:

Business data users, assisted by AI agents, deliver insights sooner, accelerating customer outcomes and internal decisions.

Trusted results:

Governance and versioned workflows restore trust by ensuring:

- A **single source of truth** for data workflows.
- A **robust change process** (tests and business-logic validation)
- **Governed datasets** on your cloud data platform

Reduced cost

- **Direct:** lower license spend
- **Indirect:** eliminate a separate Alteryx execution/automation environment (and its production support) and duplicative governance.
- **Rewrite:** avoid re-implementing Alteryx workflows in PySpark/SQL for production—build once on the data platform and deploy with tests.

When these savings are combined with AI-driven productivity gains, the case is clear.

The exact model varies by customer; we can size it with your volumes, licenses, and workload mix.

Summary

We made a case for moving from **Alteryx to AI Agent**-based data prep that enables a step-change in speed, productivity, cost, and governance.

We showed how Prophecy's **Generate → Refine → Deploy** loop enables AI agents to draft initial transforms, analysts to review and edit in a visual canvas for trusted results, and deploy governed workflows that run natively on Databricks, Snowflake, or BigQuery.

We outlined how **Prophecy uniquely delivers this**: specialized agents (Discover & Transform, Document, Harmonize, Requirements) that are metadata-aware and interface-native; Prophecy Studio for multimodal authoring with a bidirectional visual ↔ code compiler; and Prophecy Automate for connectors, native transforms, scheduling, and observability.

The outcomes are concrete: minutes/hours instead of days/weeks, roughly 2× analyst throughput, lower TCO via consolidation and no rewrites, and stronger trust through versioning, lineage, CI/CD, and policy guardrails.

Finally, we detailed a **pragmatic path to adoption that is low-cost, low-risk** using our Enterprise Express offering

See us bring magic to your team!!

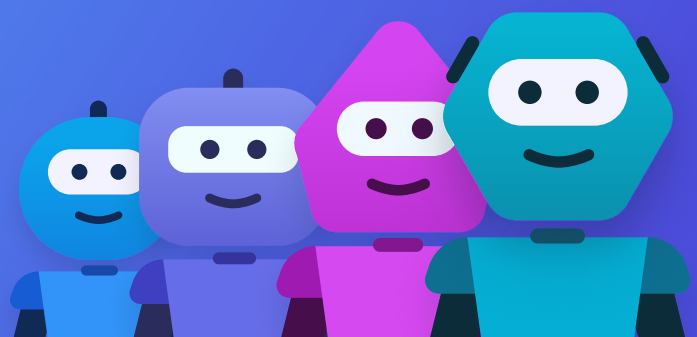

Prophecy