alteryx

E-BOOK

How to Trust AI: Revealing the Logic Behind Enterprise Agents

How an AI Data Clearinghouse helps you control and understand the decisions your enterprise LLMs and agents are making.

Table of contents

Introduction

A little more than a decade ago, machine learning made headlines. It could detect patterns better than humans and make predictions with startling accuracy, transforming everything from streaming services to data science.

Deep learning was the next major evolution. Suddenly, AI could hear and make sense of complex stimuli with far greater accuracy, enabling breakthroughs in image recognition, natural language processing, and speech understanding.

Then, generative AI flipped the business world upside down. Finally, we could talk to machines, and they could talk back. Gen AI helped us create new content, think up novel ideas, and find answers in seconds.

Today, Agentic AI is the next quantum leap forward. These digital workers are fully autonomous AI systems that can perceive the world, act with contextual insights, and learn over time. Together, LLMs and AI agents are quickly transforming modern enterprises, making real decisions that affect customers, employees, and the bottom line.

But there's a problem.

AI models are complex, and most businesses have no way to understand how these models arrive at decisions. Without explainability or lineage, business users can't verify, audit, or govern AI's outputs, creating risk and stalling adoption.

That doesn't mean risk is unavoidable. In fact, with the right solution, you can shed light on how your AI systems are making decisions and what data is fueling those decisions so you can confidently use those outputs to power business outcomes.

In this e-book, we explore how you can use an AI Data Clearinghouse to craft AI-ready data that's both traceable and understandable. You'll learn how governance and auditability across your AI systems can give you a crystal-clear lens into what's driving AI's outputs, helping you reap AI's rewards while minimizing its risks.

71%

of businesses regularly use generative AI in at least one business function, up from 65% in 2024.

The state of AI: How organizations are rewiring to capture value **McKinsey**

Welcome to the Modern Enterprise

Yesterday's digital finish line was cloud adoption. Today, it's deploying GenAI at scale, and it makes sense why. The potential ROI is too big to ignore, estimated at $4.4 trillion in added productivity. In fact, a recent study by IDC found that companies made an average of 270% ROI for every dollar invested in AI.

However, while adoption is surging, trust is plummeting. A 2025 survey from PwC showed that over half of CEOs believe GenAI will increase profitability next year, but only a third trust having AI in key business processes.

This skepticism isn't isolated to leadership. Slightly over half of workers don't trust the data used to train AI systems, and 47% of all organizations have experienced a negative consequence from using generative AI, including inaccuracy, privacy risks, and intellectual property infringement, up from 43% in 2024.

Stanford University's Center for Research on Foundation Models (CRFM) created a Transparency Index that gauges the explainability of major model providers. In 2023, when the index started, the average score among leading providers was 37 out of 100. Woefully low.

In 2024, scores improved, but most major LLM providers still hovered around 50, still too low when considering how many business-critical decisions are relegated to AI. This rising AI skepticism is the backdrop for the next technological upheaval: the rise of agentic AI and the digital workforce, and the world isn't waiting to adopt this new technology.

Respondents report increasing mitigation of inaccuracy, intellectual property infringement, and privacy risks related to use of gen AI.

Gen-AI-related risks that organizations are working to mitigate,[1] % of respondents

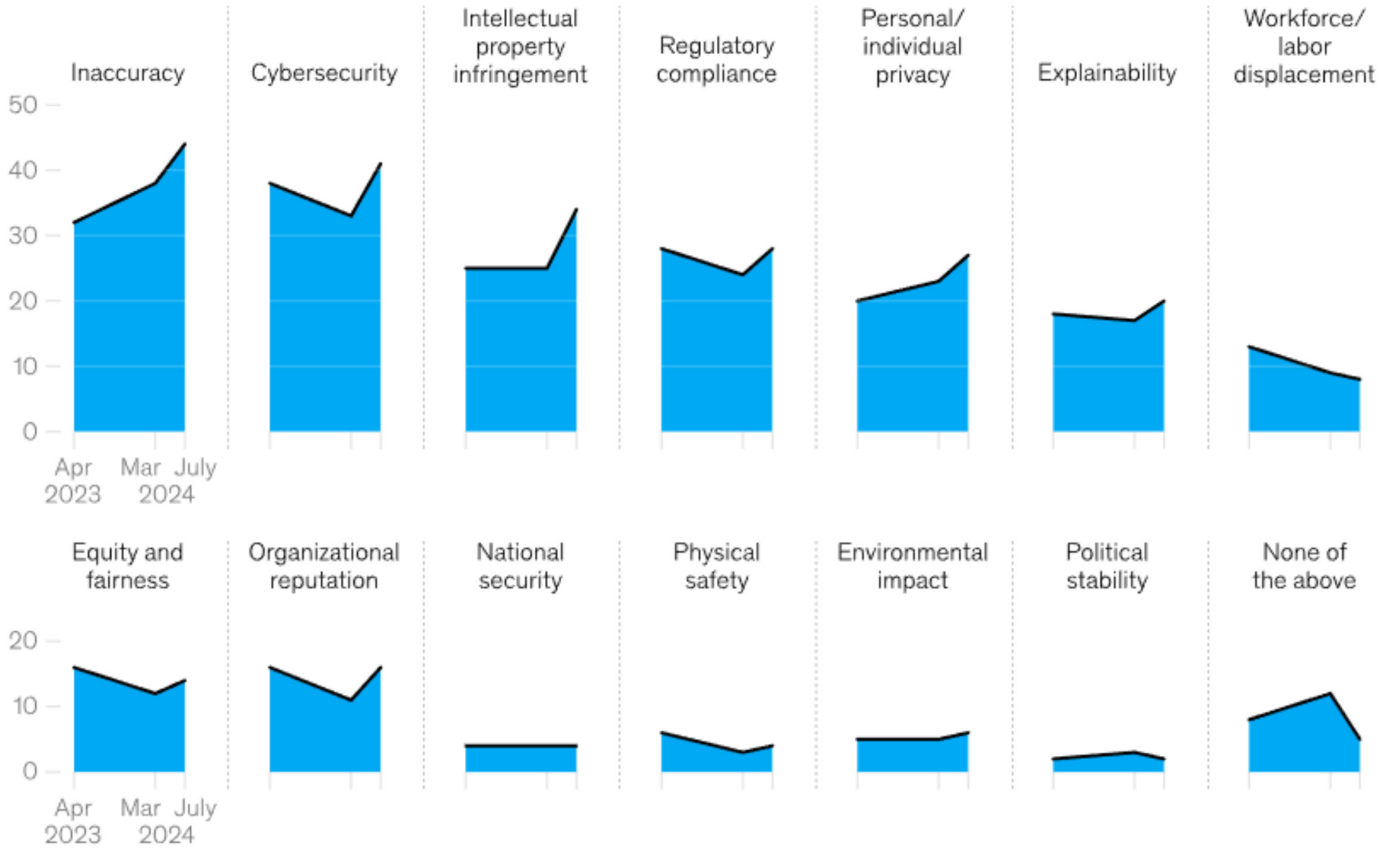

[1]Only asked of respondents whose organizations use AI in at least 1 business function. Respondents who said "don't know/not applicable" are not shown.
Source: McKinsey Global Surveys on the state of AI, 2023–24

Microsoft's 2025 study explored the idea of a new "frontier firm" that harmonizes machine intelligence with human judgment and oversight. According to the study, 81% of organizations expect AI agents to be moderately or extensively integrated into their company's AI strategy in the next 12-18 months. Separate research found that 9 in 10 leaders believe their organization will adopt AI agents faster than generative AI.

The potential productivity gains of agentic AI are impressive. In use cases like customer service and software development, AI agents can net productivity gains upward of 50%. But that doesn't matter if the outputs and decisions of these AI partners can't be trusted.

Without a logic layer that makes AI outputs transparent and auditable, most organizations will need to limit agentic AI to safe, low-risk productivity tasks. These tools may excel at summarizing meetings or drafting emails, but they won't be trusted to support decisions in areas like finance, compliance, or operations where accountability matters.

Business leaders agree: If it can't be explained, it can't be trusted.

In practice, that means agentic AI risks becoming just another productivity tool rather than a true driver of transformation. To move beyond this ceiling, businesses need a foundation that brings structure, oversight, and clarity to how AI decisions are made.

However, organizations can safely and effectively adopt the latest AI technology without exposing themselves to runaway risk. It requires deep insight into the data and logic powering these AI systems.

Understanding the Logic Behind AI's Decisions

To understand how AI systems make decisions today and how you can improve AI's transparency it's important to understand the flow of modern enterprise data systems.

Data typically flows in a linear pipeline:

- Enterprise systems like CRMs, ERPs, and HRIS send data to cloud data platforms like Snowflake and BigQuery.
- From there, data is sent to an analytics or ETL tool for cleansing, preparation, and transformation.
- The modified data is pushed back to the cloud or into embedded business dashboards and models.

While this may appear straightforward, the problem is that LLMs don't follow static rules like a chatbot running a decision tree might. They synthesize inputs, predict patterns, and produce outputs in ways that can be difficult to trace. This is especially true if you don't have visibility into the data powering these AI systems.

60%

of organizations will abandon AI projects unsupported by AI-ready data.

Gartner

There's a lack of AI-ready data

According to Gartner, through 2026, organizations will abandon 60% of AI projects unsupported by AI-ready data. The problem isn't a lack of AI use cases or even AI technology. The problem is that the data powering these AI systems lacks context, structure, and governance.

AI tools assume data is ready when they consume it, but they don't prepare it or verify its trustworthiness. However, data needs this trustworthiness — governance, auditability, and traceability — to create reliable outputs that businesses can safely use for important decisions. Otherwise, outputs could be misleading or flat-out wrong. They'll be impossible to replicate or trace, and due to this lack of credibility, enterprise AI projects will stall indefinitely.

Data lacks context

- Raw data alone (e.g., rows in a spreadsheet) isn't enough to generate high-quality outputs or results
- Business logic, relationships, and metadata are missing
- Chatbots hallucinate without relevance

Data isn't governed

- Teams haven't decided who approves what LLMs can access or say
- Sensitive or unvalidated data creates higher levels of risk
- Without an audit trail, there can't be any trust

Business and IT are disconnected

- Domain experts know the data's meaning but can't shape it in a way that's ready for AI consumption
- Data scientists spend time prepping, blending, and cleansing data instead of modeling
- AI is built in silos, misaligned with business goals

AI-ready data is rare

- Clean, unified, enriched, and semantically structured data is difficult and costly to produce
- Every team, system, and stakeholder has different needs.

The invisible calculus behind business decisions

Unclear and ungoverned data is one side of AI's murkiness. The other is the undocumented knowledge that lives inside workers' heads. For years, humans have made decisions with, well, other humans, and this paradigm persists.

The reasoning behind pricing exceptions, strategic pivots, and operational overrides is often undocumented. It lives in email approvals, hallway conversations, and ad-hoc meetings. You can't build trustworthy AI outputs on logic and undocumented institutional knowledge that lives in conversations and people's heads.

However, with the right solution, it's entirely possible to give both clean, AI-ready data and contextual insights to your AI systems. With an AI Data Clearinghouse, you gain full visibility into both the data and the logic powering your AI systems.

The AI Data Clearinghouse

Alteryx is an AI Data Clearinghouse that helps everyone, from business users to analysts and data scientists, transform raw, siloed data into explainable, compliant data ready for AI consumption. Alteryx can connect to your data wherever it lives, cleanse and prepare and transform it, then connect to a suite of AI applications, powering limitless AI use cases with pristine data.

Alteryx gives you complete visibility into what's happening to your data at every stage of the analytics and AI lifecycle. You can summarize workflows in clear, natural language with AI's help, track decisions with embedded metadata and audit logs, and monitor and adjust model and agent performance over time, all with role-based access and enterprise-grade security.

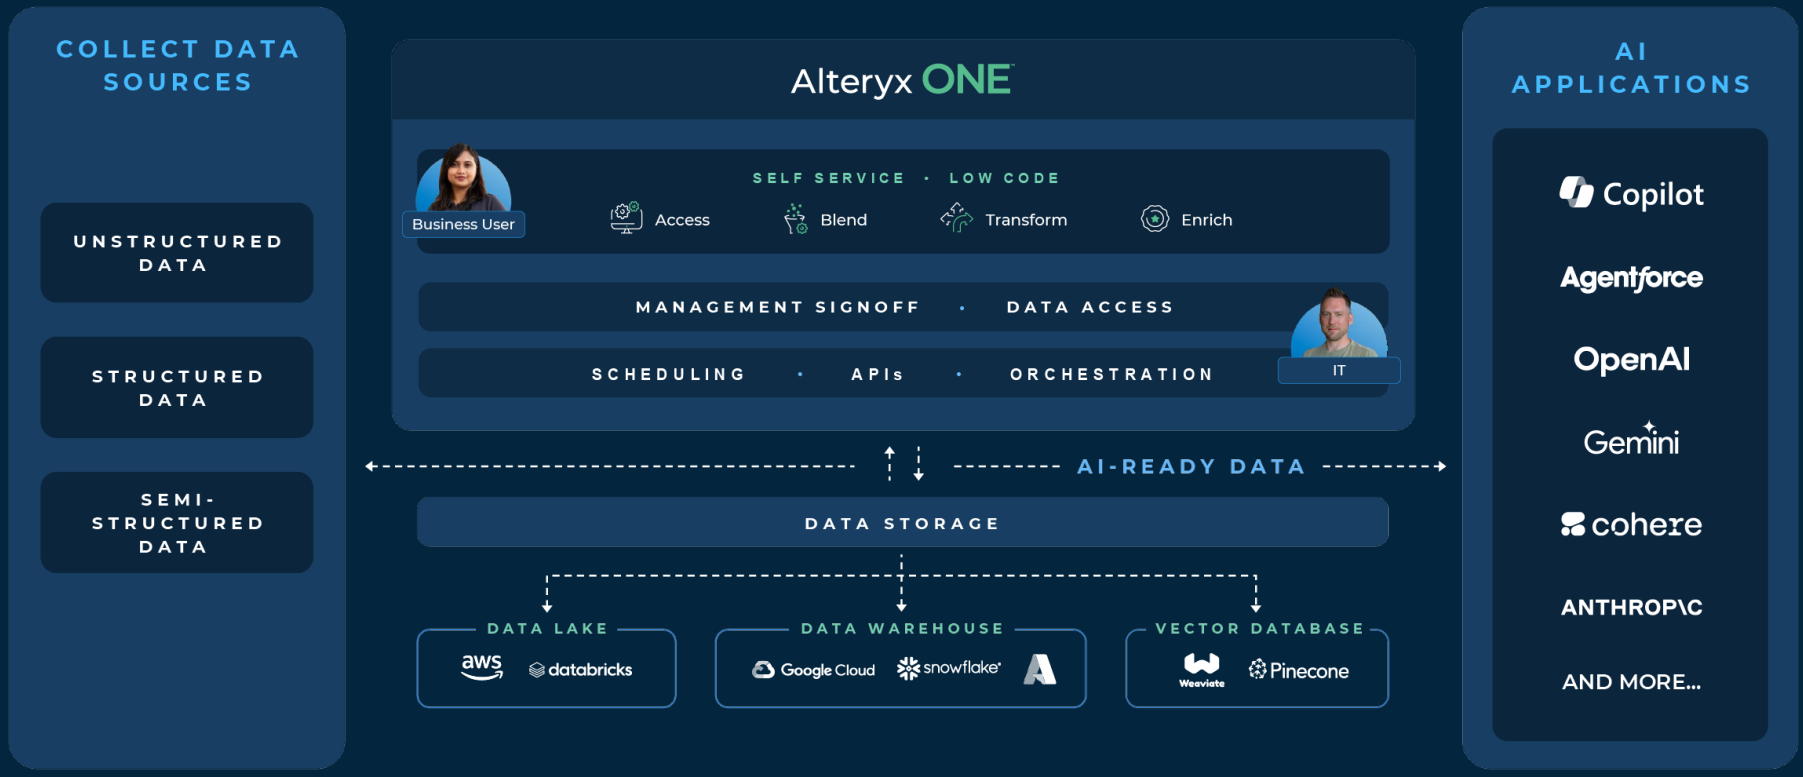

Benefits of the AI Data Clearinghouse

Connect to all your data, everywhere

With 300+ data connections and the ability to scrape web data, Alteryx empowers you to work with nearly any data source: data warehouses, ERP and cloud-based applications, flat files, MS Office applications, social media data, and even legacy analytics platforms.

Alteryx can also help you work with any type of data, pulling in structured, semi-structured, and unstructured data from anywhere. Our visual analytics canvas gives analysts and business users hundreds of tools to prepare and transform data, letting them leverage their extensive domain knowledge.

Whether your data environment is cloud-native, hybrid, or on-prem, Alteryx gives you the freedom to build, schedule, and run analytic workflows wherever it makes the most sense for your business.

Govern and approve data collaboratively

Data can't become AI-ready without oversight. Alteryx offers embedded approval steps, role-based access, version history, and traceable workflows. Governance stakeholders — including legal, compliance, and data stewards — can validate transformations and enrichments throughout the pipeline. This transparency ensures AI models use sanctioned inputs and satisfy audit requirements from day one.

Enrich with context to eliminate hallucinations

LLMs perform poorly with generic or decontextualized inputs. Alteryx enables business users to input critical metadata, logic, and business-specific signals directly to data sets and AI agents.

Adding this undocumented tribal knowledge to your AI systems enriches AI models with relevant inputs, minimizing hallucinations and surfacing insights grounded in real-world business conditions.

Orchestrate AI-ready workflows at scale

Data must be continuously updated and delivered to AI systems. With Alteryx, teams can orchestrate workflows via APIs, schedule jobs across environments, and automate end-to-end pipelines. Whether sending data to vector databases, document stores, or LLM interfaces, Alteryx ensures freshness, reliability, and repeatability across systems.

Support AI models with searchable, explainable data

To be auditable, AI inputs must be traceable. Alteryx outputs include metadata tagging, semantic structuring, and integration with vector databases that support retrieval-augmented generation (RAG), which makes LLMs less likely to hallucinate by pulling contextual information from relevant knowledge docs. Teams can easily connect downstream models to explainable, governed, and semantically searchable data, empowering trustworthy and transparent AI interactions.

Business-ready, analyst-friendly

Alteryx helps users automate time-consuming manual analytics tasks, like removing nulls and duplicate entries, grouping or summarizing data, finding unique values, and more, all without requiring extensive coding or IT support. Data professionals can build models and AI agents, working alongside LLMs in a low-code, no-code environment.

Building AI agents with Alteryx

The Alteryx One Platform can help you maximize generative AI's value in multiple ways, whether it's preparing AI-ready data for your internal and external LLMs or using integrated GenAI features to quickly build and document analytic workflows.

But Alteryx doesn't stop there. Our platform can also help you build trustworthy, explainable AI agents directly in the Alteryx One Platform.

For as complex as AI is, building an AI agent in Alteryx is extremely straightforward. Simply drag and drop the "Agent" Tool onto your Alteryx canvas. You can select the LLM of your agent, whether that's Claude, ChatGPT, or any other LLM you prefer, including internal, customized LLMs.

Then, you can specify any constraints and provide contextual information to guide your agent's decision-making and actions. Even better, after you run a workflow, you can inspect any item in your data set or output to see a detailed breakdown of an agent's reasoning and decision-making process.

Anatomy of an AI Agent

Understanding the Core Components of Modern AI Agent Architecture

Real-World Examples of Alteryx-Powered AI Agents in Action:

1. Understanding customer satisfaction surveys

2. Reducing customer churn

3. Supply chain optimization

A multinational retail company conducted a customer satisfaction survey across 15 countries and multiple channels to better understand what customers wanted, but the data they received was highly inconsistent and unstructured. It came in multiple languages and required a ton of cleansing before it was ready for analysis.

With an Alteryx-powered AI agent, the retailer automated the entire process of cleansing this data. Then, with the Agent Tool, they were able to choose which LLM powered their agent, specify its goals, and even add external information to their data set, like local currency and geographic information. At the end of the workflow, the AI agent checked its work, ensuring every action and decision was compliant with its instructions and assigned a confidence score to each outcome.

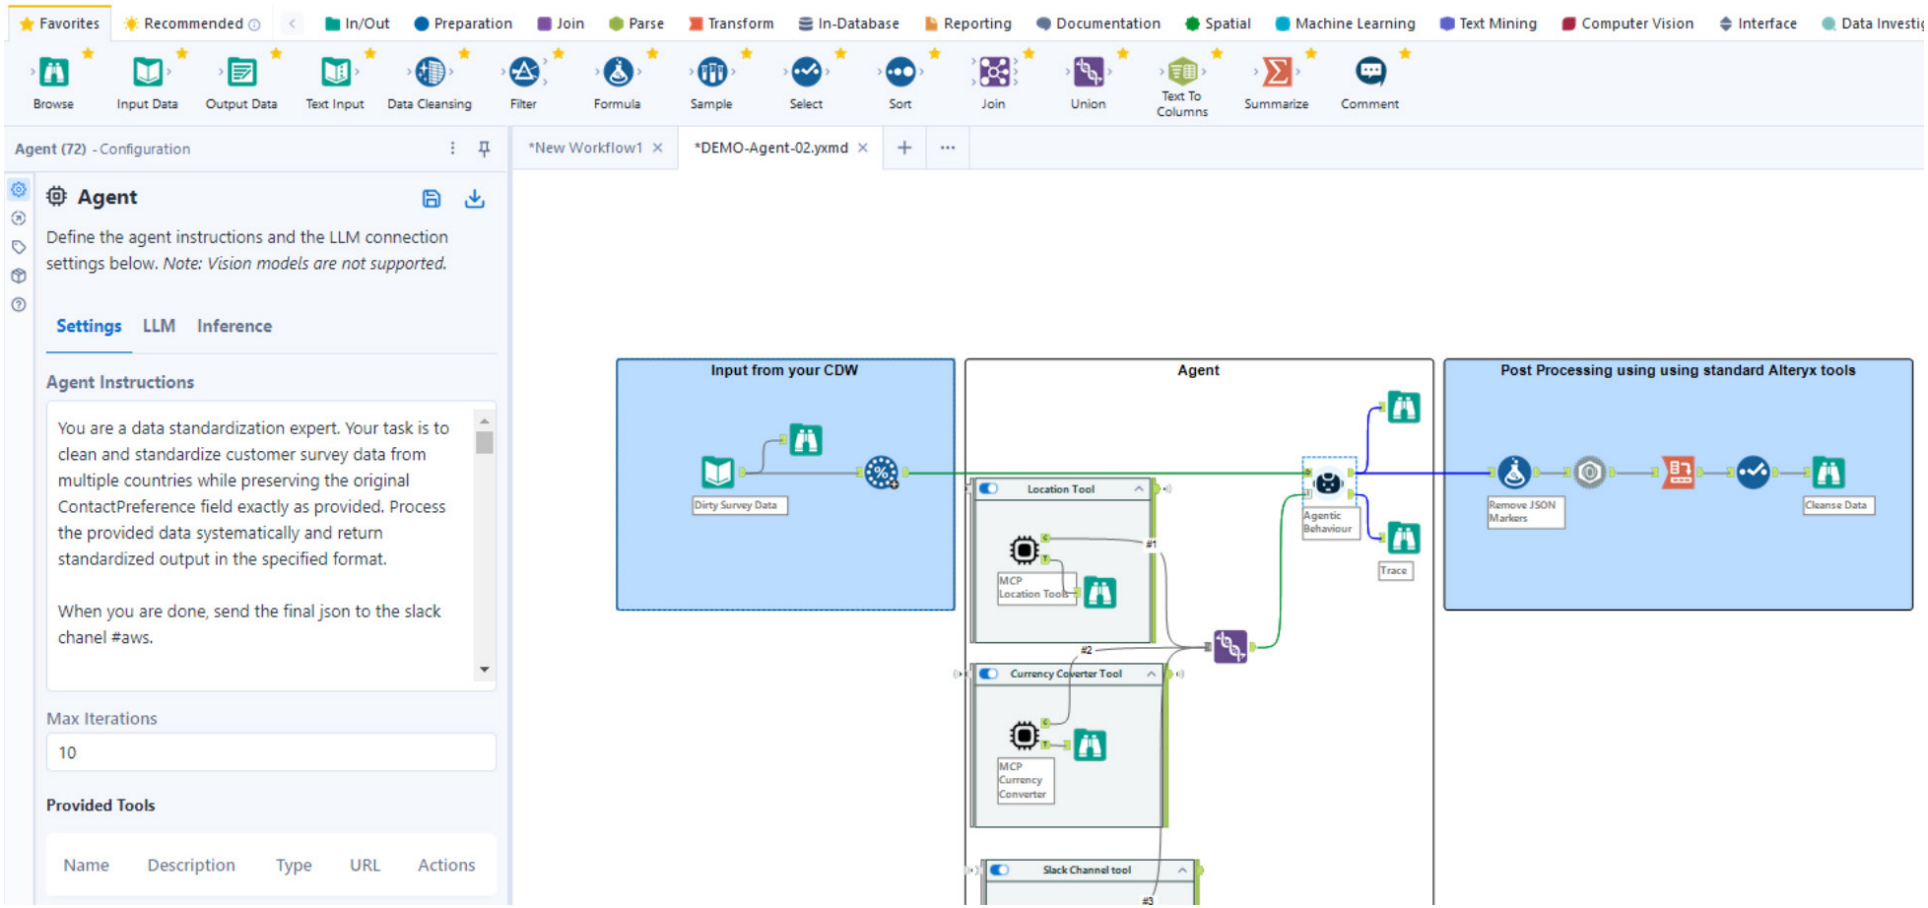

Customer Satisfaction Survey Analysis in Alteryx Designer

AI agent workflow

Perceive

The agent ingests raw data from forms in multiple languages, scanning for inconsistencies across names, addresses, date formats, and contact fields.

Act

The agent uses LLM-powered reasoning and external APIs (like Google Maps) to standardize addresses, enrich location data, and format customer preferences in structured JSON.

Decide

Based on the defined goal and constraints, the agent breaks down each record, field by field, identifying anomalies and mapping inputs to target formats.

Learn

After every transformation, the agent checks whether it met the defined output criteria, assigns confidence scores, and updates its reasoning loop, getting smarter with each iteration.

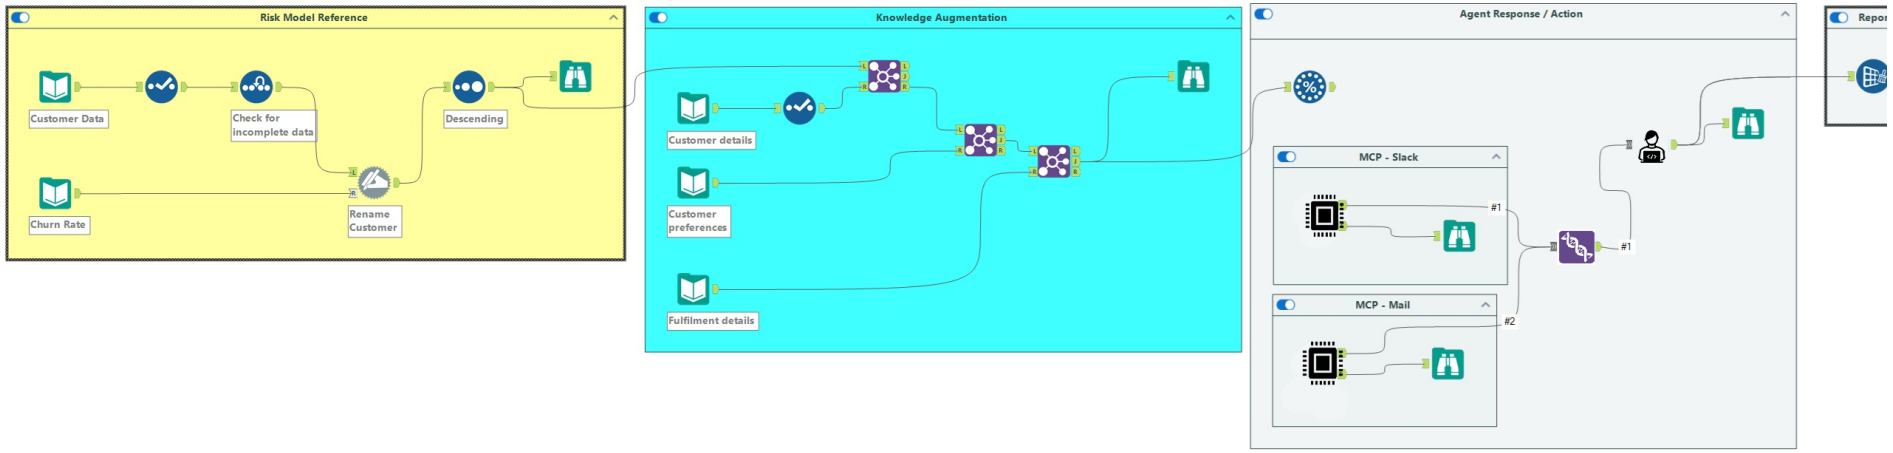

Customer Churn Agent Workflow in Alteryx Designer

Reducing customer churn

Another real-world use case is spotting early signs of customer churn, one of the most expensive and persistent challenges retailers face. Often, by the time teams detect signs of disengagement, it's too late, and the customer is lost.

Typical churn models generate risk scores for different customers, but they still require a lot of manual work, and taking corrective action can be a slow and tedious process.

With an AI agent, you have a fully autonomous digital worker that can monitor your customers 24/7 for early signs of attrition. Then, you can have the agent flag accounts for human intervention or have the agent automatically send out personalized outreach or take other corrective actions.

Customer churn AI agent workflow

Perceive
The agent monitors behavioral data from loyalty programs, app usage, purchase history, and support tickets to detect early signs of disengagement.

Decide
It combines real-time input with historical data to evaluate risk, segment customers, and determine what type of intervention is most effective.

Act
Depending on the segment and risk level, the agent triggers the appropriate action, whether sending targeted push notifications, launching loyalty offers, or drafting emails, all with the ability to route decisions through human approval if needed.

Learn
Engagement metrics and campaign outcomes are fed back into the model, allowing the agent to adjust its strategy and improve its outreach and intervention strategies over time.

Supply chain optimization

Global supply chains can be fragile, and disruptions can cost companies millions in revenue and operational headaches. Most supply chain dashboards are reactive, though. They require manual input to detect bottlenecks and assess delays. Even automated systems struggle to adapt mid-process when supply chain conditions suddenly shift.

An AI agent built in Alteryx can continuously monitor logistics data and adapt to supply chain stressors in real-time, with or without human oversight, keeping everything running smoothly across the supply chain.

Supply chain AI agent workflow

Perceive

AI monitors real-time data from IoT sensors, ERP systems, logistics feeds, and supplier networks to identify anomalies or delays.

Act

The system dynamically reroutes shipments, reallocates inventory, or adjusts procurement plans to reduce delays.

Decide

Predictive analytics assess supply chain flow, detect potential bottlenecks, and recommend mitigation strategies.

Learn

Continuous feedback improves models, helping them better anticipate future disruptions and optimize resource allocation.

Additional agentic AI use cases

Alteryx has a fleet of enterprise-grade capabilities ready-made for AI systems, including built-in MCP (model context protocol) Servers that let your AI agents safely access your business systems and data sources. You can think of MCP Servers as secure translators that help AI agents safely communicate with your databases, applications, and services while maintaining strict security and compliance controls.

Thanks to Alteryx's MCP capabilities, you can securely connect Alteryx to external agents, powering all your AI systems with AI-ready data. Even better, you can use Alteryx's MCP Servers to have multiple agents work together. Suddenly, you have an entire team of digital workers acting in real-time based on trustworthy, AI-ready data.

Making Confident AI-Powered Decisions

As enterprises race to build impactful generative AI systems and adopt AI agents, AI-ready data has become more valuable — and rarer — than ever. With so much potential upside to incorporating AI into your business, you can't afford not to be able to trust or understand AI's outputs.

With Alteryx as your AI Data Clearinghouse, you can gain a transparent view into how your AI systems are making decisions. You can see the data they ingest, the goals they're working toward, the constraints they're governed by, and the logic behind every action. Even better, you can define which actions should be automated and which need human approval.

Ready to unify your analytics, AI, and cloud solutions into a single platform experience?

Try a free trial of Alteryx today or request a demo to see how AI-ready data can help your team.

About Alteryx

alteryx