The Rise of Agentic Analytics

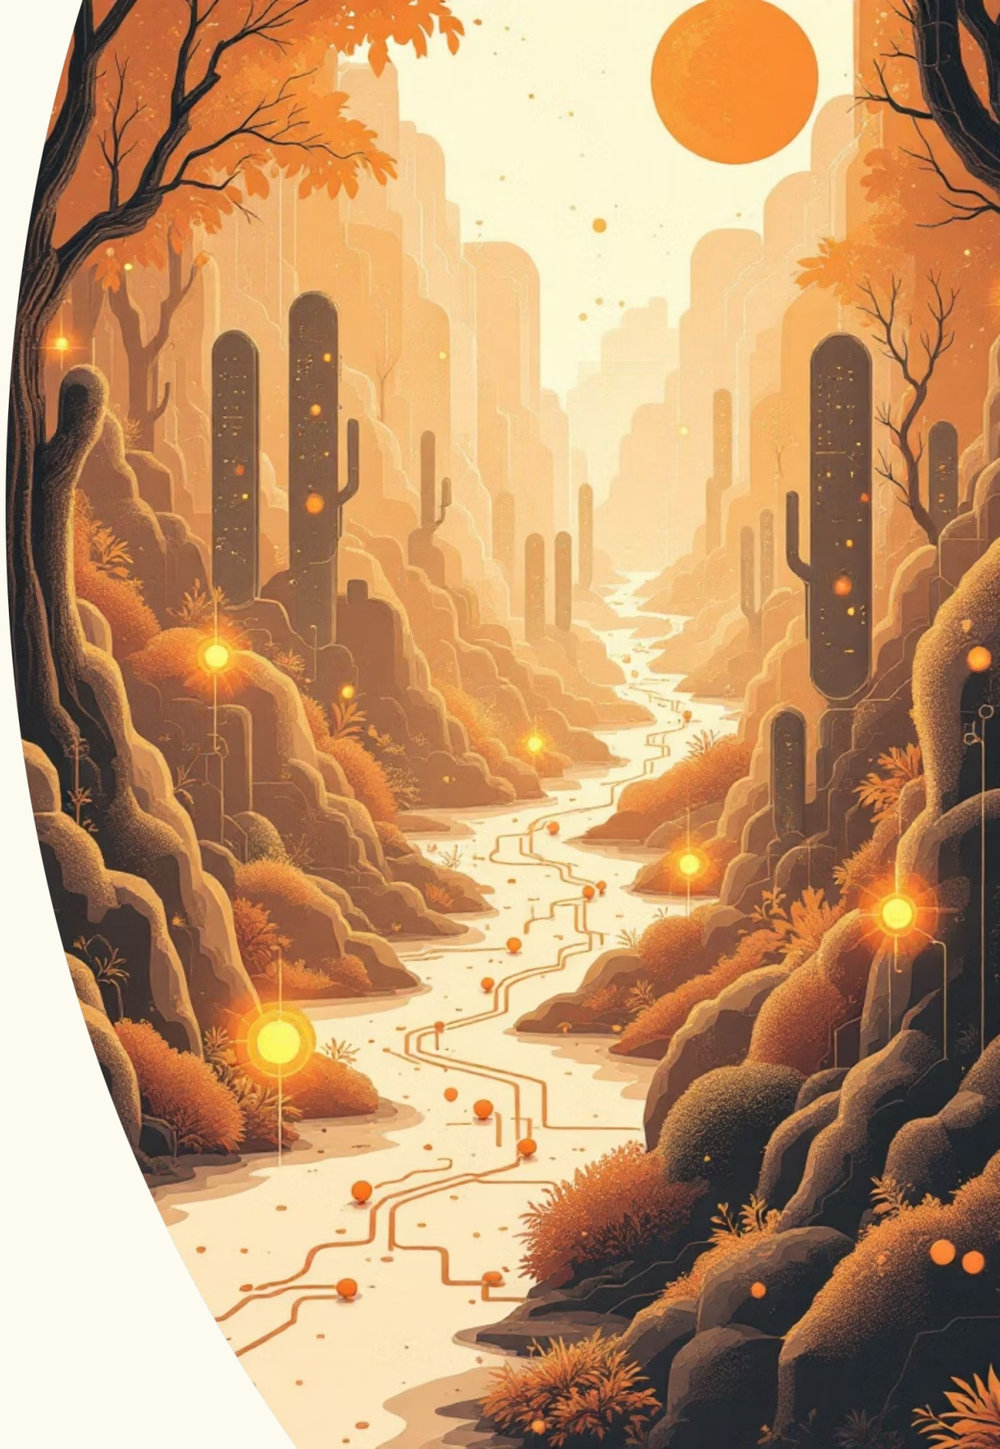

Presentation Outline

1

The Shift in Analytics

Understanding the evolution from reactive dashboards to proactive intelligent systems that transform enterprise decision-making.

2

Agentic Analytics Components

Exploring the three-tier architecture that powers autonomous analytics: diagnostic precision, cognitive reasoning, and intelligent action.

3

Applied Agentic Analytics

Real-world case study demonstrating how agentic systems optimize supply chain operations and drive measurable business outcomes.

The Last-Mile Gap in Analytics

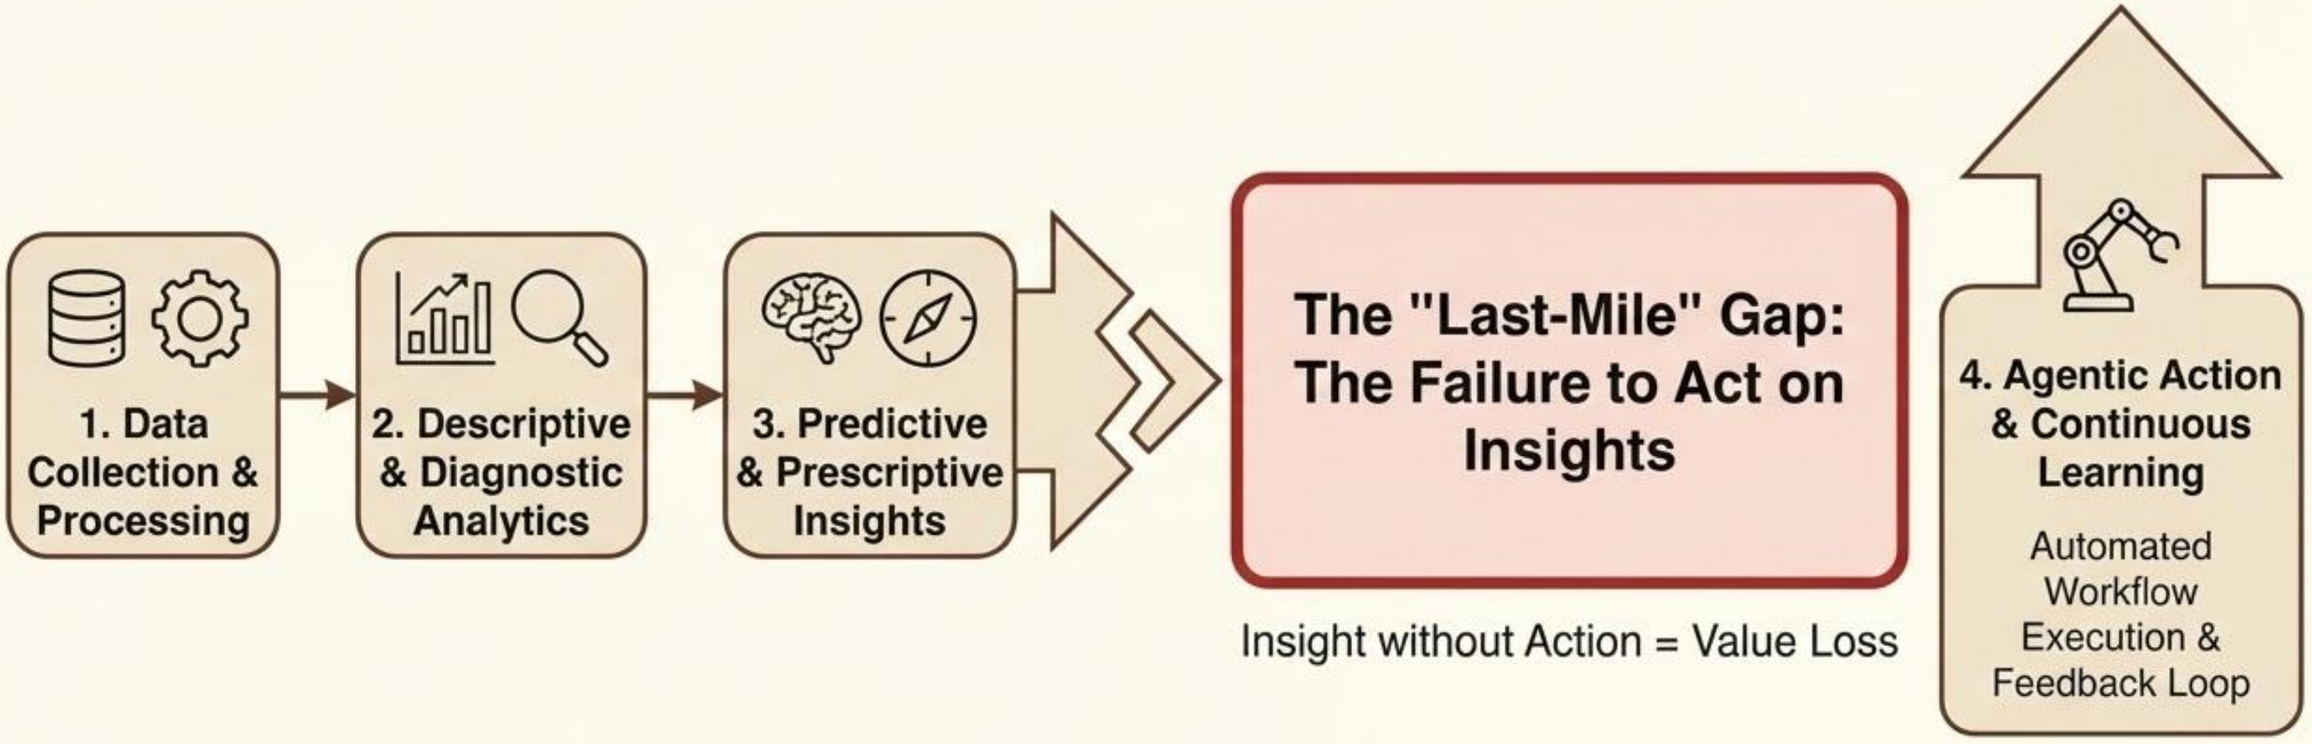

Traditional analytics value generation stops here.

The Agentic Analytics approach bridges this gap.

The Cost of Decision Latency

The silent killer of ROI in high-velocity environments.

The elapsed time between a data event and corrective action manifests as millions in inefficiencies.

The Financial Toll

70%+

of companies admit internal delays cost them revenue, yet they lack effective measurement mechanisms.[1]

- A delay of mere hours in supply chain logistics can result in missed SLAs, financial penalties, and long-term reputational damage.

The “Green Dashboard” Illusion

- Executives are often lulled into complacency by high-level dashboards that show aggregate metrics in the green.
- By the time the aggregate metric turns red, the underlying issue has often metastasized into a crisis.

Cognitive Overload

- As data volume scales, the human capacity to manually diagnose root causes across thousands of dimensions diminishes.
- The complexity of modern data ecosystems has outstripped the human ability to process it without agentic assistance.

The Rise of Agentic Analytics: Lane Assist to Full Self Driving

Current State

Meta Question: "What should I look at?"

Mechanism: Machine Learning algorithms highlight outliers or trends on dashboards, requiring constant human monitoring and interpretation.

Interface: Static visualizations and natural language query (NLQ) tools that respond to user requests.

Limitation: Fundamentally passive—requires a user to be present, attentive, and capable of interpreting suggestions. Operates in "read-only" mode regarding business state.

Future State

Meta Question: "How do we fix it?"

Mechanism: Autonomous agents continuously monitor data streams, reason through context using LLMs, and execute corrective actions via APIs.

Interface: Natural language narratives with automated notifications of actions taken and outcomes achieved.

Limitation: Fundamentally active—operates asynchronously 24/7 with "read-write" capabilities to alter business state directly and autonomously.

Three-Tier Agentic Architecture

Diagnostic Engine

The first layer of the framework is deterministic and grounded in "Hard Analytics." Before an agent can act, it must precisely understand what is breaking and where. This layer does not guess; it calculates

Cognitive Engine

A Large Language Model (LLM)-based module that transforms structured diagnostics into contextual, natural language insights and analytical narratives.

Action Engine

An autonomous, LLM-based agent that distills actionable trends and can recommend or execute corrective workflows such as rebalancing business volume, onboarding new partners, or initiating financial recovery.

The Action Engine operates under a principled, Human-in-the-Loop (HITL) governance framework that ensures safety, accountability, and continuous human oversight.

Case Study: Optimizing Damage Rate

The Analyst

Mark, a Supply Chain Data Analyst at a high-velocity e-commerce company tasked with protecting customer experience and reducing operational costs.

The Objective

Minimize the **Damage Rate** of shipments to reduce costs, avoid financial penalties, and maintain customer satisfaction through consistent SLA performance.

The Success Metric

Ensure **Customer Happiness** by eliminating product damage incidents, directly impacting brand reputation, customer retention, and long-term revenue growth.

The Challenge

The volume and complexity of data,spanning multiple carriers, regions, and warehouses,make manual root cause analysis and timely corrective action impossible, leading to high **Decision Latency** and recurring problems.

Diagnostic Engine: The Ground Truth

Precision analytics that decompose complex metrics into actionable insights, identifying exactly where problems occur and quantifying their business impact.

Carrier	Warehouse	ZIP Code	Baseline Damage	Current Damage	Baseline Share	Current Share	Δ Damage Rate	Contribution	% of Total Δ
A	WH1 (West)	90001	1.8%	4.0%	10%	15%	+2.2 pp	+0.33 pp	32.0%
A	WH2 (Midwest)	60601	2.2%	4.5%	10%	15%	+2.3 pp	+0.56 pp	54.4%
Subtotal A	—	—	**2.0%**	**4.3%**	**20%**	**30%**	**+2.3 pp**	**+0.89 pp**	**86.4%**
D	WH1 (West)	98101	1.8%	2.5%	13%	12%	+0.7 pp	+0.09 pp	8.7%
D	WH4 (East)	10001	2.2%	3.1%	12%	13%	+0.9 pp	+0.11 pp	10.7%
Subtotal D	—	—	**2.0%**	**2.8%**	**25%**	**25%**	**+0.8 pp**	**+0.20 pp**	**19.4%**
Total	—	—	**2.0%**	**3.0%**	**100%**	**100%**	**+1.0 pp**	**+1.03 pp**	**100%**

- Comprehensive metric decomposition across dimensions and time periods for complete visibility
- Full variance explainability quantifying the contribution of each factor to overall change
- Real-time data freshness ensuring decisions are based on current operational state
- Integration with internal knowledge bases for enriched contextual understanding

Diagnostic Engine: The Ground Truth

Variance Decomposition

Separates Rate Effects from Mix Effects for programmatic decision-making and determining the correct intervention.

Rate Effect

KPI change due to a segment's performance deterioration, assuming constant volume.

Example:
Carrier A's damage rate increased from 2% to 5% (vendor performance failure).

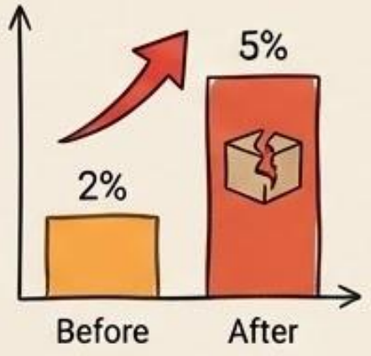

Vendor Performance Failure

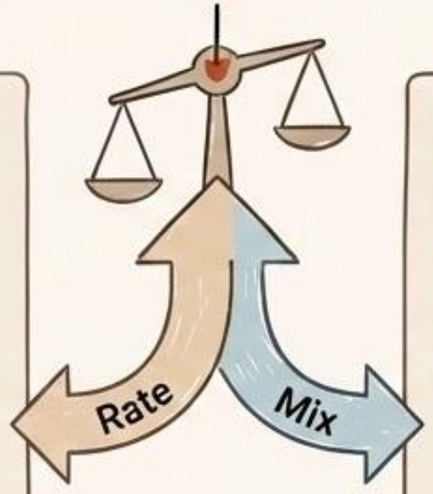

Mix Effect

KPI change due to a volume shift toward a differently performing segment.

Example: Carrier A's damage rate remained high at 5%, but we doubled volume with them because they were cheaper (error in our allocation strategy).

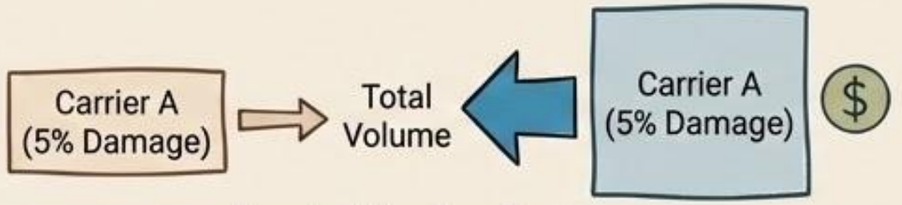

Error in Allocation Strategy

The engine calculates variance contribution for every dimension (Geography, Carrier, SKU, Time) using this decomposition logic.

Cognitive Engine : Context is all you need

An LLM agent functions as an intelligent analyst, reading metric breakdowns and producing concise narrative summaries. The agent identifies key changes, trends, and drivers, translating data into actionable knowledge.

> *"Damage rate increased from 2.0% to 3.0% this week. The rise was driven primarily by Carrier A in the Southwest region (50% of increase), and Supplier ABC Corp shipments (20%)."*

Unlike static rule-based engines, the LLM dynamically identifies top drivers by synthesizing multiple contextual signals alongside numerical decomposition data.

Natural Language Output

Plain-language summaries for quick situational awareness

Contextual Reasoning

Processes up to 512 MB of structured and unstructured input

Grounded Analysis

Data-driven reasoning reduces hallucination risk

The LLM can incorporate diverse contextual inputs: seasonality patterns, holiday calendars, macroeconomic indicators, weather alerts, historical anomalies, operational SLAs, and domain-specific text data.

Cognitive Engine: The Semantic Layer & Chain-of-Thought

The Semantic Layer & RAG

Unstructured Logs

Analyzes shipment notes, support tickets, or emails for keywords.
Example: Finds 80% of failed shipments for "Carrier A" have notes like "wet packaging" or "crushed box".

External Data

Checks weather APIs for storms or news feeds for labor disputes at the time of variance.

Knowledge Base

References internal SOPs or contracts to understand if variance is within allowable tolerances (e.g., Is a 3% damage rate acceptable?).

Chain-of-Thought (CoT) Reasoning

Prompt Step 1 (Observation)	Prompt Step 2 (Contextualization)	Prompt Step 3 (Correlation)	Prompt Step 4 (Synthesis)
Analyze diagnostic output. Identify primary driver.	Check "Shipment_Notes" vector database for Carrier A in Texas.	Check Weather API for Texas during the shipment window.	Synthesize a root cause narrative.
Agent Response	**Agent Response**	**Agent Response**	**Final Output**
Carrier A in Texas is the primary driver, contributing 0.86% to the 1.0% total increase.	Found 400 entries. Top keywords: "Rain", "Wet", "Soaked".	Confirmed. Severe thunderstorms recorded in Dallas/Fort Worth.	The spike in damage rate is primarily driven by Carrier A in the Texas region... highly correlated with severe weather events causing water damage.

This narrative output reduces the 'cognitive load' on human operators and provides the necessary context for the next stage: Action. It transforms raw data into intelligence.

Action Engine : Autonomous Decision-Making Agent

The agent ingests multiple days of summary reports to detect persistent problems and formulate corrective actions with chain-of-thought reasoning.

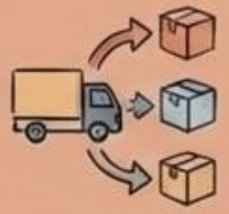

Reallocate Shipments

Shift volume from underperforming carriers to alternatives in affected regions to mitigate damages.

Identify New Partners

Search for new logistics providers with better performance records when current carriers underperform.

Engage Vendors

Draft and send emails requesting investigations or claiming refunds under service agreements.

Escalate Alerts

Post urgent messages to operations leadership when safety thresholds are crossed.

Human-in-the-Loop Governance

Before executing high-impact actions, the agent seeks human approval. For example: "*Switch 20% of volume from Carrier A to Carrier B in TX/NM region for next 2 weeks*" is sent to a logistics manager for confirmation.

This structure mirrors contemporary human–AI alignment frameworks, emphasizing traceability, constraint management, and decision auditability.

Evaluation & Future Directions

Evaluation Plan

Quasi-experimental pre/post comparative analysis within Logistics Incidence organization comparing three months before and after implementation.

Metric Category	Key Metrics
Incident Level KPIs	Damage Rate
Operational Agility	Network Changes
Financial Efficiency	Refunds, Operational & Analytical Effort (Hrs)
Responsiveness	First Response Time, Time to Detect Issues

Future Work

Advanced Tool-Calling

Enable LLMs to interface directly with enterprise systems for end-to-end execution

Enhanced Explainability

Explore Shapley value and causal inference-driven RCA for richer diagnostics

Full Autonomy

Evolve from hybrid decision-support to fully autonomous, self-reasoning platform

A Quantum Leap in Analytics Maturity

The synergy of variance decomposition and agentic LLM-based flows offers a powerful blueprint for AI-enabled business performance management, bridging data analytics and AI in action, taking insight into agentic execution.

Questions

Thank You!