
# Root Cause Analysis in the Age of Agentic AI

From Five Whys to Autonomous Intelligence

Krishna Panchal

TRANSORG ANALYTICS

# Executive Summary

For over five decades, various industry segments have relied on the Five Whys as their primary tool for root cause analysis. (A comprehensive survey on root cause analysis in (micro) services: Methodologies, challenges, and trends, 2024) The method is simple: trace the cause-and-effect chain backwards to the problem's origin. In unstable, data-sparse environments, this approach provided structure, encouraged rigour, and yielded actionable results. (Root Cause Analysis Methods for Microservices, 2025)

The RCA environment has transformed significantly. Systems are highly interconnected, events cascade across functions in milliseconds, and data volumes have surged radically. In this context, a method relying on sequential human questioning, delicate institutional memory, and manual evidence curation is slow and structurally inadequate to address such complexity.

Agentic AI has transformed investigations by removing bottlenecks analysts face: incomplete data, siloed systems, cognitive limits, and the inability to detect complex patterns. Operating on clean, governed data, AI agents generate hypotheses beyond the capacity of any individual team. With humans central to the suggestive decision-making process, the outcome is faster, more accountable, more durable, and more impactful analysis.

| CORE THESIS | Data management is not a precondition for agentic RCA. It is the practice itself. Clean data is the raw material from which patterns emerge, causes are identified, and fixes become stable. |
|---|---|

# From Five Whys to Agentic Intelligence

**The Architecture of the Traditional Approach**

The Five Whys method originated in manufacturing, where failures were physical, observable, and contained. For example, a machine stopped, the belt broke, maintenance was overdue, the schedule was not followed, and the tracking system was manual and unreliable, posing various risks. This final issue is identified as the root cause, enabling corrective action.

This approach works when problems are linear and when the data needed to answer each question is accessible in the room with evidence of historical data. It fails when causes are distributed across systems, when the relevant signals live in different databases and formats, when the people with institutional knowledge have moved on, or when the volume of contributing factors exceeds what a working group can hold in mind at once. Under those conditions, the Five Whys does not find the root cause. It finds the most plausible story a team can construct from the evidence it has available.

This distinction has serious consequences. Organisations using traditional RCA in complex, data-rich settings often address symptoms rather than causes. Problems recur, costs rise, and confidence in investigations declines. (Root Cause Analysis Solver Engine, 2024) The method is not flawed; the environment has outgrown it.

### The Emergence of Agentic AI as a Methodological Successor

Agentic AI refers to systems in which AI models are given the capacity to pursue goals across multiple steps autonomously: using tools, accessing data sources, forming hypotheses, testing them against evidence, and revising their reasoning in real time. (Sun et al., 2025) Unlike a conventional analytics dashboard or a static model, an agentic system can be tasked with finding the root cause of an event. It will independently identify which data to retrieve, which correlations to test, and which hypotheses to present for human review. (Rootly AI Auto-Detects Incident Root Causes in Seconds, 2025)

This differs from conventional automation, which follows predefined workflows assembled through coded algorithms. Agentic AI reasons under uncertainty, identifies knowledge and data gaps, navigates ambiguity, and escalates complexity as needed. It produces outputs that human teams can interpret, audit, and act on to improve outcomes for their organisations.

The evolution from Five Whys to agentic AI is therefore not a replacement for methodology. It is a completion of it. The underlying logic remains unchanged: trace the cause. What changes is the instrument. From a guided conversation conducted by a working group to a reasoning system working at machine speed across machine scale, never losing sight of the human decision that must come at the end of every investigation.

# Data Management as the Foundation of RCA

### Why Clean Data Is Not a Supporting Condition but the Core Discipline

Organisations often misframe data as mere input and analysis as value. This is mistaken with serious consequences. Data quality is not just a prerequisite for good RCA; it is the foundational discipline that enables trustworthy analysis and deliverable outcomes.

An agentic AI system investigating outages, quality failures, or compliance breaches draws on all accessible data: logs, maintenance records, transactions, sensor readings, personnel actions, environmental inputs, and more. If data is inconsistent, incomplete, duplicated, or unvalidated, the agent produces not a weaker but a confidently wrong result. Such errors, Backed by AI authority, they pose greater operational risks than having no answer at all.

This is why organisations that make serious investments in data governance, data lineage, and master data management are not just enhancing their analytics capability. They are constructing the infrastructure that every trustworthy and insightful decision-making relies on. Every hour spent cleaning historical records, establishing consistent schemas, and linking disparate systems is an investment in the reliability of every RCA that follows, as redundant records and inconsistent data can skew key performance indicators and lead to misleading analytical outcomes (2026).

| PATTERN PRINCIPLE | Patterns do not emerge from data that has not been prepared. Noise generates noise. Clean data, accumulated over time and governed with rigour, is the only substratum from which genuine patterns can surface and be trusted as a foundation for action. |
|---|---|

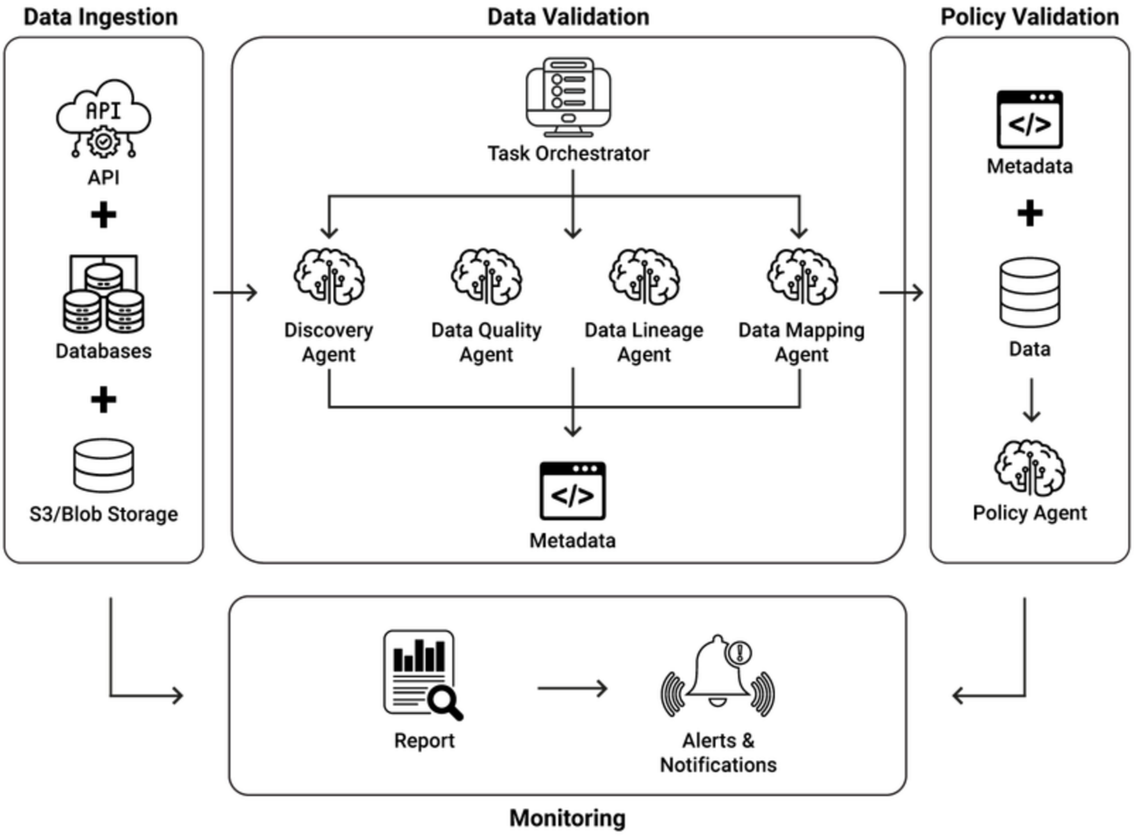


**How Patterns Emerge and What They Reveal**

Agentic RCA's greatest value lies not in single answers but in detecting patterns across time, functions, and data types that no team could observe unaided. For example, an AI agent analysing two years of maintenance records, production schedules, supplier deliveries, and shift logs can identify failure modes linked to specific, otherwise unremarkable, condition combinations.

These patterns elude traditional analysis, not because of investigator oversight, but because relevant data is scattered across sources, formats, and timeframes beyond human synthesis. Agentic AI bridges this gap by simultaneously viewing the entire governed data landscape, applying statistical and semantic reasoning, and surfacing key patterns. Importantly, it explains its findings, data sources, considerations, exclusions, and confidence levels.

After identifying a pattern and hypothesising a root cause, the system presents its insights to the human team. Decisions on action, remedies, timing, and risk remain fully with those accountable for outcomes

# Human Authority at the Centre of Agentic RCA

**Autonomy Belongs to the Organisation, Not the Algorithm**

A key feature of a well-designed agentic RCA framework is what it does not do: it does not make decisions, initiate remediation without approval, communicate findings externally, or decide which problems to investigate. These responsibilities remain with human professionals, guided by insights from data analytics and science.

This is not a constraint imposed on the technology. It is an architectural choice that reflects how trustworthy institutions operate. Executives, engineers, compliance officers, and operations leaders need to be able to account for every significant decision their organisations make. An AI system that prepares human decision-makers to make better decisions multiplies that effect. The data scientist and the data analyst occupy a critical position in this architecture. They serve as the bridge between the AI agent's raw reasoning capability and the organisation's governance requirements. They define what data the agent can access. They interpret outputs in light of the operational and regulatory context that the agent cannot fully carry. They validate that a detected pattern reflects a genuine correlation rather than a statistical artefact. And they translate the finding into a recommendation that a CXO or a board can act on with full confidence and clear accountability.

| THE HUMAN LOOP | The goal of agentic RCA is not to remove humans from the process. It is to return real decision-making power to humans who are equipped with better information, sharper patterns, and clearer evidence than any traditional method could produce on its own. |
|---|---|

**Comparison: Traditional RCA vs Agentic AI RCA**

| Dimension | Traditional RCA | Agentic AI RCA |
|---|---|---|
| Initiation | Expert recalls past incidents manually | Agent scans all historical data instantly |
| Data scope | Whatever is recalled or retrieved manually | All connected, cleaned, governed data sources |
| Speed | Days or weeks per investigation | Minutes to hours at enterprise scale |
| Pattern depth | One causal chain per session | Multi-layer cross-domain correlations in parallel |
| Human role | Primary analyst, investigator and decision maker | Decision authority, validator and strategic owner |
| Learning | Knowledge fades or leaves with people | Patterns accumulate and sharpen continuously |
| Outcome | Symptom addressed; recurrence likely | Root fixed; system learns to prevent recurrence |

# From Root Cause to Durable Fix

RCA's purpose has never been the investigation alone but the fix. Traditional methods often stop at identifying the root cause, logging corrective action, and closing the incident. Without mechanisms to monitor fix effectiveness, ongoing data capture, and agents detecting pattern recurrence, fixes remain provisional, and recurrence risks persist.

Agentic AI fundamentally changes this dynamic. The data infrastructure supporting the investigation continues to collect information post-fix, enabling the system to monitor if causal conditions have changed. If patterns reappear, the agent detects them earlier than

manual methods. The organisation learns continuously, with each investigation enhancing future speed and reliability.

This is the deliverable of a robust agentic AI RCA: not just faster answers but lasting improvements; not just identified causes but eliminated conditions; not just closed incidents but systems that are measurably more resilient than before the investigation.

| CLOSING THOUGHT | Organisations that treat data management as a strategic discipline and agentic AI as its analytical partner do not simply investigate problems faster. They build the institutional capacity to stop creating them. |
|---|---|

# References


- Wang, T. & Qi, G. (2024). A comprehensive survey on root cause analysis in (micro) services: Methodologies, challenges, and trends. arXiv preprint arXiv:2408.00803.

- (2024). Root Cause Analysis Solver Engine. Wikipedia. https://en.wikipedia.org/wiki/Root_Cause_Analysis_Solver_Engine

- Sun, J., Li, G., Zhou, P., Ma, Y., Xu, J. & Li, Y. (2025). AgenticData: An Agentic Data Analytics System for Heterogeneous Data. arXiv preprint arXiv:2508.05002. https://doi.org/10.48550/arXiv.2508.05002

- (2025). Rootly AI Auto-Detects Incident Root Causes in Seconds. Rootly. https://rootly.com/sre/rootly-ai-auto-detects-incident-root-causes-in-seconds