

WHITEPAPER

The Foundation Problem

Why 95% of Service AI Investments Underperform —
and What the 5% Do Differently

Executive Summary

Service can be the most expensive function in your business or the most profitable. The gap between the two is widening, and AI is the lever. The organizations getting it right are not chasing the \$50,000 productivity wins from summarization tools and contact center copilots. They are going after the multi-million-dollar outcomes. The ones that show up in margin, in renewal economics, in win-rate against competitors whose service model still depends on the senior technician answering the phone.

Most enterprises are not failing at AI. They are aiming too low. Gartner predicts 40% of agentic AI projects will be cancelled by 2027. A PwC study found that while 79% of companies are adopting AI agents, MIT reported fewer than 5% of AI projects achieve meaningful financial results. For C-suite leaders who have already signed the contracts and allocated the budgets, these statistics are not abstract. They are a question waiting to be asked in the next board meeting.

The reason most service AI underperforms is structural, not technical. It is not the AI model. It is not the effort your team invested in preparing data before deployment. It is that most AI architectures are designed to retrieve from clean data, not to operate on the messy, fragmented service data every enterprise actually has. AI built to retrieve from that data, rather than to structure and continuously improve it, cannot know your products, cannot learn from your resolutions, and cannot predict failures before they reach customers.


The Compounding Path to Autonomous Service

The organizations achieving the \$20M outcome are not executing transformations on day one. They are executing sequences. Each phase delivers standalone financial return. Each phase also creates the intelligence foundation that makes the next phase more accurate, more valuable, and ultimately makes autonomous service possible.

Autonomous service is where service stops being a cost center and becomes a commercial advantage. Outcome-based contracts. Uptime guarantees. Renewal economics that competitors cannot match. None of it is possible without making the right architectural decisions first.

“There’s a \$50k ROI of AI, an easy-button, and then there’s the \$5–\$20M ROI of actually solving service resolution at scale. Companies aren’t failing at ROI so much as they’re aiming too low.”

Niken Patel, Founder & CEO, Neuron7 — Emerj AI in Business Podcast



The AI Investment Reality Check

The statistics on AI adoption look impressive on the surface. Enterprises are investing. Pilots are running. Every vendor has launched an agentic AI product in the last twelve months. And yet most of those investments are producing a few minutes shaved off a call, a knowledge article surfaced faster, a contact center supervisor with a new dashboard. Productivity wins, not P&L wins.

Gartner predicts that 40% of agentic AI projects will be cancelled by 2027, primarily because organizations are deploying AI on foundations that cannot support what it promises. Gartner's same analysis paints a different picture on the vendor side: thousands of vendors are claiming agentic AI capabilities, but only about 130 are real. This is what Gartner has coined "agent washing" — traditional chatbots or RPA tools being labeled agents without the underlying capability, creating deeper confusion for buyers.

For service organizations, the consequences are visible. AI tools that surface wrong answers. Technicians who stop trusting the system and go back to calling the expert. Service leaders who cannot explain to the CFO why last year's AI investment has not moved the cost or customer satisfaction numbers.

40% of agentic AI projects will be cancelled by 2027

5% of AI pilots generated meaningful revenue acceleration

Sources: Gartner (2025), MIT NANDA (2025)

The question worth asking is not why AI is failing. It is what the 5% of companies achieving real returns are doing that the 95% are not.

The answer, consistently, is the same: they chose an AI architecture designed to operate on imperfect service data and to improve that data as it works, rather than one that requires a clean foundation before deployment can begin.



The Data Readiness Problem

The old adage of data being the new oil is not true in the AI era. Decision-ready data is. Yet most enterprises are deploying AI as a new layer atop existing CRM data, knowledge bases, and case histories. The assumption is that the underlying data is good enough for the AI to learn from, or even that AI will improve data quality over time. In almost every enterprise service organization, both assumptions are wrong.

Service data has three structural problems that most AI deployments miss.

The first is fragmentation

Service knowledge is scattered across CRM platforms, ticketing systems, device logs, engineering manuals, and the institutional memory of technicians who have spent years fixing the same issues. No single system holds a complete picture of how a product fails and how that failure gets resolved. AI that accesses only one source has a partial picture.

The second is inconsistency

The same failure described by ten different technicians produces ten different case notes. One engineer writes “replaced card reader unit.” Another documents “CR jam, full replacement required, error 265.” A third writes “fixed.” AI trained across that range of descriptions learns something incoherent, or learns just well enough to sound confident when it is wrong.

The third is context

The most valuable service knowledge is the least likely to be documented. Senior technicians close cases without recording the diagnostic reasoning, the specific part variant that mattered, or the sequence that worked. That knowledge exists in the organization. It does not exist in a form most AI can learn from.

The data readiness gap today shows wrong answers, broken trust, and AI tools that get abandoned. Tomorrow, when agents and not humans are the dominant consumers of service data, every undocumented diagnostic step becomes invisible to the systems making the decisions. Agents don't know how to phone a friend. They don't know what isn't written down.

This is why most service AI underperforms. Not because the model is wrong. Because the foundation underneath it was never built for the job. The 5% getting compounding returns from AI in service are not the ones who cleaned their data faster. They are the ones who chose an AI architecture that builds the foundation as it works.

Sources: OpenAI (2025); Elementum AI (2026); Vectara Hallucination Leaderboard (2025); SQ Magazine / Unthread (2026)

“60–70% of service data is either incomplete or inconsistent. ‘Application working as designed,’ ‘device did not function.’ That’s what most AI systems are being asked to learn from.”

Niken Patel, Founder & CEO, Neuron7.ai — Emerj AI in Business Podcast

What the Right Foundation Requires — and How to Build It

Deterministic, not probabilistic, service decision intelligence

Most AI architectures treat the foundation as a prerequisite: clean the data, then deploy the AI. In complex service environments, that sequence fails. The cleanup project stalls. The AI initiative gets pushed to the next fiscal year. The organization spends 12 to 18 months on preparation while service data continues accumulating inconsistency faster than any cleanup project can resolve it.

The organizations achieving compounding returns follow a different sequence. They select an AI architecture that does not require a clean data foundation to deliver accuracy, because the architecture itself is what builds the foundation. Three capabilities make this possible.

For high-stakes service environments such as medical devices, financial infrastructure, and industrial equipment, a suggestion is not an answer. A recommendation that is correct 70% of the time erodes trust faster than no AI at all.

Even on the narrow task of summarizing a single document, leading LLMs hallucinate 1–3% of the time. On the open-ended, multi-step reasoning that complex service actually requires, production hallucination rates climb to 15–33%. Deterministic architectures grounded in validated resolution patterns sit on the other side of that gap: by design, they cannot return an answer that is not traceable to source data. Deterministic architectures can operate accurately on incomplete data because they reason from validated patterns, not from statistical similarity to whatever happens to be in the index.

Continuous learning from every service interaction

Static knowledge bases degrade over time. One-time cleanup projects age the moment they end. As products change, new failure modes emerge, and technicians develop better approaches, that learning has to feed back automatically into the intelligence layer. Organizations achieving compounding AI returns do so because their system gets more accurate with every resolution, not despite turnover and product change, but because of it.

Active data quality management as a runtime capability

This is the capability most often missing from AI deployments, and the one buyers most often overlook when evaluating vendors. The foundation is not something the customer builds before the AI arrives. The right AI architecture scores service documentation as it is created, identifies gaps before they compound, structures unstructured case notes into reusable resolution patterns, and ensures only reliable, high-quality data influences the service decision intelligence layer. This is a continuous, runtime mechanism, not a one-time cleanup project the customer has to complete as a prerequisite.

The buyers succeeding with agentic AI in service are not the ones who invested more in data cleanup before deployment. They are the ones who chose an AI architecture that did not require the data to be clean first. That is the architectural difference separating the 5% from the 95%, and it is the single most important question to ask any agentic AI vendor under evaluation.

Niken Patel, Founder & CEO, Neuron7.ai — Emerj AI in Business Podcast



The Compounding Journey: Four Phases, Exponential Returns

The organizations achieving the strongest ROI from service AI are not executing full-scale transformations on day one. They are executing sequences. Each phase delivers standalone financial return. Each phase also creates the intelligence foundation that makes the next phase more accurate and more valuable.

Phase	Focus	What Changes for the Business	Financial Returns
Phase 1 Reactive	Teach AI your operations language. Build the Service Decision Graph. Generate high-confidence guidance on highest-frequency issues.	Escalation rates fall. Mean time to resolution drops. First-time fix rates improve. New technician ramp time shrinks.	Direct cost reduction and margin expansion. Measurable ROI within 90 days.
Phase 2 Proactive	Detect anomalies in connected and non-connected fleets before failures become customer-visible.	Emergency dispatches replaced by planned visits. Parts arrive with the technician. Repeat visits decline. Customer uptime improves.	Planned visits cost 3–5x less than reactive dispatch. Every prevented failure captures the full cost differential.

Phase	Focus	What Changes for the Business	Financial Returns
Phase 3 Predictive	Predict next likely failure per asset with confidence score and time horizon. Piggyback preventive tasks onto active service calls.	Service model shifts from cost center to competitive advantage. SLAs consistently met. Customer churn risk falls.	Eliminated truck rolls, expedited parts, overtime labor, and SLA credits compound at fleet scale into material P&L impact.
Phase 4 Autonomous	Resolve cases end-to-end without human intervention. Route high-stakes cases to humans.	Service moves from cost center to commercial advantage. Outcome-based contracts are feasible. Uptime contracts become deliverable.	Win-rate improves. Renewal economics shift from defensive to growth-oriented.

PHASE ONE:
Resolve with intelligence, not guesswork

The first phase deploys AI service decision intelligence on the service organization’s highest-frequency issues: the cases consuming the most technician time, generating the most escalations, and producing the most repeat visits. The goal is not transformation. It is to achieve 90% or greater resolution accuracy on the problems costing the most, quickly enough to build organizational trust in the system.

The financial return in Phase One is direct: reduced escalation rates, lower mean time to resolution, and measurable improvement in first-time fix rates. Organizations typically achieve measurable ROI within 90 days. Equally important, Phase One creates the Service Decision Graph, a living intelligence layer that maps the relationships between products, failure modes, and resolutions across every case the system processes. That graph is built from the customer’s actual service data as the AI operates, not from a separate cleanup effort. It is what Phases Two, Three, and Four depend on.


PHASE TWO:
Stop reacting.
Start anticipating.

The second phase activates proactive monitoring. For organizations with connected assets, this means analyzing device logs and telemetry to detect anomalies before they become customer-visible failures.

The non-connected device challenge: why most enterprises stall at predictive maintenance

For organizations with non-connected fleets, which describes most ATM networks, medical device inventories, and industrial equipment deployments, proactive monitoring means using service history, device metadata, and fleet-wide patterns to flag elevated risk before a service request is filed. The absence of real-time sensor data is not a barrier. It is a data architecture problem — and one the right foundation is built to solve.

A planned preventive visit costs between three and five times less than the same repair done reactively. Every failure detected before the customer calls eliminates that cost differential. At fleet scale, across hundreds of thousands of assets, those eliminated costs are material to the P&L.





PHASE THREE: Predict what will break next

The third phase uses fleet-level pattern analysis and validated resolution data to predict the next likely failure for specific assets, with a confidence score, a time horizon, and a resolution pathway, before any symptom is visible. Those predictions are piggybacked onto scheduled service visits: a technician already dispatched for one issue is given a second task that prevents a projected failure within the same window. One truck roll, two resolved outcomes.

Each prediction arrives with a quantified trade-off: the cost of the preventive action against the projected cost of the failure it prevents, calculated against the customer's own service economics. Leadership is not asked to trust a recommendation. They are shown the numbers behind it, in their own currency, before a single dispatch decision is made.

The financial return in Phase Three is the largest and requires measurement infrastructure to make it visible. Organizations that track predicted failures that do not happen can quantify avoided dispatches, downtime, and SLA penalties. These are real returns. The framework to capture them must be built deliberately.

PHASE FOUR:
**Service that runs
itself on what is
known. Escalates
what is not.**

The fourth phase is where the foundation built in Phases One through Three pays its largest dividend, and where it becomes a moat for service organizations. With a Service Decision Graph that knows your products, a continuous learning loop that gets sharper with every fix, and a predictive layer that sees failures before they happen, the next move is to let agents take action on the cases the system already resolves with confidence, and route high-stakes cases to humans.

This is the difference between an agent that suggests and an agent that resolves. A standard contact-center case is opened, classified, diagnosed against the Service Decision Graph, dispatched against the right part and the right technician, and closed. A predictive signal triggers a scheduled visit, the parts ship, the technician is briefed, and the asset is back online before the customer notices the symptom.

The financial outcome in Phase Four is not cost reduction. It is the inverse: service moves from cost center to commercial advantage. Outcome-based service contracts become defensible. Uptime guarantees become deliverable. Renewal economics shift, and so does win-rate against competitors whose service model still depends on expert technicians answering the phone.

Phase Four is not where the journey starts. It is what the first three phases were building toward.

Getting service AI from zero to 65% accuracy is relatively straightforward. Getting from 65% to 95% is where the economics change. For an MRI machine, that gap is the difference between meeting uptime requirements and a missed contract renewal. For a fleet of ATMs, it is the difference between a cost center and a competitive advantage. The gap is millions of dollars in ROI — and it is only reachable on a foundation built for it.

Niken Patel, Founder & CEO, Neuron7.ai — Emerj AI in Business Podcast

Proof at Scale: NCR Atleos

NCR Atleos, a \$4.1 billion financial technology company managing 600,000 ATMs across 160 countries, did not begin its AI journey by cleaning decades of accumulated service data. They began with the same data foundation every enterprise service organization starts with: fragmented across systems, inconsistent across technicians, and incomplete in exactly the places where senior experts had closed thousands of cases without documenting the reasoning behind the fix.

The Neuron7 architecture did not require them to resolve that foundation problem first. It built the foundation as it worked.

4X ROI

in 2.5 months

23%

year-over-year improvement
in service revisits and retrips

5.8M hours

of additional ATM availability in 2025

400,000 minutes

of technician time saved

Results reported publicly, March 2026. NCR Atleos NYSE-listed company.

“The Neuron7 product really gets to exactly what we needed — a product that levels the playing field for newly onboarded technicians to be just as successful as a 10-year veteran.”

Bill Girzone, SVP Global Field Services, NCR Atleos

NCR Atleos now operates its predictive maintenance program across the United States and is expanding globally. The deployment earned the Best AI Use Case and Implementation Award at Field Service Palm Springs 2025 and the ATMIA Peter Kulik Innovation Award for AI-Assisted Service Management in February 2026. These are independently validated results from a publicly traded company, delivered on service data the organization never had to pause operations to clean up before deployment.

The pattern holds across industries and fleet types. Ciena, a \$3.6 billion telecom company, achieved 44% faster case resolution and a 14-point improvement in customer satisfaction after deployment. TransLogic, a healthcare equipment provider where uptime is directly tied to patient care, reduced resolution time from three hours to three seconds with 96% accuracy and saved more than 960 warranty hours. In every case, the customer’s starting data was imperfect. In every case, the AI architecture was built to deliver accuracy on that data and to improve it as it operated.

The Question Every Executive Should Be Asking

The failure rate for service AI is not a verdict on the technology. It is a verdict on the architectural decisions most organizations are making.

Organizations deploying AI that requires a clean data foundation, then racing to clean their service data fast enough to make the AI work, are not making a technology mistake. They are making an architecture mistake. The AI is doing exactly what it was designed to do. The problem is that it was designed to retrieve from data it cannot fix, on a foundation the organization was expected to build before deployment ever began.

The difference between the 5% of companies achieving real AI returns and the 95% that are not is not budget. It is not technical sophistication. It is the decision to choose an AI architecture that can deliver accuracy on imperfect service data and improve that data continuously as a function of how it operates.

The right question is not only whether your service data foundation is ready for AI. It is whether you are aiming at the outcome that matters, and whether the AI you are evaluating can actually get you there.

The 5% of enterprises getting compounding returns from AI in service did not get there by cleaning data faster, buying more sophisticated models, or running more pilots. They got there by aiming at the outcome that justifies the investment, and choosing an architecture built to deliver it.

About Neuron7

Neuron7 is the AI agent for service that understands every product manual, support ticket, and resolution, gets smarter with every fix, and catches failures before your customers do. Purpose-built for complex service environments in medical devices, industrial equipment, high-tech, and telecom, Neuron7 delivers 90%+ resolution accuracy and measurable ROI within 90 days, on the service data customers already have, not the clean foundation they wish they had. The AI Readiness Diagnostic is available to service organizations that want to assess their current data foundation and identify where Phase One resolution intelligence will deliver ROI from their starting point. Trusted by Fortune 1000 enterprises including NCR Atleos, Ciena, TK Elevator, and TransLogic. Strategic partner of Salesforce (First Agentforce Partner for Service), ServiceNow (Strategic Investment), and Microsoft (Partner of the Year Finalist).