



The problem isn't AI. *It's operational readiness.*

An operational blueprint for AI-powered contact centers that drive tangible business outcomes

Capgemini 

Table of contents

Executive summary	03
The Readiness Amplification model	04
Section 1 – The AI rush nobody is questioning	05
Section 2 – Why most AI programs underdeliver	06
Section 3 – What operational certainty actually means	08
Section 4 – The operational certainty transformation sequence	10
Section 5 – What AI should effectively be doing in a contact center	12
Section 6 – The future operating model	13
Closing: The operational certainty imperative	15

Executive summary

The contact center industry is experiencing an AI investment cycle unlike anything before. Copilots, agent assist tools, generative AI, agentic workflows, autonomous resolution – the capabilities are real, the momentum is real, and the pressure on leaders to act is real, too.

And yet, across the industry, a quiet frustration is growing: programs are underdelivering, vendors are being replaced, boards are asking tougher questions. The technology isn't failing. Something else is.

The argument this paper makes is simple, and it runs counter to most of what the industry is currently being told:

*AI doesn't determine the outcome.
The operation does.*

When AI is introduced into a well-governed, consistently performing operation – one where demand is understood, ownership is clear, workflows are stable, and knowledge is managed as an asset – it becomes exactly what the investment promised: a lever for scale, a multiplier of capacity, and a path to genuinely better customer outcomes.

When AI is introduced into a fragmented, inconsistently performing operation – one running on tribal knowledge, unclear accountability, unmeasured failure demand, and workflows that vary by agent and by shift – it becomes something else. It scales the inconsistency. It automates the confusion, exposing gaps faster and to more customers than ever before.

This is not an argument against AI. Autonomous operations are becoming a reality, and augmentation is already delivering measurable results. The opportunity to improve efficiency, deepen personalization, and enhance customer outcomes is real and significant. But realizing capturing that value requires something the industry is systematically underinvesting in: operational certainty.

Operational certainty is the condition in which ownership, workflows, and resolution paths are visible, measurable, and governable. It is not a precondition for beginning the AI journey. It is the work itself; the foundation that



determines whether AI becomes a source of competitive advantage or an accelerant of operational debt.

This paper offers a practical framework for building that foundation. It examines why most AI programs underperform, introduces a transformation sequence that leading operations are using to move from fragmented to governed, from reactive to proactive, from handling contacts to resolving them, and describes what a genuinely AI-enabled contact center looks like when the work has been done correctly.

The winners in AI-powered customer operations will not be the companies that deploy AI first. The winners will be the companies that operationalize AI safely, consistently, and measurably on a foundation built to support it.

That foundation is what this paper is about.

The Readiness Amplification model

Before examining why AI programs underperform, one principle deserves to be stated plainly. It will recur throughout every section that follows. Every framework, every use case, and every stage of the transformation sequence that this paper describes is an expression of the same underlying logic.

*AI doesn't determine the outcome.
The operation does.*

Operational condition	AI outcome
Stable and governed	Scalable augmentation
Fragmented and inconsistent	Scaled inconsistency
Knowledge-driven	Reliable automation
Tribal and reactive	Automation instability
High operational certainty	Measurable, sustainable value



Section 1

The AI rush nobody is questioning

The market momentum behind copilots, agent assist, generative AI, agentic workflows, and autonomous resolution is real and deeply warranted. These capabilities are truly transformative when the conditions are adequate, and the investment case for moving now is not difficult to make. Investment is accelerating across every segment and every geography. The pressure on contact center leaders to act is substantial, and the consequences of being left behind are tangible.

But beneath the investment, most contact centers are still operating on fragmented workflows, inconsistent resolution paths, weak knowledge management, and broken handoffs. The question isn't whether to adopt AI; it's whether the operation underneath it can support it.

Four myths driving failed AI programs

These four assumptions appear repeatedly in organizations that are moving fast on AI deployment and obtaining disappointing results.

Myth 1: "AI reduces cost automatically." AI redirects work. It doesn't fix the work. Cost reduction follows operational improvement – it doesn't precede it.

Myth 2: "Automation maturity equals operational maturity." You can automate a broken process. You'll just break it faster. Deployment speed is not a proxy for operational readiness.

Myth 3: "Containment equals resolution." Deflecting a contact is not the same as resolving a customer's need. The demand returns – typically more frustrated and through a more expensive channel.

Myth 4: "Better AI compensates for broken operations." No model, however sophisticated, resolves what the operation itself does not control. Capability without operational certainty is capability without traction.



Section 2

Why most AI programs underdeliver

The contact center industry has never moved faster to adapt a new technology. While the urgency and potential are real, in boardrooms and operations centers alike, a quiet frustration is growing: the results aren't matching the investment.

This isn't an AI problem. It's an operational one.

Most AI programs do not fail because the AI is weak; they fail because the operating model underneath the AI is unstable, inconsistent, or poorly governed. When AI programs underperform, the instinct is to look at the technology – the model, the vendor, the implementation partner. But in most cases, the technology is functioning exactly as designed. It's

doing what it was built to do. The problem is what it's being asked to land on.

AI amplifies the reality it finds. In operations with high operational certainty – where workflows, ownership, and resolution paths are visible and governed – AI can create faster resolution, greater consistency, and meaningful capacity relief. In a fragmented operation, however, it produces something else entirely: scaled inconsistency, automated confusion, and customers who reach dead ends faster than before.

Five operational failure patterns occur consistently across underperforming AI programs.

#	Pattern	What it looks like
1	Failure demand accumulation	Repeat contacts driven by unresolved upstream failures
2	Resolution variability	The same issue handled differently depending on agent, shift, or channel
3	Fragmented ownership	No single point of accountability for customer resolution end to end
4	Tribal knowledge dependency	Operational knowledge living in people rather than governed systems
5	Metric distortion	Optimizing for speed and deflection rather than resolution consistency

Failure demand accumulation

Failure demand is the volume of contacts that exists not because customers want something new, but because something wasn't resolved the first time. "Did you get my email?" "I was told someone would call me back." "I've already spoken to three people about this." These contacts look like demand, but they aren't. They're the cost of an operation that didn't close the loop. In many operations, entire queues absorb this work – repeat inquiries and escalations generated not by new customer need but by unresolved process breakdowns elsewhere in the journey. Often 20 to 40 percent of total volume. Largely unmeasured.

AI cannot resolve what the operation itself does not control.

Resolution variability

In many contact centers, the same issue is handled differently depending on who picks up the phone. That

variation is invisible to leadership because aggregate metrics – handle time, CSAT, first contact resolution – smooth it over. AI doesn't smooth it over. AI reflects it back into the customer experience at scale. Escalations increase and trust erodes, making operation become more difficult to govern.

Consistency is not a byproduct of AI deployment. It's a prerequisite for it.

Fragmented ownership

In complex operations – particularly those that have grown through acquisition, reorganization, or years of tactical decision-making – ownership of customer issues is often unclear. A billing dispute touches finance. A delivery failure touches logistics. A service complaint touches three teams and two systems. When a customer contacts the organization, someone handles the interaction, but no one owns the resolution. AI makes this visible almost immediately: routing improves, triage becomes faster, and then the contact hits the same wall it always did – a back-office dependency

with no SLA, a queue with no owner, a process that no one has mapped end to end.

The handoff breaks. The customers repeat themselves. The contact generates another contact.

In many environments, the front office becomes a coordination layer for processes it neither owns nor controls – manually stitching together updates across logistics, finance, fulfillment, or service teams that operate on entirely different rhythms.

AI cannot restructure accountability. What it can do is expose every gap in the operating model – loudly, at scale, in front of customers.

Tribal knowledge dependency

In most contact centers, a significant portion of operational knowledge lives in people – in the heads of experienced agents, the institutional memory of team leaders, the informal networks that new hires spend months navigating. Newer agents often learn not from governed systems but from knowing which experienced person to ask. AI requires knowledge to be explicit, structured, and governed, as it cannot extract understanding from experience. When automation is introduced into an environment where knowledge is tribal, it hits a ceiling almost immediately – and that ceiling is lower than anyone expected.

Knowledge is not a support function. In an AI-enabled operation, it is the operating system.

Metric distortion

Perhaps the most consequential operational failure pattern is one that rarely appears in post-implementation reviews: the metrics driving AI investment are the wrong metrics.

For decades, contact center performance has been measured by cost per contact, handle time, and deflection rate. These are speed metrics. They reward volume throughput, not resolution quality. They treat every contact as equivalent, whether it's a customer getting a fast, accurate answer or a customer being routed through four transfers before giving up.

When AI is deployed to optimize these metrics, it does exactly that. Contacts get deflected, handle times compress, and cost per contact improves. And underneath those numbers, resolution consistency declines, failure demand grows, and customer effort increases.



The shift the industry needs – and that leading operations are already making – is from cost per contact to cost per resolution. From deflection rate to avoidable demand. From handle time to resolution consistency. These are the metrics that reflect what customers actually experience – and the ones that AI, properly deployed, can move.

AI deployed against a containment metric optimizes for the channel.

AI deployed against a resolution metric optimizes for the customer.

These five operational failure patterns share a common thread:

AI performs best in environments with high operational certainty – where ownership, workflows, and resolution paths are visible, measurable, and governable. Without that foundation, AI doesn't fix the operation. It accelerates whatever the operation already is.

Section 3

What operational certainty actually means

The phrase “operational readiness” appears often in AI implementation guidance. It typically arrives as a checklist: clean data, integrated systems, defined use cases, executive sponsorship. These things matter. But they describe the surface conditions of readiness, not the underlying state that determines whether AI delivers sustained value.

The more precise concept is operational certainty: the

condition in which an organization knows what its customers need, controls how those needs are resolved, and can measure whether resolution occurred. It exists on a spectrum – most contact centers sit somewhere in the middle, capable enough to function, uncertain enough to limit what AI can safely do. It breaks into three dimensions, each of which must reach a functional threshold before AI deployment in that area can be expected to deliver reliable value.

Demand and journey clarity	Operational system stability	Knowledge and technology readiness
Which contacts are repeatable, avoidable, or judgment-dependent?	Are workflows governed, ownership clear, failure demand visible?	Is knowledge a governed asset that AI can reliably act on?

Note: Technology readiness appears third deliberately. Operational maturity governs technology effectiveness.

Dimension 1: Demand and journey clarity


The first dimension asks a deceptively simple question: do you actually know why your customers are contacting you?

Not at the category level – most operations have that. But at the resolution level. Do you know which contact types are truly repeatable and therefore automatable? Which are avoidable – generated by failures elsewhere in the customer journey? Which require genuine human judgment? And which are hybrid – partially automatable, with a human handoff at a defined point?

Most contact centers answer these questions with assumption rather than data. Volume reports show what came in, but rarely distinguish between value demand and demand the operation created itself, rarely show why it came in, whether it was resolved, or whether it came back. Without that clarity, automation decisions are made on incomplete information, and the contacts that get automated are often not the ones that should be.

Journey clarity also means understanding where customers get stuck before they reach the contact center. Failure demand cannot be automated away. It can only be eliminated by fixing the journey that created it.

Diagnostic question: Are your automation decisions based on real resolution data, or on volume and assumption?



Dimension 2: Operational system stability

The second dimension is the one most directly connected to AI performance. Operational system stability means that workflows are defined, not improvised. That resolution paths are standardized, not agent-dependent. That ownership of customer issues is explicit, not assumed. That failure demand is visible and being actively managed – not absorbed into volume metrics where it becomes invisible. An operation that cannot distinguish between demand it generated and demand its customers generated has not yet achieved the stability that AI requires.

Diagnostic question: If you automated your most common contact type today, would customers get a consistent experience, or would the automation reflect the same variation your agents currently produce?

Dimension 3: Knowledge and technology readiness

Technology readiness matters – integration, data quality, platform architecture are real constraints and real enablers. But technology readiness without the first two dimensions is infrastructure without a purpose. Knowledge readiness is the more consequential of the two, and more frequently underestimated. In an AI-enabled contact center, knowledge is not a resource. It is the operating system. Every automated interaction, every agent assist recommendation, every guided workflow depends on knowledge that is explicit, current, validated, and accessible. When knowledge is tribal, AI hits a hard ceiling almost immediately.


Diagnostic question: Is your knowledge a governed asset that AI can reliably act on, or institutional memory that walks out the door when your best agents leave?

Operational certainty is not a starting line. It's the work.

A common misconception is that operational certainty must be fully achieved before AI programs can begin. That's not the argument this paper makes, and it's not how leading organizations approach it.

Operational certainty is built alongside AI deployment – but mature organizations expand AI only as operational stability grows. The transformation blueprint in Section 4 describes exactly how: starting with the dimensions where certainty is highest, expanding automation as stability grows, and treating each stage as both a delivery milestone and a readiness investment for the next.

AI is a multiplier. Operational certainty determines what it multiplies.



The operational certainty diagnostic below maps current state across all three dimensions and identifies where the threshold for safe AI deployment has – and has not – been reached.

Dimension	Early	Developing	Stable	Certain
Demand and journey Clarity	Volume tracked, intent assumed	Contact types categorized	Repeatable contacts identified	Avoidable demand actively managed
Operational system stability	Workflows improvised	Key paths documented	Resolution standardized	Ownership explicit, failure demand visible
Knowledge and technology readiness	Tribal, undocumented	Partially documented	Governed knowledge base	Knowledge maintained as operational asset

Section 4

The operational certainty transformation sequence

Most AI transformation programs have a sequencing problem.

They begin where they should finish. They deploy automation before workflows are stable, introduce agent assist before knowledge is governed, and scale self-service before the resolution paths underneath it have been validated. The technology works. The operation isn't ready for it. And the program stalls – not dramatically, but gradually – as escalations rise, adoption plateaus, and the

business case quietly erodes.

The Operational Certainty Transformation Sequence is a response to that pattern. It is not a phased project plan. It is a dependency chain – a logical progression in which each stage creates the operational conditions that the next stage requires. Skipping stages doesn't accelerate the program. It borrows against future stability and creates debt that compounds under the pressure of scale.

Diagnose	Stabilize	Optimize	Augment	Automate	Scale
Understand real demand	Make the operation governable	Build what automation stands on	Introduce AI where certainty is highest	Redirect repeatable contacts with confidence	Expand as operational certainty grows
Where most programs begin ↗					

Stage 1: Diagnose

Understand real demand, not assumed demand.

Every transformation begins with a diagnostic – but most diagnostics ask the wrong questions. They measure volume, handle time, and channel distribution. They produce a picture of what the operation is processing. They rarely reveal what customers are actually experiencing.

The diagnose stage asks different questions: Which contact types are effectively repeatable? Which are avoidable – generated not by customer need but by failures earlier in the journey? Where is failure demand accumulating? Which resolution paths are standardized, and which are agent-dependent? Where does ownership break down?

The output of this stage is not a recommendation. It is clarity – a data-grounded picture of demand, resolution, and operational gaps that makes every subsequent decision faster, safer, and more defensible.

Stage 2: Stabilize

Make the operation governable before making it faster.

Stabilization is the stage most AI programs skip. It feels like delay. It looks like consulting overhead. It is, in practice, the single most important investment an organization can make before deploying AI at scale.

Stabilization means addressing the operational failure patterns that the diagnostic surfaced: reducing failure demand accumulation at its source, clarifying ownership of resolution, standardizing the highest-volume resolution paths, and beginning the transition from tribal knowledge to governed knowledge. Sometimes that work is deceptively simple – clarifying ownership of shared queues, establishing internal SLAs, or eliminating the manual workarounds agents quietly built to compensate for unstable processes.

Stabilization does not require perfection. It requires a governable baseline – a known, measurable starting point from which improvement is visible and AI deployment is safe.

The instinct to skip this stage is understandable. The discipline to execute it is what separates programs that deliver sustained value from those that deliver an impressive pilot followed by a quiet plateau.

Stage 3: Optimize

Build the foundation that automation will stand on.

With the operation stabilized, the optimize stage builds the structural conditions that AI requires to perform reliably at scale. Three priorities define this stage: resolution standardization, knowledge as a governed asset, and metrics that measure resolution rather than speed.

When these three conditions are in place, the operation has achieved something specific and valuable: it has built the foundation that automation will stand on. Not perfectly. But reliably enough that what AI deploys into is knowable, and what it produces is measurable.

Stage 4: Augment

Introduce AI where the operation is ready to receive it.

The augment stage is where AI enters the operation – not broadly, and not experimentally, but deliberately, in the workflow areas where operational certainty is highest and the value of assistive AI is clearest.

Augmentation at this stage is agent-facing. The goal is not to replace human judgment but to make it faster, more consistent, and less dependent on individual experience and institutional memory: agent assist, guided workflows, interaction summarization, intelligent, triage and routing.

Each of these use cases is only as reliable as the operational foundation beneath it. This is the stage where the investment in Stages 1 through 3 becomes visible. Organizations that built the foundation find that augmentation accelerates quickly and adoption grows naturally. Organizations that skipped it find that augmentation produces inconsistent results – and reach for the technology to explain what is fundamentally an operational gap.



Stage 5: Automate

Redirect repeatable contacts with confidence – because the foundation supports it.

Automation at Stage 5 is not the same as automation attempted at Stage 1. The contact types being automated have been identified through real diagnostic data, not assumptions. The resolution paths have been standardized and validated. The knowledge assets are governed and current. The metrics to evaluate performance are in place.

This is automation deployed with operational certainty – and the difference in outcomes is significant. Customers reach resolution without agent involvement not because they've been deflected, but because the automated experience is genuinely capable of resolving their need.

Containment is not the primary measure of success. First contact resolution is.

Stage 6: Scale

Expand automation coverage as operational certainty grows.

The final stage is not an endpoint. It is an operating posture – a deliberate, governed approach to expanding automation coverage as the conditions for reliable performance are established in new areas of the operation.

Scaling means expanding horizontally – bringing additional contact types and channels into the automation framework – and deepening vertically, moving from assistive AI toward more autonomous resolution where operational certainty is highest and augmentation has performed reliably.

Organizations at this stage are not asking whether to expand AI. They are asking which parts of the operation have reached the certainty threshold that makes expansion safe, sustainable, and measurable.

The sequence as strategic doctrine

The Operational Certainty Transformation Sequence is not a methodology for slowing AI adoption. It is a methodology for making AI adoption irreversible. Programs that skip the sequence can move fast – and stall fast, when pilots don't scale, escalations rise, and the board asks why the investment isn't producing the efficiency lift that was promised. Programs that follow it move deliberately. And then they don't stop.

It multiplies what works.

Section 5

What AI should effectively be doing in a contact center

The operational certainty transformation sequence describes how to build the foundation. This section describes what becomes possible once it exists. The use cases below are organized around a simple principle: value is determined not by the sophistication of the AI but by the operational certainty of the environment it operates in.

Agent assist and guided workflows

The highest-value near-term AI deployment in most contact centers is not customer-facing – it is agent-facing. Agent assist surfaces relevant knowledge and recommends next best actions during live interactions, compressing the experience gap between tenured agents and newer ones. Guided workflows extend this into structured resolution paths that ensure high-volume contact types are handled consistently. Together they address resolution variability and tribal knowledge dependency directly. The measure of success is not adoption rate. It is resolution consistency.

Real-time knowledge retrieval

In most contact centers, agents spend a significant portion of every interaction searching. AI-powered knowledge retrieval surfaces the right information at the right moment – but only as reliably as the knowledge it retrieves. An AI that surfaces outdated guidance confidently is not an improvement over an agent who searches slowly but finds the current answer. This use case requires knowledge to be governed before it can be trusted at scale.


Interaction summarization and after-call work reduction

After-call work represents 20 to 30 percent of total interaction time in many environments and is among the most automatable. AI-powered summarization captures the substance of each interaction and produces structured documentation without agent effort – improving case continuity, data quality, and future diagnostic accuracy. Operational dependency here is relatively low, making it a strong early deployment candidate even in operations still working through the stabilize and optimize stages.

Intelligent triage and routing

Misrouting generates transfers, transfers generate repeat explanations, and repeat explanations generate escalations. AI-assisted triage identifies contact type, intent, and likely

resolution path at the point of entry – before a human agent is involved. The operational dependency is demand clarity: without a mapped understanding of what contact types exist and what resolution each requires, intelligent routing is routing toward a destination that hasn't been defined.



Case orchestration and proactive engagement

Case orchestration uses AI to manage back-office dependencies – tracking open cases, monitoring service level agreements (SLAs), and updating customers before they call back. In many operations, the underlying issue is already resolved operationally, but because visibility never reached the customer-facing layer, another inquiry is generated anyway. Proactive engagement extends this further: using AI to identify customers likely to contact and reaching out before they do. Both capabilities address fragmented ownership directly and represent the most operationally mature deployments in this section – available as certainty grows, not as starting points.

The organizations investing in operational certainty now are not moving slowly. They are moving in the only direction that leads somewhere permanent.

Section 6

The future operating model

The question is not what AI will eventually replace in the contact center. It is what a contact center becomes when AI is working the way it should. The answer is not a smaller operation. It is a more capable one – where automation absorbs what is repeatable, human judgment is elevated to handle what is complex and consequential, and the operation becomes genuinely good at resolving what customers truly need. The organizations furthest along that path share six observable characteristics.

Characteristic	What it means in practice
Resolution consistency is measurable and governed	Not just speed metrics. The operation optimizes for outcomes, not throughput.
Knowledge operates as infrastructure	Governed, maintained, continuously improved, not tribal memory.
AI governance is embedded operationally	Part of how the operation runs – not bolted on as a compliance layer.
Demand is managed proactively	Upstream conditions monitored. Contacts eliminated before they form.
Human judgment is elevated	Agents handle complexity, ambiguity, and consequence. Not search and navigation.
Automation expands deliberately	Based on operational certainty thresholds – not technology availability.

Resolution consistency is measurable and governed

The shift from speed metrics to resolution metrics – from cost per contact to cost per resolution, from deflection to avoidable demand – changes what the operation optimizes for and what AI is held accountable to deliver. Every capability is evaluated against one standard: was the customer’s need resolved?

Knowledge operates as infrastructure

Knowledge has owners, governance, and a maintenance cadence that treats accuracy as an operational obligation. It is continuously improved based on what real interactions surface. Knowledge infrastructure is what makes AI reliable over time rather than just at launch.

AI governance is embedded operationally

Governance is not a compliance layer revisited when something goes wrong. It is part of operational rhythm – AI performance reviewed with the same regularity as agent

performance, expansion conditions defined in advance, unexpected patterns surfaced and acted on. This is not bureaucratic oversight. It is operational intelligence.



Demand is managed proactively

Mature operations monitor the conditions that generate demand and act before a customer ever needs to make contact. Customers stay informed, issues are resolved proactively, and potential repeat contacts or escalations are prevented before they occur.

The shift from reactive to proactive is not a technology upgrade. It is the operational reward for having built the foundation correctly.

Human judgment is elevated

Automation handles what is repeatable. What remains – complexity, ambiguity, and critical decisions – is where human expertise matters most. Not searching for information. Not acting as middleware between disconnected systems, but handling the interactions where judgment, empathy, and expertise create the greatest value.

This is not a diminished role. It is an elevated one.

Automation expands deliberately

Automation expands not because the technology is ready but because the operation is – because the conditions for reliable performance in a new area have been established and confirmed. Organizations that internalize this discipline grow steadily and sustain performance, because every deployment was ready for what it encountered.

Resolution is the measure. Knowledge is the infrastructure. Governance is the operating rhythm. Proactive engagement is the posture. Human judgment is the differentiator. And operational certainty is what makes all of it sustainable.

That is the destination: a customer operation built to harness AI effectively – and better equipped, as a result, to do the thing it has always been asked to do:

Resolving what customers need, consistently and at scale.



Closing: *The operational certainty imperative*

The winners in AI-powered customer operations will not be the companies that deploy AI first. They will be the ones that integrate it into an operation built to use it effectively, govern it responsibly, and scale it with confidence.

That foundation is operational certainty. And like every durable operational capability, it is built through deliberate sequencing, disciplined governance, and a relentless focus on stability before scale. The technology is ready. The real question is whether the operation behind it is.

Capgemini Intelligent Customer Operations

Building the operational foundation for AI at scale.



Contact:

Tim Szymanski

Email: tim.szymanski@capgemini.com

About Capgemini

Capgemini is an AI-powered global business and technology transformation partner, delivering tangible business value. We imagine the future of organizations and make it real with AI, technology and people. With our strong heritage of nearly 60 years, we are a responsible and diverse group of over 420,000 team members in more than 50 countries. We deliver end-to-end services and solutions with our deep industry expertise and strong partner ecosystem, leveraging our capabilities across strategy, technology, design, engineering and business operations. The Group reported 2025 global revenues of €22.5 billion.

www.capgemini.com

Make
it
real.

Connect with us:



<https://www.linkedin.com/company/capgemini/>



www.capgemini.com/services/enterprise-management/intelligent-business-operations/



intelligentbusinessoperationscommunications.global@capgemini.com