The agentic service playbook for IT leaders

How to build the foundation, govern the risk, and scale AI to full autonomy



zendesk
AI masterclass

Introduction

Most IT organizations have run an AI pilot. Many have deployed a chatbot, automated a password reset, or tested a self-service workflow. And nearly all have learned the same lesson: deploying AI is the easy part. Making it work reliably, at scale, and across the business is something else entirely.

True AI-enabled IT service requires a broader shift in how service is designed, delivered, and continuously improved. Under constant pressure to do more with less, IT leaders need a proven approach that creates lasting impact without a lot of technical debt.

This playbook gives you the framework, building blocks, and self-assessment tools to get there. You'll walk away with:

- A step-by-step approach to AI adoption that helps prevent costly rework.
- A scored assessment of your organization's AI readiness across six critical dimensions.
- A way to communicate AI's business value to executives beyond feature adoption.
- A prioritized set of use cases and next steps tailored to your organization.

Let's get started.





Table of contents

04 ZEN framework: A strategic path to autonomous service

08 6 building blocks to make AI work at any phase

14 Action plan: Assessing AI readiness and building your value pipeline

17 Frequently asked questions

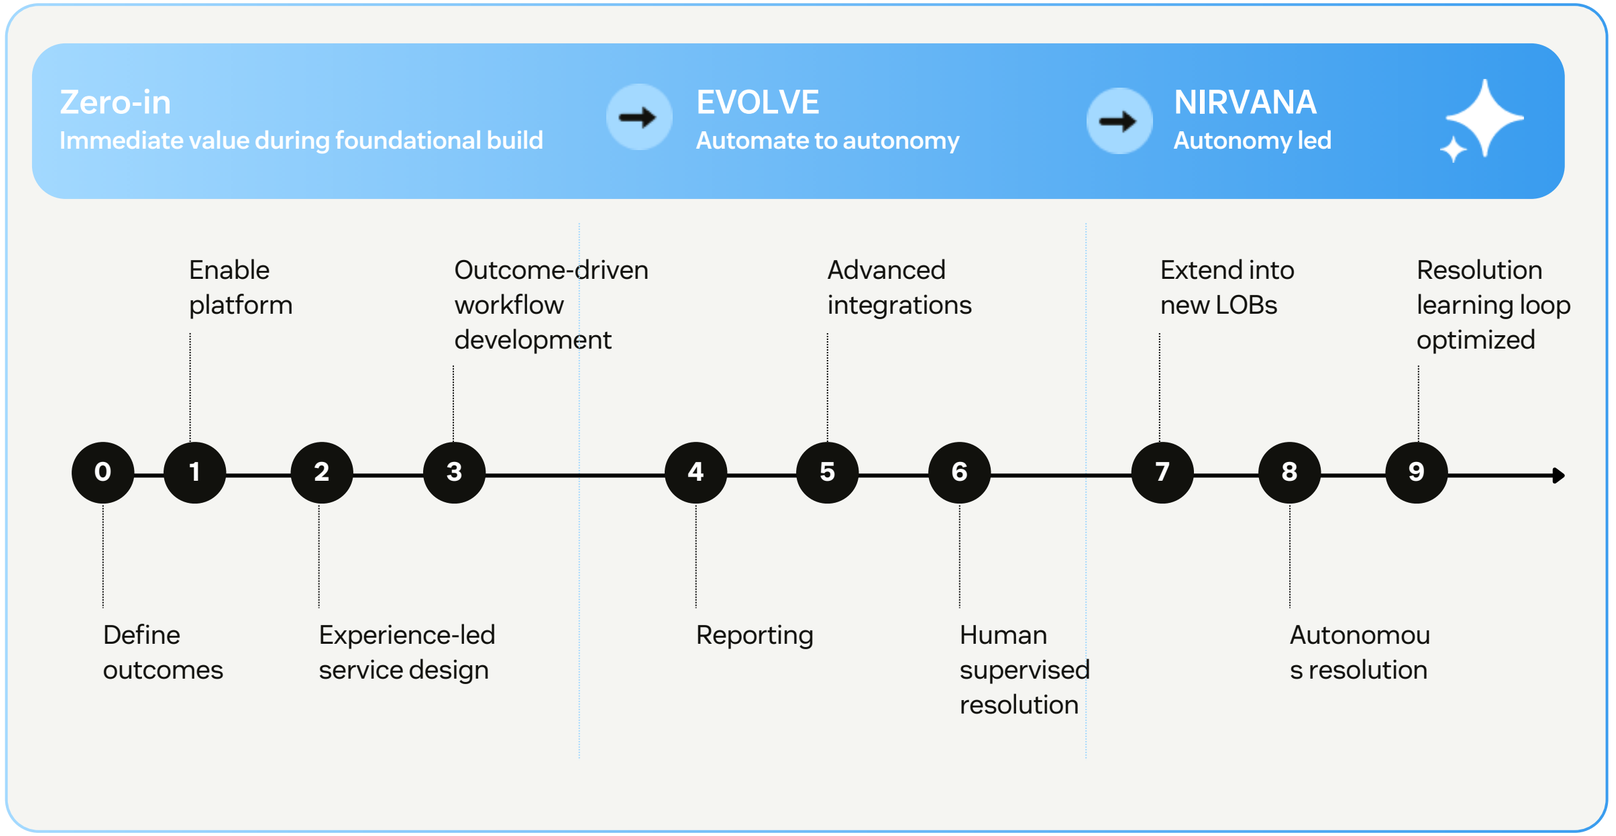
18 AI is ready. Discover whether your foundation is.

ZEN Framework: A strategic path to autonomous service

Most AI implementations fail not because the technology doesn't work, but because organizations treat AI as a series of isolated experiments.

They deploy an agent here, automate a workflow there, without grounding any of it to clear business value. The result: technical debt, stakeholder fatigue, and a growing list of AI projects that can't demonstrate ROI.

The ZEN Framework was designed to fix this. Built on industry best practices and tested across global customers in every major industry, the framework helps IT leaders approach AI as a strategic journey, not a single implementation decision. It has three progressive phases, with the end goal of delivering autonomous service resolutions with AI.



Let’s take a closer look at the three phases:

ZEN Framework: A strategic path to autonomous service CONTINUED

PHASE 1: ZERO IN

The most common mistake organizations make with deploying AI is treating the technology as the destination rather than the means.

Teams spin up agents, enable features, and automate processes, but six months later can't articulate the value they're getting from it. Zero In fixes that by forcing the outcome conversation first.
During the Zero In phase, you'll define exactly what you're trying to achieve and establish the core foundations required for AI-powered service.

This includes:

- Defining your top 3-5 business KPIs that AI should move.
- Establishing cross-functional stakeholder alignment around shared goals.
- Identifying the priority workflows and use cases connected to those KPIs.
- Building the knowledge foundation that AI will depend on.
- Connecting your AI agents and copilot to trusted knowledge sources.

Keep in mind:

Stakeholder engagement is critical here. The Zero In phase works best when cross-functional groups (like IT, operations, HR, and finance) are tied to a common, specific goal, and understand exactly what role they play.



Making Zero In possible

An employee reports a laptop overheating. Rather than emailing IT and waiting, they go to the help center. An AI agent immediately draws on the connected knowledge base to walk them through device-specific troubleshooting steps. When the issue persists, the AI recognizes that human intervention is needed and escalates seamlessly—passing the full conversation context to an IT agent so the employee never has to repeat themselves. The IT agent is guided by copilot with suggested next steps, and the issue is resolved via a remote session.
That experience—fast, context-aware, and seamless—is what Zero In makes possible. The technology isn't complex. But it requires a solid knowledge foundation, clear escalation logic, and defined outcomes to work reliably.

ZEN Framework: A strategic path to autonomous service CONTINUED

PHASE 2: EVOLVE

Once the foundation is stable, value begins to compound.

Leveraging the foundations you established in the Zero In phase, AI moves beyond knowledge retrieval and into operational execution. It can assemble context from multiple systems, apply policy, make decisions, and take actions—all within governed boundaries. This marks the shift from automation to autonomous execution.

During this phase, the most successful teams identify where AI can remove friction, not just automate the friction that already exists. The next step is to take a single, high-value service and redesign it end to end. Map the complete journey across every team, process, and system involved. Focus on the experience you want to deliver, not the constraints of today's operating model.

To do this:

- Run a value stream workshop with the teams involved.
- Select one high-value service (for example, laptop replacement or employee onboarding) rather than optimizing individual ITIL practices in isolation.
- Map the current state end to end, including who's involved, what systems are touched, where work changes hands, and how long each step takes.
- Design the ideal future-state journey before deciding where AI should automate, assist, or act autonomously.

Ask:

- Which systems can we connect to help AI gather information and take action, reducing manual effort and handoffs?
- Where can we eliminate steps—not just automate them?
- What should the end-user experience look like in the optimized state, not just the current state?
- Launch that future-state workflow, measure the results, and use it as your new best-practice blueprint for similar services.

Don't let perfection become the enemy of progress. Your first redesigned service doesn't need to solve every edge case. Build a "gold standard" operating model for one service, prove its value, then refine and scale it across additional services.

This is also where outcome-led reporting becomes visible. Because you set the destination in Zero In, you can now measure whether you're getting there.

Evolve in action

An employee needs a new laptop. They don't open a ticket or search for the replacement policy. Instead, they ask the AI agent directly. The agent checks their employee profile, looks up their current device in ITAM, notes the model and age, cross-references the equipment policy, and determines eligibility for replacement—automatically. If manual approval is required, a ticket is created and routed. Once approved, the employee receives a Slack notification that their new device is being prepared.

The entire workflow—from request to approval to fulfillment—runs without the employee or IT agent managing each step manually.

ZEN Framework: A strategic path to autonomous service CONTINUED

PHASE 3: NIRVANA

Nirvana is where AI stops being a project and becomes part of how the organization runs.

Two critical things happen at this phase:

1

Your business moves from using AI to trusting AI. This is a major organizational milestone that happens when you have enough evidence, audit trail, and consistent performance that your risk, security, and leadership teams are willing to let AI act autonomously in higher-stakes workflows.

2

You extend your proven model into new lines of business. Once the AI-led operating model works in IT, it can be replicated. The bulk of the hard work (i.e., the integrations, governance frameworks, knowledge architecture, and measurement infrastructure) is already done. Now, new use cases and new departments can be onboarded against the same foundation.

Many AI projects fail due to change management, not technology. Delivering incremental value at each phase helps your business buy into the journey. That way, by the time you reach Nirvana, leadership has already seen the results—and you've earned their trust.

Nirvana unlocked

An IT manager preparing for quarterly budget review opens their asset dashboard. They notice laptop availability is below the 5% safety benchmark. Tablet maintenance rates are above internal guidelines. Rather than building pivot tables, they query the data directly using Analyst Copilot: *"Are there any trends in tablets in repair?"*

The AI analyzes asset data, identifies that the majority of affected tablets are out of warranty, and surfaces a lifecycle insight rather than a faulty batch issue. Armed with this, the IT manager makes the decision to freeze tablet purchases this quarter and reallocate budget to rolling out new laptops.

Simultaneously, Admin Copilot is proactively surfacing a recommendation: VPN issues for Mac users have been a top ticket topic, and adding a VPN troubleshooting procedure would improve AI response quality and speed up resolution. The admin reviews an AI-generated draft of the procedure, edits it, and publishes in minutes.

The environment is built to measure outcomes and learn from results, so you can scale AI-driven service with confidence. That way, AI not only delivers service, it continuously identifies where it can do better and proactively closes those gaps. The telltale sign you've achieved Nirvana.

6 building blocks to make AI work at any phase

The ZEN Framework tells you where you're going, but you need a roadmap to get there. That's where the six building blocks of agentic AI come into play.

These six foundational elements are the infrastructure that sits underneath every phase of the framework. As your organization matures, so does the level of depth required.

Building block	Core question	Common failure mode
1 **Data & Context**	Do we have the right data, with context, for our priority use cases?	Collecting everything; understanding nothing
2 **Integrations**	Can AI see and act on the systems that matter?	Read-only integrations that deflect but don't resolve
3 **Governance**	Do we know what AI is allowed to do—and can we prove it?	No audit trail; agent sprawl
4 **Knowledge**	Is our knowledge current, owned, and connected to AI?	Treating knowledge setup as a one-time project
5 **Workflows**	Are our priority workflows defined, classified, and AI-ready?	Automating broken processes without redesigning them
6 **Measurement**	Are we measuring outcomes, not just activity?	Retrofitting reporting after the fact



6 building blocks to make AI work at any phase CONTINUED



1 Data & Context

Data is a signal. Context tells you what it means. Together, they are one of the most important building blocks. They're also where most organizations fail.

The key to collecting data is understanding the outcome you want to achieve and working backwards to capture only the data you need to get there. Less data, kept accurately, is worth more than vast amounts of data that no one can trust or interpret.

To do this:

- Prioritize data relevancy over data quality. Define what data you need for the specific outcome you're targeting. Then focus on keeping that data accurate. You'll find that the amount of data you need to maintain shrinks significantly, which makes keeping it accurate much more achievable.
- Don't confuse data and context. When evaluating your data readiness, go beyond asking, *"Do we have the data?"* to *"Do we understand what it means and what decision it should drive?"*

2 Integrations

Integrations are the extension of your data and context. They provide the flow of information between your systems and your agents—whether that's reading from a CRM, writing to an ERP, updating a CMDB, or triggering a Slack notification.

But there are two very different types of integration, and they produce very different outcomes:

Type	What it enables	Business result
Read-only	AI can see data and surface answers	Deflection: AI answers questions
Read/write (bidirectional)	AI can take action on data	Resolution: AI completes tasks

When it comes to integrations, resist the need to see everything all the time: pulling in every data source, connecting every system, and building every integration before you start. Instead, ask: *"Does the AI need to see this information to do its job, or does it need to act on it?"* The answer determines the type of integration you need—and the governance that needs to sit around it.

6 building blocks to make AI work at any phase CONTINUED

3 Governance

The moment an AI agent can take action, you have to define exactly what it's allowed to do, under what conditions, and with what level of approval. Without governance, AI becomes a black box. And a black box will never be trusted with anything important in your business.

The four components of effective AI governance are:

1. **Role-based boundaries:**
   Who can the AI act on behalf of?
   What roles or permissions are required before an action is taken?

2. **Task-specific permissions:**
   Which actions are fully autonomous?
   Which require human approval?
   Which are always human-only?

3. **Escalation rules:**
   What triggers a handoff to a human?
   What context is passed at that point?

4. **Full traceability:**
   Is every action taken by an AI agent logged and auditable?

Solving the agent sprawl problem

With AI tools becoming increasingly accessible, the governance challenge has evolved. Anyone can build an agent now. Without guardrails, you get 10 teams building 10 agents against the same process with no shared standards—the AI equivalent of shadow IT.

The solution? An AI governance board or a governed agent library. Any proposed new agent build goes through a rapid triage:

- Is this duplicative of an existing agent?
- Is it tied to a strategic initiative or does it drive meaningful operational efficiency?
- Does it meet your security, data, and compliance requirements?

If it passes: it gets built. If it doesn't: it gets deprioritized. The key is making this process fast and transparent enough that teams engage with it rather than route around it.



6 building blocks to make AI work at any phase CONTINUED



4 Knowledge

AI is only as good as the knowledge available to it. Most organizations understand this, but very few have an operational model in place to support their knowledge base as a living, governed asset.

The good news: AI has fundamentally changed what knowledge management looks like.

Good knowledge management	Poor knowledge management
Clear ownership of the knowledge base, including someone who is accountable for its health	No defined owner; knowledge becomes a maintenance burden nobody claims
A distinction between knowledge for employees (self-service) and knowledge for agents (assisted resolution), optimized for each audience	Treating knowledge setup as a one-time project rather than an ongoing practice
Regular review and publishing, not a one-time exercise	Inconsistent, contradictory, or stale articles that cause AI to deliver unreliable answers
AI-assisted drafting, with humans validating (not writing from scratch)	Measuring knowledge by volume of articles rather than by whether those articles drive resolution

6 building blocks to make AI work at any phase CONTINUED

5 Workflows

AI needs your workflows more than you do. Here's why: A person can improvise around a broken process. They can use judgment, apply context they've gathered over years, and fill in the gaps. An AI agent without a defined workflow and clear guardrails cannot. It will make decisions based on what it knows. And when it doesn't know, it will either guess or fail.

This creates a critical imperative: before you automate a workflow, define it to make it clear. What is the standard path? What are the exceptions? What triggers an escalation? Under what conditions should AI act autonomously versus requiring human approval?

Just as importantly, don't redesign workflows in isolation. Rather than optimizing practices one at a time, take a single service and map it end to end across every process, team, and system it touches. Focus on creating the ideal future-state experience—not preserving today's operating model. Once you've established that "gold standard" service journey, use it as the blueprint for similar services across the organization.

Don't let perfection become the enemy of progress. Start with one high-value service, prove the model, refine it through real-world usage, and scale what works.

We recommend taking a three-tiered approach to workflows:

The right distribution across these tiers depends on your maturity, your risk appetite, and the trust you've built with leadership and security teams. There's no universal answer, but there are definitely wrong ones. Define your workflow classification criteria upfront, and apply them consistently. Making these decisions on a case-by-case basis is where inconsistency creeps in.

Tier	Description	Who handles it
Manual	High complexity, high risk, or high sensitivity; always requires human judgment	Human agent
Automated	Structured, repeatable, rule-based; AI follows a defined path	AI with oversight
Agentic	Multi-step, cross-system, autonomous; AI applies context, makes decisions, and takes action	AI autonomously



6 building blocks to make AI work at any phase CONTINUED

6 Measurement

Measurement is how you prove the impact of everything else. And the fundamental mistake most organizations make is building their operating model first and trying to retrofit reporting into it afterward. That approach will only ever give you activity metrics, not outcome metrics.

Instead, here are the metrics to track at each phase:

Phase	What to measure
Zero In	Self-service adoption rate, knowledge article deflection rate, first-contact resolution for targeted use cases
Evolve	End-to-end automation rate, MTTR for automated workflows, manual steps eliminated per workflow
Nirvana	AI resolution rate (vs. deflection), cost per autonomous resolution, coverage across lines of business, continuous improvement cycle time

Keep in mind:

The business outcomes you're targeting don't change because you've introduced AI. What changes is how they're delivered, and sometimes what they look like in an AI-augmented environment.



Action plan: Assessing AI readiness and building your value pipeline

PART 1. EVALUATE YOUR AI READINESS AND USE CASE MATURITY

Now that you understand the framework and the building blocks, it's time to assess your own AI readiness. This is where the playbook becomes personal.

Zendesk's **AI Readiness Assessment** helps you evaluate your organization across the six building blocks, surfaces your current state, and maps it to the ZEN Framework phase you're in. That way, you know not just where your gaps are, but what to do about them in the right sequence.

1. **Define your KPIs.**
The first step of the ZEN journey is Zero-In. Here, you'll define your desired outcomes, functionality scope, line(s) of business, and industry vertical.

2. **Complete your AI Readiness Assessment**
Next, score your organization across each building block on a 0-3 scale. Then, add optional comments for context. Together, your scores and comments give you an honest baseline—one you can share with your team, your stakeholders, and your Zendesk team to build a practical plan.

Complete your AI Readiness Assessment →

Bonus step: Assess your use case maturity
Once you've scored your building blocks, the next step is to map your readiness to specific use cases based on your KPIs, your line of business, and your industry. Score your current maturity for each use case on a 1-5 scale. We recommend partnering with your CSM or AI expert to complete the exercise together.

Complete your Use Case Maturity Assessment →

The output is a prioritized view of where to focus first: the use cases where your readiness is highest and the business value is greatest. This is your starting point for building your value pipeline.



Action plan: Assessing AI readiness and building your value pipeline CONTINUED

PART 2. BUILDING YOUR VALUE PIPELINE

Your assessment gives you a scored view of where you are. Now let's talk about how to move forward.

1 Identify your top 3 use cases to start with

From your use case maturity assessment, select the three use cases where you have:

- The highest current readiness score (fewest gaps to close)
- A direct connection to a KPI your leadership cares about
- Clear data, knowledge, and workflow foundations already in place (or closeable quickly)

These are your first value streams.

2 Define the value for each For each use case, define:

- What does success look like in 90 days?
- What is the current state baseline? (e.g., How long does this take today? How many tickets does it generate? And so on.)
- What is the target state? (e.g., 60% self-service resolution, 30-minute reduction in MTTR)
- What does that translate to in cost or time saved?
-

Now you have the template for translating value for your CIO. For example: "We removed 450 hours of tier-one handling from the team in Q1, equivalent to $X in operational cost."

3 Sequence your value drops

Map your value streams across the next 12 months in a value pipeline.
Your starting structure could be:

Quarter	Phase	Focus	KPI Target
Q1	Zero In	Top 3 use cases: foundations + first deployments	Self-service adoption: +X%
Q2	Evolve	Workflow automation for priority services and cross-system integrations	MTTR reduction: -X%
Q3	Evolve	Expanded use cases	Cost per resolution: -$X
Q4	Nirvana	Analytics loop, admin optimization, second line of business	Autonomous resolution rate: X%

Action plan: Assessing AI readiness and building your value pipeline CONTINUED



PART 2. BUILDING YOUR VALUE PIPELINE CONTINUED

4 Set your governance cadence

Finally, set the cadence for reviewing your governance. We suggest:

Monthly	Review AI performance against use-case-level KPIs.
Quarterly	Run a knowledge audit and identify stale or missing content surfaced by AI. Also, review the agent library. Were there any new builds proposed? Any duplicates to consolidate? And so on.
Annually	Reassess building block readiness and refresh the value pipeline.

Frequently asked questions

Q: How do I get my CIO to trust AI enough to approve autonomous workflows?

A: Start with transparency. Choose low-risk, high-volume, highly repeatable use cases first, like password resets, FAQ resolution, or standard equipment requests. Prove the model, build the audit trail, and show leadership that every action is logged, every decision is explainable, and every outcome is measurable. When the evidence is there, autonomy becomes a logical next step—not a leap of faith.
The value pipeline is your tool here. CIOs want to see a 12-month plan with defined value drops at each stage—and to track progress against those drops—before they fund the next phase.

Q: Our processes are a mess. Do we have to fix everything before we can start?

A: No, but don't fix processes in vertical isolation either. Rather than taking each ITIL process (incident, problem, change, CMDB) and fixing it independently, pick one complete service and map it horizontally end to end. Determine what the optimized operating model looks like for that service, across all the processes it touches. Then set that as your benchmark and apply consistent principles. Processes built on real-world benchmarks from a journey-based perspective, rather than a function-by-function one, will soon start to resolve themselves.

Q: With AI tools so easy to build, how do I stop every team creating their own agent?

A: In a word: governance. That is, make the governed path faster and easier than building something unsanctioned.
Here's how to do it: Create a governed agent library. Define a simple triage process for new build requests. Make it clear that agents built outside the framework won't get IT support, won't have access to production systems, and won't be maintained. The agents that go through the process will be better—and the teams that build them will know it.



Q: How do I upskill my team as AI takes over tier-one work?

A: The best upskilling happens through role evolution, not training programs alone. The people who've been handling tier-one tickets are your highest-value knowledge validators. They understand the context of those requests better than anyone. Move them into knowledge management, AI oversight, or workflow design roles. They don't need to become developers. They need to become AI behavior architects who understand the business context and can ensure AI operates within it correctly.

Q: If AI handles the simple tickets, won't my MTTR go up?

A: Yes. And that's a signal of success, not failure. When AI is handling tier-one resolution, the tickets that reach your human agents are the ones that genuinely require expertise. A higher MTTR in that context means your team is working on harder problems, which is exactly what you want. Build this context into your reporting from day one, so leadership interprets the data correctly.

AI is ready. Discover whether your foundation is.

The organizations that lead in the AI era aren't the ones that move fastest. They're the ones that build on connected systems, clean knowledge, defined workflows, and governance that earns trust rather than restricting progress.

The ZEN Framework and six building blocks give you the foundation, structure, sequencing, and measurement model to turn AI from an experiment into a compounding, enterprise-wide capability.

The next step is completing your assessment. Know where you stand. Identify your gaps. Build your value pipeline. And take the first value stream to completion. That first proof point is what gets leadership bought in, teams energized, and the rest of the journey funded.

Take your AI Readiness Assessment now →

Already completed your assessment?

Start your free trial →



zendesk

AI masterclass