



The AI agent strategy playbook

A PRACTICAL FRAMEWORK FOR EVALUATING AND
DEPLOYING AI AGENTS ACROSS THE ENTERPRISE

PUT IT IN PLAY

WINNING MOVE

THE LINEUP

STRATEGY MAP

OPENING MOVE

Opening move:


Moving from prototype to production

Enterprises are attempting to move beyond AI prototypes into real-world deployment, but 86% of enterprises¹ require upgrades to their existing tech stack to deploy AI agents. Without a strategic approach, businesses risk wasted AI investments, rising operational costs, and governance challenges.

This playbook breaks down the four most common AI agent deployment strategies, the trade-offs of each, and how to determine which model best fits your architecture, security model, and resourcing capacity.

We'll also walk through:

- **What each model looks like in practice**
- **Where each model fits (and doesn't)**
- **How to assess deployment paths based on control, speed, scalability, and governance**
- **Why AI-ready iPaaS is uniquely equipped for cross-functional, secure, and scalable agent deployment**



Enterprises want AI agents—but their tech stacks aren't ready for game day


According to Gartner, by 2029,

80%

of enterprises² with mature automation practices will pivot to a consolidated platform that orchestrates business processes and agentic automation.

Getting there will be no easy feat for enterprises. Our recent survey of 1,000+ enterprise technology leaders and practitioners uncovered the truth behind slow AI agent adoption rates among enterprises.

The defensive line:
What's blocking
your agents?



Governance and security	Without strong governance, AI agents introduce risk of unauthorized actions, policy violations, and data leakage.
Scalability and performance	AI agents break under scale when legacy platforms can't process tasks or data fast enough to meet business needs.
Integration complexity	Poor system connectivity limits agents' ability to trigger workflows or access data, reducing their usefulness and reliability.
Talent and expertise gap	Lack of internal AI/ML expertise delays agent development and increases dependencies on external vendors or consulting.
Cost and ROI pressures	Without clear ROI, some hesitate to fund AI, while others overspend chasing momentum, risking stalled projects or costly missteps.
Deployment timelines	Lengthy deployment cycles stall momentum and cause AI projects to miss windows of business relevance.


Strategy map:

Four AI agent deployment strategies

There's no single way to deploy AI agents. Depending on priorities and capabilities, enterprises are experimenting with different models—from vendor-tied SaaS tools to fully custom builds. Here, we break down the four primary deployment strategies, how they work, where they fit best, and the trade-offs of each.

We'll explore these four strategic plays:

- **Play 1: Internal custom agents**
- **Play 2: Off-the-shelf agents**
- **Play 3: SaaS app agents**
- **Play 4: AI-ready iPaaS**



What they are

Custom-built AI agents developed in-house by engineering teams.

How it's done


Internal teams design and build agents from the ground up, including data pipelines, orchestration logic, models and integration. Every component is tailored for internal systems and requirements.

Where it fits

Best for highly specialized use cases where full control over architecture, data, and execution is critical—and where internal AI expertise and engineering resources are available.

Trade offs

- Requires deep in-house AI/ML expertise
- Long build cycles delay business impact
- High upfront and ongoing costs
- Custom integration increases technical debt
- Difficult to scale across teams and use cases
- Governance, security and compliance must be built from scratch



SUMMARY	
Speed to deploy	Slow
Governance	Build from scratch
Customization	Full control
Scalability	Hard to scale
Integration depth	Custom for every system



Run this when:

You need full control over architecture and data



Watch out for:

High cost, long timelines, deep AI/ML expertise required

What they are

Pre-built AI agents designed for a specific function, like chatbots, scheduling assistants, or basic Q&A agents.

How it's done


Enterprises license or subscribe to these agents from a vendor and configure them within preset parameters. Most offer limited customization or extensions.

Where it fits

Useful for simple, purpose-built tasks where speed of deployment is more important than control, scale, or system integration.

Trade offs

- Limited integration reduces agent effectiveness
- Rigid design makes it hard to adapt to new use cases
- Lacks enterprise governance and compliance
- Difficult to scale across workflows or departments
- Creates vendor lock-in
- Leads to fragmented oversight and policy gaps



SUMMARY	
Speed to deploy	Fast
Governance	Limited
Customization	Minimal
Scalability	Not scalable
Integration depth	Low



Run this when:

You want fast deployment for a simple use case



Watch out for:

Limited control, integration and scalability

What they are

Agents embedded within SaaS platforms (e.g., CRM or ERP systems) that automate tasks specific to that ecosystem.

How it's done


Delivered as part of a SaaS subscription, these agents are pre-configured by the vendor to support platform-specific tasks like data entry, routing, and recommendations.

Where it fits

A quick option for automating repetitive tasks within a single platform — ideal for teams that already operate heavily within one vendor ecosystem.

Trade offs

- Confined to a single vendor ecosystem
- Minimal customization beyond predefined workflows
- Limited interoperability with external systems
- Governance depends on the SaaS vendor's controls
- Ongoing licensing and switching costs reduce ROI
- No cross-functional agent orchestration



SUMMARY

Speed to deploy	Fast (within platform)
Governance	Vendor-defined
Customization	Limited
Scalability	Confined to vendor
Integration depth	Platform only



Run this when:

You need quick automation inside a SaaS tool



Watch out for:

No cross-functional orchestration, limited reach

What they are

An enterprise integration platform designed to support secure, scalable, cross-functional AI agent deployment.

How it's done


Combines low-code development, orchestration tools, composability, and universal integration for IT and business teams to build, deploy, and monitor agents across systems—while enforcing centralized governance and observability.

Where it fits

Ideal for enterprises scaling agents enterprise-wide, needing unified governance, pricing, and security across teams and systems.

Trade offs

- Upfront onboarding required to scale adoption
- Requires centralized ownership to avoid vendor sprawl and fragmented agent strategies
- Scaling agents across teams adds pressure to replace legacy tools and unify integration



SUMMARY	
Speed to deploy	Moderate
Governance	Centralized
Customization	High
Scalability	Enterprise-wide
Integration depth	Full stack



Run this when:

You're scaling agents across systems and teams




Watch out for:

Requires platform onboarding and central ownership

The lineup:

Comparing AI agent deployment models

The right deployment model depends on your goals, maturity and appetite for control and scale. Here's how each model addresses the top enterprise challenges.



Choose your play



	Internal custom agents	Off-the-shelf agents	SaaS app agents	AI-ready iPaaS
Governance and security	⚠️ Enterprise must build controls from scratch	❌ Limited oversight, inconsistent access, audit, and security controls	⚠️ Depends on SaaS vendors' policies	✅ Centralized guardrails, auditability, and controls
Scalability and performance	⚠️ Possible, but requires significant investment	❌ Rigid and single-purpose, not scalable	⚠️ Limited to vendor ecosystem	✅ Built to scale across systems and data volumes
Integration complexity	❌ Custom integration required for every system	❌ Minimal integration options	⚠️ Works within platform, limited outside	✅ Hundreds of pre-built connectors, APIs and orchestration
Talent and expertise gap	❌ Requires in-house AI/ML and engineering expertise	✅ No technical expertise required	✅ Easy to adopt within existing platforms	⚠️ Low-code reduces build effort; onboarding still needed
Cost and ROI pressures	❌ High build cost, long ROI horizon	✅ Low cost for basic use cases	⚠️ Ongoing costs add up over time	✅ Efficient build + reuse for exponential ROI over time
Deployment timelines	❌ Long development cycles delay outcomes	✅ Fast to deploy	✅ Fast (within platform)	⚠️ Faster than custom builds; setup and governance planning needed


Winning move:

Why AI-ready iPaaS wins

Enterprise AI adoption is moving fast, but success depends on more than just getting agents live. The real challenge is building a strategy that works not just for one team or one use case, but across departments, systems, and data boundaries.

This requires more than off-the-shelf tools or point solutions. It requires a platform that combines connectivity, control, and flexibility without introducing governance gaps or scalability limits.

AI-ready iPaaS offers the most complete foundation—but not all integration platforms are designed for the demands.



Why Tray is built for enterprise AI agents

AI agents can't succeed in silos. They need to connect to the systems your business runs on, follow your governance model, and scale without breaking. Tray's AI-ready iPaaS was built to do exactly that. Here's how:

Build	Merlin Agent Builder	• Use natural language to design and configure agents • Leverage pre-built agent templates, reusable logic, and guided workflows • Test, iterate, and deploy without heavy dev resources
Govern	Merlin Guardian	• Apply granular access control and role-based policies • Enforce compliance guardrails, input validation, and auditing • Monitor agent behavior and enforce secure, observable execution
Work with data	Vector Tables & IDP	• Vector Tables let agents retrieve and reason on structured and unstructured data • Intelligent Document Processing (IDP) automates knowledge ingestion • Bring your own LLM or use Merlin for generation, classification, and summarization
Connect	700+ integrations	• Pre-built connectors to SaaS apps, databases, and internal tools • REST, GraphQL, SOAP, and file-based integrations supported natively • Event-driven triggers and webhook support for real-time orchestration
Scale	Serverless, cross-system orchestration	• Orchestrate agents across multiple departments and systems • Zero-provisioning, serverless infrastructure handles billions of tasks • Elastic scalability without added infrastructure overhead

Inspiration for cross-functional agents

Tray makes it possible to go beyond isolated use cases and build a governed, repeatable model for scaling AI agents across your business. Here are just a few ways teams are putting agents into play—use them to spark ideas for the agents you can build next.



IT & internal support agents

ITSM – Automates IT ticket triage and resolution
Knowledge – Retrieves and applies organizational knowledges across systems and teams
Policy monitoring – Checks system activity against internal rules and thresholds



HR agents

HR service desk – Answers policy, onboarding, PTO questions
Recruiting assistant – Screens resumes, schedules interviews
Onboarding – Triggers provisioning and comms for new hires



Finance agents

Finance – Answers budget, spend, or ERP-related questions
Data & analytics – Builds dashboards, extracts performance data
Procurement approval – Automates requests and approval routing
Invoice reconciliation – Extracts and routes invoice data



Sales agents

CRM – Retrieves sales insights (e.g., opps, win rates)
Sales – Drafts outreach, call prep, follow-up
Sales pipeline – Flags stuck deals, recommends next steps
RFP – Fills proposals using existing content and logic



Marketing agents

Marketing campaign – Coordinates multi-channel campaign launches
Product feedback – Mines internal notes and call logs
Content search – Recommends best-fit assets
Event follow-up – Routes leads and triggers follow-ups



Customer experience agents

Customer support – Deflects tickets, provides contextual answers
Customer risk monitoring – Flags high-risk accounts
Account briefing – Compiles insights before customer check-ins

From idea to impact:

It's not about launching one agent. It's about building a foundation that scales, adapts, and stays in your control.

Tray gives you the platform to operationalize your AI agent strategy—across teams, systems, and use cases.



Talk to a strategist →

Get direct, technical input on how this model will work in your stack.



Join a hands-on workshop →

Work side by side with our team to build a real agent—from idea to deployment—with your goals and systems in mind.



Benchmark your strategy →

See how 1,000+ IT leaders are tackling adoption—based on Tray's enterprise AI agent survey.





About Tray.ai

Tray.ai is the creator of Merlin Agent Builder — the way to build and deploy high-value AI agents without constraints. Unlike single-purpose SaaS apps, Merlin fits your business needs. Whether you start from scratch or use our pre-built accelerators, configure, deploy and manage horizontal AI agents — all in one place, without coding or complexity. Built on the Tray composable AI integration platform, it includes the connectivity, data integration, guardrails, governance controls, and scalability needed to deploy agents enterprise-wide. Learn more at [Tray.ai](https://tray.ai).

Citations

- 1 Tray.ai (2024). Enterprise AI agent survey report. <https://tray.ai/enterprise-ai-agent-survey-report>
- 2 Gartner. (2023, October). Strategic roadmap for automation. Gartner, Inc.