

genpact

The enterprise buyer's guide to agentic AI





Table of contents

1. Setting the stage: What agentic AI is, why this decision is different, and whether you're ready	4
2. The decision matrix and total cost of ownership: Build vs. buy	9
3. Your due diligence checklist: How to evaluate, compare, and make the call	16
4. Move: What to expect from a partner and your next steps	23

A structured decision guide for **senior leaders evaluating, justifying, and deploying agentic AI in enterprise operations**

Agentic AI doesn't just change what work gets done – it changes who owns execution. That shift carries real accountability. And it makes one question unavoidable for every leadership team: do you build this capability yourself, or do you buy it?

This guide gives you the tools to answer that with confidence. Inside, you'll find a build vs. buy decision matrix, a total cost of ownership framework, a due diligence checklist, a readiness self-assessment, and a clear set of red flags to help you separate a capable partner from a costly one.

It shows how enterprise teams can move from problem identification to purchase, including stakeholder alignment, requirements building, proof points, and decision-readiness checkpoints.



1.

Setting the stage: What agentic AI is, why this decision is different, and whether you're ready

Who this guide is for and how to use it

This guide is for the leaders who shape and sign off on agentic AI:

- CIOs and CTOs aligning architecture, integration, and risk
- COOs and shared services heads who own end-to-end execution
- CFOs weighing cost, payback, and defensible ROI
- Transformation leaders and AI program owners sequencing the roadmap
- Risk, compliance, and people leaders who shape and stress-test the decision



Recap

Agentic AI shifts who owns execution. The buying bar is higher because you're delegating action, not output. Trust, governed through four leadership decisions, is what separates pilots from scale.



From prototype to autonomy

Most enterprises have already experimented with gen AI. Chatbots, copilots, summarizers. Useful, but assistive by design. A human still owns every decision, every action, every consequence.

Agentic AI works differently. It reasons across steps, moves across systems, triggers workflows, and acts within defined guardrails. It works beyond suggesting the next move

to making the next move. An agent can read a document, validate the numbers, post the result, and escalate the edge case autonomously.

That shift rewrites how you evaluate, fund, and govern the technology. And here's what leaders often miss: agentic AI doesn't hide your weaknesses. It amplifies them. Broken processes scale and bad data is replicated.



Stat spotlight

92% of senior executives believe agentic AI will fundamentally change how work gets done. The conviction is near-universal. The organizational readiness isn't.

The operational discipline required

Autonomy without structure is just risk at speed

To run agents safely, you need capabilities most enterprises haven't yet built:

- Clear guardrails that define exactly what an agent can and can't do
- Escalation paths so the right human steps in at the right moment
- Evidence capture that logs every action, input, and decision for audit and review
- Permission control that starts supervised and expands only as trust is established

This isn't automation with a smarter engine. It's a new execution model, and it asks more of your operating model than previous technology waves.



The new enterprise blueprint

The destination is an operating model where agents handle exceptions, route work, resolve queries, and close tasks end to end – while people focus on judgment, strategy, and the edge cases that genuinely need human input.

Getting there isn't about bolting agents onto existing systems. It requires redesigning how work flows, where decisions sit, and who's accountable when an agent acts. The architecture that matters isn't the model. It's the orchestration, governance, measurement, and workforce design around it.

Get those right and autonomy scales. Skip them and you join the roughly 80% of enterprises still running agents in supervised, low-autonomy modes.



**Board-level
question**

Are you redesigning workflows for autonomy or layering AI on top of processes that are undefined and unreliable?

Why the buying bar is higher for agentic AI

Buying traditional software requires specific criteria: Does it work? With agentic AI, the question shifts to: How safely and reliably does it act?

Model quality is baseline. The real evaluation is about orchestration, integration, auditability, and accountability. This is not a tool that waits for instructions. Buyers are delegating execution to a system that acts on their behalf. And buyers own the outcome.



Stat spotlight

92% believe agentic AI will change how work gets done. Only 22% are comfortable operating agents with domain-level or broad autonomy. That gap is the challenge. Trust, not technology, is the barrier to scale.

The trust-first principle: Four leadership decisions

Autonomy scales only when it's governed.

Genpact's study with HFS Research, *Autonomy requires trust in AI*, surveyed 545 senior executives across 11 industries. The conclusion is consistent: the hardest part of scaling agentic AI isn't the technology. It's organizational readiness, and it comes down to four decisions.

Leave any one of these unresolved and scale stalls.



Accountability

Who owns outcomes when agents act?
Identifying an owner is required to expand autonomy



Measurement

What does agent success look like? Metrics must reflect what the system actually does, or investment decisions will misfire



Workforce clarity

Is the human impact defined? Decision rights, oversight responsibilities, and escalation paths need to be set before launch, not after



Process design

Do your workflows support autonomy? Agents running inside broken processes produce inaccurate outcomes. Rebuild the process first



What good looks like

Agentic workflows need named accountable owners, agent-native KPIs defined before launch, documented decision rights and escalation triggers, and a process that was redesigned for autonomy – not retrofitted.

How enterprise buyers move from problem to purchase

Enterprise buying teams rarely make the agentic AI decision in one step. They move through a sequence of tasks: proving the problem, exploring solution paths, defining requirements, comparing vendors, and securing purchase approval. Treating the guide as a journey – not only as an evaluation checklist – helps buyers know what evidence, stakeholders, and decisions are required at each stage.



2. The decision matrix and total cost of ownership: **Build vs. buy**

What the matrix covers: Six decision dimensions

Score each dimension for one priority workflow. The pattern of your scores points toward build or buy.

1. Workflow complexity and differentiation

What it measures

How unique and strategic the workflow is

Build

☐

1

☐

2

☐

3

☐

4

☐

5

Buy

Highly differentiated, proprietary logic

Common, well-understood process

2. Internal capability

What it measures

Strength across product, engineering, security, governance

Build

☐

1

☐

2

☐

3

☐

4

☐

5

Buy

Deep, proven, full-stack

Gaps in one or more areas

3. Speed to value

What it measures

How fast you need measurable outcomes

Build

☐

1

☐

2

☐

3

☐

4

☐

5

Buy

Time is flexible

Speed is critical

4. Risk appetite and governance maturity

What it measures

Tolerance for risk and maturity of controls

Build

☐
1☐
2☐
3☐
4☐
5

Buy

High maturity, defined controls

Governance still developing

5. Foundation readiness

What it measures

Process, data, technology, and talent readiness

Build

☐
1☐
2☐
3☐
4☐
5

Buy

Strong across all four

Visible foundation gaps

6. Governable autonomy readiness

What it measures

Accountability defined, agent-native KPIs, decision rights, escalation paths

Build

☐
1☐
2☐
3☐
4☐
5

Buy

All in place


Still being established

Total score

How to score:

- **Below 18:** A credible build path is open for this workflow
- **18–24:** Buy, with a partner who closes your gaps
- **Above 24:** Buy, and fix the foundation as you go

Run the matrix per workflow. You may build in one differentiated area and buy everywhere else.



When building may make sense

Building can be the right call when:

- The workflow is highly differentiated and central to your competitive position
- You hold unique internal data or decision logic worth protecting
- You have strong in-house capability across every matrix dimension, not just engineering
- Governance, workforce clarity, and process design for autonomy are already mature

Be honest with yourself. Building doesn't just mean owning integration. It means designing the accountability structures and measurement frameworks that make autonomous systems safe to expand. Most enterprises clear that bar in a small number of workflows.



Red flag

Building solely to avoid vendor dependency is a preference, not a strategy. Building means staffing product, engineering, security, and governance at full strength.

When buying is the smarter path

Buy covers a range of options: packaged platforms, managed services, implementation partners, and platform-led ecosystems. Most enterprises blend these.

Buying makes sense when:

- Speed matters and you can't wait to build capability
- The workflow is common rather than competitively differentiated
- Governance maturity is still developing
- You need help closing foundation gaps while delivering value, not before

For most enterprises, buying is the faster and lower-risk path – particularly with a partner who brings process expertise and governance design, not just model access.



What good looks like

A buy decision that closes a capability gap and leaves you more capable over time: cleaner processes, better measurement, and governance you can operate independently.

The full cost picture: Hidden costs and five real cost drivers

The first business case is rarely complete. Both build and buy carry costs that don't make it onto the opening slide.

Foundation matters here. In Genpact's recent study: [*How four enterprise debts will make or break your AI future*](#), surveyed over 2,000 enterprise leaders and found roughly \$18 trillion in trapped value across process, data, technology, and talent debts. The research found 85% of leaders say enterprise debt actively limits their AI value. Build agents on weak foundations and costs inflate on either path.

The five cost drivers most business cases miss:



Integration and identity

Connecting systems, permissions, and access controls cleanly and securely.



Escalation and oversight design

Building human-in-the-loop mechanisms that keep autonomy safe.



Accountability frameworks

Defining and maintaining who owns what.



Agent-native measurement

Building KPIs that reflect autonomous performance, not just task speed.



Workflow redesign and workforce transition

Rebuilding processes and shifting decision rights.



More than half of enterprises have no funded debt resolution plan, which raises costs on both sides of the decision. Map your enterprise debts before committing to a path.

Cost category	If you build	If you buy
Initial setup	High	Moderate
Integration and identity	You own it	Shared with partner
Governance and accountability design	You own it	Partner-supported
Measurement infrastructure	You build it	Offer included
Ongoing maintenance	High, in-house	Bundled in service
Time to first value	Slower	Faster



Red flag

A business case showing only license cost and projected savings. Accurate numbers must include costs for integration, governance design, measurement, and change management.



Measurement and ROI: The metrics problem

Most enterprises are measuring agentic AI with metrics built for a different era.



Stat spotlight

71% of executives expect agentic AI to deliver ROI faster than previous technology waves. Yet 53% don't have KPIs that reflect autonomous system performance, and 67% still rely on legacy productivity metrics. High expectations, but inaccurate measurements.

Move beyond productivity metrics. Track agent-native measures instead:

- Workflows completed end to end without escalation
- Decisions removed from human queues
- Independent exception handling rate
- Reduced handoffs across functions
- Cycle time for outcomes, not just tasks

The business case, the board update, and the decision to expand autonomy all depend on measuring the right things.



Board-level question

Do you have agent-native KPIs before you scale – or are you measuring a new system with old numbers?

The cost of delay

Inaction has a cost too. Every quarter spent deliberating is another quarter of AI investment landing on unstable foundations. It's also another quarter the leading enterprises use to widen their advantage.



Stat spotlight

Only 6% of enterprises are proven debt resolvers – the ones treating debt resolution, AI, and governance as a single program. They're not waiting for perfect conditions.

Quantify the cost of inaction before comparing paths

Before deciding whether to build or buy, leadership teams should first name the hidden cost of operational friction. A strong business case should quantify the cost of manual work, delayed cycle times, rework, exception handling, weak controls, supplier dissatisfaction, and enterprise debt. This prevents a technology-first comparison and anchors the decision in outcomes the business can defend.

Use this prompt in the business case: What is the dollar value of doing nothing for another year, and which process, data, technology, or talent debts make that cost worse?



Recap

Score the six dimensions. Account for hidden costs. Address your metrics. Build only where you're genuinely differentiated and ready. Otherwise, buy – and close the foundation gaps as you go.

3.

Your due diligence checklist: How to evaluate, compare, and make the call

Once you've settled the build vs. buy question and mapped the true cost of ownership, the next decision is sharper: which partner can you actually trust to act on your behalf? That call shouldn't rest on a demo.

This section gives you the structured tools to get it right – eight evaluation criteria, the governance and trust requirements to insist

on, the questions every stakeholder will raise, a way to define your requirements before you talk to a single vendor, and a readiness self-assessment.

Work through them, and you'll reach a decision you can defend to your board, your auditors, and yourself.

Eight enterprise evaluation criteria

For each criterion, here's what to look for – and its importance by buyer type.

✓ **Business fit:**

Solves a meaningful workflow problem with clear operational or commercial value.

✓ **Time to value:**

Moves from pilot to measurable production outcomes without unnecessary delay.

✓ **Integration readiness:**

Connects to your systems, data, identity, and workflow tools without requiring a major rebuild.

✓ **Governance and compliance:**

Controls, logging, auditability, and policy enforcement are built in from day one.

✓ **Human oversight:**

Handles approvals, exceptions, escalations, and decision rights safely.

✓ **Scalability and resilience:**

Performs across functions, geographies, volumes, and real-world complexity.

✓ **Measurement and ROI:**

Defines and tracks success with both business metrics and agent-native KPIs.

✓ **Partner maturity:**

Has the operational experience, support model, and process expertise to help you scale.



Specifics for mid-market buyers:

Speed and business impact take priority. Weight time to value, integration readiness, and business fit highest. Look for partners who move quickly without requiring a long rebuild. Just don't let speed become an argument against an audit trail.

Specifics for large enterprises:

Compliance and operational scale are the primary concerns. Weight governance and compliance, scalability and resilience, and human oversight. Ask specifically what happens when agents produce incorrect outputs – how errors are detected, contained, and recovered from.

Specifics for highly regulated industries:

Human oversight and regulatory compliance come first. Weight human oversight, governance and compliance, and measurement. Require documented evidence that the solution meets applicable regulations and that humans retain clear control over consequential decisions. Governance isn't a feature here – it's a requirement.



Red flags

No audit trail. If you can't reconstruct what an agent did and why, you can't govern it or defend it. That's a fundamental gap.

No explainability. A system that can't account for its decisions can't be defended to a regulator, a board, or a customer.

The model only learns against itself. Without human feedback in the loop, quality drifts silently and errors compound. Require evidence of human-in-the-loop learning.



Governance and trust by design, not bolted on

Logging, accountability, policy enforcement, and audit readiness need to be present from day one. Escalation paths, human decision rights, and intervention triggers should be defined explicitly. When decision rights are unclear, accountability breaks down – usually at the moment it matters most.

Only 22% of enterprises are comfortable with broad agent autonomy. The ability to begin in supervised modes and expand permissions deliberately, as trust is earned, is a practical requirement.

For mid-market buyers:

Keep governance proportionate and aligned to needs. Supervised modes, clear escalation, and basic logging are the minimum. Don't trade governance for speed.

For large enterprises:

Expect enterprise-grade controls, role-based access, and full evidence capture. Accountability belongs in the deployment design, not as an after-thought.

For highly regulated industries:

Governance is the core requirement. Demand proven experience in regulated environments, documented control points, and the ability to demonstrate compliance on demand.



Red flag

Governance described as a phase-two deliverable. If it isn't built in from the start, it won't be added later without a significant incident as the trigger.



What good looks like

Graduated permissions you control, complete logging, named accountability for every workflow, and compliance you can demonstrate in minutes.

Questions every stakeholder will ask

Getting to a decision requires alignment across multiple functions. Here's what each will want answered.

What's the total implementation cost, all-in? Which agent-native metrics will we use to track ROI? What's the cost of not acting this year?

Finance

How does this fit our existing architecture and identity model? How does it affect our current integration and legacy debt? How do we monitor, update, and roll back agents safely?

CIOs and IT leaders

How complete is the audit trail? Who's accountable when an agent acts incorrectly? What are the escalation triggers and policy controls?

Risk and compliance

How do decision rights change when agents are in the workflow? How does it perform in real, messy processes – not just controlled demos? How are exceptions returned to human handlers?

Operations

How do roles and oversight responsibilities change? What's the plan for reskilling and workforce transition? How do we keep people clear on what they own?

HR and people leaders



Who joins the buying process, and when

The buying committee expands as the decision matures. Early conversations often start with the business owner and finance leader. Later stages require IT, security, procurement, legal, operations, executive sponsorship, project management office, and implementation leadership. Mapping this progression early reduces late-stage stalls.



Board-level question

Have finance, IT, risk, operations, and HR all reviewed this – or are we one function away from a stall?

Requirements building before vendor selection

Buying teams should define requirements before vendor selection, not during vendor negotiation. The goal is to create a shared baseline for what must be true operationally, technically, financially, and from a control standpoint before any partner is evaluated.

Requirement area	Questions to answer
Baseline	Which process metrics, exception volumes, service levels, cost drivers, and pain points define the starting point?
Controls	Which controls must remain human-governed, and which can move to graduated autonomy?
Integration	Which ERP, workflow, identity, reporting, and data environments must the agentic workflow connect to?
KPIs	Which business and agent-native metrics define success before launch?
Security and compliance	Which security, privacy, audit, and regulatory requirements are non-negotiable?
Operating model	Which roles, decision rights, escalation paths, and governance cadences need to change?



Readiness self-assessment

Section 2 told you what to choose. This section tells you how to make it work. Here you'll find the practical requirements to move from a smart decision to a system that actually delivers.

#	Readiness dimension	Score (1–5)
1	Process maturity	
2	Use-case clarity	
3	Data readiness	
4	System and integration environment	
5	Governance framework	
6	Executive sponsorship	
7	Cross-functional alignment	
8	Internal resourcing	
9	Change-management capacity	
10	Measurement readiness	
11	Accountability clarity (named outcome owners)	
12	Workforce clarity (defined decision rights, escalation paths, and oversight)	
	Total	

Reading your score:

- **48–60:** High readiness. A build path is credible for differentiated workflows
- **30–47:** Moderate readiness. Buy, with a partner who closes gaps as you deliver
- **Below 30:** Lower readiness. Buy, and treat foundation work as part of the program – not a prerequisite to it



Recap

Evaluate the eight criteria, weight them for your context, watch the red flags, run the requirements checklist, and score the 12 readiness dimensions. Lower readiness tilts toward buy. Higher readiness opens a build path in select workflows.

4.

Move: What to expect from a partner and your next steps

What to expect from a vendor or buying partner

The strongest partners don't just implement AI – they help redesign the workflows that need to run it, and they stay accountable for the outcomes.

Look for a partner who brings:

- Domain and process expertise in your specific operations, not just AI capability in the abstract
- Enterprise orchestration that coordinates agents across systems and handles real-world exceptions
- Regulated-environment experience if compliance is part of your operating context
- A single integrated program covering AI, governance, and operating-model change together
- Post-launch accountability – measurement frameworks, workforce transition support, and accountability for the results

If a partner leads with the model and skips the process and governance design, to keep looking. Deployment is the easy part.




What good looks like

A partner who can demonstrate a governed agent operating in a high-exception, real-world workflow – with a complete audit trail, escalation logic, and agent-native metrics – and clearly explain what happens when something goes wrong.

Proof to expect at each buying stage

The evidence a buyer needs changes by stage. Use this as a practical filter for partner conversations and asset requests.



Moving forward: What real change looks like

Here are two examples of how enterprises have moved from challenge to measurable outcome with agentic AI.

Case study: From 7% to 65% touchless AP for a global supply chain company



The challenge:

A global distribution and supply chain provider was processing close to 3.5 million invoices per year across a fragmented set of ERP systems inherited after a merger. Accounts payable was largely manual –

team members were sifting through hundreds of invoices by hand, resolving exceptions over email, and struggling to keep pace with volume. On-time payment rates sat at 60%.



The solution:

The company partnered with Genpact to deploy Genpact AP Suite, combining AP Capture for invoice data extraction and AP Advance for three-way matching, exception handling, and non-PO processing. Twenty-one

pretrained AI agents handle intelligent routing and prioritization, adapting to regional and business-specific rules. Human experts retain oversight of complex exceptions and validate outcomes throughout.



The results:

Touchless invoice processing rose from 7% to 65%. On-time payments went from 60% to 95%. Data extraction accuracy improved from 40% to 92%. Invoice processing cycle times dropped from 18–29 days to 9–14 days, and first-pass yield climbed from 67% to 85%. Discount capture grew from \$35 million to \$44 million. The team also identified approximately \$381 million in duplicate invoices.

[Read the full case study](#)





Case study: Reducing manual AP work by up to 99% for a global financial institution



The challenge:

A leading retail and institutional bank was struggling with accounts payable operations built on legacy infrastructure. OCR accuracy was poor, employees spent significant time manually resolving invoice mismatches,

and duplicate payment risk was a persistent concern. The manual workload was slowing the team and affecting supplier and customer relationships.



The solution:

Genpact AP Suite was implemented across three integrated modules: AP Capture for accurate invoice data extraction with multilingual support; an intelligent invoice matching module built on Salesforce

to automate exception handling for non-PO invoices; and a smart anomaly detection module using machine learning to flag duplicate payments and unusual billing patterns.



The results:

Manual work cut by up to 99%, duplicate payments reduced to near-zero, working capital leakages down by 30%, and a projected 40% cost optimization. OCR

accuracy improved significantly, exception resolution accelerated, and employees reported higher job satisfaction as the volume of manual intervention dropped.

[Read the full case study](#)

Both examples share the same pattern: a specific operational problem, a structured agentic solution with human oversight built in, and outcomes that are measurable and defensible. That's the standard to hold any deployment to.



Board-level question

Can you define the operational problem precisely enough to project outcomes like these – or are you still describing a general ambition?

The right partner

Agentic AI can't scale without process intelligence and governance. These underpin every section of this guide.

Look for a partner that brings AI together with deep process and operations expertise, experience in complex and compliance-sensitive environments, and accountability across the full program: strategy, workflow redesign, deployment, governance, change management, and ongoing performance measurement.

The enterprises that pull ahead in this next period won't just have better models. They'll have more reliable execution – governed, accountable, and built to scale.



What good looks like

A single partner accountable for strategy, redesign, deployment, governance, and change management – with agent-native measurement demonstrating value throughout.



Decision-readiness checklist

Before moving into purchase or implementation, use this checklist to confirm the buying team has moved from interest to readiness.

- ☐ We quantified the cost of inaction
- ☐ We agreed on the target workflow and scope
- ☐ We defined baseline metrics and success criteria
- ☐ We identified governance, control, and security requirements
- ☐ We mapped required integrations and data dependencies
- ☐ We involved the right stakeholders at the right stage
- ☐ We compared vendors using agreed upon evaluation criteria
- ☐ We clarified ownership, escalation paths, and accountability
- ☐ We approved a commercial model, implementation plan, and proof milestones



Your next three moves



Step 01

Run the readiness self-assessment. Score it honestly and map your foundation gaps



Step 02

Use the decision matrix for one priority workflow. Pick one, score it, and decide




Step 03

If you decide to buy, start partner conversations with process expertise, accountability design, and governance at the center – not model specs

Agentic AI will reward the enterprises that move forward with discipline.

Start with one high-value workflow, define the outcomes you need, and choose a build path or a partner who can help you redesign, govern, and scale execution safely.





genpact

About Genpact

Genpact (NYSE: G) is the Agentic Operations company, where applied AI meets context-rich process intelligence. We run and transform mission-critical operations for global enterprises. Genpact's Agentic Operations are grounded in decades of operating core business processes at scale across finance, supply chain, banking, insurance, and more. Our flywheel of advanced technology, trusted data foundations, and deep ecosystem partnerships moves enterprises rapidly from experimentation to scale. AI agents bring speed and precision, human experts bring judgment, and together they deliver outcomes clients can measure.

Get to know us at genpact.com and on [LinkedIn](#), [X](#), [YouTube](#), and [Facebook](#).