Progress® Data Platform

The Agentic AI Maturity Model: Four Stages of Enterprise AI Growth

A practical guide to identifying where you are on your AI journey and how to grow AI's impact and value in your business using Progress® Agentic RAG solution and the Progress® Data Platform

WHITEPAPER





Executive Summary

Agentic AI is moving fast; faster than most organizations' data, governance and operating models can keep up with. Many teams start with a successful pilot and assume the hardest part is behind them. In practice, the opposite is true: the moment AI proves useful, usage expands, more data sources get pulled in, security and compliance expectations rise and leaders begin asking whether results can be trusted enough to automate decisions.

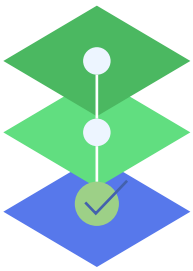
This is the point where "AI that demos well" and "AI that runs the business" diverge. And we're not the only ones that think this way. Gartner's recently released analysis showed that:

> **". . . organizations that report successful AI initiatives invest up to four times more (as a percentage of revenue) in these four foundational areas: data quality, governance, AI-ready people and change management."**

This whitepaper introduces four practical growth stages that customers typically move through as they adopt agentic AI. Each stage describes what customers are trying to accomplish, what usually breaks and what "good" looks like when you're ready to progress. It is designed as a self-assessment and a roadmap, so you can identify where you are today and plan the most reliable next step.

The guiding idea is simple—as your AI becomes more impactful, it must also become more *boring:* governed, grounded, explainable and repeatable. That's how organizations move from pilots to production, and from assistive experiences to operational outcomes.

© 2026 Progress. All Rights Reserved.



Why a “Growth-Staged Model” for Agentic AI?

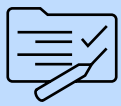
Most AI adoption guidance is either too abstract (with dictums like “be responsible”) or too technical (à la “choose a vector database”). But what customers actually need is a practical view of *how adoption changes the problem.*

Early on, a team wants speed and value. Later on, the organization needs consistency, shared context, policy enforcement, lineage—and an approach that scales across teams, data types and business-critical workflows. These are not separate initiatives; they are the natural progression of making AI useful.

This model helps you:



Diagnose your current stage without over-engineering too early

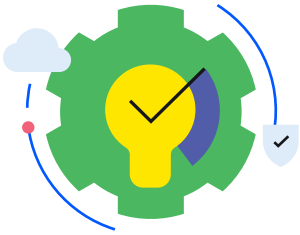


Align stakeholders (business owners, security, IT, data teams) around the next step



Invest in foundations at the right time, when AI success creates demand for reliability

© 2026 Progress. All Rights Reserved.



The Four Stages of Enterprise AI Growth

STAGE 1 — AI ACTIVATION:

Start Fast, Prove Value
Getting started with AI and making data AI-ready, focusing on a specific use case to produce quality outputs by leveraging unstructured data

STAGE 2 — AI EXPANSION:

Break Down Silos
Getting answers not only from unstructured data, but also from other sources using a reasoning engine

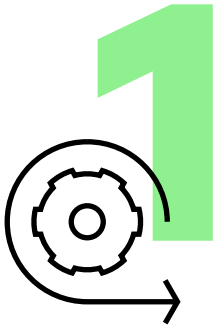
STAGE 3 — AGENTIC DATA FOUNDATION:

Reliable & Repeatable AI
Connecting your knowledge applications to create an agentic data foundation for specific use cases

STAGE 4 — SCALED AGENTIC INTELLIGENCE

Sustainable & Operational AI
Connecting your knowledge applications into an enterprise-level agentic knowledge platform

© 2026 Progress. All Rights Reserved.



AI Activation

What's happening in the business

At this stage, AI is being adopted to deliver immediate value, usually by one team or function. The use case is narrow and well defined; common examples include discovery, Q&A, summarization and decision support. The priority is proving AI can produce useful results with the data you already have.

This is often the moment when AI becomes "real" inside an organization. Thus, a working pilot changes the conversation from *"Should we do AI?"* to *"How quickly can we do more?"*

What customers need

Customers in the AI activation stage need fast deployment with minimal platform overhead. They want clear business outcomes and a short path to credibility, without the need for heavy re-architecture, long data migration projects or a large operational burden.

What "good" looks like

When AI activation is working well, the team can point to a small set of measurable outcomes (e.g., time saved, fewer searches, faster answers, higher satisfaction) and can explain how responses were produced at a basic level. People trust the system enough to use it regularly, but they still validate results.

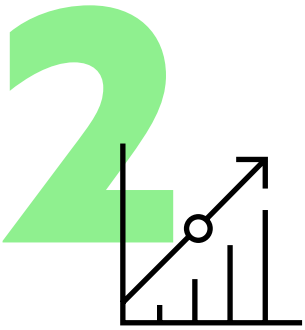
What it proves

AI can work with your data.

GUIDANCE FOR MOVING FORWARD

In the AI activation phase, lead with speed. The goal is not to build the perfect enterprise foundation, but to create tangible value quickly while learning where trust breaks as usage grows. If the pilot succeeds, plan early for what happens next: more users, more content, more sources, more scrutiny.

© 2026 Progress. All Rights Reserved.



AI Expansion

What's happening in the business

After early success, more teams want AI. Additional data sources are introduced, often in a siloed way: with new document repositories, shared drives, wikis, ticketing systems, CRM notes and structured systems that contain high-value context.

This expansion creates a new set of expectations. Security teams want stronger controls. Compliance leaders want defensibility. IT leaders want the system to be stable as traffic grows. Business leaders want consistent answers, not "It depends . . ."

What customers need

At this stage, customers need flexible, extensible integrations and a way to share context across agents and use cases. They need governance that is automated and policy-driven, plus reusable logic and retrieval orchestration capabilities so teams don't rebuild the same patterns repeatedly.

What "good" looks like

When AI expansion is working well, AI can draw from multiple sources without collapsing into inconsistency. You can apply basic policies (e.g., who can see what, which sources are allowed, which rules apply) without hand-coding everything per app. Teams can scale use cases without multiplying complexity.

What it proves

The need for a data foundation.

GUIDANCE FOR MOVING FORWARD

As scope expands, treat governance and retrieval orchestration as product capabilities, not custom one-offs. This is the point to introduce a data foundation as the mechanism that keeps agentic AI working as complexity increases.

© 2026 Progress. All Rights Reserved.



Agentic Data Foundation

What's happening in the business

At this stage, AI is now proven, and the conversation shifts from *"Can it work?"* to *"Can we trust it?"* because trust, compliance and lineage become non-negotiable. The organization needs knowledge to be shared across teams, not trapped in single apps or within single projects.

Leaders begin asking harder questions:

- Can we explain why the model answered the way it did?
- Can we show the sources and policy controls behind outputs?
- Can we reuse knowledge and logic across multiple applications?
- Can we demonstrate that results are consistent and defensible?

What customers need

Customers now require unified access to structured and unstructured data, with governance, security and policy enforcement built in. They need semantic enrichment to improve relevance and enable deeper responses. And they need a reusable knowledge layer that multiple agents and applications can rely on.

This is where the organization stops thinking in terms of an "AI app" and starts thinking in terms of a "shared knowledge layer for AI."

What "good" looks like

When an agentic data foundation is working well, teams can build new AI experiences faster because the hard parts, including access, permissions, governance, enrichment and reusable context, already exist. AI becomes more repeatable—not perfect, but consistent enough to support real workflows.

What it proves

A data platform is not mere infrastructure; it makes agentic AI enterprise-grade.

GUIDANCE FOR MOVING FORWARD

After proving AI works, make it reliable and repeatable. This is the stage where "boring AI" becomes a competitive advantage: governed, grounded, explainable and reusable.

© 2026 Progress. All Rights Reserved.



Scaled Agentic Intelligence

What's happening in the business

At this stage, AI is embedded into workflows. Agents don't just answer questions, they reason, orchestrate and act. AI outcomes now affect customers, revenue and operations.

This is where agentic AI becomes operational. Organizations require confidence that AI behavior is predictable under policy, that knowledge is current and trusted and that actions are auditable.

What customers need

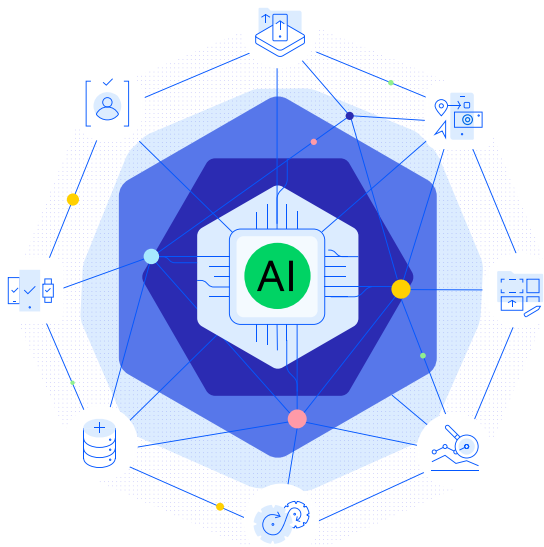
Customers need agentic AI operating on trusted, governed knowledge at scale. The knowledge layer becomes the backbone—increasingly invisible, but still the source of durability.

What "good" looks like

When scaled agentic intelligence is working well, teams can onboard new workflows without reinventing governance and data access. The organization can monitor outcomes, enforce policy, measure quality and evolve safely. AI is no longer an assistive experiment; it's part of how the business runs.

GUIDANCE FOR MOVING FORWARD

1. Prove value with agentic AI in a focused use case
2. Expand and scale by investing in the shared data and governance foundation



© 2026 Progress. All Rights Reserved.

Where Are You Today? A Practical Self-Assessment

Use the questions below to quickly identify your current stage. Most organizations will recognize themselves in more than one—choose the stage that best describes your dominant reality today:



AI Activation Signals

- One team owns the use case and the success metrics
- Data scope is limited (often a single content source)
- You are focused on speed and business value, not standardization
- Users validate results manually before acting



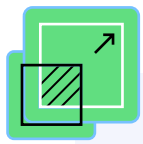
AI Expansion Signals

- Multiple teams are now requesting AI use cases
- You are adding new sources quickly and inconsistently
- Security, privacy and access control requirements are increasing
- You are starting to see inconsistent answers, duplicated logic or “prompt sprawl”



Agentic Data Foundation Signals

- AI is considered valuable, but trust is now the gating factor
- You need governance, policy enforcement and lineage
- You need a shared knowledge layer across teams and apps
- You need semantic enrichment to make outputs more accurate and explainable



Scaled Agentic Intelligence Signals

- AI is embedded into workflows and operational processes
- Agents coordinate actions across systems
- You require auditable outcomes, monitoring and repeatability
- Failure has business impact (i.e., revenue, compliance, customer experience)

© 2026 Progress. All Rights Reserved.

How to Move from One Stage to the Next



Activation to Expansion

The key shift is from proving value to managing scope. Your next step is to reduce the "cost of adding a new source" by standardizing how retrieval works and how policies are applied by:

- Defining a consistent pattern for ingestion and source selection
- Introducing baseline access controls and content segmentation
- Capturing feedback loops so quality improves over time



Expansion to Agentic Data Foundation

The key shift is from building more AI experiences to building a shared knowledge layer that multiple experiences can rely on. This is where governance and repeatability become foundational to:

- Unifying structured and unstructured access under consistent policies
- Establishing semantic enrichment (i.e., metadata, entities, relationships, meaning)
- Building reusable retrieval orchestration and knowledge services



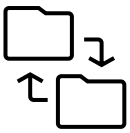
Foundation to Scaled Intelligence

The key shift is from reliable answers to reliable outcomes. This requires operationalizing agent behavior with monitoring, policy enforcement, auditability and lifecycle management through:

- Establishing outcome monitoring and quality measurement
- Enabling safe action (i.e., approval workflows, guardrails, audit trails)
- Scaling knowledge and governance across business functions

© 2026 Progress. All Rights Reserved.

A Simple 30/60/90 Plan



Stuck in AI Activation

30 hours Define your success criteria and make your pilot measurable. Keep the scope tight, but formulate an explanation of how responses are generated and what sources are used.

60 hours Add one additional source deliberately and observe where consistency breaks. This will reveal your biggest trust risks.

90 hours Document a repeatable pattern for ingestion, retrieval and evaluation so the pilot can become a template, rather than a one-off.



Stuck in AI Expansion:

30 hours Inventory your sources and identify the highest-risk content (i.e., sensitive data, regulated content, stale knowledge).

60 hours Standardize your retrieval approach and introduce policy-driven governance so every new use case doesn't require custom controls.

90 hours Define the design for a shared knowledge layer that supports multiple teams, and begin implementing semantic enrichment.

© 2026 Progress. All Rights Reserved.

A Simple 30/60/90 Plan Continued

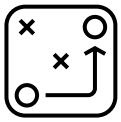


Stuck in Agentic Data Foundation

30 hours Establish trust requirements with what must be explainable, auditable and enforceable.

60 hours Implement unified access and policy enforcement across structured and unstructured sources.

90 hours Operationalize reuse with shared knowledge services, shared retrieval orchestration and shared evaluation.



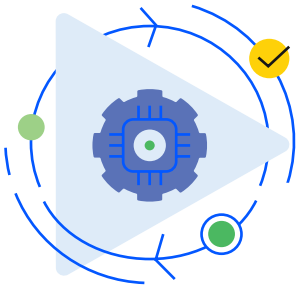
Stuck in Scaled Intelligence

30 hours Map agent workflows to business outcomes and define how you will measure success and failure.

60 hours Implement outcome monitoring, approval gates where needed and auditable trails for AI-driven actions.

90 hours Expand your platform model by onboarding additional workflows and teams using the same knowledge and governance foundation.

© 2026 Progress. All Rights Reserved.



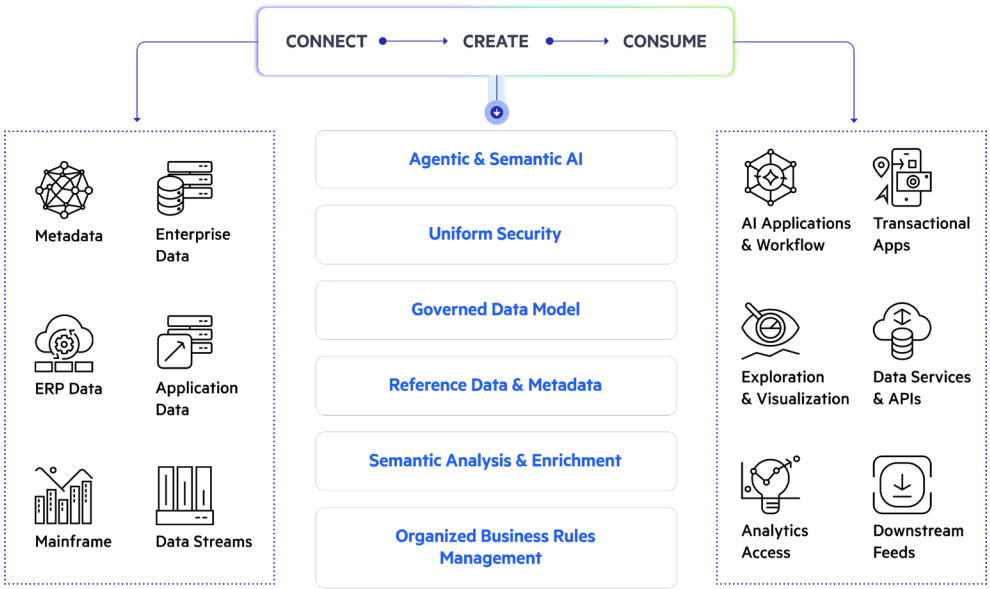
Architectural View: From AI App to Agentic Knowledge Platform

Early-stage agentic AI is often an application pattern: a model, a prompt and a narrow set of sources. However, as AI expands, this approach strains—because every new use case becomes a new mini-platform.

A more sustainable model is to separate knowledge foundation capabilities from front-end AI experiences. In this approach, applications and agents consume a shared layer that provides:

- Unified access to structured and unstructured knowledge
- Policy enforcement and security controls
- Semantic enrichment (e.g., meaning, relationships, domain context)
- Reusable retrieval orchestration
- Evaluation and measurement loops

When this layer is in place, scaling agentic AI becomes less about building more bespoke apps and more about onboarding new workflows onto a stable foundation.



© 2026 Progress. All Rights Reserved.



KEY TAKEAWAYS

Agentic AI maturity is not about adopting more AI features, it's about making AI dependable as it becomes more important and prevalent.

If your AI is going to scale, it needs to get more boring: **governed, grounded, explainable and repeatable.**

Glossary

Agentic AI — AI systems that can plan, reason, orchestrate steps and take actions across tools and workflows.

Agentic RAG — Retrieval-augmented generation designed for agentic behavior, where retrieval and reasoning support multi-step tasks and workflow execution.

Agentic Data Foundation — A shared knowledge layer that unifies access, governance, enrichment and reuse for multiple AI applications and agents.

Lineage — The ability to trace outputs back to sources, transformations and applied policies.

Policy Enforcement — Automated application of access controls, rules and governance to data and AI interactions.

Speak to an expert to assess where you are today and build a practical roadmap with the **Progress Agentic RAG solution** and the **Progress Data Platform.**

About Progress Software

Progress Software (Nasdaq: PRGS) empowers organizations to achieve transformational success in the face of disruptive change. Our software enables our customers to develop, deploy and manage responsible AI-powered applications and personalized digital experiences with agility and ease. Businesses of all sizes get a trusted provider in Progress, with the products, expertise and vision they need to turn AI disruption into a competitive advantage. Millions of developers and technologists at hundreds of thousands of organizations depend on Progress every day. Learn more at www.progress.com

/progresssw
/progresssw
/progresssw
/progress-software
/progress_sw_

© 2026 Progress Software Corporation and/or its subsidiaries or affiliates.
All rights reserved. Rev 2026/05 | RITM0364882

Progress