Progress® Data Platform

Practical AI

From AI-ready data to defensible answers, and the new economics that decide which enterprise AI scales and which quietly fails.

A PROGRESS® DATA PLATFORM FIELD GUIDE | MAKE AI PRACTICAL SERIES



Practical AI: From AI-Ready Data to Defensible Answers

First edition, 2026. Published by Progress Software Corporation, 15 Wayside Road, Burlington, MA 01803, United States. www.progress.com

© 2026 Progress Software Corporation and/or its subsidiaries or affiliates. All rights reserved.

Progress, the Progress logo, Progress Data Platform, Progress MarkLogic, Progress Semaphore, Progress Corticon and Progress Flow Conductor are trademarks or registered trademarks of Progress Software Corporation and/or its subsidiaries or affiliates in the United States and other countries. Other product or service names mentioned herein may be the trademarks of their respective owners.

The information, statistics, modeling assumptions, customer outcomes, market figures and pricing referenced in this book are drawn from publicly cited sources and Progress internal modeling. They are presented as directional analysis, not as commercial quotes or warranties. Real-world outcomes will vary by data, governance posture, vendor agreements, workload patterns and organizational maturity.

The numerical illustrations and modeling assumptions presented in this book are directional planning examples, not commercial quotes or guarantees of financial outcomes. Progress Data Platform pricing and implementation costs depend on the customer's specific requirements and are subject to enterprise quotation.

The document-processing illustration uses Claude Sonnet 5.5 at published base API pricing of $2 per million input tokens and $10 per million output tokens. The example measures model-processing compute costs only. Model pricing changes frequently, and results should be recalculated against current pricing and workload assumptions before being used in a business case. See Appendix B for the methodology.

Authors: Philip Miller, AI Strategist & Senior Product Marketing Manager, with Hinal Patel, Senior Product Marketing Specialist · Progress Software

How to cite this book: Miller, P. and Patel, H. (2026). Practical AI: From AI-Ready Data to Defensible Answers. Progress Data Platform · Make AI Practical Series.

Editor: Monica Delyani, Copywriter · Progress Software

Series: Make AI Boring (2025) · Agentic AI Growth Stages (2025) · Practical AI (2026, this volume)

Table of Contents

FOREWORD

A letter to the reader / 5

PART I | THE NEW REALITY OF ENTERPRISE AI

Why enterprise AI needs to become practical / 7

From AI activity to AI economics / 11

PART II | THE DEFENSIBLE ANSWER

The cost-per-defensible-answer formula / 13

Compute is the visible bill; trust is the real bill / 17

PART III | THE THREE HIDDEN COST CENTERS

Retrieval: Where plausible becomes defensible / 20

Remediation: The silent cost center of AI / 22

Human review is not a failure / 25

PART IV · BUILDING PRACTICAL AI

Agentic coding, MCP servers, skills and tools / 27

The neuro-symbolic advantage / 30

Practical AI workflows: Where the value shows up / 32

PART V | FROM MATURITY TO EXECUTION

How Progress® Data Platform changes the economics / 34

From agentic growth stages to practical execution / 37

© 2026 Progress. All Rights Reserved.

Table of Contents

PART VI | THE FIELD REPORTS

Three customer stories that show the thesis in practice / 39

PART VII | THE PRACTICAL AI PLAYBOOK

A 30/60/90 plan and how to talk about ROI / 42

The knowledge store that compounds / 46

The last word: Answers the business can defend / 50

APPENDICES

About the authors / 52

Glossary & key concepts / 53

Sources, methodology & further reading / 55

About the Progress Data Platform / 62

© 2026 Progress. All Rights Reserved.

Foreword

If your AI looks impressive, that's a problem.

A letter to the reader on the difference between AI that demonstrates well and AI the business will actually stand behind.

The first wave of enterprise AI was a wave of demos—brilliant ones. We saw models that could draft contracts, summarize research, generate code, answer questions and reason across documents we had never bothered to digitize. Boards were impressed. Pilots were funded. Pilots multiplied. Slides were made.

Three years on, the slides still look great. But it's reality that looks different.

Stanford HAI's 2026 AI Index, drawing on McKinsey survey data, reports that 88% of respondents said their organization used AI in at least one business function in 2025. Separately, McKinsey's State of AI (November 2025) found that only 39% of respondents reported earnings before interest and taxes (EBIT) impact at the enterprise level, and most of those said AI was responsible for less than 5% of it. And around one-third of respondents reported that their companies have begun to scale their AI programs.

The dots take a moment to connect. AI works. AI is everywhere. And yet, as a portfolio, it is failing to do the one thing the business actually pays for: produce trusted, defensible answers at a price that scales.

This book is about that gap.

Two earlier Progress whitepapers laid the groundwork: the Make AI Boring theme argued that the AI worth building is the AI you can govern, audit and rely on at 3 a.m.; the Agentic AI Growth Stages framework mapped how organizations move from first AI pilots to scaled agentic intelligence. And this Practical AI: From AI-Ready Data to Defensible Answers volume picks up where those left off, with a deceptively simple question:

> Once your data is AI-ready, what does it actually cost to produce an answer your business would defend?

Throughout these pages, you will meet a metric we have come to believe is the most useful single number in the enterprise AI conversation: cost per defensible answer. You will find the formula behind it, the four cost centers it exposes, the public evidence and adjustable assumptions that inform it and the customer

© 2026 Progress. All Rights Reserved.

stories that show it in motion. You will also find a working playbook, including a 30/60/90 plan, for moving from AI readiness to AI execution.

We will not pretend that AI will never get things wrong, that platforms remove governance work, or that one customer's outcome is everyone's benchmark. The point of practical AI is the opposite: lower the cost of trust, raise the share of outputs your business can stand behind and stop measuring AI by how much it generates.

If your AI looks impressive but the business still will not commit to its answers, this book is for you. If your AI is boring but you are ready to make it productive, this book is also for you. And if you are asked in the next quarter to explain to a CFO why the AI bill keeps growing while business outcomes do not, then read this twice.

> The first generation of enterprise AI was about possibility. The second was about reliability. The next one is about practical value.
>
> **THIS IS THE FIELD GUIDE FOR THAT TURN**

Philip Miller | AI Strategist & Senior Product Marketing Manager

Hinal Patel | Senior Product Marketing Specialist

Progress Software · October 2026



PART I

The New Reality of Enterprise AI

Adoption is settled, value is not. Why the enterprise conversation has to move from what AI does to what AI costs to be trustworthy.

1. **Why enterprise AI needs to become practical**
2. **From AI activity to AI economics**

Why enterprise AI needs to become practical.

AI feels transformative. So did the corporate intranet in 1999. The difference between technology that excites and technology that scales is not how impressive it looks; it's how dependable it becomes when nobody is watching.

CHAPTER 1 | ADOPTION VS. VALUE

Adoption is not the same as value.

Enterprise AI has crossed the adoption threshold. The question is no longer whether organizations will use AI because they already are. Stanford HAI's 2026 AI Index, drawing on McKinsey survey data, reports that 88% of respondents said their organization used AI in at least one business function in 2025, up from 78% in 2024. On the narrower question of generative AI (GenAI) specifically, 79% of respondents reported regular GenAI use in at least one business function. The AI Index notes that this survey data is self-reported and should be treated as directional. The trend is still unambiguous: AI is part of enterprise-operating reality, not a side experiment.



© 2026 Progress. All Rights Reserved.

Yet, adoption is not the same as value.

McKinsey's State of AI (November 2025) found that while AI use is widespread, only 39% of respondents reported EBIT impact at the enterprise level, and most organizations had not embedded AI deeply enough into workflows and processes to realize material enterprise-level benefits. The same research found that approximately one-third of respondents reported their companies had begun to scale AI programs.

88%
of respondents said their organization used AI in at least one business function

STANFORD HAI | 2026 AI INDEX

39%
reported EBIT impact at the enterprise level

MCKINSEY | STATE OF AI 2025

~1/3
said their companies have begun to scale AI programs

MCKINSEY | STATE OF AI 2025

That is the gap this book addresses.

The previous Progress® Data Platform campaign, Make AI Boring, argued that enterprise AI has to become reliable, governed, explainable and repeatable before it can safely scale. The Agentic AI Growth Stages whitepaper described the maturity journey from early AI activation to scaled agentic intelligence, where agents reason, orchestrate and act across business workflows. This volume continues that journey. It assumes the reader has accepted the need for boring AI and has started building, buying or evaluating an AI-ready data foundation.

The next question is more commercial: What is the business value of that foundation?

The answer is not simply "better AI;" it is better economics, better trust and better operational scale.

From the "wow" moment to the work

Enterprise value rarely comes from the moment something feels impressive. It comes later, when the system becomes dependable enough to change how work is done. That was the central point of Make AI Boring. The argument was

© 2026 Progress. All Rights Reserved.

not that AI should become dull, limited or less ambitious. It was that the most valuable enterprise systems do not feel like a breakthrough day to day. They feel dependable. They work the same way, for the same reasons, under the same controls, and can be trusted when outcomes matter. Practical AI is the next step.

Boring AI earns trust. Practical AI turns that trust into measurable outcomes.

This distinction matters. Many organizations can now create AI outputs. Far fewer can create AI-assisted decisions that are repeatable, auditable and economically viable. A chatbot answer may be useful. A legal summary may be plausible. A coding agent may produce working code. A research assistant may produce a convincing report. But enterprise leaders eventually ask harder questions.

01

Can we trust it?

Is the the output reliable enough to act on without a second opinion?

02

Can we prove it?

Is there evidence, provenance and lineage behind the claim?

03

Can we repeat it?

Does the same question produce the same quality of answer tomorrow?

04

Can we govern it?

Do permissions, policy and retention travel with the answers?

05

Can we audit it?

Can we reconstruct what the system saw and why it answered that way?

06

Can we afford it?

Does the unit cost hold up when usage multiplies by fifty?

FIGURE 1.1 | THE SIX QUESTIONS OF DEFENSIBILITY

A pilot proves that AI can produce an answer. A production workflow proves whether the business can defend it. That is where practical AI begins.

© 2026 Progress. All Rights Reserved.

THE SHORT ANSWER

Enterprise AI adoption is effectively universal, but financial impact is not. The gap is not model capability, it is the cost and reliability of the system that produces the answer: the data, the retrieval, the rules, the review and the evidence. Practical AI is the discipline of managing that system as an economic asset.

A pilot proves that AI can produce an answer. A production workflow proves whether the business can defend it or not.



© 2026 Progress. All Rights Reserved.

From AI activity to AI economics.

Most enterprise AI dashboards count what the system does: prompts run, tokens consumed, agents deployed. However, none of those numbers tell you whether the business can use what came out the other end. The metric we need is one the CFO can read aloud.

CHAPTER 2 | THE SHIFT

From activity metrics to outcome economics.

Most AI metrics are activity metrics. How many prompts were run? How many users adopted the assistant? How many documents were summarized? How many agents were deployed? How many tokens were consumed? How much did inference cost?

These numbers are useful, but they are incomplete. They describe the movement of the machine, not the value of the outcome. The danger is that organizations start measuring AI like a demo environment. They count what the system generates, not what the business can use.

This is why AI ROI remains difficult. McKinsey's 2025 survey found that AI use is now common, but most organizations are still navigating the transition from experimentation to scaled deployment; many may be capturing value in pockets without yet seeing enterprise-wide financial impact.

Gartner's April 2026 analysis reinforces the same point from a data and analytics perspective. Gartner reported that organizations with successful AI initiatives invest up to four times more, as a percentage of revenue, in foundational areas such as data quality, governance, AI-ready people and change management, than organizations experiencing poor AI outcomes. In the same research, only 39% of technology leaders said they were confident their enterprise's current AI investments would have a positive impact on financial performance.

TWO DIFFERENT 39% FIGURES—DO NOT CONFLATE THEM

McKinsey's 39% describes respondents reporting EBIT impact from AI use. Gartner's 39% describes technology leaders expressing confidence in the financial return of their current AI investments. Two unrelated surveys, two different populations, one coincidental number. They point the same way, but neither corroborates the other.

© 2026 Progress. All Rights Reserved.

THE HIDDEN WORK IS THE WORK

AI value is not created only by the model—it's created by the full answer-production system around the model. That system includes data quality, retrieval, access control, semantic context, rules, workflow design, review routing, auditability, monitoring, change management and human judgment.

In other words, the hidden work is the work and it is precisely where the invisible share of AI cost lives.

The question has changed.

Practical AI requires a different metric because the problem has changed. The question is no longer:
Can AI produce an answer?

But instead:
What does it cost to produce an answer the business would actually defend?

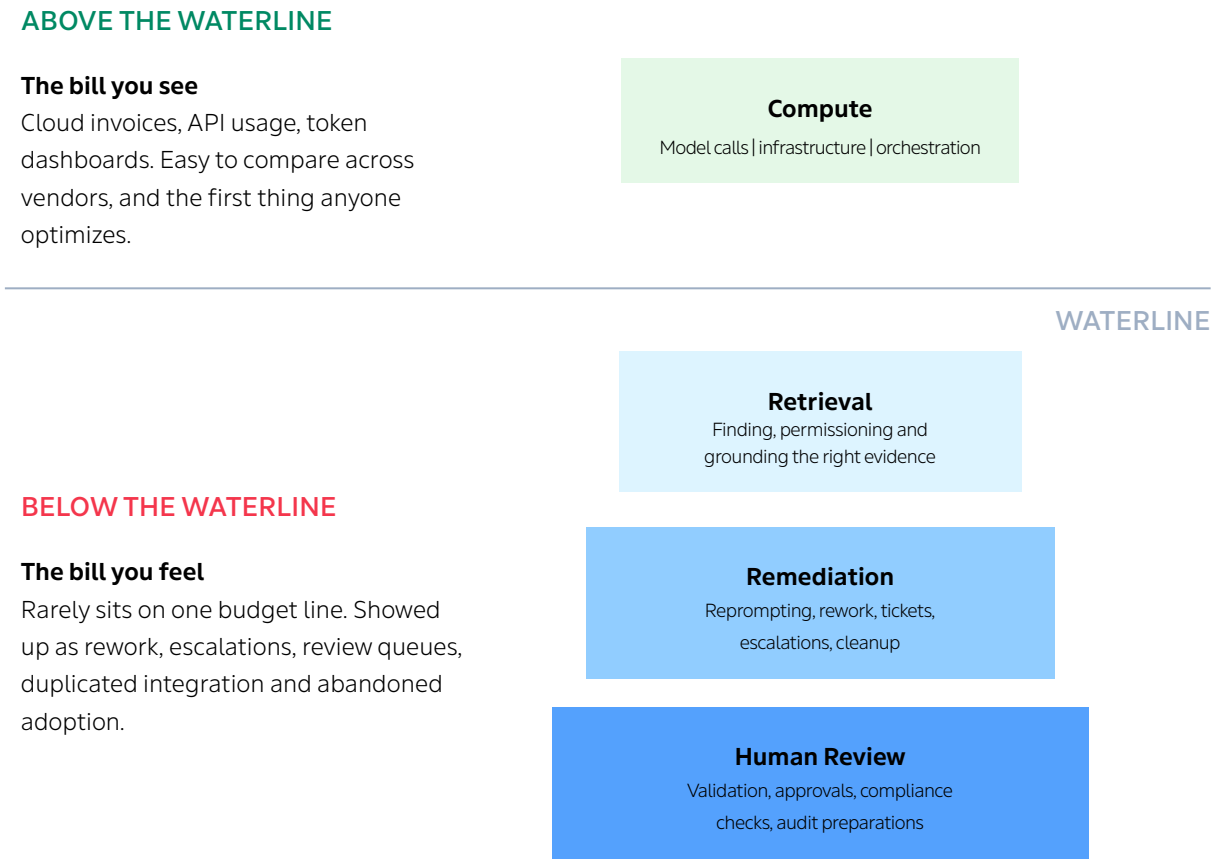

FIGURE 2.1 | THE AI COST ICEBERG

Compute is the visible bill. Retrieval, remediation and human review are where the real operational cost hides and where practical AI either earns its money or quietly bleeds it.

© 2026 Progress. All Rights Reserved.

THE SHORT ANSWER

Activity metrics measure the machine. Outcome economics measure the business. The shift from one to the other is what turns AI from a line item into an investment case, and it starts by counting only the answers you would actually defend.

The AI bill you see is compute. The AI bill you feel is retrieval, remediation and review.

PART II

The Defensible Answer

One formula, four cost lines and a denominator that changes what the CFO is being asked to fund.

3. **The cost-per-defensible-answer formula**
4. **Compute is the visible bill; trust is the real bill**

The cost-per-defensible-answer formula.

One formula, four cost lines and one denominator that changes the conversation, built deliberately simple so a CFO and a CTO can argue about it on a single sheet of paper.

CHAPTER 3 | THE FORMULA

One metric for the next phase of enterprise AI.
The formula is deliberately simple:

THE PRACTICAL AI EQUATION

compute + retrieval + remediation + human review

÷

answers you would actually defend

Both the naive AI path and the Progress Data Platform–grounded path use the same equation. They differ only in the inputs.

© 2026 Progress. All Rights Reserved.

This is not intended to replace every ROI model, business case or financial analysis. It is intended to expose the hidden economics of enterprise AI in a way business and technical leaders can discuss together. It works because it separates AI activity from AI usefulness.

What the four cost lines mean

1. Compute
Model calls, infrastructure, orchestration and processing: the cost most organizations can already see, that appears in cloud bills, API invoices, internal platform budgets and usage dashboards. Compute matters, but it is rarely the full story. If the same model, prompt and token volume produce different business outcomes depending on the quality of the data and retrieval layer, then compute is not the differentiator—the system around the model is.

2. Retrieval
Finding, accessing, contextualizing and grounding the right information: whether the system can connect to the right sources, apply the right permissions, understand business meaning and return evidence that is actually relevant. Retrieval is where many AI systems either become useful or become expensive. A weak retrieval layer gives the model more room to guess. A strong retrieval layer gives the model better context, better boundaries and better evidence.

3. Remediation
The cost of fixing what AI gets wrong: reprompting, rework, data fixes, support tickets, escalations, manual corrections, workflow exceptions and downstream cleanup. Bad AI does not just produce bad answers, it produces work.

4. Human review
These include expert validation, approvals, legal or compliance review, governance checks, risk review and audit preparation. In high-value workflows, review is not a failure, but part of the control system. The question is not whether humans should be involved, but whether they're reviewing everything because the AI cannot be trusted, or reviewing the right things because the system knows when to escalate.

The denominator that changes the conversation
A defensible answer is an output the business would stand behind: accurate enough, explainable enough, governed enough and useful enough to support a real decision or workflow. This is the denominator that matters.

© 2026 Progress. All Rights Reserved.

WHY THE DENOMINATOR IS THE WHOLE ARGUMENT

If AI produces 100,000 outputs but only 60,000 are trusted enough to use, the organization should not divide cost by 100,000. It should divide cost by 60,000. Suddenly, the unit economics of every productivity story look different.

WHAT CHANGES BETWEEN NAIVE & PROGRESS DATA PLATFORM-GROUNDED

COST LINE	NAIVE AI PATH	PROGRESS DATA PLATFORM-GROUNDED PATH
Compute	Compute costs depend on model selection, token consumption and workload complexity	Prepared, reusable context can reduce unnecessary input-token consumption and associated compute costs
Retrieval	Basic vector lookup	Governed retrieval, plus amortized platform fee
Remediation	High cost per error × high error count	Low cost per error × low error count
Human review	High share reviewed	Low share reviewed—Progress® Corticon® rules engine routes only the genuinely uncertain
Defensible answers	N × (1 − ungrounded hallucination rate)	N × (1 − grounded hallucination rate)

For this illustrative model, defensible-answer yield is approximated using one minus the assumed hallucination rate. This serves as a simplified proxy and does not independently measure governance, compliance, provenance, auditability or business usefulness. Organizations should validate defensibility against their own workflow requirements and acceptance criteria.

© 2026 Progress. All Rights Reserved.

The illustrative scenario: Reducing compute costs through AI-ready data

Preparing enterprise content for AI can reduce the number of tokens required to process information, creating measurable savings in one component of the cost-per-defensible-answer equation: compute.

In an illustrative document-processing scenario, a 20-page document processed with Claude Sonnet 5.5 has an estimated model-processing cost of $0.215 when read as a raw PDF, compared with $0.017 when processed as prepared Markdown. Under the modeled assumptions, this represents approximately 92% lower compute costs per document.

At a volume of 10,000 documents per month, the estimated annual savings are approximately $23,700.

These figures are illustrative and based on adjustable assumptions, including token consumption and published model pricing. They measure model-processing costs only and do not account for the full costs of retrieval, remediation, human review or document preparation. They should not be interpreted as a measured reduction in cost per defensible answer.

RAW PDF

$0.215 compute cost per document

PREPARED MARKDOWN

$0.017 compute cost per document

Illustrative scenario: 10,000 documents per month, 20 pages per document, Claude Sonnet 5.5 at published pricing of $2 per million input tokens and $10 per million output tokens. Assumes 4,000 PDF tokens per page, 200 Markdown tokens per page, 500 output tokens per call and a directional 1.30 tokenizer multiplier. Caching, batch discounts and document preparation costs are excluded. See Appendix B for methodology and assumptions.

The strategic point is larger than any single number.

When AI is weakly grounded, economics depend on more review. When AI is strongly grounded, economics depend on better context.

THE SHORT ANSWER

Cost per defensible answer divides the full operational cost of AI—compute, retrieval, remediation and human review—by the number of outputs the business would actually stand behind. It's a single number that exposes whether an AI program is getting cheaper to trust or simply cheaper to run.

© 2026 Progress. All Rights Reserved.

Compute is the visible bill; trust is the real bill.

Token prices are falling. AI bills are growing. The two facts are not in conflict, but describe an industry that has solved one cost line and not started on the others.

CHAPTER 4 | THE ECONOMICS BEHIND THE HEADLINES

Why the AI bill keeps growing while inference gets cheaper.

The market tends to talk about AI cost as if it is mostly an inference problem. Which model is cheaper? Which provider has the lower token price? Which prompt uses fewer tokens? Those questions matter. But they are not the whole economic picture. Token prices can fall while the cost of trustworthy AI remains high.

That happens when the organization still has to pay people to validate outputs, fix bad answers, reconcile sources, check permissions, explain decisions and clean up the operational consequences of weak AI. It can also happen because multistep agentic workflows make repeated model and tool calls. Lower per-token prices do not necessarily mean lower total inference cost.

This is why practical AI economics are not solved by choosing a cheaper model alone. A cheaper model can reduce compute cost while leaving retrieval, remediation and review untouched.

The better question

Where does the work go after the model responds?

- If every answer has to be checked, human review cost rises with usage
- If every poor answer creates a ticket, remediation rises with usage
- If every new use case requires a custom integration, retrieval cost rises with scope
- If every team builds its own policies, governance fragments as adoption expands

That is how AI becomes expensive even when compute looks manageable.

© 2026 Progress. All Rights Reserved.

RAW PDF

$0.215

compute cost per document

PREPARED MARKDOWN

$0.017

compute cost per document

Compute	Prepared content can reduce input-token consumption and associated processing costs
Retrieval	Governed, reusable context can improve access to relevant evidence
Remediation	Better grounding can reduce the effort required to correct errors
Human review	Policy-aware workflows can focus expert review on consequential decisions
Defensible answers	Yield must be evaluated against workflow-specific acceptance criteria

The figures above model document-processing compute only, not the full cost per defensible answer.

The other cost categories are directional opportunities, not measured savings.

Source : Progress PDF vs. Markdown calculator v3 · Claude Sonnet 5.5 · 20 pages per document

FIGURE 4.1 · WHERE THE COST ACTUALLY LIVES

Document preparation illustrates one measurable opportunity to reduce compute costs. Retrieval, remediation and human review represent additional opportunities to improve the economics of enterprise AI, but their financial impact is not quantified in the accompanying PDF vs. Markdown calculator. A complete cost-per-defensible-answer analysis requires workload-specific measurement across all four cost categories and defensible-answer yield.

© 2026 Progress. All Rights Reserved.

A NOTE ON MODEL PRICING

The document-processing example holds the model and assumed output-token volume constant while varying the input-token requirements of raw PDF and prepared Markdown content.

Worked examples use Claude Sonnet 5.5 at published pricing of $2 per million input tokens and $10 per million output tokens. The calculator applies an adjustable 1.30 tokenizer multiplier to Claude models to account for differences in tokenization. Because image-based PDF processing depends on factors such as page resolution, this adjustment is directional rather than a universal measurement.

Lower input-token consumption can reduce compute costs, but these savings represent only one component of the broader cost-per-defensible-answer equation.

Actual results will vary with document complexity, token usage, pricing agreements, preparation costs and workload characteristics. Recalculate scenarios against current vendor pricing before incorporating results into a business case.

THE SHORT ANSWER

Falling token prices reduce one of four cost lines. Retrieval, remediation and human review all scale with usage rather than with model price, so an AI program can get cheaper per call and more expensive per trusted answer at the same time. That is the paradox practical AI is designed to resolve.

Token prices can fall while the cost of trustworthy AI remains high.

© 2026 Progress. All Rights Reserved.

PART III

The Three Hidden Cost Centers

Retrieval, remediation and human review—three lines that never appear on an AI invoice and decide whether the program scales.

5. **Retrieval: where plausible becomes defensible**
6. **Remediation: the silent cost center of AI**
7. **Human review is not a failure**

Retrieval: where plausible becomes defensible.

Large language models (LLMs) are good at fluency; that is part of their power and part of their risk. A fluent answer can feel right before it is proven right. Retrieval is the discipline that turns fluency into evidence.

CHAPTER 5 | RETRIEVAL

Not all retrieval is equal.
Retrieval-augmented generation, or RAG, is widely used to connect GenAI models to external knowledge sources so they can produce more relevant and higher-quality responses. IBM describes RAG as an architecture for optimizing AI model performance by connecting models with external knowledge bases; it lists benefits including access to current domain-specific data, lower hallucination risk and increased user trust, while noting that RAG does not make a model error-proof. NVIDIA similarly describes RAG as a technique for improving the accuracy and reliability of GenAI models with information fetched from specific and relevant data sources, and notes that source citation can help users check claims.

© 2026 Progress. All Rights Reserved.

But not all retrieval is equal.

Vector search can find similar content, but that does not always mean it has found the right evidence. In enterprise settings, relevance depends on more than similarity. It depends on:

- ✓ Is this source authoritative?
- ✓ Is it current?
- ✓ Is the user allowed to see it?
- ✓ Does the answer require structured data, unstructured content or both?
- ✓ Does the retrieved passage reflect the right policy, geography, product, customer segment or regulatory context?
- ✓ Does the system understand the business meaning of the terms being used?
- ✓ Does it know when sources conflict?
- ✓ Does it know when not to answer?

This is where an AI-ready data foundation changes the economics. When enterprise data is unified, enriched with business meaning and governed consistently, retrieval improves. Better retrieval reduces guessing; less guessing reduces remediation; and clearer evidence reduces review burden. More outputs become defensible.

The flow of meaning

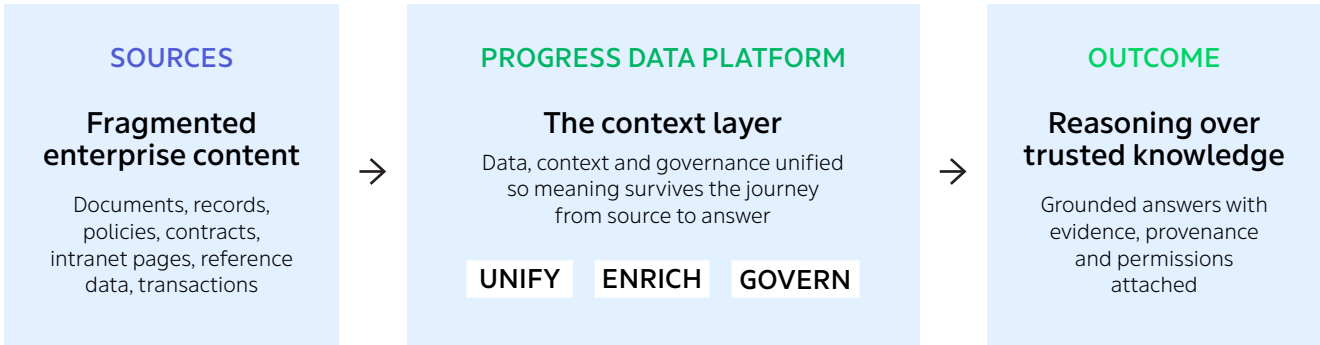

FIGURE 5.1 | THE FLOW OF MEANING

The Progress Data Platform unifies data, context and governance into a flow of meaning so AI models can reason over trusted enterprise knowledge rather than retrieve disconnected fragments.

© 2026 Progress. All Rights Reserved.

The Progress Data Platform frames this as unifying data, context and governance into a flow of meaning, so AI models can reason over trusted enterprise knowledge rather than simply retrieve disconnected fragments. Semantic enrichment, metadata, business rules, lineage and governed access are the core capabilities for explainable and defensible AI outcomes.

THE SHORT ANSWER

Vector similarity finds content that looks like the question. Governed retrieval finds evidence the business is allowed to use, can prove the origin of and can defend later. The difference between the two is where most of the hidden cost in enterprise AI is created or avoided.

Practical AI does not start with the prompt. It starts with what the system can know, retrieve and prove.

Remediation: the silent cost center of AI.

Bad AI is not a quality issue; it is an operating cost, one that hides in lost productivity, abandoned adoption and risk mitigation. The most expensive AI is that which quietly produces work for everyone else.

CHAPTER 6 | REMEDIATION

What happens after trust breaks.
Remediation is what happens after trust breaks. It is the analyst checking the source because the answer sounded odd. It is the support agent correcting a chatbot response. It is the developer fixing code the assistant generated confidently, but incorrectly. It is the compliance team reviewing a claim that should never have been made. It is the data team cleaning up source content because the AI exposed the mess. It is the business user who stops trusting the tool and goes back to manual work.

This is why bad AI is not merely a quality issue, it is an operating cost.

Remediation tends to be undercounted because it rarely sits on a single budget line. It appears as lost productivity, duplicated effort, escalations, review queues, manual reconciliation, abandoned adoption and risk mitigation. It is also one of the most important areas where practical AI can create value.

© 2026 Progress. All Rights Reserved.

Hallucination rates by use case

The Progress methodology models remediation as the cost of chasing down a bad answer, including reprompting cycles, support tickets, data fixes and escalations. For illustration, Progress uses adjustable modeling assumptions for each use case. Public research informs the direction of those assumptions, but published evaluations use different tasks, datasets and definitions and should not be treated as one comparable benchmark. Replace these assumptions with results from your own workload wherever possible. your own workload wherever possible.

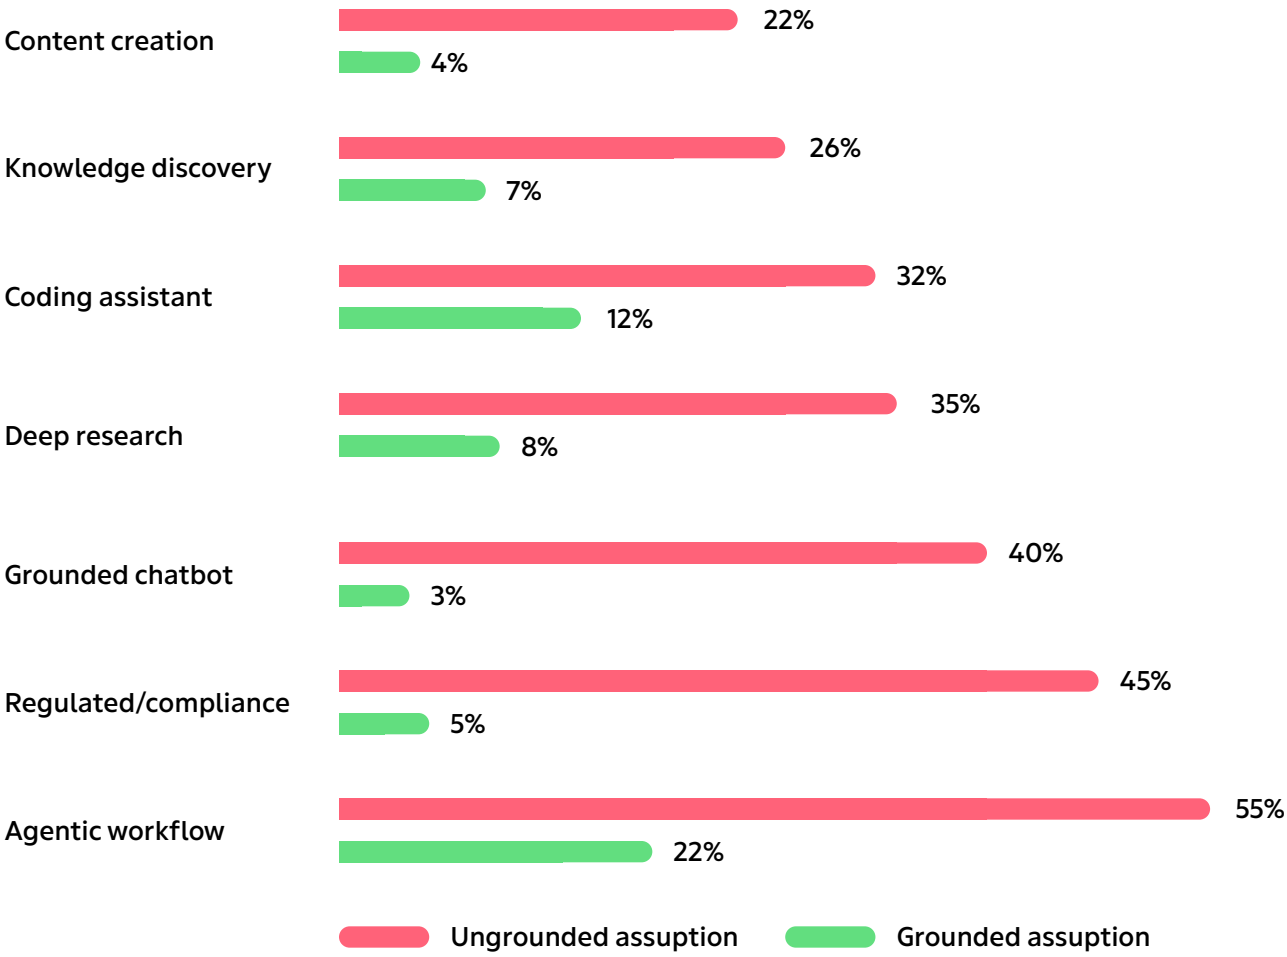

FIGURE 6.1 | HALLUCINATION RATES · UNGROUNDED VS. GROUNDED

Illustrative Progress modeling assumptions for ungrounded and grounded AI workflows, plotted on a shared scale. Public evaluation research supports workload-specific testing of factuality, retrieval and groundedness, but does not provide one universal rate across these use cases. Validate the assumptions against your own sources, tasks, review rules and risk tolerance before using them in a business case.

© 2026 Progress. All Rights Reserved.

Where remediation costs compound

In low-stakes content workflows, remediation may be a copy edit. In regulated workflows, remediation may require legal review, audit preparation or customer correction. In agentic workflows, remediation can be more expensive because errors compound across steps. A wrong answer is one thing, but a wrong action taken by an agent inside a workflow is another.

This is why agentic AI raises the stakes. As the Agentic AI Growth Stages whitepaper argues, when AI becomes more embedded in workflows, agents do not just answer questions; they reason, orchestrate and act, and outcomes can affect customers, revenue and operations. Practical AI reduces remediation by improving the system before the answer reaches the user.

1. Better data quality

Trusted sources, reduced inconsistency and lineage that explains every transformation

2. Better retrieval

Permission-aware, policy-aware and context-aware retrieval over a curated knowledge graph

3. Better escalation

Rules and confidence signals that route uncertainty to the right human reviewer instead of all reviewers

The aim is not to pretend AI will never be wrong, but to catch more issues earlier, route uncertainty appropriately and reduce the cost of correction, turning remediation from a back-of-house tax into a controlled flow.

THE SHORT ANSWER

Remediation is the operational cost of correcting AI outputs: reprompting, rework, tickets, escalations and cleanup. It is usually the largest uncounted line in an AI program because it never appears on a single budget. Grounding, rules and targeted escalation are the three levers that reduce it.

© 2026 Progress. All Rights Reserved.

Human review is not a failure.

There is a persistent myth in enterprise AI that the goal is to remove humans from the process. In high-value work, the goal is the opposite—to scale human expertise, not eliminate it.

CHAPTER 7 | HUMAN REVIEW

Targeted judgment, not blanket validation.
In some workflows, automation can and should reduce manual effort dramatically. In high-value, regulated or decision-heavy work, the goal is not to remove humans from decisions, but to scale human expertise.

Human review is not a sign that AI has failed; it is often what makes AI usable. The issue is not human involvement, but instead, undifferentiated human involvement.

If every AI output needs the same level of review, AI becomes a labor amplifier.

It produces more material for humans to check. It increases throughput without reducing bottlenecks. It looks productive at the front end and expensive at the back end.

Practical AI changes the pattern. It uses data, rules, confidence signals, policy checks, audit trails and workflow design to decide which outputs can move forward, which need review and which should not be produced at all.

What high performers actually do
This aligns with McKinsey's 2025 findings about high-performing AI organizations. McKinsey reported that high performers are more likely than others to say their organizations have defined processes to determine how and when model outputs need human validation, and that is one of the management practices that differentiates them.

This is the practical version of human-in-the-loop. Not humans as a blanket safety net, but humans as targeted judgment, escalation and accountability. The Progress Data Platform supports this by providing capabilities for governed data access, policy-aware workflows, business rules, auditability and controlled integration with downstream AI applications and agents.

© 2026 Progress. All Rights Reserved.

THE REVIEWER-ROUTING PATTERN

A well-designed Progress Data Platform-grounded workflow can route outputs that require expert judgment to a human reviewer and let lower-risk outputs follow approved paths. For illustration, the Progress Data Platform-grounded path assumes 6% to 18% of outputs are reviewed, compared with 20% to 60% for the naive path. These are adjustable Progress modeling assumptions, not universally observed rates. Reviewers can spend less time on routine checks and more time on consequential judgment calls, exceptions, policy interpretation and accountability decisions.

NAIVE PATH | ASSUMPTION

20–60%

of outputs reviewed

- Review applied broadly because confidence is unknown
- Expert time absorbed by routine verification
- Review cost rises in step with adoption

PROGRESS DATA PLATFORM-GROUNDED PATH | ASSUMPTION

6–18%

of outputs reviewed

- Rules and confidence signals decide what escalates
- Expert time concentrated on genuine judgment calls
- Lower reviewed share helps moderate review costs as output volume grows

Reviewed-share figures are adjustable. These figures are illustrative Progress modeling assumptions, not observed rates across a population of organizations. Set them from your own workflow before using them in a business case.

The better the system, the more valuable human review becomes. Experts spend less time checking obvious things, and more time on judgment, exceptions, policy interpretation and decisions where accountability matters.

THE SHORT ANSWER

The measure of a mature AI workflow is not how little human review it uses, but how well it targets the review it does use. Blanket validation makes AI a labor amplifier; rule-driven escalation makes it a capacity multiplier.

The goal is not to remove humans from decisions. The goal is to scale expertise and confidence.

© 2026 Progress. All Rights Reserved.

PART IV

Building Practical AI

Agents, tools, symbols and workflows—what practical AI is actually made of once the economics case is accepted.

8. **Agentic coding, Model Context Protocol (MCP) servers, skills and tools**
9. **The neuro-symbolic advantage**
10. **Practical AI workflows: where the value shows up**

Agentic coding, MCP servers, skills & tools.

The business value of agentic coding is not that AI can write code faster, it's that AI can move teams through repeatable delivery workflows with more context, fewer handoffs and stronger controls.

CHAPTER 8 | AGENTIC CODING

From code generation to governed delivery.
AI-assisted software development is often discussed in terms of speed. Can AI write code faster? Can it generate tests? Can it explain a codebase? Can it modernize old applications? Those are useful questions, but they are still too narrow.

The business value of agentic coding is not simply that AI can generate code, it's that it can help teams move through repeatable delivery workflows with more context, fewer handoffs and stronger controls. That shift depends on three things:

1. Tools
Allow agents to inspect, run, test, search, edit or interact with approved systems

2. Skills
Package repeatable workflows, standards and domain-specific guidance

3. Controlled connectivity
Gives agents access to the right context and capabilities without turning every integration into a one-off security exception

© 2026 Progress. All Rights Reserved.

Why MCP-style connectivity matters

This is where MCP-style connectivity comes in—the Model Context Protocol documentation describes it as an open-source standard for connecting AI applications to external systems, including data sources, tools and workflows. The MCP specification describes servers that can offer resources, prompts and tools, and stresses that these capabilities introduce important security and trust considerations because they can involve arbitrary data access and code execution paths.

That last point is critical: Agentic coding becomes practical only when the agent's power is matched by control. The MCP specification explicitly calls for user consent and control, data privacy, tool safety, explicit consent before tool invocation and appropriate access controls and data protections.

Skills provide another part of the pattern. OpenAI's Codex documentation describes agent skills as packages of instructions, resources and optional scripts that help Codex follow workflows reliably, using progressive disclosure so the agent loads the full skill only when needed. Anthropic's documentation similarly describes skills as reusable resources that give an agent domain-specific expertise.

What this looks like in practice

The practical value is that engineering teams can turn repeated work into governed patterns via the:

- ✓ Validation of a code change
- ✓ Running of required tests
- ✓ Following of security review requirements
- ✓ Preparation of a release
- ✓ Updating of documentation
- ✓ Application of repository standards
- ✓ Handling of exceptions and approvals

© 2026 Progress. All Rights Reserved.

AN ILLUSTRATIVE DATA POINT

OpenAI's published example from its Agents SDK repositories reported that repo-local skills, repository policy files and GitHub Actions helped turn recurring engineering work—verification, release preparation, integration testing and pull-request review—into repeatable workflows. The report cited 457 merged pull requests across the two Agents SDK repositories combined between December 1, 2025 and February 28, 2026, compared with 316 in the prior three-month period.

That is an illustrative example from OpenAI's own repositories, not a general result. It shows the kind of workflow-level value agentic coding is moving towards, in a codebase whose maintainers also wrote the tooling.

The pattern, not the toolchain

For enterprise leaders, the point is not the specific toolchain, but the pattern. AI becomes more useful when it can operate with approved data, tools, workflows, rules, review paths and a trusted platform foundation.

This is where the Progress Data Platform fits—it provides the governed data, semantic context, business rules, policy-aware connectivity and workflow foundations that help agentic coding move from ad hoc assistance to controlled delivery.

THE SHIFT IN VALUE

The value is not that AI writes more code, but instead:

- ✓ Developers spend less time searching for context
- ✓ Teams reuse approved patterns
- ✓ Reviewers see evidence and rationale
- ✓ Security and governance are embedded earlier
- ✓ Testing and validation become part of the workflow
- ✓ Delivery becomes faster without bypassing control

That is practical agentic coding.

THE SHORT ANSWER

Agentic coding pays off at the workflow level, not the keystroke level. Tools give agents reach, skills give them repeatable method and MCP-style connectivity gives them governed access. Without the third the first two become a security exception, rather than a delivery capability.

© 2026 Progress. All Rights Reserved.

The neuro-symbolic advantage.

The enterprise does not need AI that only sounds intelligent, it needs AI that can reason with business meaning: flexibility from models, explicit meaning from semantics and discipline from rules.

CHAPTER 9 | NEURO-SYMBOLIC AI

Models bring fluency; the platform brings meaning.
In simple terms, neuro-symbolic AI combines the flexibility of statistical AI, such as machine learning and language models, with symbolic knowledge and reasoning: rules, metadata, taxonomies, ontologies, knowledge relationships and policy logic. IBM Research describes neuro-symbolic AI as augmenting and combining statistical AI with symbolic knowledge and reasoning.

For enterprise audiences, the useful point is not the academic label, but the operating model. Language models are good at language, synthesis, pattern recognition and flexible reasoning. Symbolic systems are good at explicit meaning, rules, relationships, constraints, provenance and repeatability. Practical AI needs both.



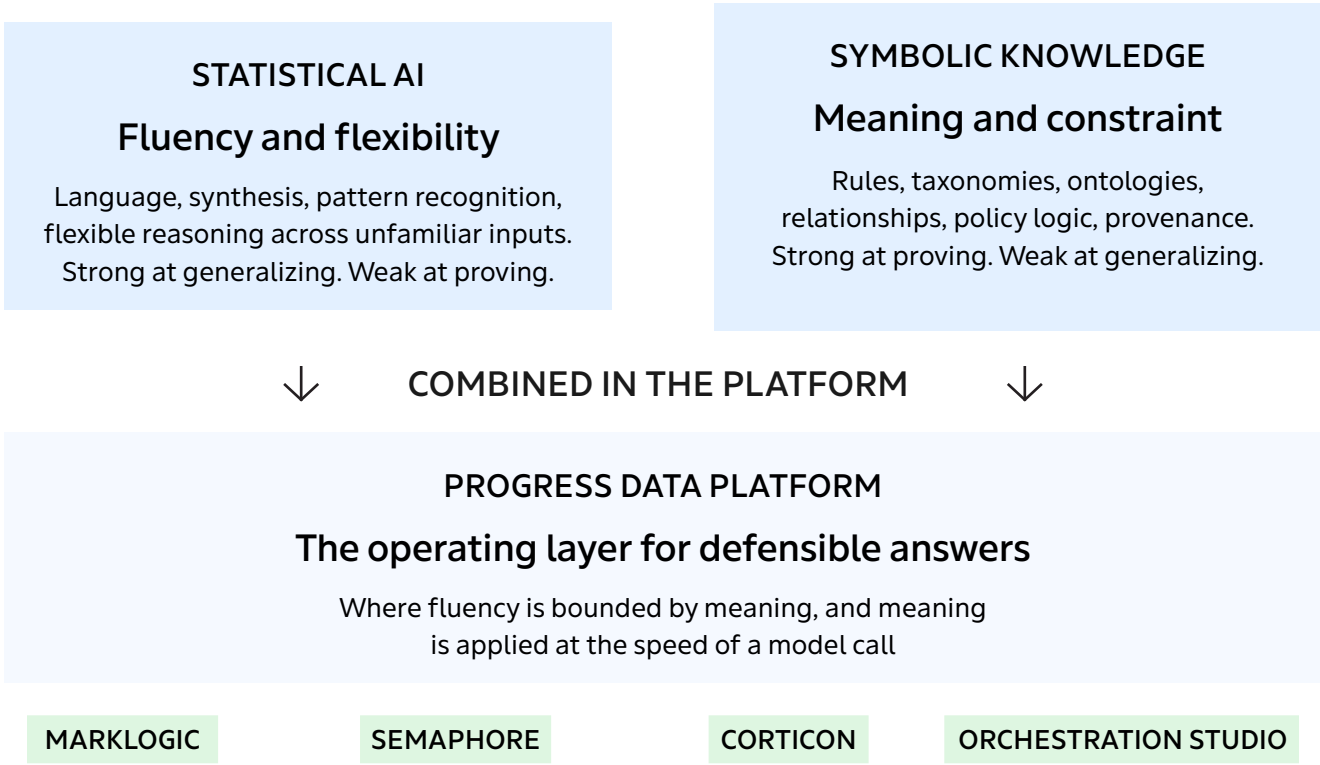

FIGURE 9.1 · THE NEURO-SYMBOLIC STACK

IBM Research describes neuro-symbolic AI as augmenting and combining statistical AI with symbolic knowledge and reasoning. The Progress Data Platform is positioned as the operating layer that makes that combination practical for enterprise workflows.

Why this matters in the enterprise

This is especially important in enterprise settings because business knowledge is not just text. It is structured data, unstructured content, policies, reference data, definitions, relationships, prior decisions, permissions, regulatory constraints and tacit domain logic.

- If AI only sees words, it guesses.
- If AI sees words plus meaning, rules and context, it can produce outputs that are easier to trust.

The Progress Data Platform is positioned around this combined foundation: data, context and governance unified into a platform that helps AI systems produce explainable, auditable and defensible outcomes.

This is the practical neuro-symbolic story:

- The model brings fluency and reasoning.
- The platform brings meaning, control and trust.
- Together, they produce answers the business can defend.

© 2026 Progress. All Rights Reserved.

THE SHORT ANSWER

Neuro-symbolic AI pairs statistical models with explicit symbolic knowledge, including rules, taxonomies, ontologies and policy logic. In enterprise terms, it is the difference between a system that generalizes confidently and a system that can also show it's working. Practical AI needs both halves.

Practical AI workflows: where the value shows up.

Practical AI is not a technology category, but a list of workflows. The most valuable ones share a pattern: knowledge-intensive, context-dependent and expensive to validate manually. .

CHAPTER 10 | WHERE IT SHOWS UP

Seven workflows where practical AI earns its money.
Practical AI should not be described as a technology category alone; it should be described through the work it improves. The most valuable use cases share a pattern: knowledge-intensive, context-dependent and expensive to validate manually. They are exactly the kinds of workflows where better grounding, governance and review design can reduce the cost per defensible answer.

A. Knowledge discovery
From search to evidence-based answer production. Faster time to answer, less duplication and more reuse of institutional knowledge, when retrieval respects permissions and context.

B. Deep research
For analysts, legal, product and strategy teams. Practical AI preserves evidence and shows provenance, so experts spend their time on judgment rather than source reconciliation.

C. Customer and employee assistants
Assistants fail when they answer from stale or unauthorized context. They become practical when grounded in current, governed enterprise knowledge and able to escalate uncertainty.

© 2026 Progress. All Rights Reserved.

D. Regulated question answering

For legal, financial services, healthcare, life sciences and public sector. Practical AI combines source-grounded answers, policy-aware retrieval, human review, audit trails and controlled access to produce faster answers, with evidence.

E. Content creation

The enterprise problem is not generating content; it's holding to approved claims, current messaging, brand requirements, legal constraints and source evidence. Faster operations with less compliance rework.

F. Agentic coding and DevOps

AI coding assistants become more valuable when they operate within approved delivery practices, repository rules, test workflows, security standards and governed data access.

G. Agentic workflows

Agents recommend, route, draft, transform, validate and, in bounded cases, act. They need trusted data, policy checks, tool boundaries, human approvals and monitoring. The outcome is operational AI that can scale without becoming operational risk.

Across all seven use cases, the common pattern holds: practical AI lowers cost per defensible answer by attacking retrieval, remediation and review, not just compute.

THE SHORT ANSWER

The workflows where practical AI pays back fastest are the ones where an answer is expensive to verify by hand: regulated question answering, deep research, knowledge discovery, grounded assistants, governed content, agentic coding and bounded agentic workflows. If verification is cheap, so is the upside.

© 2026 Progress. All Rights Reserved.

PART V

From Maturity to Execution

Where the platform enters the formula, and how the maturity journey changes the question a leader should be asking at each stage.

11. How Progress Data Platform changes the economics
12. From agentic growth stages to practical execution

How Progress Data Platform changes the economics.

It's not a magic layer, a replacement for models or a claim that AI becomes risk-free. It's a foundation that helps reduce the cost and risk of producing answers the business can defend.

CHAPTER 11 | THE PLATFORM VALUE

Across the formula: six ways Progress Data Platform contributes.
The Progress Data Platform provides a trusted data and platform foundation for practical AI execution. Its value maps directly to the formula.



PROGRESS DATA PLATFORM CONTRIBUTION	WHAT CHANGES	BUSINESS VALUE
Lower retrieval friction	Connects structured and unstructured data, enriches with business meaning, provides governed access to context	Faster onboarding of AI workflows, less repeated integration work
Higher answer quality	Semantic enrichment, shared metadata, reference data and governed data models so AI understands context, not just keywords or vectors	Fewer plausible-but-wrong outputs, more answers with useful evidence
Reduced remediation	Better grounding, policies, rules and workflow checks reduce outputs that need correcting after the fact	Less rework, fewer escalations, lower hidden cost
Targeted human review	Business rules, policy-aware workflows and governed escalation patterns route the right outputs to human experts	Review capacity focused on exceptions and judgment, not blanket validation
Stronger defensibility	Lineage, provenance, access control, metadata, semantic context and auditability help explain what happened and why	Confidence in regulated, high-value, customer-facing AI workflows
Better scale	Shared knowledge layer so teams build new AI experiences faster; access, permissions, governance, enrichment and reusable context already exist	Reuse: teams stop rebuilding the same controls and integrations for every new use case

© 2026 Progress. All Rights Reserved.

Both sides of the equation

In practical AI terms, Progress Data Platform changes both sides of the formula.

NUMERATOR ↓

Less cost per answer

compute + retrieval + remediation + human review

- Lower retrieval waste through reusable governed context
- Lower remediation through better grounding and rules
- Less unnecessary review through targeted escalation

DENOMINATOR ↑

More answers you would defend

- Evidence and provenance attached to the output
- Policy and permissions enforced at retrieval time
- Semantic meaning, so context is understood, not guessed

That is the economic story—not a replacement for models or a promise—a foundation.

THE SHORT ANSWER

A platform changes AI economics in two directions at once by reducing the numerator by removing repeated retrieval, remediation and review work and raising the denominator by increasing the share of outputs the business would stand behind. Point solutions typically move only one.

© 2026 Progress. All Rights Reserved.

From agentic growth stages to practical execution.

Maturity is not about adopting more AI, but about reducing the cost and risk of trusted outcomes at every stage of the journey.

CHAPTER 12 | MATURITY MEETS PRACTICAL AI

Four stages, four practical questions.
The Agentic AI Growth Stages framework describes four stages of enterprise agentic AI growth: AI Activation; AI Expansion; Agentic Data Foundation; and Scaled Agentic Intelligence. Practical AI fits across those stages, but it matters most after the first signs of success.

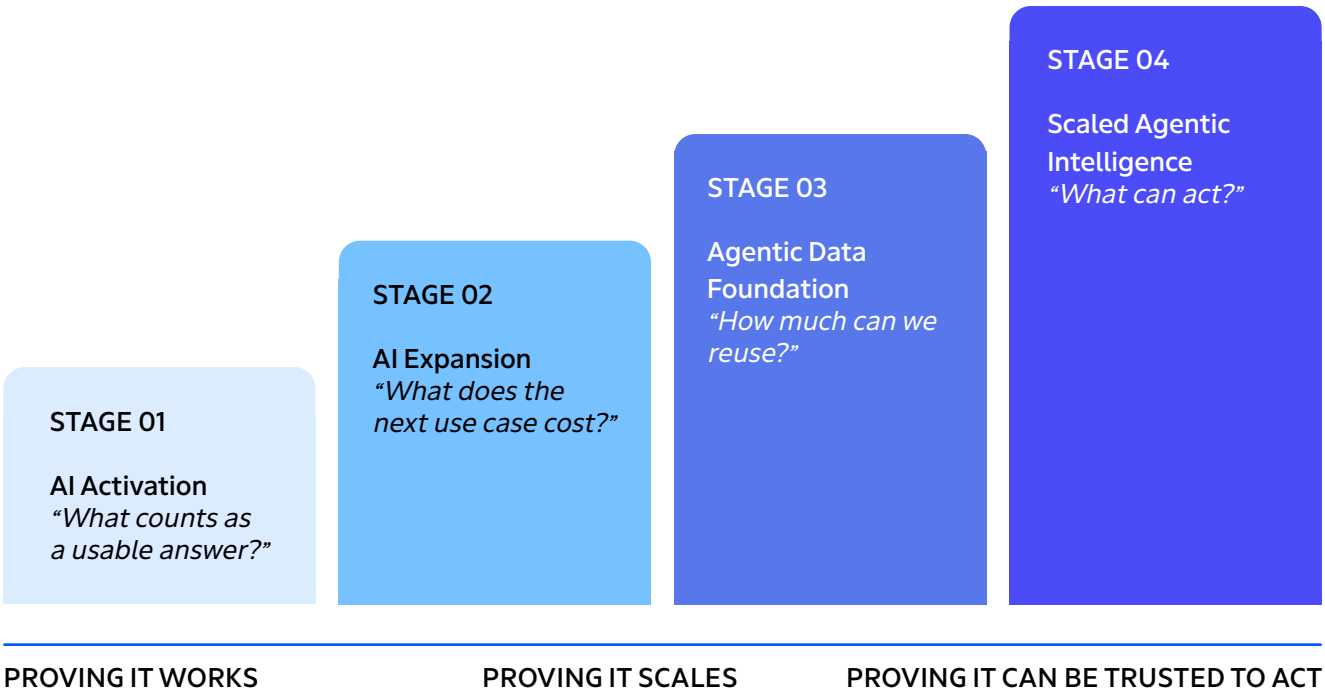

FIGURE 12.1 · THE FOUR STAGES OF AGENTIC AI MATURITY

The Agentic AI Growth Stages framework. Practical AI sharpens the question at each stage from "Does it work?" to "What does it cost to defend?"

© 2026 Progress. All Rights Reserved.

The four practical questions

1. AI Activation: “What counts as a usable answer?”
The organization proves AI can work with its data. Define defensibility early:

- What sources are allowed?
- What review is required?
- What accuracy is acceptable?
- What must be logged?

2. AI Expansion: “What does the next use case cost?”
More teams want AI. More data sources enter the system. Context starts to fragment. If each new AI workflow requires custom retrieval, custom policy, custom review and custom integration, the cost per defensible answer rises.

3. Agentic Data Foundation: “How much can we reuse?”
The organization builds shared access, governance, context and enrichment. This is where Progress Data Platform becomes central. Teams build on trusted data, shared meaning and governed workflows instead of rebuilding the same foundations.

4. Scaled Agentic Intelligence: “What can act?”
AI is embedded in operational workflows. Agents coordinate actions. Failure has business impact. The goal is not maximum autonomy, but policy-driven autonomy, proving reliability before expanding scope.

THE SHORT ANSWER

Each maturity stage has its own economic question. Activation asks what a usable answer is. Expansion asks what the next use case costs. The data foundation stage asks how much can be reused. Scaled agentic intelligence asks what the system can be trusted to do without asking. Autonomy is earned through readiness.

AI maturity is not about adopting more AI. It is about reducing the cost and risk of trusted outcomes.

© 2026 Progress. All Rights Reserved.

PART VI

The Field Reports

Three published customer stories, read through the cost-per-defensible-answer lens. Different industries, same unglamorous foundation.

13. Three customer stories that show the thesis in practice

Three customer stories that show the thesis in practice.

A global financial services firm, a state Medicaid agency and a national news agency—different industries with different starting points. The same lesson: practical AI rests on an unglamorous foundation.

CHAPTER 13 | CASE STUDIES

Practical AI in motion.
Practical AI is not a slide-deck claim. It shows up, or fails to, in real customer workflows. The three field reports that follow are drawn from published Progress customer stories and summarized here as illustration, not as outcomes any other organization should expect to replicate.

CASE STUDY | FINANCIAL SERVICES · LEGAL RESEARCH

Legal research that gets faster and more defensible.
A global financial services firm used Progress Data Platform to build an AI-powered legal assistant. According to the Progress customer story, associates had been spending about four hours compiling dossiers of relevant legislation, case law, articles and analysis. The AI-driven platform can complete the compilation and analysis of data in minutes, and responses are systematically measured against human-created responses to maintain quality and trust.

The important lesson is not simply that AI made research faster, but that speed came from a governed knowledge foundation. The firm needed authoritative sources, semantic context, relevant retrieval and a way to measure outputs against human-created references.

© 2026 Progress. All Rights Reserved.

Read through the cost-per-defensible-answer lens, the published outcome maps onto several cost centers at once: retrieval that is faster and more relevant, human effort shifted from compiling information towards higher-value work and review anchored to a measured comparison rather than a spot check. Those are our inferences from the customer story, not figures the customer reported. What the story does report is the time saving and the systematic measurement against human-created responses.

~4 hrs → minutes
Reported time to compile a legal dossier

Measured
Responses systematically compared against human-created responses

CASE STUDY | PUBLIC SECTOR · HEALTH & HUMAN SERVICES

Modernization that protects continuity and future AI readiness.
The Mississippi Division of Medicaid (MS-DOM) faced a modernization challenge involving siloed systems, regulatory requirements and the need to consolidate administrative, claims, eligibility and clinical data. Progress reports that the agency implemented the Progress Data Platform as a cloud-based operational data store, delivered key requirements in less than 12 months and achieved procurement and deployment at roughly one-third the cost of the original RFP path, avoiding multiyear delays. The same case study notes readiness for future AI-driven use cases such as policy analysis and eligibility scenarios.

This is relevant because practical AI often depends on unglamorous foundations. Before AI can answer policy questions, analyze eligibility or support public-sector service delivery, data has to be consolidated, governed, secure and reusable. That is not an AI demo; it is AI readiness becoming operational infrastructure.

The business value is continuity, compliance, cost avoidance and future optionality. In cost-per-defensible-answer terms, the foundation reduces retrieval friction before the AI use case even starts.

< 12 months
To deliver key requirements

~⅓ cost
vs. the original RFP path, procurement and deployment

© 2026 Progress. All Rights Reserved.

CASE STUDY | MEDIA · CONTENT & METADATA

Content discoverability across languages and formats.
A leading Canadian news agency used Progress Data Platform capabilities to improve content discoverability across multilingual content in formats including text, photos, video and graphics. The Progress customer story states that the agency replaced a mostly manual and inconsistent tagging process with auto-tagging, richer metadata and relationships between tags, topics and content, and that this transformed the search experience. The customer describes the opportunity ahead in their own words: with this technology they can create new products and new revenue streams by exploiting metadata to its fullest. That is a stated opportunity, not a reported result.

This is a useful example because it shows that practical AI is not only about generation. Sometimes the most valuable AI foundation is better discovery. Better metadata means content is easier to find, reuse, personalize, govern and monetize. It also means downstream AI systems have better context to retrieve from.

The lesson is simple: if the content foundation is inconsistent, AI inherits that inconsistency; if the content foundation is semantically enriched, AI can use that meaning. That is why practical AI starts with AI-ready data.

Auto-tagged
Multilingual content across text, photo, video and graphics

Stated opportunity
New products and revenue streams from richer metadata—customer's forward-looking view, not a measured outcome

Sources: Progress.com customer stories: AI-powered legal assistant; Mississippi Division of Medicaid; and Canadian news agency. Outcomes are summarized from published Progress customer stories and are illustrative. Results in any other organization will depend on data, governance and workload conditions. Where a customer describes an opportunity, rather than a delivered result, it is labeled as such above.

THE SHORT ANSWER

Across three very different sectors, the reported gains came from the same place: a governed data foundation built before the AI use case, not a better model applied after it. Speed, cost avoidance and discoverability are all downstream of the same work.

© 2026 Progress. All Rights Reserved.

PART VII

The Practical AI Playbook

Ninety days, one workflow, with one number to move. Plus, the knowledge store that makes the second workflow cheaper than the first.

14. A 30/60/90 plan and how to talk about ROI
15. The knowledge store that compounds
16. The last word: answers the business can defend

The 30/60/90 plan and how to talk about ROI.

Practical AI does not begin with a platform selection; it begins with choosing one workflow where trust, cost and value are visible, then working backwards from the answer the business needs to defend.

CHAPTER 14 | THE PLAYBOOK

Seven steps, three months, with one defensible workflow.
A good practical-AI candidate workflow has four characteristics:

- ✓ The answer matters
- ✓ The current process is slow or expensive
- ✓ The organization can define what a defensible answer looks like
- ✓ There is enough volume that better economics matter

Once that workflow is selected, leaders should work backwards from the answer

© 2026 Progress. All Rights Reserved.

The seven-step practical AI workflow

1. Define the defensible answer
What does the business need to trust? Is it a source citation; confidence score; review status; policy check; rationale; lineage record; or human approval? A defensible answer should be defined before the system is measured.

2. Map the answer-production system

- Where is data found?
- Who checks it?
- What is reworked?
- What is escalated?
- Which systems are touched?
- Which policies apply?
- Where does review happen?

Mapping exposes the real cost.

3. Estimate cost per defensible answer
Use the cost-per-defensible-answer formula to estimate the full operational cost of your selected workflow, including compute, retrieval, remediation and human review. Then divide that cost by the number of outputs meeting your organization's defensibility criteria. Finally, use workload-specific measurements and treat assumptions as directional planning estimates rather than commercial quotes.

4. Identify the biggest cost center
In some workflows, retrieval. In others, remediation. In regulated workflows, review may dominate. In agentic workflows, the cost may come from exception handling. Do not assume compute is the main problem until the full system is visible.

5. Design the Progress Data Platform-grounded workflow
Trusted data access; semantic enrichment; metadata and reference data; governed retrieval; policy enforcement; business rules; human-in-the-loop routing; auditability; controlled tool and data connectivity; workflow monitoring. Tie every capability to a cost line or a trust requirement.

6. Measure defensible-answer yield
Track how many outputs are accepted, corrected, escalated, rejected or reused. Yield matters more than output volume. If yield improves, the economics improve.

7. Scale only after trust mechanisms scale
The earlier *Agentic AI Growth Stages* whitepaper argued that autonomy should expand only after reliability is proven, moving from grounded answers to suggested actions, bounded reversible autonomy and eventually policy-driven autonomy. That remains the right path. Do not ask the organization to trust autonomy. Prove reliability first, then expand scope.

© 2026 Progress. All Rights Reserved.

A practical 30 / 60 / 90 plan

Day 1 – 30
BASELINE

Find the number you are actually trying to move:
- Select one workflow where the answer matters and volume is real
- Define what a defensible answer requires: evidence, policy check, approval
- Map the answer-production system end to end
- Estimate the current cost per defensible answer

DELIVERABLE

One-page cost -per-defensible answer baseline

Day 31 – 60
DESIGN

Show how the grounded workflow changes each cost line:
- Identify the dominant cost center – retrieval, remediation or review
- Design the Progress Data Platform-grounded workflow against that cost center
- Tie every capability to a cost line or a trust requirement
- Set the escalation rules that decide what reaches a human

DELIVERABLE

A grounded workflow design showing how retrieval, remediation and review will change

Day 61 – 90
PROVE

Measure yield, not output volume:
- Run the workflow and track accepted, corrected, escalated and rejected outputs
- Recalculated cost per defensible answer against the baseline
- Document what governance and audit evidence now exists
- Decided what earns the right to expand scope

DELIVERABLE

A business case based on before-and-after answer economics, not model performance

FIGURE 14.1 · THE 30/60/90 PRACTICAL-AI PLAN

Three months with three deliverables: a one-page cost-per-defensible-answer baseline; a Progress Data Platform-grounded workflow design showing how retrieval, remediation and review will change; and a practical-AI business case based on before-and-after answer economics, not just model performance.

© 2026 Progress. All Rights Reserved.

How to talk about ROI without overpromising

Enterprise AI ROI should be commercially sharp, but defensible, and should not:

- × Promise that AI will remove all human review
- × Claim that grounding removes hallucinations
- × Imply that a platform alone solves governance
- × Use one customer result as a universal benchmark

Instead, make the value case in terms of controllable levers: reduce retrieval waste, remediation and unnecessary review, increase defensible-answer yield, improve auditability, accelerate repeatable workflows and scale expertise.

This is a more credible ROI story because it reflects how AI actually becomes valuable in the enterprise. It also aligns with the broader market direction. Gartner's April 2026 analysis reports that organizations with the highest maturity of AI-ready data and analytics capabilities are achieving up to 65% greater business outcomes, such as revenue growth and cost optimization, and frames context—including semantics and metadata—as mission-critical for AI agents to deliver trusted intelligence.

THE BUYER MESSAGE

DO NOT SAY

"AI will transform everything overnight."

Unfalsifiable, unbudgetable and the first thing a skeptical CFO will discount

SAY

"Start with one workflow, measure the cost of defensible answers and then use the foundation to reduce that cost and scale what works."

Specific, measurable and a business case a buyer can defend

THE SHORT ANSWER

Ninety days is enough to baseline one workflow, design a grounded version of it and measure the change in defensible-answer yield. That sequence—baseline, design, prove—produces a business case grounded in answer economics, rather than model benchmarks.

© 2026 Progress. All Rights Reserved.

The knowledge store that compounds.

AI should not have to rediscover your organization every time it answers a question. It needs a living, independent memory that gets richer as the business learns, while keeping evidence, permissions and human accountability intact.

CHAPTER 15 | A LIVING KNOWLEDGE STORE

Prepare knowledge once improve it continuously; reuse it everywhere.
Most enterprise AI still behaves as if every question is the first question. It searches again, rebuilds context, resolves the same terms and asks people to check the same facts. Teams then repeat that work in separate assistants, projects and departments. The organization pays for rediscovery instead of benefiting from what it has already learned.

A better pattern is a living knowledge store: an independent layer that brings together approved enterprise knowledge, adds business meaning, records where each fact came from and keeps useful context ready for the next workflow. New material can enrich what is already known. Conflicts can be surfaced instead of silently flattened. Approved answers can become reusable knowledge, rather than disappearing into a chat history.

> The store does not replace systems of record; it gives AI a governed memory of what the enterprise knows, where it came from, what changed and who may use it.

CHAPTER 15 | THE KNOWLEDGE THAT COMPOUNDS

Independent means reusable, not disconnected
The word *independent* matters. The knowledge store should not belong to one model, one chatbot or one vendor's interface. Models will change, applications will multiply, but the trusted context behind them should remain reusable. Original documents, databases and content systems remain the sources of record; the knowledge store sits between those sources and AI experiences as a governed context layer.

© 2026 Progress. All Rights Reserved.

That makes the knowledge store more than a personal notebook or a folder of summaries. It becomes shared business infrastructure that's connected to structured and unstructured information, aware of relationships, protected by access rules and able to show the evidence behind an answer.

How a living knowledge store works

1. Connects without replacing
Brings together documents, data and content from the systems the business already trusts. The originals remain the source of record.

2. Adds business meaning
Classifies information, recognizes important people, products, policies and topics and links related knowledge so AI sees context, not isolated files.

3. Preserves the evidence
Keeps source links, history and relationships. Flags contradictions, stale material and gaps rather than presenting uncertain knowledge as settled fact.

4. Governs every update
Applies permissions, policies, version history and approval steps so people can see what changed, why it changed and who approved it.

5. Is reusable across experiences
Serves the same governed context to search, assistants, analytics and agents. Teams can change models or interfaces without rebuilding enterprise memory.

6. Improves through use
Feeds checked answers, new evidence and expert corrections back into the store. Reviews it for contradictions, staleness, orphaned knowledge and unanswered questions so quality compounds over time.

Why Progress Data Platform changes the pattern
Progress Data Platform can provide this living knowledge store at enterprise scale. Progress® MarkLogic® software can hold documents, data, metadata and relationships together and make them searchable with their history and provenance. Progress® Semaphore® semantic AI platform can add shared business meaning by classifying content and identifying important facts and relationships. Progress® Corticon® decision management system can apply business rules, policy checks and approval requirements. Progress® Flow Conductor can connect sources and coordinate repeatable flows so knowledge is prepared once and made available to multiple AI workloads.

© 2026 Progress. All Rights Reserved.

One context layer
Structured and unstructured knowledge can be brought together without forcing every source into the same shape.

Meaning with evidence
Terms, topics and relationships are enriched while permissions, provenance and lineage remain attached.

Control at reuse
Rules, approvals and access controls can travel with the knowledge into new assistants, agents and workflows.

The result is not a frozen repository, but an agile knowledge layer that can accept new sources, improve its structure, support new questions and remain independent of the model that happens to use it today.

The write-back rule: improve the store without corrupting it
A living store should learn, but it should not treat every model response as fact, including keeping raw sources unchanged; treating AI-produced summaries, links and conclusions as proposals; checking them according to the risk of the workflow; then promoting approved material as reusable knowledge with its evidence and decision history attached.

SITUATION	HOW THE STORE SHOULD HANDLE IT
Low risk	Automated checks and sampled human review may be enough
Higher risk	A named expert or policy owner approves the update
Conflicting evidence	Preserves both sources and makes the disagreement visible
Outdated knowledge	Retires it from active use without erasing the audit trail

© 2026 Progress. All Rights Reserved.

The economic effect: knowledge that compounds

This pattern changes AI economics because useful preparation is no longer paid for one question at a time. Connections, classifications, relationships, permissions and policies can be reused across workloads. Retrieval becomes less repetitive. Source reconciliation becomes easier. Remediation and unnecessary review can fall. More answers arrive with context the business can evaluate.

That is the deeper value of an independent knowledge store—it allows every approved source, correction and answer to improve the next one. The organization stops accumulating disconnected AI conversations and starts building a governed memory that compounds.

THE SHORT ANSWER

A living knowledge store is a governed context layer that sits between systems of record and AI experiences. It prepares knowledge once, keeps its evidence and permissions attached, improves through approved updates and serves every downstream model and application. It is what makes the second AI use case cheaper than the first.

Platform capabilities summarized from current Progress Data Platform, MarkLogic, Semaphore and Corticon product materials at publication. The personal-scale pattern that helped inspire this chapter is credited in Appendix B. Product fit, configuration and outcomes depend on the organization's sources, controls and workflows.



The last word: answers the business can defend.

The first phase of enterprise AI was about possibility. The second was about reliability. The next is about practical value and the answers we would actually defend.

CHAPTER 16 | THE LAST WORD

The most valuable AI will be the AI the business can defend.
Organizations have invested in data platforms, AI-ready foundations, copilots, agents, pilots and proofs of concept (PoCs). Many have seen enough to believe AI can work. Now they need to know where it can scale, where it can be trusted and where it produces economic value.

The answer will not come from model performance alone, but from the full system that produces it.

THE FOUNDATION

- Data
- Context
- Retrieval

THE DISCIPLINE

- Rules
- Workflow
- Review

THE DEFENSE

- Evidence
- Governance
- Human judgment

That is why cost per defensibe answer is the right metric for this moment. It forces the organization to stop measuring AI by how much it generates and start measuring AI by what it can actually use.

Progress Data Platform helps organizations move from AI-ready data to practical AI execution by providing the trusted foundation for governed access, semantic context, retrieval, rules, policy-aware workflows, auditability and agentic execution.



The goal is not more impressive AI; it's AI that lowers the cost of trusted outcomes since it can be:

- ✓ Used
- ✓ Scaled
- ✓ Defended

Practical AI starts when the answer is no longer just plausible, but defensible too.

CLOSING ARGUMENT

RUN THE COST-PER-DEFENSIBLE-ANSWER CALCULATOR

Model your own workload. See where AI cost is really coming from. Compare a naive AI path with AI grounded by Progress Data Platform.

CALCULATE YOUR COST

Plug in your use case, model and volume. The calculator shows compute, retrieval, remediation and review side by side.

progress.com/data-platform/insights/ai-calculator

SPEAK TO A PROGRESS DATA PLATFORM SPECIALIST

Identify one practical AI workflow where better data, trusted retrieval, governed workflows and human-in-the-loop controls can reduce cost and improve confidence.

progress.com/data-platform | Contact

© 2026 Progress. All Rights Reserved.

ABOUT THE AUTHORS

The people behind the practical AI argument.
A product marketer who thinks in systems and a campaign architect who turns positioning into pipeline. Together, they are the voices behind the Progress Data Platform Make AI Boring · Make AI Practical series.



Philip Miller
AI Strategist & Senior Product Marketing Manager | Progress Software

Philip serves as an AI Strategist at Progress, where he oversees the messaging and strategy for data and AI-related initiatives across the Progress Data Platform. A frequent writer, he contributes to blogs and helps present and moderate product and community webinars. He is dedicated to advocating for customers and works to drive innovation and improvement within the Progress AI-powered platform.

Outside his professional life, Philip is a devoted father of two daughters, a dog enthusiast with a mini-dachshund, and a lifelong learner always eager to discover something new.

> The most valuable enterprise AI does not feel like a breakthrough day to day. It feels reliable, governed and accountable, and the organizations that win with AI will be the ones that can defend the answers their systems produce.
> **Connect with me:** Linkedin | Progress



Hinal Patel
Senior Product Marketing Specialist | Progress Software

Hinal Patel is a Senior Product Marketing Specialist at Progress Software focused on the Progress Data Platform and enterprise data and AI solutions. Her work centers on communicating the value of data, AI and decision automation technologies, helping organizations understand how trusted data, business rules and AI capabilities can support operational efficiency, better decision-making and digital transformation.

Some of her favorite things to do outside of work are hiking, reading and spending time with family and friends.

> The best technical stories are the ones a non-technical reader can pick up and understand on the first pass, without losing the substance underneath.
> **Connect with me:** Linkedin | Progress

Appendix A

Glossary & key concepts

A pocket dictionary for the practical AI conversation. Use these definitions to keep meetings honest and slides defensible.

APPENDIX A | GLOSSARY

The practical AI vocabulary

TERM	DEFINITION
Practical AI	AI that is repeatable, governed, trusted, measurable and economically useful in real business workflows.
AI-Ready Data	Data that is accessible, governed, semantically enriched and usable by AI systems with appropriate context, permissions and lineage.
Defensible Answer	An AI-assisted output that is accurate, explainable, governed, trusted and useful enough to support a business decision or workflow.
Cost per Defensible Answer	A metric that divides the full operational cost of producing AI-assisted answers by the number of answers the business would actually defend.
Retrieval	The process of finding, accessing, contextualizing and grounding the right data or evidence for an AI response.
Remediation	The cost and effort of correcting poor AI outputs, fixing data issues, resolving inconsistencies, rerunning workflows and handling exceptions.
Human-in-the-Loop	A workflow design where people review, approve, reject or escalate AI outputs where judgment, risk or accountability requires it.

© 2026 Progress. All Rights Reserved.

Agentic AI	AI systems that can plan, reason, orchestrate steps and take actions across tools and workflows.
MCP-Style Connectivity	A standards-based way for AI applications to connect to external systems, tools, data and workflows, with appropriate consent, controls and security considerations.
Skills	Reusable packages of instructions, resources and optional scripts that help an AI agent perform repeatable workflows reliably.
Neuro-symbolic AI	A practical combination of statistical AI capabilities, such as language models, with symbolic knowledge and reasoning: rules, metadata, relationships, policies and knowledge structures.
Living Knowledge Store	An independent, governed context layer that brings together enterprise knowledge, preserves its sources and permissions and improves through approved updates so it can be reused across AI models and workflows.
Defensible-Answer Yield	The proportion of AI outputs the business would stand behind. Yield, not volume, is the denominator that matters.
Grounding	Anchoring AI responses to authoritative sources via retrieval, semantic context and policy controls, not just vector similarity.
Policy-Aware Retrieval	Retrieval that enforces who can see what, under which conditions and with full audit lineage.
Context Layer	The governed layer between systems of record and AI experiences where data is unified, semantically enriched and served with its permissions and provenance intact.

© 2026 Progress. All Rights Reserved.

Appendix B

Sources, methodology & further reading

Where the numbers come from, what we modeled and what we deliberately did not. Treat this section as the paper trail.

APPENDIX B | SOURCES & METHODOLOGY

The receipts.

Market & adoption data

Stanford HAI · 2026 AI Index Report:
88% of respondents said their organization used AI in at least one business function in 2025, up from 78% in 2024; 79% reported regular generative AI use in at least one business function. Survey-based and self-reported.
hai.stanford.edu/assets/files/ai_index_report_2026.pdf

McKinsey & Company · The State of AI (November 2025):
39% of respondents report EBIT impact at the enterprise level; approximately one-third report their companies have begun to scale AI programs; high performers are more likely to say their organizations have defined processes to determine how and when model outputs need human validation.
mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

Gartner press release, 16 April 2026:
Organizations with successful AI initiatives invest up to four times more, as a percentage of revenue, in foundational data and analytics areas; only 39% of technology leaders are confident their enterprise's current AI investments will have a positive impact on financial performance; organizations with the highest maturity of AI-ready data and analytics capabilities are achieving up to 65% greater business outcomes.
gartner.com/en/newsroom/press-releases/2026-04-16-gartner-says-organizations-with-successful-ai-initiatives-invest-up-to-four-times-more-in-data-and-analytics-foundations

A NOTE ON THE TWO 39% FIGURES

McKinsey's 39% and Gartner's 39% are unrelated statistics from different surveys with different populations: respondents reporting EBIT impact in one case, technology leaders expressing confidence in future financial impact in the other. They are not corroborating figures and should never be presented as though one supports the other.

© 2026 Progress. All Rights Reserved.

RAG, retrieval & grounding

- IBM: Retrieval-Augmented Generation overview. ibm.com/think/topics/retrieval-augmented-generation
- NVIDIA: What is Retrieval-Augmented Generation. blogs.nvidia.com/blog/what-is-retrieval-augmented-generation
- IBM Research: Neuro-symbolic AI. research.ibm.com/topics/neuro-symbolic-ai

Evaluation research and illustrative modeling assumptions

The percentages below are adjustable Progress modeling assumptions for distinct use cases. They are informed by the direction of public research, but they are not results from a single benchmark and should not be compared as though they were measured under the same conditions. NIST and recent RAG evaluation research support task-specific evaluation, source and provenance checks and ongoing monitoring. Replace these defaults with workload-specific measurements wherever possible.

USE CASE	UNGROUNDED ASSUMPTION	GROUNDED ASSUMPTION
Content creation	22%	4%
Knowledge discovery	26%	7%
Coding assistant	32%	12%
Deep research	35%	8%
Grounded chatbot scenario	40%	3%
Regulated / compliance	45%	5%
Agentic workflow	55%	22%

NIST AI 600-1:

Recommends task-specific evaluation of retrieval-augmented systems, documenting sources and provenance and monitoring performance. nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf

RAGEval and MIRAGE:

Scenario-specific and component-aware RAG evaluation frameworks that illustrate why one configuration or metric does not fit every task. aclanthology.org/2025.acl-long.418 · aclanthology.org/2025.findings-naacl.157

© 2026 Progress. All Rights Reserved.

Model pricing used in worked examples

The document-processing illustration uses Claude Sonnet 5.5 at published base API pricing of $2 per million input tokens and $10 per million output tokens.

The example holds the model constant while comparing the estimated input-to-ken requirements of raw PDFs and prepared Markdown. A directional 1.30 tokeniz-er multiplier is applied to account for changes in tokenization. The comparison measures model-processing compute costs only.

Prompt caching, batch-processing discounts and negotiated enterprise rates are excluded from the illustration. Vendor pricing and processing characteristics may change, so organizations should recalculate estimates using current pricing and workload-specific assumptions before incorporating them into a business case.

Source: https://platform.claude.com/docs/en/about-claude/pricing

Document-processing cost illustration

The following illustration compares the estimated compute cost of processing enterprise documents as raw PDFs vs. AI-ready Markdown. It uses published AI model pricing and adjustable token-consumption assumptions to demonstrate how preparing content for AI can reduce processing costs.

This comparison focuses specifically on compute costs, one component of the broader cost-per-defensible-answer equation.

© 2026 Progress. All Rights Reserved.

Modeling assumptions

The illustrative scenario uses the following assumptions:

INPUT	ASSUMPTION
AI Model	Claude Sonnet 5.5
Input Pricing	$2 per million tokens
Output Pricing	$10 per million tokens
Documents processed monthly	10,000
Pages per document	20
PDF input tokens per page	4,000
AI-ready Markdown input tokens per page	200
Output tokens per call	500
Tokenizer multiplier	1.30

Calculation methodology

Input-token consumption is calculated by multiplying the number of pages by the estimated tokens per page and the tokenizer multiplier. The same multiplier is applied to the assumed output-token volume.

Model-processing costs are calculated using the following formula:

Cost per document = (input tokens ÷ 1,000,000 × liput-token price) + (output tokens ÷ 1,000,000 × output-token price)

The estimated savings per document are calculated by subtracting the Markdown processing cost from the PDF processing cost. Annual savings are calculated by multiplying the difference by the monthly document volume and 12 months.

© 2026 Progress. All Rights Reserved.

Illustrative results

Under these assumptions, processing a 20-page document using Claude Sonnet 5.5 produces the following estimated costs:

- Raw PDF: $0.2145 per document
- AI-ready Markdown: $0.0169 per document
- Reduction in model-processing costs = approximately 92%
- Estimated annual savings at 10,000 documents per month = approximately $23,700
- The comparison demonstrates how reducing the input-token requirements of enterprise content can lower AI processing costs without changing the underlying model.

Limitations and considerations

These figures are illustrative estimates based on adjustable assumptions, not measured customer outcomes or guaranteed savings.

The comparison focuses exclusively on model-processing compute costs. It does not include the costs of document conversion, preparation, implementation or ongoing maintenance. It also does not quantify retrieval, remediation, human review or defensible-answer yield.

Actual token consumption depends on document complexity, content format, processing method and model behavior. Image-based PDF processing may vary by resolution, and the 1.30 tokenizer multiplier is a directional assumption rather than a universally applicable rate.

The example uses published base API pricing and excludes prompt caching, batch-processing discounts and negotiated enterprise rates. These factors may change the absolute costs and relative savings.

A complete cost-per-defensible-answer analysis requires workload-specific measurements across all four cost categories and the proportion of outputs meeting the organization's defensibility criteria.

Source: Anthropic Claude API pricing, https://platform.claude.com/docs/en/about-claude/pricing. Calculations based on illustrative Progress modeling assumptions.

© 2026 Progress. All Rights Reserved.

Persistent knowledge, governance & platform capabilities

- Progress Data Platform: governed context layer spanning enterprise data, semantics, policy and orchestration. progress.com/data-platform/insights/practical-ai
- Progress practical-AI methodology: source systems remain the source of record while the platform prepares governed context that can be reused across AI workloads. progress.com/data-platform/insights/ai-calculator
- Progress MarkLogic, Semaphore and Corticon platforms contain model-independent knowledge, hybrid search, provenance, semantic enrichment, classification and traceable business rules. progress.com/marklogic/solutions/generative-ai · progress.com/semaphore/platform · progress.com/corticon
- Andrej Karpathy, LLM Wiki: a personal-scale pattern for turning raw sources into a persistent, interlinked wiki that records contradictions, strengthens synthesis and improves through human curation. It inspired the living-store framing; Chapter 15 extends the idea to governed enterprise knowledge. gist.github.com/karpathy/442a6bf555914893e9891c11519de94f

Standards & agent infrastructure

- Model Context Protocol: open-source standard for connecting AI applications to external systems; security and trust considerations, including user consent and control, data privacy and tool safety. modelcontextprotocol.io · /docs/getting-started/intro · /specification/2025-11-25
- OpenAI Codex Skills: agent skills as packages of instructions, resources and scripts. developers.openai.com/codex/build-skills
- OpenAI Agents SDK Skills: repo-local skills, repository policies, GitHub Actions; 457 merged pull requests across the two Agents SDK repositories combined, December 2025 to February 2026, vs. 316 in the prior three months. developers.openai.com/blog/skills-agents-sdk
- Anthropic Claude · Skills: domain-specific expertise loaded on demand. platform.claude.com/docs/en/managed-agents/skills

© 2026 Progress. All Rights Reserved.

Progress customer stories

- AI-powered legal assistant, global financial services firm. progress.com/customers/ai-powered-legal-assistant
- Mississippi Division of Medicaid (MS-DOM) modernization case. progress.com/customers/ms-dom
- Canadian news agency, content discoverability. progress.com/customers/canadian-news-agency

What was deliberately not modeled in the document processing illustration:

- × **Retrieval, remediation and human review:** These cost categories are part of the broader cost-per-defensible-answer framework, but are not quantified in the PDF vs. Markdown calculator.
- × **Defensible-answer yield and model accuracy:** The comparison does not measure accuracy improvements, governance compliance or the proportion of outputs a business would accept.
- × **Document preparation and implementation costs:** The cost of converting, preparing, validating and maintaining AI-ready content is excluded from the calculation.
- × **Caching, batch discounts and enterprise pricing:** These may change actual processing costs. The illustration uses published base API rates without these adjustments.
- × **Strategic governance benefits:** Explainability, audit readiness, reusable context and stronger control can provide additional business value, but their financial impact is not calculated in this example.

Disclaimer

This book and its illustrative modeling scenarios are directional planning tools, not commercial quotes or guarantees of financial outcomes. The PDF vs. Markdown calculator estimates model-processing compute costs using adjustable token assumptions and published model pricing. It does not calculate the full cost per defensible answer. Hallucination and human-review assumptions presented elsewhere in this book are separate modeling illustrations and are not inputs to the PDF vs. Markdown calculator. Actual results depend on data, workload characteristics, governance requirements, implementation costs and vendor pricing agreements.

© 2026 Progress. All Rights Reserved.

Appendix C

About the Progress Data Platform

A trusted foundation for practical AI: data, context, governance, rules, retrieval and agentic execution, within one platform.

APPENDIX C | THE PLATFORM

The components behind the story.

The Progress Data Platform brings together four products designed to work as one platform for trusted, governed, practical AI:

M. **Progress MarkLogic Multimodel Database and Context Engine**

Unifies structured and unstructured data, supports knowledge graphs and rich metadata and provides governed access to enterprise context.

S. **Progress Semaphore Semantic Enrichment and Classification**

Adds business meaning to content with taxonomies, ontologies, auto-tagging and disambiguation, turning text into structured semantics.

C. **Progress Corticon Business Rules and Policy Engine**

Encodes the rules, policies and decision logic that govern when answers are allowed, escalated or blocked and makes them auditable.

O. **Progress Flow Conductor Workflow Automation**

Coordinates the full answer-production system: retrieval, ruleset application, escalations, integrations and human-review routing.

Next steps

- **AI ROI Calculator —** Compares the estimated cost of processing PDFs vs. AI-ready Markdown at: progress.com/data-platform/insights/ai-calculator
- **Make AI Boring —** The first whitepaper in this series on governed, dependable enterprise AI: progress.com/data-platform/insights/if-your-ai-isnt-boring-it-isnt-ready-reliable-governed-ai
- **The Agentic AI Growth Stages —** The maturity framework that complements this book: progress.com/data-platform/insights/agentic-ai-maturity-model
- **Progress Data Platform —** Product overviews, customer stories and architecture at: progress.com/data-platform

© 2026 Progress. All Rights Reserved.



Explore Progress Data Platform

About Progress Software

Progress Software (Nasdaq: PRGS) provides the context and control organizations need to reliably extract value from AI—context drawn from an organization's data, content and workflows, and control over the security, governance and cost of their AI initiatives. Learn how hundreds of thousands of businesses, powering the work of tens of millions of professionals worldwide, realize value from trusted, enterprise-ready AI at www.progress.com

© 2026 Progress Software Corporation and/or its subsidiaries or affiliates.
All rights reserved. Rev 2026/09 | 1217474390383978

Worldwide Headquarters

Progress Software Corporation
15 Wayside Rd, Suite 400, Burlington, MA 01803, USA
Tel: +1-800-477-6473

facebook.com/progresssw
x.com/progresssw
youtube.com/progresssw
linkedin.com/company/progress-software
progress_sw_

Progress