

Sema4.ai

The enterprise AI maturity model

A practical guide to moving from
experimentation to scale

Table of contents

Executive summary	3
Defining AI maturity	4
The five pillars of AI maturity	5
AI maturity vs. AI development: why the confusion hurts	6
The five stages of AI maturity	7
How to benchmark your AI maturity	9
The role of an AI Center of Excellence	10
Your roadmap to AI at scale	15
What's blocking AI maturity	16
Learnings from our customers	18
Moving forward	20
Recommended next steps	21
Further reading	22
Appendix 1 – AI maturity self-assessment worksheet	24
Appendix 2 – Links and Sources	27



Executive summary

Most enterprises have AI projects. Few have mature AI capabilities. According to RAND Corporation's 2025 analysis of more than 2,400 enterprise AI initiatives, 80% fail to deliver their intended business value, with a third abandoned before ever reaching production. For agentic AI specifically, recent enterprise surveys show that only 12% of proofs of concept have reached production deployment (Prosigns, State of Enterprise AI 2026).¹ Similar feedback is presented by more major industry analyst firms.

In 2026, AI maturity, not experimentation, separates winners from laggards. Organizations that thrive won't be those with the most AI projects, but those that have systematically embedded AI into decision-making, workflows, and value creation.

This white paper introduces a comprehensive AI maturity model to help you benchmark your organization's readiness and build a clear roadmap toward scalable AI that delivers measurable business impact.



80%

of enterprise AI initiatives
fail to deliver their intended
business value



Defining AI maturity

AI maturity isn't measured by how many AI tools you've deployed or how much budget you've allocated. Real AI maturity is defined by how deeply AI is embedded into your organization's decision-making, operational workflows, and value streams. It's the difference between **having AI** and **being AI-first**.


Mature AI organizations don't just run successful pilots. They systematically identify high-value use cases, deploy solutions rapidly, govern them effectively, and continuously improve based on measurable outcomes.

The challenge is well-documented: most enterprises remain stuck at the experimental stage. They've proven AI can work in controlled environments but struggle to operationalize successes across departments, geographies, and use cases. Data remains siloed, governance frameworks lag behind deployment speed, and technical teams become bottlenecks as demand outpaces their capacity.



The five pillars of AI maturity

Think of AI maturity as a structure supported by five essential pillars. Each must be strong for the framework to support enterprise-scale AI. Weaknesses in any area create vulnerabilities that prevent scaling and put your AI investments at risk.



Strategy and alignment

AI initiatives must align with clear business objectives and have strong executive sponsorship. This ensures AI investments support strategic priorities rather than becoming technology projects in search of problems.

Data and integration

AI systems require access to quality data across the enterprise. This addresses data governance, integration architecture, and the ability to connect AI to the applications and data assets that power your business.

Technology and tooling

The right platform capabilities enable rapid development and secure deployment – the infrastructure, development tools, and operational capabilities needed to build, run, and manage AI at scale.

Talent and culture

AI maturity requires both technical expertise and organizational readiness – skills development, change management, and fostering a culture that embraces AI-augmented work.

Governance and risk

Enterprise AI demands robust frameworks for security, compliance, ethical use, and risk management – ensuring AI systems operate within appropriate boundaries while maintaining audit trails and accountability.

AI maturity vs. AI development: why the confusion hurts

A critical misconception slows many organizations: confusing stages of AI **development** with organizational AI **maturity**.

AI development refers to business process owners designing solutions for their teams and the technical process of building solutions, i.e. developers assembling frameworks, coding integrations, and stitching together tools. This is a project-focused process centered on individual implementations.

Organizational AI maturity is fundamentally different. It measures how effectively your entire enterprise leverages AI across operations – encompassing strategy, governance, security, culture, and systematic value delivery. You can have technically sophisticated agents while maintaining low organizational maturity if those agents remain siloed, ungoverned, or dependent on scarce developer resources.

The solution is clear role separation. Business users should build and maintain agents using natural language – defining how work gets done in plain English. Developers should focus on enterprise integrations and platform capabilities. And governance should ensure both operate within enterprise standards. This division accelerates the journey from technical projects to true organizational maturity.


The pilot trap

The dangerous pattern: Companies celebrate deploying impressive custom-coded agents while ignoring whether business users can maintain them independently and have the right business context, whether the company can tackle complex team-based work, or whether the organization can scale AI without creating unsustainable technical dependencies.

The result: Technically impressive DIY projects where business users remain dependent on developers for every modification, teams are slowed down due to limited context or inability to work on complex processes, and AI initiatives stall due to a lack of governance.

The five stages of AI maturity

Understanding where your organization sits helps you identify specific challenges and chart a clear path forward.



Stage 1: Ad hoc

AI projects emerge organically without central coordination. Individuals and teams experiment with consumer-grade tools. There's no standardization, limited knowledge sharing, and no governance. Organizations remain stuck here when they lack executive sponsorship or a clear strategic vision.

Stage 2: Opportunistic

Success stories create pockets of proven value, but remain isolated and difficult to replicate. Different departments pursue AI independently, often duplicating effort. Technical debt accumulates as teams build one-off integrations. Organizations stall here when they fail to establish centralized leadership or shared platforms.

Stage 3: Systematic

Leadership establishes an AI Center of Excellence. The organization develops an AI roadmap aligned and prioritized with business strategy, selects enterprise platforms, and implements governance. This stage represents the critical transition from experimentation to enterprise capability.

Stage 4: Integrated

AI becomes embedded across the organization. The CoE enables business units while maintaining standards. AI systems integrate seamlessly with enterprise applications through semantic understanding, eliminating data silos. Teams leverage platform optimizations to improve accuracy and reduce LLM costs. Business users are empowered to build and deploy agents for their domains.

Stage 5: Optimized

AI agents drive autonomous operations and continuous innovation. Agents handle end-to-end processes 24/7, only escalating exceptions that require human judgment. The organization captures knowledge from operations to drive continuous improvement.

Where most organizations sit today:

Between Stages 2 and 3, with pockets of success but lacking the systematic approach and enterprise platforms needed to scale. Few are operating at Stage 4 and almost none at Stage 5.



How to benchmark your AI maturity

Rather than treating maturity as a single score, evaluate your organization across the five pillars using a simple low/medium/high framework. See Appendix 1, the AI maturity self-assessment worksheet, to score your organization.


Pillar	Key Questions
Strategy and alignment	Do we have a clear AI strategy supported by executive leadership? Are initiatives tied to measurable business outcomes?
Data and integration	Do we have consistent data governance? Can AI systems access enterprise data securely?
Technology and tooling	Do we have enterprise AI platforms or rely on fragmented point solutions?
Talent and culture	Do we have both technical AI capabilities and organizational readiness?
Governance and risk	Do we have frameworks for AI security, compliance, and ethical use?

Apply this framework department by department. Marketing might rate high on adoption but low on governance, while IT rates high on security but struggles with business alignment. These gaps reveal where focused investment delivers maximum impact.

The role of an AI Center of Excellence

Organizations plateau at Stage 2 without establishing an AI Center of Excellence. A CoE serves as the central hub that coordinates governance, establishes standards, provides platforms, and enables business units to deploy AI solutions aligned with enterprise requirements.

The CoE's role isn't to gatekeep – that creates bottlenecks. A mature CoE enables federated development by providing the platforms, frameworks, and support that allow business units to build independently while maintaining enterprise standards.



Getting started: from mandate to momentum

The most common mistake organizations make is treating a CoE as a long-term planning exercise. Successful CoEs launch with a narrow focus supporting operational goals and expand from demonstrated results.

Start with three foundations:

- **Executive sponsorship with authority.** The CoE needs a mandate from leadership that goes beyond advisory. It must have authority to select enterprise platforms, establish standards, and allocate resources. Without this, the CoE becomes a suggestion box that teams ignore.
- **A single high-value use case.** Don't try to govern everything on day one. Select one use case that meets the three qualifying criteria (failed with traditional automation, significant resource cost, unmet SLAs) and prove the model works. This first success becomes your internal case study for expansion.
- **A cross-functional founding team.** The initial CoE team should include a platform architect, a governance specialist, a business process owner from your first use case, and a change management lead. Avoid staffing exclusively with AI researchers or data scientists – you need people who understand enterprise operations, not just algorithms.

The 90-day launch sequence:

Days 1-30

Secure mandate, select founding team, identify first use case, evaluate and select enterprise AI platform

Days 31-60

Deploy platform, build first agent with cross-functional pair (business + technical), establish initial governance policies

Days 61-90

Validate results, document patterns and learnings, present outcomes to leadership, identify next 2-3 use cases for expansion

Key insight:

Your CoE doesn't need to be fully mature to start delivering value. It needs to be mature enough to govern one successful deployment – then scale its governance as it scales its portfolio.

Structure: what to centralize, what to federate

The most effective CoE structure follows a “hub and spoke” model – centralizing what requires consistency while federating what requires domain expertise.

Centralize these (the hub):

- **Platform and infrastructure.** One enterprise AI platform, consistently deployed and managed. This eliminates shadow AI, reduces vendor sprawl, and ensures every agent operates within the same security and governance boundary.
- **Security, governance, and compliance.** Standards for data access, audit requirements, model selection policies, and deployment approval processes. These non-negotiables must be consistent across every business unit.
- **Shared configuration and integrations.** Approved LLM configurations, data connections, action packages, and MCP servers distributed from a central authority. Teams shouldn’t independently negotiate data access or build redundant integrations.
- **Knowledge management.** Patterns, learnings, and reusable components captured from every deployment and made available to the entire organization. One team’s discovery should benefit every subsequent project.

Federate these (the spokes):

- **Use case identification and prioritization.** Business units know their processes best. They should identify where AI agents deliver the most value within their domain.
- **Agent creation and maintenance.** Business process owners define how agents work using natural language. They own the logic, the exception rules, and the domain-specific knowledge.
- **Outcome measurement.** Each business unit defines and tracks the metrics that matter for their use cases – time savings, error reduction, throughput, cost per transaction.

Function	Centralized	Federated
Platform selection and deployment	✓	
Security policies and audit standards	✓	
LLM and data connection configuration	✓	
Use case identification		✓
Agent creation (natural language Runbooks)		✓
Integration development (SDK/custom code)	✓	
Outcome measurement and reporting		✓
Knowledge sharing and pattern documentation	✓	

Scaling: the steering committee and expansion playbook

The CoE delivers its first successful deployment. Now what? Scaling from one use case to an enterprise-wide program requires deliberate mechanisms for knowledge transfer, pattern reuse, and systematic expansion.

Establish a recurring cadence. Meet bi-weekly or monthly with representation from each business unit to explore or deploy agents. The agenda is simple:

Share results

What's working, what metrics are improving, what's been learned

Socialize patterns

Reusable approaches that other teams can adopt without starting from scratch

Identify expansion opportunities. New use cases that leverage existing integrations, data connections, or agent patterns

Remove blockers

Governance, access, or resource issues that are slowing teams down

The expansion flywheel:

Each successful deployment should make the next one faster and cheaper. This happens through three mechanisms:

- **Pattern reuse.** Document what works – the agent architecture, the data access approach, the exception handling logic – and make it available to the next team. The second invoice reconciliation agent shouldn't take as long as the first.
- **Shared integrations.** Every enterprise integration built for one use case becomes available to all. A Snowflake connection configured once serves every agent that needs that data.
- **Compounding intelligence.** As agents operate, they build institutional knowledge – vendor quirks, data format variations, exception patterns. This knowledge carries over to new deployments in the same domain and accelerates time-to-value.

Common mistakes that prevent scaling:

- **Over-centralizing agent development.** If every agent modification requires a CoE ticket, you've recreated the IT bottleneck you were trying to eliminate. Business users must own their agents.
- **Insufficient authority.** A CoE that can recommend but not mandate platform standards will watch shadow AI proliferate around it.
- **Ignoring change management.** Technical readiness without organizational readiness produces unused agents. Invest in adoption, training, and demonstrating how agents augment (not replace) expertise.
- **Measuring activity, not outcomes.** "Number of agents deployed" is vanity. Measure business impact: hours saved, error rates reduced, SLAs met, cost per transaction.

The economics of maturity

The cost equation changes dramatically as you advance through maturity stages. At Stage 2, every new use case requires net-new development effort. By Stage 4–5, intelligence compounds across the organization:

- Each new deployment is faster than the last because institutional knowledge carries over
- The 500th transaction costs a fraction of the first as proven patterns promote to deterministic execution
- Business users handle agent modifications independently, freeing technical resources for strategic work

The bottom line: Mature AI organizations don't just get better outcomes – they get better economics with every deployment.

Your roadmap to AI at scale

Building a practical roadmap requires translating the maturity model into concrete actions.

1. Identify high-impact use cases

Don't start by trying to transform everything. Identify specific processes where AI agents can deliver clear ROI – complex document processing, multi-system reconciliation, help desk automation, or procurement workflows. Prioritize use cases that combine high business value with feasibility, given your current maturity level.

2. Stand up your AI CoE

Establish central leadership with a clear mandate. Staff with diverse capabilities – platform architects, governance specialists, and change management experts alongside AI practitioners. Give the CoE authority to select platforms and establish standards while empowering business units to innovate within guardrails.

3. Adopt a secure AI agent platform

Select an enterprise platform that supports your journey to Stage 4-5. The platform must enable business users to build agents through natural language while providing IT with security, governance, and integration capabilities. Key requirements:

- Deploys within your own cloud infrastructure (AWS, Azure, Google Cloud, or Snowflake)
- Empowers business users to define processes without code
- Provides centralized governance and lifecycle management
- Enables zero-copy data access without moving sensitive information
- Supports comprehensive observability and audit trails

4. Deploy, measure, and iterate

Launch with focused pilots that prove value and establish patterns. Measure both technical success (accuracy, reliability, performance) and business outcomes (time savings, cost reduction, quality improvement). Use insights to refine your approach and identify the next wave. Scale systematically.

What's blocking AI maturity

Even organizations with clear frameworks face persistent barriers.

Fragmented tools



Many enterprises have accumulated dozens of AI and automation tools. This creates integration nightmares, duplicated effort, and governance blind spots. Each tool requires separate expertise, creating knowledge silos that make consistent standards impossible.

Siloed data



AI systems require cross-application data access, but legacy architecture isolates data in disconnected systems. Traditional approaches that copy data create security risks and synchronization problems. Zero-copy architectures solve this by enabling agents to access data where it lives – without moving sensitive information outside your security boundaries.

Compliance risks



Regulatory requirements for data privacy, model explainability, and algorithmic fairness create legitimate concerns. Without governance frameworks and audit capabilities, organizations face genuine legal and reputational risks. Comprehensive audit trails, transparent reasoning, and deployment within your own infrastructure address these requirements directly.

Talent gaps



The shortage of AI expertise forces organizations to compete for scarce technical talent while business users with deep domain knowledge remain unable to build solutions themselves. Platforms that enable natural language agent creation eliminate this bottleneck by empowering process experts to build directly.

Organizational resistance



Teams fear AI will expose inefficiencies or reduce their value. This manifests as passive roadblocks – delayed approvals, reluctance to share knowledge, insistence that “our work is too complex for AI.” Overcoming this requires demonstrating how agents augment rather than replace human expertise.

Signs you're ready to advance

From Stage 2 to Stage 3

- You have at least 2-3 proven AI success stories
- Executive leadership is asking "how do we scale this?"
- Multiple teams are requesting AI capabilities independently
- You're experiencing governance gaps or shadow AI concerns

From Stage 3 to Stage 4

- Your CoE is operational with clear standards
- Business users are requesting more direct involvement in agent creation
- You need cross-system data access for AI use cases
- Initial pilots have demonstrated measurable ROI

From Stage 4 to Stage 5

- Agents are running reliably across multiple business units
- You're ready for 24/7 autonomous operations with exception-based escalation
- Your organization has a track record of capturing operational knowledge systematically

Learnings from our customers

The maturity model describes the destination. The following lessons are drawn from enterprises actively advancing through the stages and offer guidance on that journey.

1. Ask “How good is our data for this use case?”

Agents are far more flexible than traditional software at handling data variations, but “good data” is always a “good thing”. As organizations moved beyond pilots, a consistent pattern emerged: use cases built on structured or semi-structured data in governed warehouses were more successful, while those built on fragmented data scattered across spreadsheets and mangled ERP exports faced more difficulty. The most mature organizations now treat data readiness as an explicit criterion in use case prioritization. If the data foundation exists in Snowflake, Redshift, or a similar environment, they pull that lever first. A useful side effect: this discipline drives better data pipelines across the enterprise, strengthening the Data and integration pillar for every future deployment.

2. Scope in weeks, not quarters.

The most successful customers do not attempt to automate an entire end-to-end process in one initiative. They isolate the highest-value slice of a use case and prove it quickly, typically in two to three-week development, pilot, and iterate cycles. In a technology landscape moving this quickly, a three-month build is obsolete before it ships. This applies at the strategic level too. Organizations that declare themselves “all in on AI” and interpret that as rationalizing every tool and framework on the market get stuck in a permanent figuring-out phase. Going all in should mean empowering every contributor to build, learn, and deploy. Start small, ship, learn, repeat. That flywheel is what builds internal talent and compounds results.

3. The line of business must own the use case.

The old pattern is familiar: business defines requirements, IT builds in isolation, and review meetings reveal that four or five requirements were misinterpreted. Customers who break this cycle by making business sponsors the owners of agent development, with IT as a partner handling integrations and platform capabilities, consistently outperform. This is the role separation the maturity model calls for in practice. When skilled developers can offload lower-complexity projects to business users on a governed platform, developers are freed for the complex, customized work only they can do, and domain experts stop being spectators to their own processes.


4. Executive-sponsored cadence is what makes a CoE work.

The organizations deploying the most production use cases share a pattern: an executive sponsor who takes ownership from the top, holds use case champions accountable, and sticks to a weekly or bi-weekly cadence for knowledge sharing and status updates. That rhythm does more than drive accountability. It surfaces reusable skills, tools, and patterns that every subsequent project inherits, which is exactly the expansion flywheel described earlier. Guardrails matter here too: platforms with central governance let business users experiment safely, so a small group of experts never becomes the bottleneck that limits learning and adoption.

5. Dedicate real capacity and prioritize against enterprise KPIs.

No one produces enterprise-ready AI solutions in their spare time. Everyone has a day job, deadlines, and SLAs, and organizations that “throw AI at everyone” and hope for side-of-desk innovation move slowly. The fastest movers either hire dedicated AI development and experimentation resources or intentionally free up the teams that own use cases. As backlogs grow, which they will, successful organizations pair this investment with a prioritization framework requiring every AI proposal to identify KPIs tied to enterprise priorities. Defining those success metrics early, with broad consensus, gives every line of business a shared north star and makes prioritization a business conversation rather than a political one.

Moving forward



The AI maturity model reveals a fundamental truth: competitive advantage won't come from simply using AI for personal productivity – it will come from having matured your AI capability to the point where agents operate autonomously, governance is systematic rather than reactive, and business users drive innovation without technical bottlenecks.

Maturity isn't about accumulating more tools or running more experiments. It's about building governed, scalable systems that deliver ROI at speed while operating within enterprise compliance and risk frameworks.

The framework we've outlined – five pillars supporting five stages of progression – provides your roadmap. But frameworks alone don't create maturity. You need enterprise platforms that embody these principles, enabling advancement while removing the barriers that keep most organizations trapped in pilot purgatory.

What mature organizations achieve:

- Business users own, define and modify agent processes, without IT dependencies
- Agents access enterprise data securely across systems without data movement or duplication
- Every agent decision is traceable and auditable, satisfying compliance requirements by design
- Intelligence compounds across deployments – each new use case deploys faster than the last
- Operations scale with volume, not headcount

Sema4.ai is the enterprise AI agent platform purpose-built for organizations advancing to Stage 4–5 maturity. Deploy securely within your own infrastructure across AWS, Azure, Google Cloud, or Snowflake. Empower business users through natural language Runbooks. Govern with complete lifecycle management and audit trails. Scale with an architecture where every transaction makes the system smarter, faster, and more cost-effective.

Recommended next steps

1. Assess your current state

Use the five-pillar framework to benchmark your maturity. Involve stakeholders across IT, finance, operations, and compliance. Document gaps between where you are and where you need to be. See the Appendix 1 worksheet.

2. Select your first use case

Look for processes that have failed with traditional automation, consume significant resources, and have unmet SLAs. Start with one – not ten.

3. Build your cross-functional team

Pair a business process expert (who defines what the agent does) with a technical expert (who builds how it connects to systems). Dedicate real capacity, not side-of-desk time for both.

4. Stand up your CoE

Establish central AI leadership with authority to select platforms, set standards, and enable business units. Staff with platform architects, governance specialists, and change management expertise – not just AI enthusiasts.

5. Choose an enterprise platform

Evaluate against five requirements: deploys in your cloud (AWS, Azure, Google, Snowflake), empowers business users without code, provides centralized governance and audit trails, enables zero-copy data access, and compounds intelligence across deployments.

6. Deploy, measure, scale

Launch focused pilots. Measure technical success (accuracy, reliability) and business outcomes (time savings, cost reduction). Use learnings to identify the next wave. Scale systematically.

Further reading

Resource	Focus
Building an AI agent Center of Excellence	Structuring and operationalizing your CoE
Enterprise AI agents for finance operations	Use cases across AP, AR, reconciliation, and close
Sema4.ai platform overview	Architecture, deployment, and capabilities
Security and governance documentation	Deployment architecture and compliance frameworks

Connect with our team to see how enterprises are scaling with AI agents on Sema4.ai. Move beyond experimentation to true AI maturity with the platform built for organizations that demand security, governance, and business user empowerment without sacrificing operational sophistication.

Case study

Global manufacturer transforms invoice reconciliation

Industry: Manufacturing | Stage advancement: Stage 2 → Stage 4
Timeline: 12 weeks to production

The challenge

A Fortune 500 manufacturer faced a complex and time-consuming invoice reconciliation process across its energy procurement operations. The team of eight analysts processed up to 350 invoices per month, with approximately 80% containing discrepancies requiring detailed manual review. A single invoice could take up to three hours to reconcile – requiring analysts to cross-reference complex PDF invoices against financial data stored across multiple enterprise systems.

Previous automation attempts with traditional tools had failed due to the variability of invoice formats across hundreds of unique suppliers and the complexity of multi-step reconciliation logic that required human-like reasoning.

The approach

The organization paired a senior finance analyst who understood the reconciliation process intimately with a developer who built integrations to their enterprise data warehouse. Together, they defined the reconciliation workflow in plain English, describing how the agent should read invoices, extract structured data, access financial records, and identify discrepancies.

The team started with a focused pilot of 50 invoices, validated accuracy against manually reconciled results, and expanded to full production within 12 weeks.

The outcomes

Metric	Before	After
Processing time per invoice	Up to 3 hours	~2 minutes
Invoices requiring human review	100%	~5% (true exceptions only)
Monthly throughput	350 invoices across 8 analysts	350+ invoices processed autonomously
Format coverage	Manual review of each format	90%+ of varying formats handled automatically

Appendix 1 - AI maturity self-assessment worksheet

Use this worksheet to benchmark your organization's current AI maturity. Rate each dimension as Low (1), Medium (2), or High (3), then total your score to identify your stage.

Pillar 1: Strategy and alignment

Dimension	Low (1)	Medium (2)	High (3)	Score
Executive sponsorship	No AI mandate from leadership	Informal support; no dedicated budget	Active C-suite sponsor with funded roadmap	_____
Strategic alignment	AI projects unconnected to business goals	Some projects tied to objectives	All initiatives mapped to measurable outcomes	_____
Use case prioritization	No formal selection process	Ad hoc evaluation of opportunities	Systematic framework ranking value vs. feasibility	_____

Pillar total: ____ / 9

Pillar 2: Data and integration

Dimension	Low (1)	Medium (2)	High (3)	Score
Data access	Siloed; manual extraction required	Some integrations; significant gaps remain	Agents access data across systems seamlessly	_____
Data quality	No standards; inconsistent across sources	Basic governance in some areas	Enterprise-wide quality standards enforced outcomes	_____
Cross-system connectivity	Point-to-point; custom code per integration	Shared middleware; some standardization	Semantic data layer with zero-copy access	_____

Pillar total: ____ / 9

Pillar 3: Technology and tooling

Dimension	Low (1)	Medium (2)	High (3)	Score
Platform consolidation	Fragmented tools; shadow AI prevalent	Preferred tools identified; partial adoption	Single enterprise platform for build, run, manage	_____
Business user enablement	All AI work requires developers	Some low-code options available	Business users build and modify agents via natural language	_____
Deployment and operations	Manual; no lifecycle management	Some automation; basic monitoring	Full CI/CD with observability, versioning, rollback	_____

Pillar total: ____ / 9

Pillar 4: Talent and culture

Dimension	Low (1)	Medium (2)	High (3)	Score
AI skills	Concentrated in a few individuals	Dedicated AI team; growing capabilities	Cross-functional teams; business users are fluent	_____
Organizational readiness	Resistance or indifference to AI	Pockets of enthusiasm; uneven adoption	Culture embraces AI-augmented work across functions	_____
Knowledge capture	Tribal knowledge in people's heads	Some documentation of processes	Institutional knowledge systematically codified in agents	_____

Pillar total: ____ / 9

Pillar 5: Governance and risk


Dimension	Low (1)	Medium (2)	High (3)	Score
Security framework	No AI-specific security controls	Basic controls; not comprehensive	Enterprise-grade: VPC deployment, RBAC, secrets management	_____
Audit and compliance	No audit trail for AI decisions	Partial logging; manual evidence gathering	Full decision traceability; automated compliance documentation	_____
Risk management	No formal AI risk framework	Policies exist but inconsistently applied	Systematic governance with automated enforcement	_____

Pillar total: ____ / 9

Scoring

Grand total: ____ / 45

Score	Stage	What it means
5-10	Stage 1: Ad hoc	Uncoordinated experimentation; no foundation in place
11-18	Stage 2: Opportunistic	Pockets of success but can't replicate or scale
19-27	Stage 3: Systematic	Coordinated approach with platform and governance; scaling begins
28-36	Stage 4: Integrated	Business users empowered; intelligence compounds across deployments
37-45	Stage 5: Optimized	Autonomous operations with continuous improvement



Appendix 2 - Links and Sources

“More than 80% of AI projects fail – twice the rate of failure for IT projects that don’t involve AI”

https://www.rand.org/pubs/research_reports/RRA2680-1.html

(Published Aug 13, 2024 by RAND Corporation report based on 2,400+ enterprise AI initiatives)

“82% of enterprise AI initiatives never reach production” and “Only 12% of agentic AI proofs-of-concept have reached production”

<https://prosigns.io/us/resources/state-of-enterprise-ai-2026/>

(Published April 1, 2026 – “State of Enterprise AI 2026” report based on survey of 1,200 enterprise leaders conducted December 2025 – February 2026)

Sema4.ai

Sema4.ai provides a comprehensive enterprise AI agent platform that enables organizations to build, run, and manage AI agents at scale. Our platform empowers business users to create intelligent agents using natural language, connects agents to enterprise applications and data through pre-built actions and universal connectivity, and provides complete lifecycle management with enterprise-grade security and governance. Sema4.ai is trusted by leading enterprises to transform business operations through autonomous, intelligent agents that work 24/7 with complete transparency and control.

For more information or to [get a demo](#), visit www.sema4.ai