

Sema4.ai

Building an AI agent Center of Excellence to accelerate agent adoption



Table of contents

1: Executive summary	3
Establishing a CoE	3
2: Vision and focus	4
2.1 Objectives	4
2.2 Deliverables	4
3: Methodology and approach – for Sema4.ai Agents	5
3.1 Identify use cases	6
3.2 Build agents – Rapid Agent Deployment (RAD)	7
3.3 Deploy	8
3.4 Train	8
3.5 Support	9
3.6 Educate	9
4: Staffing the CoE	10
4.1 AI CoE responsibilities	10
5: Governance and communication	11
5.1 Governance structure	11
6: Performance metrics and success criteria	12
7: Resources and tools	13
7.1 CoE–provided resources	13



1 | Executive summary

The past eighteen months have seen a tremendous interest from enterprises to harness the power of Generative AI to build more efficient, more precise, data driven businesses. As with any new technology, enterprises that follow a deliberate prescriptive approach towards adopting Generative AI are likely to emerge as leaders. The rest are more likely to dabble with different technologies, spend R&D dollars and fail to demonstrate meaningful ROI to their stakeholders. For the past year, Sema4.ai has been working closely with large customers to help them on their Generative AI adoption journey. This document captures our learnings from this period and provides a blueprint for how enterprises can adopt Generative AI successfully.

Establishing a CoE

AI agents can automate complex processes and as enterprises start the process of identifying the parts of the business that can benefit the most from agentic adoption, pattern recognition is key. In order to build the pattern recognition muscle, what works best is to have a center of excellence that is composed of business stakeholders and technology experts. The business stakeholders can provide insights into the expected ROI from solving a particular problem with agents and the technology experts can provide insights on the suitability of the problem for agentic automation. This addresses the two failure modes in adopting this technology i.e. solving the wrong problem and solving the problem wrong, both of which have disastrous downstream consequences and can stymie other efforts.

The CoE serves as a scalable framework for mutually identifying and accelerating AI use case adoption. It will focus on training AI agents that drive business transformation and innovation. The CoE is tasked with delivering cost efficiencies and value generation across your company by deploying Sema4.ai Conversational and Worker Agents. Once deployed, teams will be able to measure the value generation and outcomes, documenting them into a shared use case library to be leveraged by the entire organization.

2 | Vision and focus

The vision of the CoE is to establish a well-defined and repeatable way to empower teams through AI enablement across the organization. Through the CoE, you can manage the creation and maintenance of a library of repeatable use case patterns and a deployed library of reusable agents. By documenting each AI agent use case background, creation process, and outcome, teams from each division will be enabled to rapidly deliver on AI agent use cases, further accelerating value generation.

2.1 | Objectives

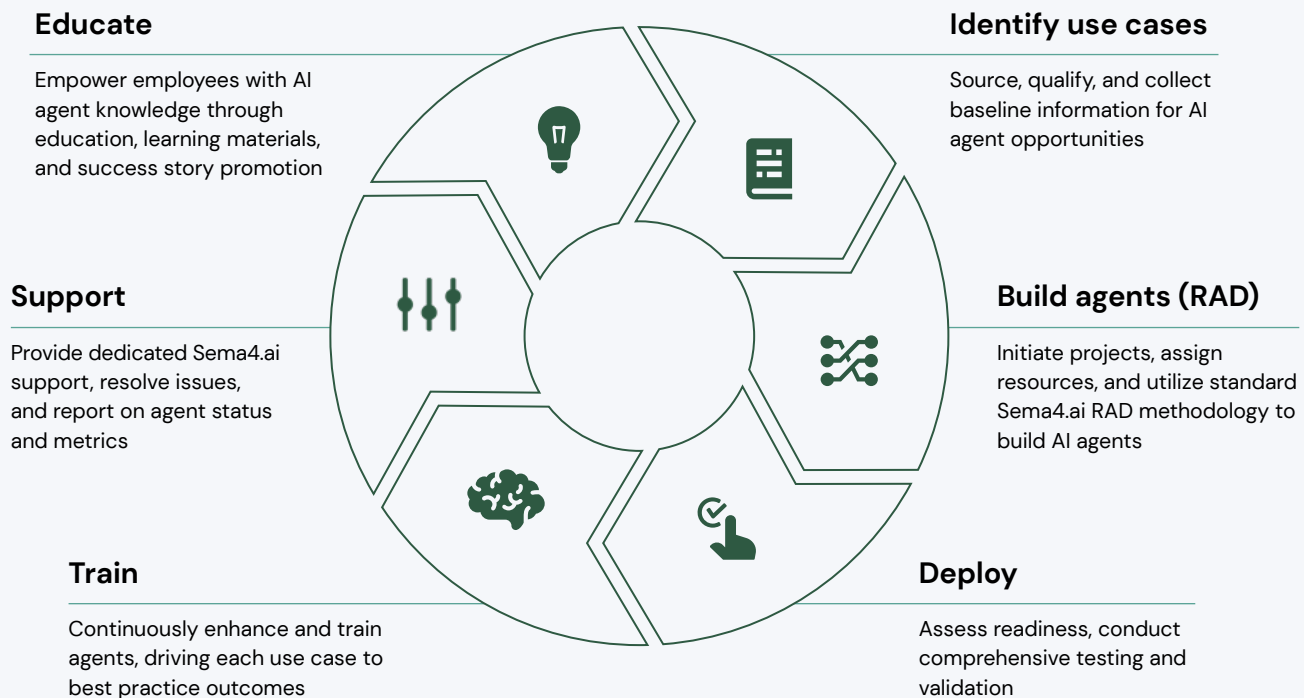
- Build a library of repeatable AI agent use cases that help transform your existing business processes and applications.
- Implement new AI agent use cases across your organization.
- Document success stories to share globally to foster a strong AI and AI agent talent pipeline.
- Educate company leadership and executive teams on the benefits of adopting agent-first mindsets to maximize efficiency and value creations.

2.2 | Deliverables

- Production deployed agents across all divisions.
- Continuous improvement via updated Runbook and Agent training – targeting improved accuracy and automation.
- Library of use cases, Sema4.ai Runbook, Action, and data access templates.
- Comprehensive documentation, training materials, and certification programs.
- Established support channels and processes for case management and expedited resolution.
- Regular performance reporting cadence and improvement recommendations.
- Internal case studies and content for internal showcases.

3 | Methodology and approach – for Sema4.ai Agents

The methodology and approach section outlines the strategic framework for implementing Sema4.ai Agents. It covers the entire lifecycle of Sema4.ai Agents adoption, from use case identification to ongoing support and integration, ensuring a systematic and effective approach to leveraging Sema4.ai's products and platform.



3.1 | Identify use cases

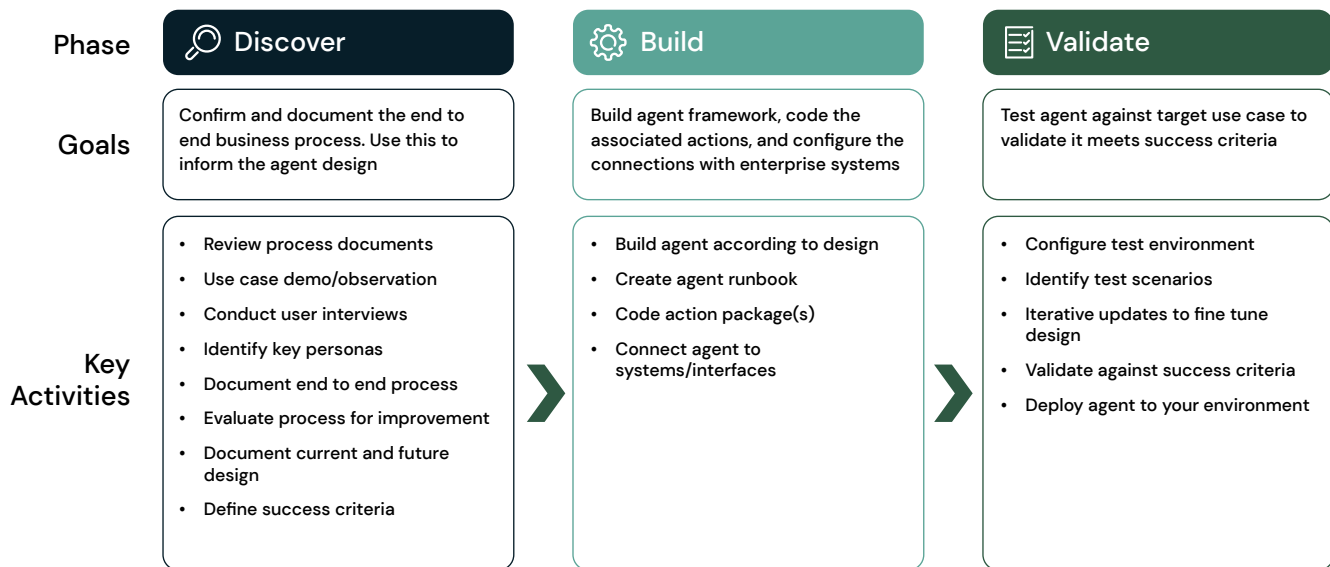
- Engage with stakeholders across various departments to identify potential areas where AI agents can add value.

Parameters to consider	Evaluation criteria
Does it solve a real urgent problem vs a nice to have issue?	Unless you are saving money, improving efficiency and solving a well understood acknowledged problem, even a successful solution will not translate into adoption gains.
How frequently are humans required to execute this business process?	Processes (like compliance at banks) have to be done continuously. Infrequent processes are less likely to translate into real gains.
Does 24*7 execution make the business more responsive?	If you can resolve a problem faster or get paid quicker or make decisions faster using agents, then it is a good candidate. Also if demand fluctuates and you need to auto scale, that is an even better candidate.
Is the problem repeatable across business divisions?	Solving repeatable mainstream problems will amplify adoption as subsequent efforts become easier.
Does the problem align with key imperatives of your organization?	If your CEO can talk about this on their earnings call, then it is definitely worth automating.

- Gather comprehensive baseline data, including current process metrics, pain points, and desired outcomes, to establish a clear starting point for the AI agent implementation. Because ROI requires you to understand the as-is state so that you can compare it to the new agentic implementation.

3.2 | Build agents - Rapid Agent Deployment (RAD)

- Conduct discovery sessions to document end to end business processes and key agent requirements.
- Develop Sema4.ai Runbooks, Actions, and Data Access for Agents working with line of business process owners.
- Use the Sema4.ai Rapid Agent Deployment (RAD) process and methodology for development, testing, and deployment into your Sema4.ai environments:



3.3 | Deploy

- Conduct a thorough readiness assessment, evaluating all aspects of the AI agent's performance to ensure it meets the required standards for production deployment, including reviews of functionality, scalability, and security measures.
- Execute a rigorous testing protocol and validate the AI agent's performance across a wide range of scenarios to confirm accuracy, reliability, and compliance with both technical specifications and business requirements.
- Ensure that approval processes, approval exceptions, and reasoning guard rails are in place and tested before you deploy the solution.
- Document key metrics to be used in Runbook and Agent training (sample of possible measures below):
 - > **Task completion rate:**
 - Percentage of tasks successfully completed by the agent
 - Quality and correctness of task outcomes
 - > **Contextual understanding**
 - Ability to maintain context over long conversations
 - Relevance of responses to the given context
 - > **Reasoning and problem-solving**
 - Ability to break down complex problems
 - Logical consistency in multi-step reasoning
 - > **Adaptability**
 - Performance across different domains or types of tasks
 - Ability to understand and follow new instructions
 - > **Policy compliance**
 - Adherence to company guidelines and policy constraints
 - Ability to recognize and act based on provided company policies



3.4 | Train

In the spirit of continuous improvement, it is important to use feedback and reflection techniques to deliver more accurate and consistent results, driving best practice outcomes.

- Track workflow exception paths in each use case and create a feedback loop to refine Runbooks, training agents to deliver better results with each process run.
- Track and analyze key performance indicators (KPIs) for each AI agent. These metrics will provide quantitative insights into the agent's effectiveness and areas for potential refinement.
- Establish channels for collecting ongoing feedback from both internal company employees and end users, ensuring we capture diverse perspectives on the agent's performance and user experience.

3.5 | Support

- Establish and adhere to mutually agreed-upon Service Level Agreements (SLAs) that outline response times, resolution targets, and escalation procedures for different types of support issues.
- Ensure that you are working with vendors who can demonstrate the same level of urgency as your enterprise in resolving issues.

3.6 | Educate

- Create and maintain a shared repository of use case materials, reference architectures, and best practices based on successful implementations.
- Document cost efficiencies, value generation and lessons learned from each successful implementation.
- Develop and offer certifications for key AI agent personas. Sema4.ai offers the following certifications which we think are essential for successful adoption of Generative AI across your organization.

Sema4.ai Certified Agent **Architect**

Sema4.ai Certified Agent **Developer**

Sema4.ai Certified Agent **Operator**



4 | Staffing the CoE

It is important to start small but ensure that you have folks with the right business domain experience and the right technical expertise as part of the CoE. It is also important to ensure that you have an influential sponsor who can create visibility for the CoE and can be a voice to amplify the success of the CoE. Having an enterprise architect who understands and can ask the right questions around security, accuracy, availability requirements of the agentic solutions is critical, so that there are no dropped balls by the time you get to deployment.

This section delineates the specific roles and responsibilities of both Sema4.ai and the AI Operations team in the CoE. It ensures clear accountability, promotes efficient collaboration, and establishes a framework for shared success in the implementation and ongoing operation of the CoE.

4.1 | AI CoE responsibilities

- **Assign ongoing key personnel, including:**
 - CoE sponsor (Senior Leadership)
 - Enterprise architect (Subject Matter Expert on Sema4.ai)
 - Program manager
- Participate in agent discovery, planning, testing, and feedback sessions.
- Manage internal change and adoption processes, when appropriate.
- Hold internal education sessions on successful use cases.
- Collaborate on internal case studies and white papers.
- Provide timely feedback on new and experimental product features.



5 | Governance and communication

This section establishes the structure for oversight, decision-making, and information sharing within the CoE. It ensures that the COE maintains alignment with internal teams, can quickly address challenges, and key stakeholders have clear channels for collaboration and progress tracking.

5.1 | Governance structure

- Establish a steering committee.
- Clearly defined escalation procedures.
- Weekly project status reports via project collaboration tools.
- Quarterly in-person CoE sponsor meetings to discuss progress, lessons learned.



6 | Performance metrics and success criteria

The performance metrics and success criteria provide a comprehensive framework for measuring the impact and effectiveness of the CoE. These metrics will guide ongoing optimization efforts, demonstrate value to stakeholders, and inform strategic decisions about the future direction of Sema4.ai Agent initiatives within your company.

The CoE will define and measure the below KPI targets. Once a quarter, these metrics will be reviewed with key stakeholders and CoE sponsors. This quarterly meeting will track progress, report lessons learned, and set targets for the next quarter.

- Number of conversations with Conversational Agents.
- Number of work items (eg. documents) processed by Worker Agents.
- Number of use cases onboarded.
- Use case pipeline by name and division, contact name and next step.
- Number of people certified to work on Sema4.ai.
- System performance benchmarking.
- End user satisfaction scores.




7 | Resources and tools

This section outlines the key resources and tools that a CoE needs to maintain. It ensures that the necessary infrastructure, expertise, and support systems are in place to facilitate successful Agent adoption and ongoing operations.

7.1 | CoE-provided resources

- Access to all necessary internal systems and data.
- Development and testing environments.
- Real-time, digital end-user feedback mechanisms.
- Real-time project management and collaboration tools.





Sema4.ai

At Sema4.ai, we're helping customers build, run, and manage intelligent AI agents that combine actions, intelligence, and enterprise context to automate complex human workflows.