

Executive Director, Agentic Lab and Architecture Lead - Remote

Job ID

REQ-10086322

Aug 21, 2026

LOC_US

About the Role

Reporting to the ED, Head of Agentic Factory, the Executive Director, AgentidLab and Architecture Lead defines the technical architecture, standards, and innovation roadmap for enterprise agentic AI capabilities. The role leads the Agentic Lab to prototype, validate, and industrialize next-generation AI agent technologies, while establishing reusable architectural patterns that enable secure, scalable, governed deployment across the enterprise.

Responsibilities:

Enterprise agentic AI architecture and standards

- Define the technical architecture, reference patterns, and engineering standards for enterprise agentic AI capabilities, ensuring alignment with approved platform, security, Responsible AI, and data governance requirements.
- Establish reusable patterns for agent orchestration, memory, evaluation, observability, tool use, RAG, APIs, and integration with enterprise systems.

Agentic Lab strategy, prototyping, and validation

- Lead the Agentic Lab innovation roadmap, focused on rapid experimentation, technical validation, and controlled incubation of high-potential agentic AI technologies.
- Evaluate foundation models, agent frameworks, orchestration technologies, developer tooling, and emerging technical approaches for enterprise applicability.

Reusable capabilities, accelerators, and reference implementations

- Build reusable proof-of-concepts, accelerators, architectural blueprints, and reference implementations that shorten time-to-value for the Agentic Factory and broader SPT product teams.
- Drive adoption of shared technical assets and standards that prevent one-off builds, reduce duplication, and improve scalability across AI agent solutions.

Industrialization and transition to enterprise scale

- Partner with AI Engineering, Platform Engineering, Product, Architecture, DDIT, security, and business teams to move validated prototypes into production-ready enterprise capabilities.
- Ensure successful transition from lab validation to scaled implementation, including architecture readiness, technical documentation, risk assessment, and operating model handoff.

Responsible AI, risk, and compliance by design

- Embed Responsible AI, security, privacy, accessibility, data governance, regulatory, and quality expectations into agentic AI architectures and lab validation methods from the start.

- Define evaluation and monitoring expectations for reliability, explainability, performance, user safety, human oversight, and business value realization of AI agents.

Technical leadership and capability building

- Lead, coach, and develop technical talent across agentic architecture and lab activities, fostering disciplined experimentation, engineering excellence, and continuous learning.
- Communicate technical direction, architectural decisions, trade-offs, risks, and investment needs to senior stakeholders in clear, executive-ready language.

Education:

Bachelor's or advanced degree in Computer Science, Artificial Intelligence, Engineering, or a related technical discipline.

Required experience:

- 12+ years of experience designing enterprise software platforms, AI architectures, data/AI products, or cloud-native engineering capabilities, including executive-level technical leadership responsibilities.
- Deep architecture expertise across LLMs, agent frameworks, multi-agent orchestration, RAG, vector and graph databases, tool/function calling, memory, evaluation, observability, APIs, event-driven patterns, and cloud-native deployment.
- Proven ability to define enterprise reference architectures, standards, reusable patterns, technical guardrails, and validated blueprints for secure, scalable, governed agentic AI deployment.
- Hands-on credibility with modern software engineering and AI delivery practices, including versioning, CI/CD, testing, release readiness, model/prompt evaluation, telemetry, cost/performance optimization, and operational support models.
- Demonstrated ability to move emerging technologies through disciplined technical scouting, lab validation, architecture review, production readiness, and handoff into engineering roadmaps.
- Strong fluency in Responsible AI, security, privacy, data governance, human oversight, access control, prompt/model risk, auditability, and compliance requirements for regulated enterprise AI.

Preferred experience and skill set:

- Experience building enterprise agentic AI platforms, developer ecosystems, or AI architecture practices in pharma, healthcare, life sciences, financial services, or another regulated environment.
- Familiarity with Microsoft Azure AI Foundry/OpenAI, Semantic Kernel, Copilot Studio, LangChain/LangGraph, MCP/A2A protocols, knowledge graphs, GraphRAG, and enterprise search architectures.
- Experience influencing senior technology and business leaders on architecture trade-offs, investment choices, platform reuse, operational risk, and scaling strategy.
- Track record coaching principal engineers, architects, AI engineers, and product teams on reusable patterns and disciplined experimentation.
- Delivery of agreed agentic architecture, lab, technical validation, governance, and capability milestones within planned timelines.
- Quality, reuse, and adoption of agentic AI reference architectures, standards, accelerators, proof-of-concepts, and reusable engineering patterns.
- Successful transition of validated lab concepts into production-ready enterprise capabilities with clear documentation, handoffs, and measurable value potential.
- Alignment with enterprise architecture, Responsible AI, security, data governance, privacy, accessibility, and applicable

regulatory requirements.

- Stakeholder confidence, technical risk transparency, executive-ready reporting, and continuous improvement of agent performance, reliability, and value realization.

The salary for this position is expected to range between \$225,400.00 and \$418,600.00 annual per year. The final salary offered is determined based on factors like, but not limited to, relevant skills and experience, and upon joining Novartis will be reviewed periodically. Novartis may change the published salary range based on company and market factors. Your compensation will include a performance-based cash incentive and, depending on the level of the role, eligibility to be considered for annual equity awards. US-based eligible employees will receive a comprehensive benefits package that includes health, life and disability benefits, a 401(k) with company contribution and match, and a variety of other benefits. In addition, employees are eligible for a generous time off package including vacation, personal days, holidays and other leaves.

Role Requirements

Division

DIV_IU

Business Unit

General Management

Location

LOC_US

Site

Remote Position (USA)

Company / Legal Entity

U014 (FCRS = US014) Novartis Pharmaceuticals Corporation

Functional Area

FCT_DD

Job Type

Full time

Employment Type

Regular

Shift Work

No

[Apply to Job](#)

Job ID

REQ-10086322

Executive Director, Agentic Lab and Architecture Lead - Remote

[Apply to Job](#)

Source URL: <https://jobapi.novartis.com/req-10086322-executive-director-agentic-lab-and-architecture-lead-remote>

List of links present in page

1. <https://jobapi.novartis.com/req-10086322-executive-director-agentic-lab-and-architecture-lead-remote>
2. https://novartis.wd3.myworkdayjobs.com/en-US/Novartis_Careers/job/Remote-Position-USA/Executive-Director--Agentic-Lab-and-Architecture-Lead_REQ-10086322-1
3. https://novartis.wd3.myworkdayjobs.com/en-US/Novartis_Careers/job/Remote-Position-USA/Executive-Director--Agentic-Lab-and-Architecture-Lead_REQ-10086322-1