Agentic AI in Life Sciences Discovery

Reinventing Life Science R&D with Agentic AI

Agentic AI refers to intelligent systems that understand their industry, can devise plans on their own, set goals, and execute strategies without constant human oversight. Unlike generative AI, which relies heavily on human prompts to generate content, agentic AI operates independently, making decisions and taking actions based on real-time data and predefined objectives.

A study by the Wharton School (Wharton AI & Analytics Initiative 2025) and Accenture assessed the potential impact AI agents and robots may have on the life sciences industry. By performing a bottom-up analysis, the study found that approximately 50% of hours spent in research & development (R&D) will be impacted by the emergence of AI agents and robots. This whitepaper takes a closer look at how adoption of agentic AI can accelerate and improve early discovery in life sciences by empowering scientists and democratizing access to data and information, and looks at how Accenture and AWS can jointly help unlock that potential, specifically focusing on how science is conducted and supported in life sciences R&D.

An AI-enhanced, multi-agent drug discovery process has the potential to outperform the traditional approach on key metrics, like return on investment (ROI) on R&D investments or probability of technical regulatory success (PTRS). It can deliver new drug candidates faster (compressing timelines by automating analysis and design), at lower cost (fewer failed experiments and more efficient trials), and with improved success rates (better decisions upfront lead to fewer late-stage failures). AI agents achieve this by not just extracting knowledge, but by applying reasoning capabilities and acting on complex workflows with dynamic data sources and tools. By lowering technological barriers and democratizing access to complex analyses and insights, they improve speed, cost efficiency, and decision accuracy at each phase of discovery. Thus, the agentic AI approach not only accelerates the path to first-in-human (FIH) trials, but also increases the probability that those trials will ultimately yield safe and effective new therapies for patients.

Shifting Imperatives

The pharmaceutical industry is facing mounting pressures that strain existing R&D models, slowing innovation, and inflating costs. Despite scientific advances, development timelines have plateaued, with costly late-stage failures and low success rates (BIO 2021). The ROI for R&D has significantly declined (Deloitte Global 2024), and drug development

costs are escalating due to protocol complexity, patient dropout and operational inefficiencies. To address these challenges, industry leaders are rethinking their operating models to deliver pipelines with higher productivity and lower costs. Specifically, they are looking for models that enable a large volume of targets with high PTRS and rapid clinical validation; prioritizing targets with multiple indications greatly expands the potential for success and higher ROI in early discovery and clinical trials.

To enable this new R&D operating model, pharma must leverage an integrated and global workforce of humans and technology. This R&D workforce of the future combines physical and digital agents with a human workforce that brings together specialized roles, strategic decisions, and global operations.

Agentic AI: A New Paradigm

Agentic AI represents the next level of evolution for AI systems, injecting reasoning capacity into workflows. This advanced AI can autonomously perform tasks, inform decisions, and interact with users or systems. Agents can analyze structured and unstructured biomedical data, plan multistep workflows under scientific or regulatory constraints, act by invoking downstream services, and learn continuously via reinforcement or active learning feedback loops.

The paradigm shift from generative AI to agentic AI usually involves a multitude of agents – from orchestrators to super and utility agents. By leveraging agent-to-agent communication and deploying agents in a multitude of models (e.g., structured flow, swarm, hierarchical, collaboration), it can help drive automation and enhance decision-making.

Establishing governance and setting up operating models is essential for maintaining the integrity and reliability of agentic AI systems. It involves setting policies, standards, and procedures that guide the development and deployment of AI technologies across the value chain. Effective governance ensures that AI systems are used ethically, comply with regulations, and deliver consistent results.

An Outline of a Future-State, Early Discovery Agentic Architecture

In early discovery and pre-clinical development, agentic AI can accelerate workflows, reduce costs, and improve success rates. Hierarchies of AI agents are assigned to each major task and stage of drug discovery. Each agent, or set of agents, isn't just automating a task, it contributes to a distributed intelligence system, where discovery emerges from

agent-agent and agent-human interaction. Shared utility agents (e.g., a literature search agent) also contribute across the value chain.

Target Identification and Validation Agents: Scan multi-modal data (genomics, proteomics, literature, clinical data) to identify disease-relevant targets. Target identification agents use natural language processing on publications and knowledge graphs to find novel target candidates. They interact with validation agents for feedback and rank targets by druggability and disease linkage confidence, passing top targets to molecule design agents. Where available, agents interact with existing predictive and statistical models, leveraging established tools and metrics rather than reinventing the wheel. These agents accelerate target selection by processing huge data volumes quickly, uncovering hidden patterns humans might miss, thus saving time and focusing resources on the most promising biological targets.

Molecule Design and Screening Agents: Generate or identify candidate compounds and molecules for a given target. This may involve a generative AI model (to propose novel molecules) and a virtual screening AI that evaluates large libraries computationally (in-silico). These agents filter out compounds unlikely to bind the target or meet basic criteria (using predictive scoring) and work closely with predictive modeling agents, e.g., sending generated molecules to an AI predictive model that evaluates binding affinity and drug-likeness. Promising hits are passed to the lead optimization stage. By automating compound generation and testing millions of candidates virtually, these agents dramatically speed up hit discovery and reduce cost. What once took months of lab screening can happen in days.

Lead Optimization Agents: Refine the hit molecules into optimized leads, using predictive models (QSAR, protein-ligand docking, ADMET predictions) to iteratively modify structures for better potency, selectivity, and safety. These agents apply a multi-objective optimization (e.g., balancing efficacy and toxicity), advancing only molecules that meet learned thresholds. They use generative models and property prediction tools to suggest new analogs, and provide integration with lab-in-the-loop systems to drive directed evolution for Design-Make-Test-Learn cycles – if a design fails to improve a property, the agent requests alternative modifications. AI agents shorten the design cycle, exploring chemical modifications faster than manual chemistry could, and improving decision accuracy by quantitatively predicting properties (e.g. binding affinity, toxicity) before synthesis. As a result, a more diverse set of candidates with better chances of success are pursued, reducing risk of late-stage failure.

Preclinical Testing and Risk Assessment Agents: Analyze and plan preclinical studies with AI support, using predictive safety models for toxicity, off-target effects, and pharmacokinetics and helping to design in vitro/in vivo experiments. Agents evaluate whether a lead is likely safe enough to enter clinical trials, forecasting risks (e.g. toxicity or efficacy issues) using learned patterns from past data. They interface with laboratory automation systems to prioritize experiments with a risk prediction AI that flags high-risk candidates. If a compound shows concerning predictions (e.g. potential liver toxicity), agents can recommend dropping it or looping back to the optimization agent for structural changes. By predicting many failure modes in-silico, these agents reduce costly late-stage failures and ensure only the most promising and safe leads advance.

Path to Reinvention

Agents can synthesize literature, internal omics data, and preclinical results to suggest novel hypotheses, shortening the time required from days to minutes. Similarly, where a biomarker lead would struggle to unify genomic, imaging, and IHC data across cohorts, an agent provides immediate feedback on robustness, cohort relevance, and translational potential. Agentic workflows seamlessly integrate and analyse multi-modal data from key trials to identify new predictive markers. Agentic AI handles all the labor-intensive work, not only speeding up the process but allowing many more possibilities to be considered, improving the likelihood of clinical success.

The integration of agentic AI into R&D processes is transforming the way pharmaceutical companies operate. By leveraging an integrated and global workforce of humans and agents, organization can optimize testing, decrease time to market, and increase the number of disease-relevant targets. Agentic AI enables a new R&D operating model that combines digital agents, human expertise, and physical agents to drive efficiency and scale. However, the path to reinvention involves preparing the human workforce for new ways working, rethinking workflows, and identifying existing internal assets for integration into utility agents and agentic workflows. Companies must plan for the skillsets of the future, emphasize AI and agentic training programs, and leverage global innovation to sandbox the right processes for "agentification". By establishing optimizing mechanisms and continuously monitoring value and performance, organizations can ensure the successful integration of agentic AI into their R&D processes.

Outlook

Agentic AI is no longer experimental: organizations are deploying agents in production, and it is a transformative technology that is beginning to reshape the pharmaceutical R&D landscape. By enabling faster, smarter decisions, and enhancing collaboration between humans and AI agents, agentic AI is driving the future of R&D. As companies embrace this technology, they will be better positioned to navigate the complexities of drug development, improve productivity, and deliver innovative treatments to patients more efficiently.

Laying the Foundation for Agentic AI

To truly benefit from the promise of Agentic AI, organizations must ensure several foundational elements are in place. These span strategy, technology, governance, people, and process alignment.

Holistic AI Strategy and Executive Alignment

A holistic, enterprise-wide approach is essential to realize the full benefits of agentic AI integration. Gone are the days of isolated, siloed use cases. Instead, a comprehensive strategy that aligns with an organization's vision and business objectives is needed to ensure that the removal of data silos does not result in agentic silos instead. To secure enduring executive alignment, organizations need to balance short-term wins with long-term transformation goals.

It is important to identify and prioritize use cases where Agentic AI can deliver the greatest business value – from automating complex workflows to enhancing decision-making to personalizing patient experiences. A readiness assessment can pinpoint gaps and opportunities, ensuring that resources are focused on initiatives with clear ROI and the potential to scale. Often, this means starting with smaller projects to gain understanding and trust, proving ROI before moving on to tackle larger opportunities.

Robust Digital Core – Technology and Data Foundations

To unleash the full potential of agentic AI, organizations must have a scalable, secure, and flexible *digital core*. This includes the ability to support autonomous, intelligent agents across multiple workflows and systems with governance, as well as organized access to data.

Life sciences organizations are no strangers to vast, heterogeneous datasets – from clinical to genomic to real-world evidence. But to truly harness the power of agentic AI, leaders must prioritize data readiness – break down those silos, standardize data, and enable secure, real-time access to high-quality information. Lakehouse architectures combining data lakes, data warehouses and databases with near-time data transfer and zero ETL are key to achieving these objectives.

Businesses also need to be ready to consider a future where the agent is the primary actor and therefore prepare data and integration architectures that cater more to agent interaction rather than human navigation. Data curation and integration become mission-critical, as AI agents rely on rich, accurate datasets for complex tasks like drug discovery, diagnostics, and regulatory submissions.

Finally, agents need to access this data seamlessly, whether it's in the cloud or on-premises. Seamless integration with existing platforms, such as customer relationship management (CRM) and enterprise resource planning (ERP), is essential. Moreover, the ability to orchestrate AI agents, humans, and robotics on a unified platform is key for maximizing productivity and control.

Trust: AI Governance, Risk Management and Security

Trust and accountability are paramount in the life sciences domain. As a highly regulated industry, agentic AI in life sciences organizations must operate within strict compliance frameworks, such as the EU AI Act, FDA regulations, and GxP standards. This means integrating AI governance and risk assessments into organizational structure, with a focus on transparency, interpretability, and documentation of AI models.

It is also important that in the use of these AI tools, scientific veracity is protected. This means that avoiding and detecting hallucinations is critical, as is the ability to go to the source data and/or code to reproduce findings.

In summary, guardrails and monitoring mechanisms are essential to manage autonomous agents, consistently and transparently assessing the impact on patient safety, data privacy, and ethical use.

Focus on People: Value-Focused Transformation and Change

Agentic AI should be positioned to augment, not replace, human workers. It needs to become a catalyst for business transformation rather than just an assistive tool. Most applications of AI today are assisting with current processes, and not deeply integrated into

revised workflows. While this may be appropriate for simple repetitive tasks, the work performed by scientists in research presents a higher degree of complexity and cross-functional collaboration.

Unlocking the full potential of agentic AI in research therefore requires more than plugging agents into existing workflows. AI initiatives should no longer be scoped around a single or small group of use cases, but instead around the end-to-end reinvention of a full process.

Leading organizations reimagine workflows entirely, considering agents as a core part of their workforce and preparing for a future where agents perform the majority of the tasks. This allows scientists to focus on higher-value, creative, and strategic tasks.

Genentech built an agentic solution that automates the time-consuming manual search process, enabling its scientists to focus on high-impact research and accelerate drug discovery innovation. Accepting that curent processes could be challenged, the system uses autonomous agents that can break down complicated research tasks into dynamic, multi-step workflows. Unlike traditional automation systems that follow predetermined paths, these agents adapt their approach based on information gathered at each step, thus enabling a new process workflow.

"Agents can't replace our scientists, but they actually boost our scientists. The really cool part happens when you let the agent and our top performing scientists work together. They then perform together far better than our best scientist, and that means that our best scientist is now a better version of themselves."

Aviv Regev, EVP and Head of Genentech Research and Early Development

---

However, life sciences workflows are complex and cross-functional. Agentic AI must be integrated in a way that supports regulated workflows – from clinical trial management to pharmacovigilance to commercial planning – while ensuring robust human oversight at critical decision points. Organizations must ensure clear processes for human-AI handoffs, defining where AI agents act autonomously and where human oversight is required.

Change management is crucial to prepare employees for these new ways of working. It is important to address resistance, build trust, foster a culture of innovation and continuous

learning, and upskill the human workforce to work alongside AI agents. AstraZeneca, presenting at the AWS Life Science symposium in 2025, described how it is seeking to address needs like this by using showcases, playgrounds, events, tips and testimonials, complemented with learning modules and accreditation. Many other organizations are investing time and resources into similar motions.

Attracting and retaining talent with hybrid skills that combine deep domain expertise and technical ability is paramount for life sciences organizations, and employees must be empowered to leverage evidence-based, AI-augmented decision-making.

Continuous Improvement and Adaptability

It is imperative to acknowledge that generative AI and agentic AI are fast moving fields of active research with many breakthroughs yet to materialize. As such, organizations must embrace a mindset of ongoing optimization, leveraging agentic AI ability to learn from data and adapt processes in real-time. Regular review and refinement of AI systems, workflows, and outcomes is necessary to stay responsive to changing business needs and market dynamics, as well as the breakneck speed at which AI capabilities are evolving.

Agentic AI has the power to transform the life sciences landscape. By aligning strategy, technology, governance, and human-centric approaches, organizations can unlock unprecedented efficiency, innovation, and patient-centricity.

Building the Agentic Future of Science

Accenture helps organizations navigate the complexities of AI reinvention with agility and precision, focusing on the human aspect of change and accelerating it using AI and behavioral science. Fundamental to any successful reinvention are strategy and alignment across the enterprise. To that end, Accenture created a new business unit called Reinvention Services (Accenture 2025) to consolidate its strategy, consulting, creative, technology, and operations units into a single entity focused on delivering end-to-end solutions and outcomes for clients. The goal is to make it easier to embed data and AI into solutions and accelerate value delivery.

AWS Life Science-Specific Agentic Capabilities

Today, 9 out of the top 10 global biopharma companies use AWS for generative AI, fueling a wide range of innovations from advancing protein design to accelerating clinical trial research and reducing false-rejects. Many of these life sciences leaders, including Roche, Lilly, and

AstraZeneca are now taking their innovations further by embracing agentic AI capabilities on AWS to drive greater efficiencies and accelerate discoveries at-scale.

To help organizations get started faster, AWS built the open-source **Life Sciences Agents Toolkit**, available on GitHub (AWS 2025). With over 20 starter agents, purpose-built for healthcare and life sciences use cases, the toolkit helps life sciences organizations accelerate the development and deployment of agents. Key features of the toolkit include an open-source code repository, conversational interface for rapid testing, multi-agent collaboration examples, fine-grained guardrails and identity policy examples, infrastructure-as-code templates for rapid deployment, and tools for public APIs, AI model endpoints, database queries, and code interpretation.

Robust Digital Core – Technology and Data Foundations

AWS' relentless drive for innovation spans the complete agentic AI architecture, with end-to-end investments at every layer of the AI stack, from AI assistants to foundation models to chips. With nearly twice the production-ready capabilities of other cloud providers, AWS takes care of the stability and reliability of agentic AI, making it possible for customers to focus on innovating.

Robust infrastructure and data management are crucial for AI agents as they require seamless access to diverse data types (structured, unstructured, and multimodal) and sources to function effectively. AWS provides a comprehensive set of services for the entire end-to-end data journey for all workloads, all types of data and all desired business outcomes, ensuring that data is optimized for agentic AI. For example, **Amazon Bedrock Knowledge Bases** provide one of the first managed, out-of-the-box, RAG solution that enables clients to natively query structured data wherever it resides. This capability helps break data silos across sources and can reduce the time to build generative AI apps from more than a month to just days.

Moving from agentic AI experiments to production-ready systems requires addressing fundamental infrastructure challenges around security, reliability, identity management, and operational complexity. AWS recently announced a comprehensive set of enterprise-grade services that eliminate the undifferentiated heavy lifting typically required to deploy and operate AI agents at scale. This capability, known as **Amazon Bedrock AgentCore,** helps customers deploy and operate AI agents securely at scale using any framework and model.

Data for Generative AI and Agentic AI

To address the challenges of getting the data landscape ready for generative AI and agentic AI, AWS and Accenture have created the **Accenture Data Modernization Factory for AWS**. This structured, repeatable framework modernizes data infrastructure to overcome organizational and technical barriers preventing the at-scale adoption of generative AI and agentic AI. The offering leverages a business-focused methodology with four key phases and has flexible entry points based on the customer's needs:

1. *Assess and Identify:* Identify where the high-value data and business use cases are.
2. *Architect and Migrate*: Design modern architectures and apply migration patterns.
3. *Activate and Augment*: Curate data into a single catalog and establish governance.
4. *Evaluate and Scale*: Initiate agentic solutions to feed business outcomes.

Cost

Cost is an important consideration when scaling agentic AI workflows. In addition to supporting model training and inference with its own chips, AWS has launched Amazon S3 Vectors to help increase the price performance of Vector stores. Vectors are the building blocks of generative AI applications and the new service is the first cloud object store with native vector support, reducing vector storage costs by 90% while maintaining sub-second query performance.

Choice

With the rise of agentic models, choice matters more than ever when it comes to model selection. AWS was the first to embrace a diversity of LLMs, with hundreds of foundational models available from leading companies through Amazon Bedrock. In late 2024, AWS announced the Amazon Nova family of models that are optimized for agentic applications. ("100% of the generative AI for as little as 25% of the cost.") AWS has continued to optimize these models with the release of a suite of Amazon Nova customization capabilities and the Nova Act SDK, which includes new security and access controls that help developers take their prototypes to production and build agents that can reliably take action in web browsers.

AI implementation is not one-size-fits-all. With AWS, clients can choose to buy, build, or partner to build AI solutions best suited for their organizations: deploy an AWS pre-built solution like Amazon Q, choose from the broadest and deepest set of cloud capabilities, or discover over 4,000 AI Solutions available in AWS Marketplace from ISVs and partners. With "AI Agents and Tools" in AWS Marketplace, customers have streamlined access to a

one-stop shop with over a thousand AI agent solutions and tools, allowing them to fast-track AI initiatives. Customers using AI Agents and Tools will build faster with these ready-to-integrate solutions, and can develop their AI strategies with professional services that specialize in building, maintaining, and scaling agents.

To effectively leverage agentic AI, organizations should align their strategy with the desired level of customization and control. AWS offers three main options: specialized agents, fully managed agents and do-it-yourself (DIY) agents.

Specialized Agents with Amazon Q

Amazon Q provides ready-to-deploy agents that leverage company content and data to automate tasks, provide answers, and take actions. Amazon Q Business and Q Developer are examples that allow immediate deployment of agentic experiences with minimal technical overhead, helping streamline processes and boost productivity across the organization.

Fully Managed Agents with Amazon Bedrock

A comprehensive toolset for developers to build custom AI agents in a secure, managed environment, Bedrock agents automatically break down tasks, create orchestration plans, and securely connect to company data through APIs without manual coding. It includes features like action groups, knowledge bases, and multi-agent collaboration with enterprise-ready security and guardrails.

DIY Agents

For organizations that have started with open-source agentic frameworks like LangGraph and CrewAI, this option offers maximum flexibility but requires more technical expertise. It helps address common challenges around awareness of Bedrock agents' value proposition and bandwidth constraints for migration. The DIY approach typically requires more time for technical validation, but provides the most customization options.

Additional specialist capabilities recently released include AgentCore Browser which provides a fast, secure, cloud-based browser to enable AI agents to interact with websites at scale for tasks like form completion or navigating a website.

Speed and Scale

Amazon Bedrock AgentCore Serverless Agent Runtime delivers low-latency, serverless environments with complete session isolation for added control, supporting any agent framework including popular open source frameworks like CrewAI, LangGraph,

LlamaIndex, and Strands Agents. The runtime automatically handles scaling from zero to thousands of concurrent session, while providing the longest running workload capability available today. Built-in checkpointing and recovery mechanisms ensure graceful recovery from unexpected interruptions, making it suitable for complex, long-running analytical workflows that span hours or days.

Interoperability

As highlighted, agent interoperability is paramount to a successful deployment. AWS has demonstrated its commitment to supporting agentic AI open standards by joining the steering committee for Model Context Protocol (MCP), an open protocol that standardizes how applications provide context to LLMs. By supporting multiple protocols, AWS ensures developers can build breakthrough agentic applications without being tied to one standard. The new Amazon Bedrock AgentCore gateway transforms existing APIs, AWS Lambda functions, and AWS services into agent-compatible tools with minimal code, eliminating months of integration work. The gateway provides a unified MCP interface across diverse tools while supporting cross-cutting features like authentication, authorization, throttling, and multi-tenancy. Tool selection capabilities help agents discover the most relevant resources for specific tasks, while runtime discovery enables dynamic workflow adaptation.

Trust: AI Governance, Risk Management, and Security

When it comes to building enterprise AI applications and agents responsibly, clients must stay vigilant to industry threats and evolving industry guidelines. AWS recently became the first major cloud provider to announce ISO 42001 accredited certification across AI services. ISO 42001 is an international standard that outlines requirements and controls for organizations to promote the responsible development of AI systems.

AWS has been investing in capabilities specific to Agentic AI to ensure clients can maintain compliance in regulated environments like life sciences. An example of this is the Enterprise Memory Management capability in the AWS Bedrock AgentCore Memory service that ensures agent context preservation, critical for agents performing complex, multi-step workflows. The service manages short-term (session) and long-term memory and supports memory sharing across collaborating agents, offering configurable retention policies.

Security and Responsible AI

To provide more trust and transparency into a model's decision-making process, AWS added automated reasoning safeguards in Amazon Bedrock Guardrails. AWS is the first

and only major cloud provider to integrate automated reasoning into their generative AI offerings for hallucination prevention. This is the first and only generative AI safeguard that helps prevent factual errors due to model hallucinations, pioneering a new scientific approach and opening new use cases that demand the highest levels of precision.

Observability of agentic AI is key to ensuring trust and transparency. AWS's recently released AgentCore Observability enables customers to track and trace every action of an agent and AgentCore Observability powered by Amazon CloudWatch gives developers real-time visibility through built-in dashboards and telemetry for key metrics, while integrating with existing observability systems.

Management of identity is even more important with agentic AI. Traditional identity models must evolve as agents act on behalf of users across multiple systems. AWS AgentCore Identity delivers secure agent access with temporary, fine-grained permissions and standards-based authentication, integrating with leading identity providers including Amazon Cognito, Microsoft Entra ID, and Okta.

To help tackle the wider need for responsible AI, Accenture and AWS have collaborated to release the **Accenture Responsible AI Platform, powered by AWS,** a fully automated, always-on solution that vigilantly monitors AI 24x7. With a focus on governance and risk management, it swiftly identifies and neutralizes any attempts to bypass AI safeguards from hallucinations to adversarial attacks. Leveraging insights from thousands of customers' generative AI experience, the platform provides industry-specific guidance for tailored support.

Focus on People: Value-Focused Transformation and Change

To help customers ensure a value-led transformation of key business processes, Accenture has created "Accenture AI / Agentic AI Enterprise Process Reinvention Services", available in the AWS marketplace. Delivered over three phases, this eight-to-18-week modular program is designed for industry-specific applications and successful delivery of business value powered by generative AI and agentic solutions on AWS. The key phases are:

- *Foundation* (2–4 weeks): Assess enterprise readiness, evaluate technology maturity, design compliance and security frameworks, define operating model, and estimate initial business value.
- *Automation Architecture* (3–6 weeks): Identify and redesign key processes, design orchestration and integration frameworks, build reusable components, and develop a refined value estimate.

- *Implementation and Scale* (3–6 weeks): Deploy pilot processes, implement compliance and change management, measure value delivery, and create a roadmap for scaling automation.

Continuous Improvement and Adaptability

AWS continues to lead the industry in generative AI and agentic AI innovation, demonstrating an unwavering commitment through substantial investments, rapid product development, and the launch of transformative technologies.

In 2025, AWS expanded its generative AI portfolio with over 100 new foundation models in Amazon Bedrock, introduced advanced agentic AI capabilities for multi-agent collaboration, and launched Amazon Transform—a service that leverages agentic AI to automate complex infrastructure and application modernization tasks.

These efforts, combined with continuous enhancements to custom silicon (like Trainium and Inferentia chips) and the integration of new AI agents in Amazon Q, highlight AWS' relentless pace of innovation and its role in shaping the future of AI for enterprises worldwide.

For over fifteen years, the powerful collaboration between Accenture and AWS has merged innovative technologies with human ingenuity to transcend business challenges, ignite transformation and cultivate boundless growth opportunities for our clients.

Jennifer Jackson, Accenture AWS Business Group (AABG) Global Lead

---

Bibliography

Accenture. 2025. "Accenture Changes Growth Model to Reinvent Itself for the Age of AI."

AWS. 2025. *Life Sciences Agents Toolkit*. https://aws-samples.github.io/amazon-bedrock-agents-healthcare-lifesciences/.

BIO. 2021. "Clinical Development Success Rates and Contributing Factors 2011-2020." https://go.bio.org/rs/490-EHZ-999/images/ClinicalDevelopmentSuccessRates2011_2020.pdf.

Deloitte Global . 2024. "Measuring the return from pharmaceutical innovation."

MCP. 2025. https://modelcontextprotocol.io/.

Wharton AI & Analytics Initiative. 2025. "Humans, AI and Robots: The economics of reinventing work and the workforce." https://www.accenture.com/us-en/insights/strategy/humans-ai-robots.

Jens Hoefkens, Ph.D.
Global Life Sciences Lead
Accenture AWS Business Group
jens.c.hoefkens@accenture.com

Jan-Paul Groeneboom
Global Life Sciences Lead
Accenture AWS Business Group
jpboom@amazon.co.uk

Dinesh Agaram
European R&ED Lead
Accenture Life Sciences
dinesh.agaram@accenture.com

Lisa McFerrin, Ph.D.
WW Lead- Research, Discovery, & Translational Medicine
lisamcfe@amazon.com