

Six Capability Pillars of the **AI-Era** Life Sciences Consultant

From Six Capabilities to
Evidence-Backed Deployment Readiness



Executive Summary

AI is no longer a specialist function in life sciences consulting; it is a baseline expectation across client engagements. Yet organizations often frame the shift primarily as a skills challenge – adding AI literacy to domain knowledge, communication, problem solving, outcome orientation, and business acumen. That is necessary, but it is not sufficient.

The real challenge is determining whether these capabilities converge in practice for a specific person, role, and client engagement. AI judgment, domain fluency, communication, problem framing, outcome discipline, and business context must work together on real problems. Course completion and self-assessed skills matrices alone cannot establish that readiness.

This paper introduces Life Sciences Capability Intelligence: an organization's ability to identify the gap between what a role requires and what an individual has demonstrated, develop against that gap, and build the evidence leaders need to make defensible deployment decisions. It distinguishes domain readiness, role readiness, project readiness, and deployment confidence.

Alexandria AI is the platform Agilisium built to operationalize this intelligence. Agilisium's Forward Deployment Experience (FDXe) – spanning Forward Deployed Experts (FDX), Forward Deployed Agents (FDA), and Forward Deployed Services (FDS) – puts this model into practice across its people and client work.


“

The harder problem is not naming six capabilities. It is knowing, for a specific person, a specific role, and a specific client engagement, whether those six capabilities actually converge.



1. The Consulting Model Is Changing

For much of the past two decades, life sciences consulting operated along two relatively distinct lines: domain specialists who understood science, regulation, and clinical or commercial workflows, and technology specialists with engineering, analytics, or data expertise. The two disciplines overlapped, but the separation largely held.



Generative and agentic AI are rapidly eroding that divide. McKinsey's research on scaling generative AI in life sciences highlights the growing need for capabilities such as prompt design, model evaluation, and translating AI outputs into defensible business or clinical recommendations. These capabilities sit between traditional domain and technology roles – an area where many organizations remain underprepared.

Recent Global Human Capital Trends research, based on more than 9,000 business and HR leaders across 89 countries, points to the same execution gap. Organizations redesigning work around human-AI collaboration report stronger financial results, yet only 6% of leaders say they are making meaningful progress. The intent to change is widespread; execution remains uneven.

For life sciences consulting, the implication is immediate. Client challenges rarely fall neatly into “domain” or “technology” categories. Accelerating clinical trial site selection, for example, requires an understanding of operational workflows, regulatory constraints, data quality, and the limits of AI-driven feasibility models. Increasingly, clients need professionals who can connect science, workflow, technology, and business consequence – and organizations that can demonstrate that capability rather than assume it.

2. Six Capabilities Now Have to Converge


The instinctive response to this shift has been to train more AI. Recent life sciences industry research found that 69% of life sciences companies plan to invest in digital fluency and AI upskilling – up 51% since 2023 – and that 94% of life sciences leaders now see AI as crucial to building scalable operations.

69%

Consider AI and GenAI skills critical for their future workforce

94%

Believe domain expertise remains essential



None of these six capabilities is new to consulting. What is different is that they now need to operate together, in the same person, on the same problem – not distributed across a larger team, and not developed as six unrelated training tracks that happen to share a syllabus.



It is not enough to have these capabilities separately. The real differentiator is the ability to bring them together – in context, on real work – to **create meaningful client impact**.

3. The Hidden Enterprise Problem: Capability Does Not Equal Readiness

Naming six capabilities is only the starting point. The harder enterprise challenge is determining readiness when professionals begin from different levels of experience, domain depth, and AI fluency. A capability framework alone cannot resolve those differences.

A technology-first consultant may be comfortable with AI tools but need deeper clinical or regulatory context. A domain specialist may bring years of market access experience yet need stronger AI fluency to judge whether a model's output is trustworthy. Senior leaders may require less foundational training and more cross-functional, AI-specific capability, while early-career professionals may need broader development across several pillars.

A single curriculum cannot address these starting points equally. Organizations need to know where each person stands today through assessment, what their intended role requires, where the gaps are, and what development should come next. They also need opportunities to apply capability to genuine business problems and evidence that it has been demonstrated – not simply inferred from course completion.



The question is no longer whether consultants have these capabilities in isolation, but whether they can apply them together – and whether organizations can prove it.

Common Proxies Used Today



Course completion



Internal assessments



Manager intuition



Past experience

Each of these signals has value. But none, on its own, proves whether a consultant can apply all six capabilities at the required level in the context of a specific client problem.

4. From Skills to Life Sciences Capability Intelligence

Life Sciences Capability Intelligence is an organization's ability to continuously determine what a role requires, what an individual has demonstrated, where the capability gap lies, how that gap is being closed, and whether evidence from real or realistic work supports a specific deployment decision.

Readiness is not a single state. A consultant may be domain-ready in clinical operations but not yet role-ready as an AI-enabled consultant. They may be role-ready for one type of engagement but not project-ready for a specific client context. Deployment confidence sits above all three: an evidence-based organizational judgment, not a single automated score.

When readiness is assumed rather than evidenced, gaps often surface after staffing-in client conversations, delivery quality, escalation risk, slower ramp-up, or increased dependence on senior specialists. The impact extends beyond learning effectiveness to deployment quality and confidence in delivery.

From
potential
to proven
impact.



This is why "ready" must be defined more precisely. In practice, it represents four distinct levels:

01



Domain-ready

Can this person operate credibly within the required life sciences context-clinical, regulatory, commercial, or access?

02



Role-ready

Can this person perform the specific capabilities required by the intended role, beyond domain knowledge alone?

03



Project-ready

Can this person apply those capabilities within the context, constraints, and stakeholders of a specific engagement?

04



Deployment confidence

Does the organization have sufficient evidence, from multiple sources, to make and defend a staffing decision?

5. FDX Readiness: Where Six Capabilities Converge

FDX readiness is demonstrated when six capabilities come together in the context of real client work. Consider an AI/data consultant preparing to support a Clinical Operations engagement focused on AI-enabled site feasibility – helping a sponsor determine which investigator sites can realistically meet enrollment targets for a new protocol.

- 01 AI Literacy**
Can they interrogate the feasibility model's confidence and limitations rather than present its output as fact?
- 02 Domain Literacy**
Do they understand site feasibility, patient recruitment dynamics, and the operational constraints under which investigator sites work?
- 03 Problem Solving**
Are they solving the sponsor's actual site-selection problem, rather than simply applying an available AI tool?
- 04 Communication**
Can they explain the recommendation to Clinical Operations stakeholders in a way that supports a real decision, not just a technical readout?
- 05 Business Acumen**
Do they understand the cost, timeline, and risk trade-offs the sponsor is weighing?
- 06 Outcome Orientation**
Can they connect the recommendation to outcomes the client cares about, such as enrolment velocity or study execution risk?

FDX readiness is not demonstrated through six capabilities in isolation. It emerges when all six are applied together to a real client problem.

A consultant may be strong in AI literacy, for example, but if they cannot connect the recommendation to the domain context, business decision, and intended outcome, that capability remains isolated. An FDX-ready professional demonstrates all six capabilities together at the level required for the engagement – and the organization has evidence to validate that readiness. This is the convergence that Life Sciences Capability Intelligence is designed to make visible.

6.

Alexandria AI: Operationalizing Capability Intelligence

Alexandria AI is Agilisium's Life Sciences Capability Intelligence Platform, built to translate capability requirements into individualized development, applied evidence, and defensible readiness decisions. It operationalizes the Capability Intelligence model as one connected system – not a course library or a single AI-generated score.




Together, Alexandria AI connects baseline assessment, gap identification, contextual development, applied evidence, and multi-source readiness decisions – turning Capability Intelligence into an operational system for workforce readiness.

7. FDXe: The Applied Deployment Model

The six-pillar model defines the capability standard. Alexandria AI operationalizes and verifies readiness against that standard, while Agilisium's Forward Deployment Experience (FDXe) puts verified capability to work in real life sciences engagements.

FDXe brings together three connected forms: Forward Deployed Experts (FDX), who combine domain knowledge, AI-enabled execution, and business judgment in client work; Forward Deployed Agents (FDA), which extend that capability into execution; and Forward Deployed Services (FDS), the delivery model through which FDX and FDA are applied to client engagements.

FDXe is not separate from the capability model – it is its applied expression. Agilisium uses the same six-pillar standard and readiness evidence to assess its own forward-deployed professionals before deployment. In this way, Capability Intelligence moves from a framework for development to a standard applied in real client delivery.



Same capability. Real-world impact.

FDXe brings together verified people, AI agents, and delivery services to solve real client challenges in life sciences – with speed, quality, and measurable outcomes.



Conclusion

The AI-era life sciences consultant needs six capabilities working together: AI literacy, domain literacy, communication, problem solving, outcome orientation, and business acumen. The challenge is not simply defining them, but determining whether they converge in a specific person for a specific engagement.

Life Sciences Capability Intelligence provides that discipline-comparing required capability with demonstrated capability, closing the relevant gaps, and building multi-source evidence to support readiness decisions. Alexandria AI operationalizes this approach, while Agilisium's FDXe model applies verified capability in real client work.

For life sciences leaders, the question is therefore not whether the workforce has completed enough AI training. It is whether the organization can demonstrate that its people are domain-ready, role-ready, project-ready, and ultimately FDX-ready when a real deployment decision is made.



References

- Clarkston Consulting. "2026 Life Sciences Industry Trends."
clarkstonconsulting.com/insights/2026-life-sciences-industry-trends/
 - Deloitte. "2026 Global Human Capital Trends."
deloitte.com/us/en/insights/topics/talent/human-capital-trends.html
 - McKinsey & Company. "Scaling gen AI in the life sciences industry."
mckinsey.com/industries/life-sciences/our-insights/scaling-gen-ai-in-the-life-sciences-industry
- 



About Agilisium

Agilisium is the life sciences industry's AI-powered data innovation partner, working with pharmaceutical, biotech, and medical device companies to build data-driven, AI-enabled solutions across the value chain. Alexandria AI and the Forward Deployment Experience (FDXe) model – spanning Forward Deployed Experts (FDX), Forward Deployed Agents (FDA), and Forward Deployed Services (FDS) – described in this paper reflect Agilisium's own approach to building an AI-era consulting workforce – one Agilisium is applying internally before extending the conversation to clients.

To explore how Alexandria AI can assess, develop, and validate capability against specific life sciences roles and project requirements.

Contact sales@agilisium.com or visit <https://www.agilisium.com/>

About Alexandria AI

Alexandria AI is a capability assurance platform that certifies readiness rather than only delivering learning. Built from an organisation's own material through Agilisium's consulting-led model, each pathway applies adaptive assessment, a personalised pathway, and a scored business capstone across four capability levels: Explorer, Practitioner, Orchestrator, and Mastermind. Every learner receives a Go/No-Go verdict with a documented rationale, giving leaders auditable evidence of readiness before deployment.

Contact sales@vidalis.ai or visit <https://vidalis.ai/>