

GraphRAG Pilot FAQs for AI Product Teams

Execution-readiness answers for Product Managers and AI Product Leads evaluating a GraphRAG pilot.

Quick answer

What does a GraphRAG pilot actually let us test?

GraphRAG pilots let product teams test answer quality, retrieval performance, traceability, and orchestration in a controlled scope before committing to a wider rollout. A well-scoped pilot proves whether GraphRAG improves product experience, supports more reliable answers, and fits realistic delivery timelines — before it becomes a roadmap commitment.

FIT CHECK

1

How do we know if GraphRAG is right for our AI product?

GraphRAG is worth piloting when simple retrieval or keyword search is falling short. Good signals include:

- Users ask complex, multi-step questions
- Answers need citations or provenance
- Domain context and relationships between entities matter
- Vector-only RAG is producing shallow or incomplete answers
- Your team needs more confidence before scaling an AI feature

For a deeper self-assessment, including how this plays out across regulated and complex-data workflows, see [GraphRAG for AI-Powered Products](#).

PROOF POINTS

2

What should a GraphRAG pilot actually prove?

A good pilot is judged against a short, specific list — not a general impression of "it works":

- [Better answer relevance and multi-hop reasoning](#) — validated against a small set of objective metrics agreed with your team, chosen to fit the specific use case rather than one fixed benchmark
- [Clear citations and traceability](#)

- [Manageable latency for the target workflow](#)
- [Realistic integration effort](#)
- [User value in a defined, real workflow — not a demo scenario](#)
- [Governance and access-control feasibility \(at a pilot-appropriate level\)](#)

This also covers whether a full ontology is needed up front: it isn't. A pilot can start with a lightweight relationship model scoped to the use case and mature it as value is proven.

TECH STACK FIT

3

Does GraphRAG integrate with our existing tech stack?

Yes, and integration is usually additive rather than a replacement. In practice:

- GraphRAG sits alongside your existing vector store rather than requiring you to rip it out
- It connects to common data sources — document stores, APIs, structured databases — through standard connectors
- It works with major LLM providers, so it doesn't lock you into a specific model
- It can be exposed through your existing product API layer, so the frontend and app logic don't need a rebuild

The pilot is the right place to confirm this for your specific stack: it validates integration effort against your real systems rather than a generic checklist.

SCOPE & TEAM

4

How is a GraphRAG pilot scoped, and who needs to be involved?

Pilots typically run 6–8 weeks across use-case definition, data selection, retrieval setup, orchestration, and evaluation. For the full phase-by-phase breakdown, see the 6–8 Week GraphRAG Pilot Roadmap.

From a staffing perspective, plan for:

- A product owner to define success criteria and workflow scope
- IT support for source access and integration
- Subject-matter input to validate answer quality and relationships

Datavid runs the technical delivery; your team's main commitment is scoping time up front and evaluation time at the end.

USER EXPERIENCE

5 Does GraphRAG change what users see in the product?

Not necessarily. Many teams keep the exact same look and feel and simply get stronger foundations underneath — better answers, more reliable reasoning, and cleaner traceability, with no visible change to the interface.

If you do want the improvement to be visible, all of the following can be added:

- Inline citations or source references alongside answers
- Follow-up or related-question suggestions grounded in the graph
- Higher answer confidence on questions that span multiple sources or topics


Whichever direction you choose, this is a product decision rather than a technical constraint — the pilot can validate the underlying answer quality either way, and any UI changes can be layered on afterwards if and when they're wanted.

PERFORMANCE

6 What performance questions should we be asking?

Framing performance in product terms, not just infrastructure terms, means:

- Defining an acceptable response time for this specific workflow before the pilot starts, so there's a clear target to measure against
- Identifying which step adds the most latency — retrieval, graph traversal, ranking, or generation — so effort goes where it actually matters
- Routing simpler queries down a faster path, reserving full graph traversal for the queries that need it
- Using caching or pre-computation for common questions to cut repeat latency
- Agreeing upfront on how performance will be measured and reported during the pilot, so results are comparable at the end



Lean, predictable step



Where latency accumulates

PERFORMANCE

7 How does GraphRAG handle latency?

GraphRAG can introduce additional steps compared with simple vector retrieval, but latency is managed through query routing, caching, retrieval optimisation, and scoping the pilot to high-value workflows rather than every possible query. The pilot's job is to test whether the gain in answer quality, citations, and traceability justifies the response time for that specific user experience.

ORCHESTRATION

8 How does GraphRAG support AI agents or agentic workflows?

GraphRAG provides a trusted context layer for agents: connected, governed knowledge that constrains and grounds what an agent retrieves, reasons over, and acts on. For product teams, the value isn't automation in the abstract — it's giving agents context they can use to produce more reliable, traceable outputs on a specific task. Orchestration should be piloted around one or two real user tasks, not tested as a general capability.

DELIVERY RISK

9 What are the main delivery risks in a GraphRAG pilot, and how are they reduced?

Most pilot risk comes from scope and definition, not the technology itself:

- Unclear use case — fixed by defining one workflow and one user group before kickoff
- Poor data quality or weak success criteria — fixed with a short readiness check in week 1
- Overly broad scope — fixed by capping the pilot to a single high-value workflow
- Unclear ownership between product, data, and engineering — fixed with named owners at kickoff
- Governance considered too late — fixed by scoping access control from day one, even at pilot scale
- Cost overrun — fixed by capping the pilot to a fixed budget and timeline upfront, with any extension tied to a defined trigger such as proven value, not open-ended iteration

SECURITY & COMPLIANCE

10 What are the security and compliance considerations for a GraphRAG pilot?

At pilot scale, security and compliance stay practical and scoped down, but they're not skipped:

- The pilot runs on a defined, limited data set rather than a full production feed

- Access is restricted to the people evaluating the pilot, using role-based permissions that match your existing policies
- Sensitive or regulated data — PII, financial, health — is identified upfront and either excluded or handled under your existing compliance requirements
- Datavid works within your security and data residency requirements rather than introducing new infrastructure
- A pilot-appropriate audit trail is in place so any answer can be traced back to its source

This is deliberately lighter than a full production security review — the pilot's job is to prove the approach works without taking on production-level risk. A complete governance, RBAC, and audit model is scoped separately when moving toward production.

SUCCESS METRICS

11

How should we measure whether the pilot succeeded?

Useful, PM-ready metrics include:

- Improvement in answer relevance vs. current baseline
- Reduction in incomplete or shallow answers
- Citation and provenance coverage
- User task completion rate
- Response time against an agreed threshold
- SME or user confidence score
- Reduction in manual lookup effort
- Feasibility of integrating the approach into your existing roadmap

The aim is to measure both technical performance and real user value. A pilot should show not only that GraphRAG can produce better, more traceable answers, but that it improves a defined workflow enough to justify the next stage of investment.

See it in practice: Roche worked with Datavid to deliver a functional semantic knowledge prototype in six weeks, helping users get faster answers, reduce manual advisor effort and validate a path towards broader AI-powered support.
[Read the Roche case study.](#)

NEXT STEPS

12

What happens after the pilot — and what does Datavid deliver?

At the end of the pilot, Datavid delivers a pilot report and roadmap recommendation, covering what worked, what to adjust, and what production would require. From there, typical next steps are to:

- Proceed to MVP with the validated use case
- Narrow the use case if results were mixed

- Improve data readiness before scaling further
- Expand into additional workflows once the first is proven
- Build the production architecture and investment case

WHY DATAVID

13

Why should we consider Datavid over building in-house or using another vendor?

Each option carries a different trade-off:

- Building in-house means finding and retaining specialised graph and retrieval engineering talent, then carrying the full timeline and risk of a first attempt yourself
- A generic AI vendor may move fast but often lacks specific experience with knowledge graphs or regulated, complex data — the exact conditions GraphRAG is built for
- Datavid has delivered GraphRAG and knowledge graph pilots in regulated, complex-data environments — including Syngenta, Roche, and ABN AMRO — so the pilot draws on patterns that already work, not a first attempt
- Delivery is fast without cutting corners: a functional pilot in 6–8 weeks, not a multi-quarter build
- Governance and traceability are built in from day one rather than retrofitted once something breaks
- The engagement is pilot-first and low-risk: you validate value before committing to a larger build, in-house or otherwise
- It doesn't have to be all-or-nothing: Datavid can build the solution with your team in the loop from the start, then hand it over cleanly so you manage it in-house — you keep the knowledge internally while still de-risking the initial build

In short: Datavid removes the build-it-yourself timeline and risk, and brings graph-specific, regulated-data delivery experience that most generalist vendors don't have.

Ready to reduce delivery risk before it hits your roadmap?

Use the discovery pilot to validate whether GraphRAG can improve answer quality, traceability, and product experience — in a controlled 6–8 week scope, before it becomes a roadmap commitment.

Book a GraphRAG Discovery Pilot