



Working with Graph Research Labs

GRL helps organisations turn knowledge graph and agentic AI work into production-grade, governed infrastructure that can perform at the level regulated industries and organisations require.

Five engagement shapes. Each is bespoke; quoted against your specific scope, environment, and outcomes.

01

GRL Engagement - Information Modernisation

Led by Dougal Watt, GRL connects any system in your business - legacy or modern, structured or unstructured - into a unified, real-time, auditable intelligence layer with full provenance. Information once stranded across ageing machines, document stores, and disconnected applications becomes usable in weeks, not years, at a fraction of typical cost and risk.

Our tooling and methodology bridge old and new tech, and are stack-agnostic across the semantic graph ecosystem. They produce data assets your engineers already know how to operate, run inside your secure environment as an asset to your security rather than a liability, and adapt as the landscape shifts. New databases, new ontologies, new AI platforms - GRL acts as the bridge for data to move seamlessly across open-standard technology. With GRL you're never locked into a 2026 runtime or a vendor's platform.

Whether you're bringing legacy into a modern data fabric, extracting information to retire a system, or building knowledge graph and agentic AI capabilities, GRL compresses work that normally takes years.

Engagements are scoped to your situation: architecture assessment, modernisation programmes, ontology engineering, governance design, knowledge transfer, production deployment, and ongoing advisory - including GRL tooling licences where applicable.

Best for organisations with mission-critical information challenges that require a strategic view of their entire information landscape - fragmented systems, legacy data, stalled modernisation, or KG and agentic AI initiatives that need senior expertise.

02

GRL Agent Acceleration Engagement

GRL works alongside your team to take your existing agentic AI infrastructure to production-grade accuracy (90%+) and meaningfully reduced latency. Through governed semantic context and the GRL Semantic AI Agent Harness, we close the typical 50–60% production-accuracy gap while reducing per-query cost. The trade-off between speed, accuracy, and compliance, which most enterprise AI initiatives accept as inevitable, is structurally eliminated when the agent has a governed substrate to reason against.

Best for organisations with deployed or in-development agent infrastructure that needs to clear the production-readiness threshold for regulated environments.

03

GRL Strategy Workshop

A focused day with Dougal Watt and your senior team. Structured problem-solving on your specific knowledge graph, ontology, or agentic AI challenges. Working session with your senior team, ending with a written report containing directional recommendations.

Best for organisations weighing approach options or wanting senior practitioner perspective on a specific question, before committing to a larger engagement.

Bookable directly via Stripe, or via standard procurement for organisations that prefer formal procurement.

04

GRL Assessment Workshop

GRL tools applied to your existing ontology and production results, with senior practitioner analysis from Dougal Watt and team. We deploy our tools against your data, run conformance and capability analysis, and produce a written assessment of your current state, governance posture, and capability uplift opportunities, including specific, costed recommendations for closing the gap to production-grade reliability.

Best for organisations with deployed ontology or agent infrastructure who want a concrete tools-based diagnostic before committing to a larger engagement, or who need formal documentation of their current production-readiness posture.

Bookable directly via Stripe or via standard procurement for organisations that prefer formal procurement.

GRL AI Safety & Governance Engagement

Led by Dougal Watt, GRL helps organisations make their AI deployments safer through architecture and model-level design, with a custom Semantic Agent Harness - provably auditable, compliant with internal and external regulation, and strategically aligned with the organisation's compliance needs.

The most consequential AI incidents to date have not been failures of the model. They have been failures of architecture: poor provenance, weak access controls, missing audit trails, brittle integration, and policy gaps between systems. AI safety in regulated environments is both a model and an architectural discipline. Model-level controls are necessary, but they are not sufficient. The architecture of the Agent Harness surrounding the model is where compliance lives or dies.

GRL works with your model risk, AI leadership, and compliance teams to convert AI from an ungoverned, unstructured activity into a governed, defensible capability, anchored in the organisational architecture rather than tacked on.

Engagement scope (modular): architectural safety review, agent and model-risk posture assessment, provenance and lineage architecture, internal AI policy authoring, regulator-readiness preparation, board-level governance briefings, ongoing advisory, including GRL tooling licences where applicable.

Best for organisations in highly regulated, mission-critical industries whose AI ambitions require safety and governance built in, not tacked on.

TO ENGAGE

To learn more, contact: dougal.watt@graphresearchlabs.com.
