

GraphRAG, but with Governance: Where Microsoft's Architecture Stops Short

GraphRAG validates the thesis that knowledge graphs beat naive retrieval for entity-rich reasoning. It uses a language model to extract entities and relations into a structured graph that is then queried. The substrate goes further: the graph is not just a query target, it is a governed computational medium where every traversal step is admissibility-gated.

What GraphRAG Is

GraphRAG is Microsoft Research's approach to retrieval-augmented generation over a knowledge graph rather than over a flat vector store. As described in its public materials, the method uses a language model to read a source corpus and extract entities and the relationships among them into a structured graph, derives summaries over clusters of that graph, and then answers a query by retrieving from the graph and its summaries instead of from isolated text chunks. The motivation is sound and well evidenced: naive vector retrieval struggles with questions that require connecting many entities or reasoning over a whole corpus, and a graph gives the system structure to traverse where a bag of embeddings gives it none. GraphRAG has been published as research and graph-based retrieval is being integrated into Microsoft's Azure offerings, which makes it the most visible validation of a thesis the discovery substrate also holds: structure beats undifferentiated retrieval for entity-rich reasoning.

The agreement is real and worth stating plainly. GraphRAG and the discovery substrate both reject the premise that retrieval should be a flat nearest-neighbor lookup, and both treat a structured, navigable representation of knowledge as the right medium for reasoning. The difference is not about whether to use a graph. It is about what the graph is permitted to be.

Where It Stops Short

In GraphRAG as described, the graph is a passive data structure. A language model builds it, and a language model queries it, but the graph itself does not govern the traversal that crosses it. There is no per-step admissibility evaluation on a transition from one node to another, no typed object that carries the query's intent, accumulated memory, policy, and lineage as it moves, and no record of why each step was taken that travels with the result. Hallucination is therefore handled the way it is handled everywhere in the retrieval-augmented family: as a statistical risk to be reduced by giving the model better context, not as a category of output that the architecture can refuse. The graph improves what the model is shown; it does not constrain what the model is allowed to commit.

This is the structural ceiling. A graph that is only a query target inherits the governance properties of the model that queries it, which is to say it has none of its own. The provenance of an answer is whatever the model reports, the constraints on an answer are whatever the prompt manages to hold, and the line between a supported step and an unsupported one is drawn by attention rather than by rule.

The Governance Increment

In the discovery substrate the graph is not a query target; it is a governed computational medium, and every step across it is a single fused operation. At each anchor the traversal narrows the candidate set, updates the discovery object's structured semantic state, and evaluates the proposed transition for admissibility,

simultaneously and inseparably. The inference engine proposes; the execution substrate disposes. A proposed transition is admitted, rejected, or decomposed against policy, lineage continuity, entropy bounds, and temporal validity before the traversal advances, and that determination is recorded in the discovery object's lineage whether the step was taken or refused. Because the model only proposes and the substrate decides, an unsupported or non-compliant step is a category failure that the architecture rejects, not a probability the architecture tries to lower. The result arrives with the full sequence of transitions and admissibility determinations that produced it, an epistemic lineage rather than a model's after-the-fact explanation.

The increment over GraphRAG is therefore not a better extraction step or a richer graph. It is the addition of governance at the granularity of the traversal step: a typed object carrying intent and constraint, an admissibility gate at every transition, and a lineage that makes the path auditable. The graph stops being something a model reads and becomes something a model is held to.

Complementary Positioning

GraphRAG and the discovery substrate are not competitors for the same slot.

GraphRAG is a strong retrieval augmentation: it is a better way to give a model context for entity-rich questions, and an organization already invested in it has built exactly the kind of structured representation the substrate operates over. The substrate is a governance upgrade on top of that premise. An extraction pipeline like GraphRAG's can populate the index the substrate governs; the substrate then adds the per-step admissibility, the carried typed object, and the lineage that GraphRAG does not provide. The path from one to the other is additive, not adversarial: keep the graph, govern the traversal.

Disclosure Scope

The fused search, inference, and execution traversal step, the structural separation between the inference engine as proposal generator and the execution substrate as authority, and the treatment of unsupported output as a category failure rejected at the admissibility gate rather than a statistical risk, are disclosed in the cognition filing (U.S. Application No. 19/647,395 and its international counterpart) at Sections 10.3, 10.5, and 10.8. This article compares those disclosed mechanisms with Microsoft's publicly described GraphRAG approach and positions the substrate as a governance layer composable with graph-based retrieval. References to GraphRAG and Azure are to Microsoft's public materials and are used for comparison only; no relationship, endorsement, or infringement is asserted.

Semantic Discovery (</semantic-discovery>)

[All 36 steps → \(/inventive-steps\)](/inventive-steps)

Search, inference, and execution as one governed step.

PRIMARY TECHNICAL DISCLOSURE

- [Governed Semantic Discovery: Search, Inference, and Execution Through Adaptive Traversal \(/articles/governed-semantic-discovery-search-inference-and-execution-through-adaptive-traversal\)](/articles/governed-semantic-discovery-search-inference-and-execution-through-adaptive-traversal)

SECONDARY TECHNICAL

- [The Adaptive Index as Unified Search-Inference-Execution Substrate \(/articles/semantic-discovery/unified-substrate\)](/articles/semantic-discovery/unified-substrate)
- [Three-in-One Traversal: Search, Inference, and Execution in a Single Step \(/articles/semantic-discovery/three-in-one-traversal\)](/articles/semantic-discovery/three-in-one-traversal)
- [The Discovery Object: A Traversal-Native Semantic Agent \(/articles/semantic-discovery/discovery-object\)](/articles/semantic-discovery/discovery-object)
- [Post-PageRank Semantic Ranking: Relevance Through Governed Traversal \(/articles/semantic-discovery/post-pagerank\)](/articles/semantic-discovery/post-pagerank)

- [Persistent Semantic State: Eliminating Prompt Reconstruction \(/articles/semantic-discovery/persistent-state\)](/articles/semantic-discovery/persistent-state).
- [Traversal Lineage as Index Evolution Signal \(/articles/semantic-discovery/traversal-lineage\)](/articles/semantic-discovery/traversal-lineage).
- [Anchor Semantic Neighborhood Publication \(/articles/semantic-discovery/semantic-neighborhoods\)](/articles/semantic-discovery/semantic-neighborhoods)
- [Inference-Time Execution Control as Traversal Primitive \(/articles/semantic-discovery/inference-governance\)](/articles/semantic-discovery/inference-governance).
- [Anchor Self-Organization Under Entropy and Load Pressure \(/articles/semantic-discovery/anchor-self-organization\)](/articles/semantic-discovery/anchor-self-organization)
- [Alias Resolution as Navigational Traversal \(/articles/semantic-discovery/alias-resolution\)](/articles/semantic-discovery/alias-resolution)
- [Three Discovery Operating Modes: Human Search, Agent Reasoning, Answer Synthesis \(/articles/semantic-discovery/operating-modes\)](/articles/semantic-discovery/operating-modes)
- [Model-Agnostic Semantic Discovery \(/articles/semantic-discovery/model-agnostic\)](/articles/semantic-discovery/model-agnostic)
- [Affect-Modulated Discovery Traversal \(/articles/semantic-discovery/affect-modulated-traversal\)](/articles/semantic-discovery/affect-modulated-traversal)
- [Confidence-Gated Discovery Traversal \(/articles/semantic-discovery/confidence-gated-traversal\)](/articles/semantic-discovery/confidence-gated-traversal)
- [Integrity-Tracked Traversal Drift Detection \(/articles/semantic-discovery/integrity-tracked-drift\)](/articles/semantic-discovery/integrity-tracked-drift)
- [Biological Identity-Scoped Access During Discovery \(/articles/semantic-discovery/biological-access\)](/articles/semantic-discovery/biological-access).
- [Rights-Grade Anchor Governance for Content Discovery \(/articles/semantic-discovery/rights-grade-anchors\)](/articles/semantic-discovery/rights-grade-anchors)
- [Forecasting-Shaped Discovery Traversal \(/articles/semantic-discovery/forecasting-shaped\)](/articles/semantic-discovery/forecasting-shaped)
- [Capability-Constrained Anchor Accessibility \(/articles/semantic-discovery/capability-constrained\)](/articles/semantic-discovery/capability-constrained)
- [Collaborative Multi-Object Discovery Traversal \(/articles/semantic-discovery/collaborative-traversal\)](/articles/semantic-discovery/collaborative-traversal)
- [Discovery-Driven Sensor Invocation Closed Loop \(/articles/semantic-discovery/sensor-invocation-loop\)](/articles/semantic-discovery/sensor-invocation-loop)
- [Cross-Platform Credentialed Reader Activation \(/articles/semantic-discovery/credentialed-reader-activation\)](/articles/semantic-discovery/credentialed-reader-activation)
- [LLM-as-Bootstrap: Why Anchor Inference Engines Shrink as the Lineage Matures \(/articles/semantic-discovery/maturation-engine-substitution\)](/articles/semantic-discovery/maturation-engine-substitution)
- [Personal Cognitive Asset: How Per-User Lineage Re-Weights the Same Substrate \(/articles/semantic-discovery/personal-lineage-layer\)](/articles/semantic-discovery/personal-lineage-layer)
- [Loki, the Dog, and the Symbol Grounding Problem \(/articles/semantic-discovery/hybrid-symbol-grounding\)](/articles/semantic-discovery/hybrid-symbol-grounding)

APPLICATIONS · GENERAL

- [Enterprise Knowledge Management Through Governed Traversal \(/articles/semantic-discovery/enterprise-knowledge-management\)](/articles/semantic-discovery/enterprise-knowledge-management)
- [AI-Native Search That Replaces PageRank With Contextual Relevance \(/articles/semantic-discovery/ai-native-search\)](/articles/semantic-discovery/ai-native-search)
- [Semantic Discovery for Scientific Research \(/articles/semantic-discovery/scientific-research-discovery\)](/articles/semantic-discovery/scientific-research-discovery)
- [Semantic Discovery for Legal Case Research \(/articles/semantic-discovery/legal-case-research\)](/articles/semantic-discovery/legal-case-research)
- [Semantic Discovery for Patent Landscape Analysis \(/articles/semantic-discovery/patent-landscape-analysis\)](/articles/semantic-discovery/patent-landscape-analysis)
- [Semantic Discovery for Medical Literature Search \(/articles/semantic-discovery/medical-literature-search\)](/articles/semantic-discovery/medical-literature-search)
- [Semantic Discovery for Competitive Intelligence \(/articles/semantic-discovery/competitive-intelligence\)](/articles/semantic-discovery/competitive-intelligence)
- [Semantic Discovery for Regulatory Compliance Search \(/articles/semantic-discovery/regulatory-compliance-search\)](/articles/semantic-discovery/regulatory-compliance-search)
- [Discovery-Coordinated Multi-Sensor Perception \(/articles/semantic-discovery/coordinated-perception\)](/articles/semantic-discovery/coordinated-perception)
- [Post-AirTag Cross-Platform Object Tracking \(/articles/semantic-discovery/post-airtag-tracking\)](/articles/semantic-discovery/post-airtag-tracking)
- [Use the World as Memory: The Brain Strategy for AI \(/articles/semantic-discovery/world-as-memory\)](/articles/semantic-discovery/world-as-memory)

APPLICATIONS · SPECIFIC

- [Google Search Retrieves Results, Not Understanding \(/articles/semantic-discovery/google-search\)](/articles/semantic-discovery/google-search)
- [Perplexity Answers Questions Without Discovery State \(/articles/semantic-discovery/perplexity\)](/articles/semantic-discovery/perplexity)
- [Elasticsearch Indexes Documents, Not Discovery \(/articles/semantic-discovery/elasticsearch\)](/articles/semantic-discovery/elasticsearch)
- [Algolia Optimizes Relevance Without Discovery State \(/articles/semantic-discovery/algolia\)](/articles/semantic-discovery/algolia)
- [Pinecone Finds Vectors, Not Understanding \(/articles/semantic-discovery/pinecone\)](/articles/semantic-discovery/pinecone)
- [Weaviate Stores Semantics Without Discovery Governance \(/articles/semantic-discovery/weaviate\)](/articles/semantic-discovery/weaviate)
- [You.com Answers Questions but Does Not Govern Discovery \(/articles/semantic-discovery/you-com\)](/articles/semantic-discovery/you-com)
- [Brave Search Built an Independent Index Without Governed Traversal \(/articles/semantic-discovery/brave-search\)](/articles/semantic-discovery/brave-search)
- [Kagi Charges for Better Results, Not Governed Discovery \(/articles/semantic-discovery/kagi\)](/articles/semantic-discovery/kagi)

- [Metaphor Systems Predicts Links but Does Not Govern Traversal \(/articles/semantic-discovery/metaphor-systems\)](/articles/semantic-discovery/metaphor-systems).
- [Glean Indexes Enterprise Knowledge Without Governing Its Discovery \(/articles/semantic-discovery/glean\)](/articles/semantic-discovery/glean).
- [Coveo Personalizes Retrieval, Not Discovery Governance \(/articles/semantic-discovery/coveo\)](/articles/semantic-discovery/coveo)
- [Apple Find My Lacks Cross-Authority Reader Activation \(/articles/semantic-discovery/apple-find-my\)](/articles/semantic-discovery/apple-find-my).
- [Google Find My Network Needs Credentialed Cross-Activation \(/articles/semantic-discovery/google-find-my\)](/articles/semantic-discovery/google-find-my)
- [IETF DULT Specifies Behavior, Not Architecture \(/articles/semantic-discovery/ietf-dult\)](/articles/semantic-discovery/ietf-dult)
- [Glean Enterprise Search and Work AI \(/articles/semantic-discovery/glean-enterprise-search\)](/articles/semantic-discovery/glean-enterprise-search)
- **[GraphRAG, but with Governance: Where Microsoft's Architecture Stops Short \(/articles/semantic-discovery/microsoft-graphrag\)](/articles/semantic-discovery/microsoft-graphrag)**
- [Memory Layers for Agents: Why Mem0, Zep, and Letta Get Close \(/articles/semantic-discovery/memory-for-agents\)](/articles/semantic-discovery/memory-for-agents).

[Semantic Discovery overview → \(/semantic-discovery\)](/semantic-discovery)