Knowledge Graph Industry Benchmarking Survey

Data and Analysis on Industry Maturity

October 2022

EXECUTIVE SUMMARY

The adoption of knowledge graph has been steadily rising across several information-intensive industry sectors including those in life sciences, financial services, health care and manufacturing. This sign of maturity is one of the conclusions from the 2nd Annual Knowledge Graph Benchmarking Study – conducted jointly by the Enterprise Knowledge Graph Foundation (EKGF) and the Knowledge Graph Conference (KGC).

But growth and progress are not the only plot of the story. In the face of continuing maturity - two issues remain as fundamental challenges. The first is a skill set gap that inhibits the capacity of organizations to operationalize their initiatives. The second is internal resistance and the challenge of convincing senior stakeholders to adopt knowledge technologies. Companies report they are still swimming upstream against the current due to the lack of internal support from executive stakeholders.

Most of the respondents to the study are still in the experimental stage of their journey. Some are initiating pilots and POCs to demonstrate value while others are moving last year's experiments out of the innovation lab and into production. We are excited to note that a healthy percentage of participants are now at the point where they are extending activities across multiple use cases. In terms of organizational structure, most firms are managing their knowledge graph initiatives with a small set of specialists often embedded in technology or in the Office of Data Management. We do observe that organizations running advanced knowledge graph initiatives uniformly lean heavily toward sponsorship by executive management.

The goal of most of these initiatives is to enhance capability. Across the board, the goals of adopting semantic standards and knowledge technologies are to enhance data discovery, explore 'what if' questions and give analytical flexibility to users. And while enhancing capability is the prime motivation, the initial use cases relate more to activities associated with getting the data house in order in the form of data integration, aggregation across diverse sources and enforcing data quality rules.

The bottom line tells a story of progress and an industry at a pivot point. 'Information literacy' among senior stakeholders remains an issue but the industry is continuing to mature. One participant put it succinctly – "we have a plan per our product roadmap to leverage knowledge graphs in innovative ways." This represents a strategic commitment and recognition of the knowledge graph as a core service for the enterprise.

TABLE OF CONTENTS

BENCHMARKING PROFILE

BUSINESS SYNOPSIS

The Knowledge Graph Industry Benchmarking Survey was conducted in 2022 as a joint project by the Knowledge Graph Conference (KGC) and the Enterprise Knowledge Graph Foundation (EKGF). This is the second study. The first was conducted in 2021. The objective of the study was to capture baseline statistics on the size, nature and direction of data management using semantic standards (i.e., the capabilities that result in a knowledge graph).

The profile of the companies (regardless of size and industry sector) was similar. Almost all have complex and interconnected environments. Most are looking to find hidden connections between entities, address the challenges of data silos, integrate data across linked processes and gain efficiencies across their chain of supply.

The adoption of knowledge graph is in its formative stage. Adoption has been rising across information-intensive industry sectors such as life sciences, financial services, manufacturing conglomerates, technology companies, governments, and professional services. We collected data from over 150 companies representing organizations from across the classification spectrum. There was an equal mix of participants spanning from very large (diversified) entities to those under 100 employees.

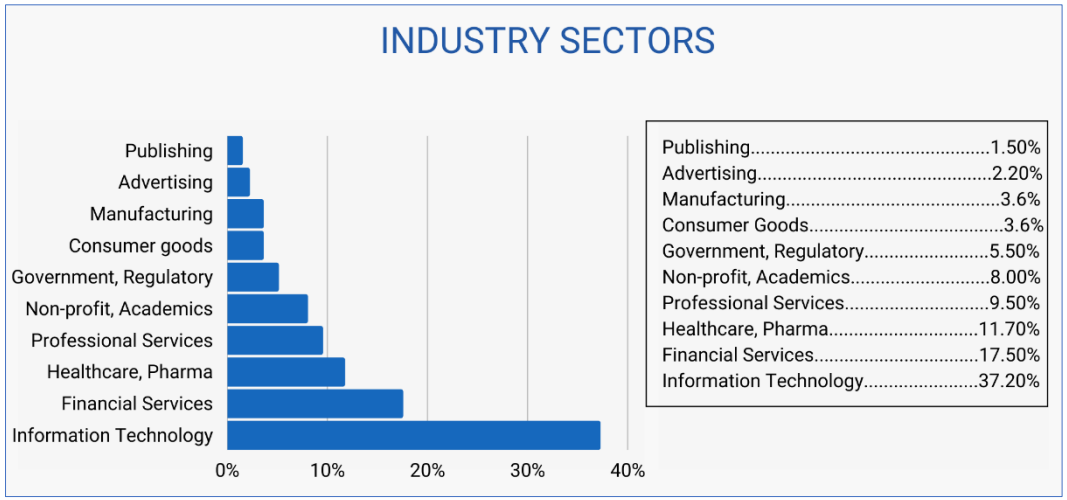

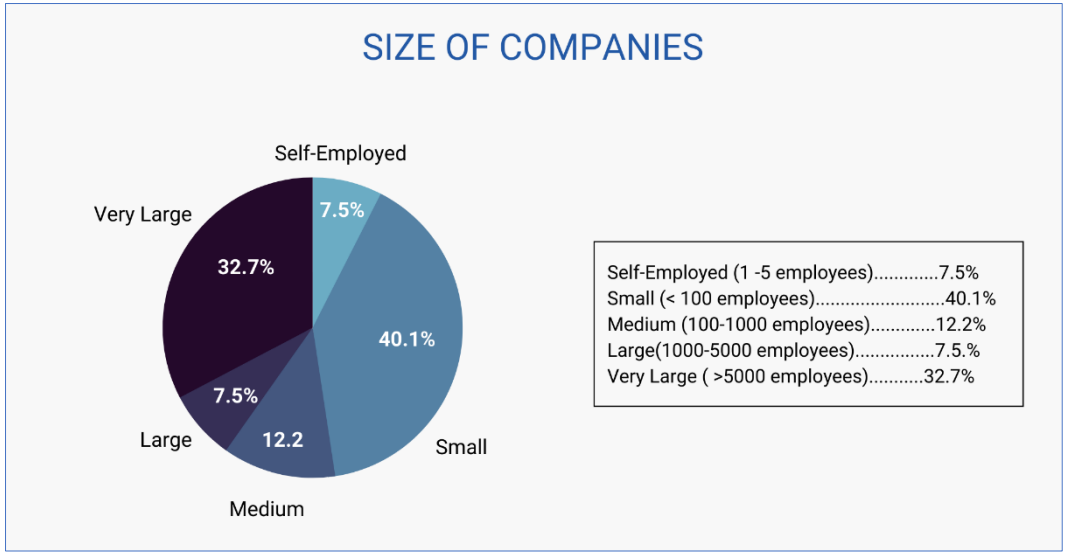

Most participants should be considered as early adopters of knowledge graph technology. Many are still doing proofs-of-concept to demonstrate capability – but they are all focused on turning data from "a problem to be managed" into data "as a resource to exploit." We have seen progress over the past year. Companies are continuing to mature at a healthy pace. Proof of Concepts (POCs) are becoming operational. Pilots are being extended. And a growing number have made knowledge graph the foundation of their data management operations.

PARTICIPANT PROFILE

Most of the participants in this year's study could be characterized as "engaged practitioners" who are already familiar with knowledge graph technology. The majority work directly with the data - primarily in modeling, precision of meaning and mapping. Mapping of data from one repository to another and dealing with the nuances of proprietary schemas is still an important data pipeline management requirement.

The number of participants that were part of technology departments declined over the past year. This is a recognition of the complementary roles between "technology" (acquire, store, distribute) and "data" (precision and contextual meaning). We see this as progress for the industry.

The share of respondents that are aligned with revenue objectives continues to rise across most industry sectors. This mirrors the rationale for the investment in knowledge graph as companies seek to digitize their internal knowledge and document business rules as well as uncover hidden insights and connections in data.

We see this as a step forward for "information literacy" as business units are beginning to recognize the causes and liabilities of data that is both out of synch with meaning and structured in rigid environments that inhibit the flexibility in its use.

BUSINESS FUNCTIONS

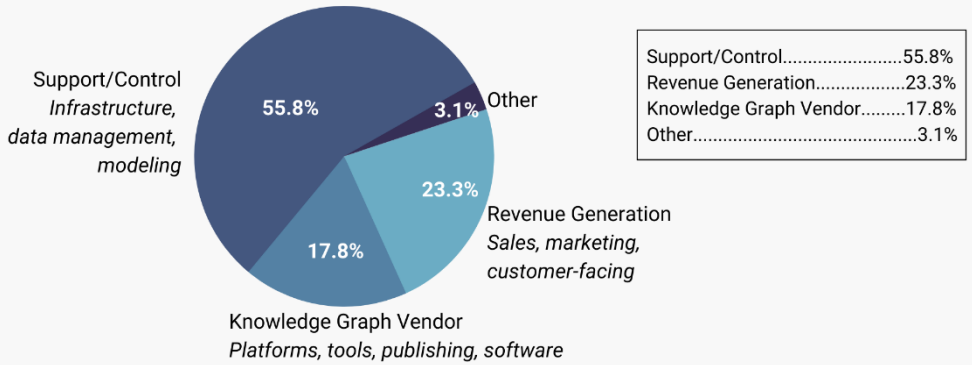

PRIMARY ROLES

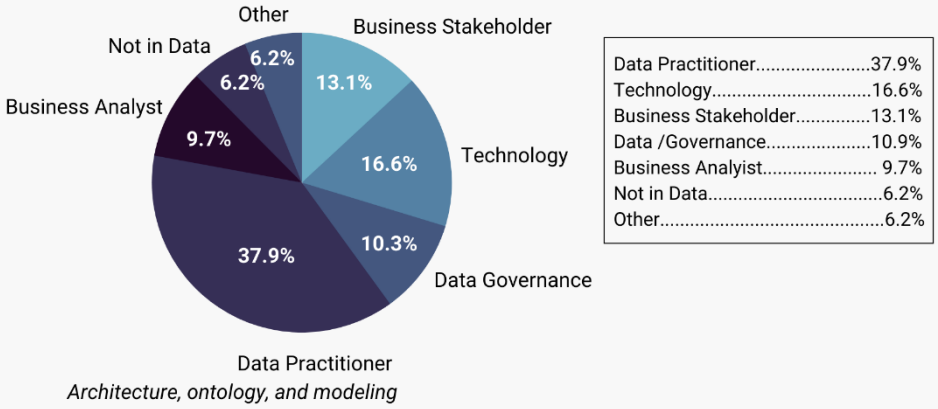

One additional note of promise is the elevation in the level of participant responding to the study. We see more senior level respondents including CEOs, chief scientists, lead ontologists, knowledge engineers and enterprise architects. This is critical because leadership is essential to facilitate adoption of knowledge graph across organizational boundaries and in competition with scarce resources.

ORGANIZATIONAL STRUCTURES

Ensuring that the organization supporting the knowledge graph is appropriately positioned within the company is an essential ingredient to success of these initiatives. It is important because this is about getting a diverse set of stakeholders to both collaborate and to adopt alternative approaches to the goal of data management.

One of the core challenges is bureaucracy. It is necessary but can also be debilitating. One thing we know for certain is that it is hard to synchronize strategy across operational boundaries. Even when users are excited, it is often hard to get the knowledge graph on the roadmap. This is particularly true in circumstances where staff are freely (and frequently) moved around the organization.

Most organizations in our study are managing the knowledge graph with a small set of specialists. These can be viewed as "heroes" who understand the nuances of semantic architecture and are performing all the tasks associated with onboarding, modeling, mapping, architectural engineering, and delivery.

A well-staffed team normally consists of between 5 and 15 depending on the skills of the people. There is a requirement for an experienced architect to design the approach and lead the team. Ontologists are required for content engineering and mapping. Knowledge engineers are needed to model concepts and manage the data pipeline. And the organization needs a project manager to advocate for the team and the development process.

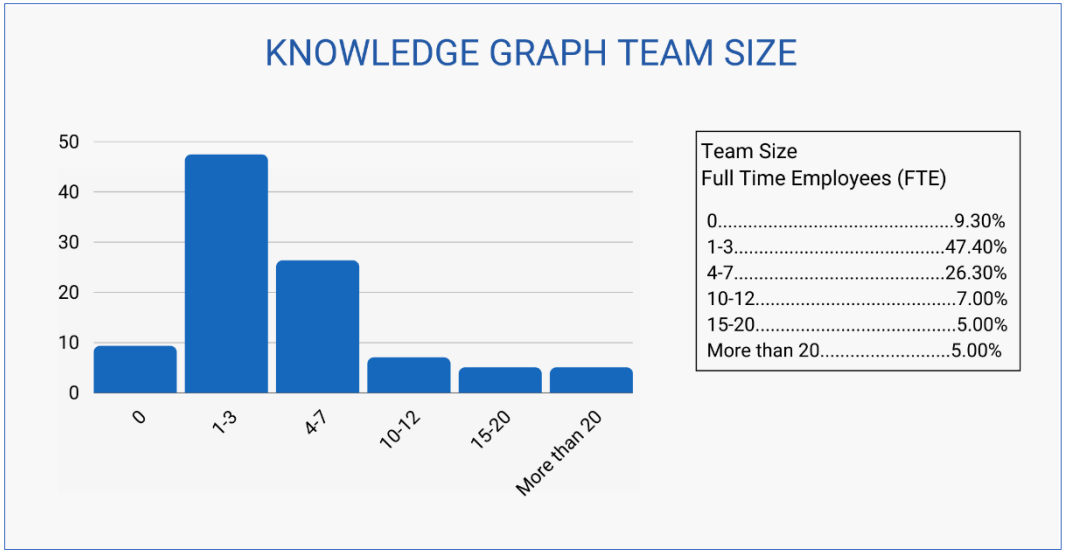

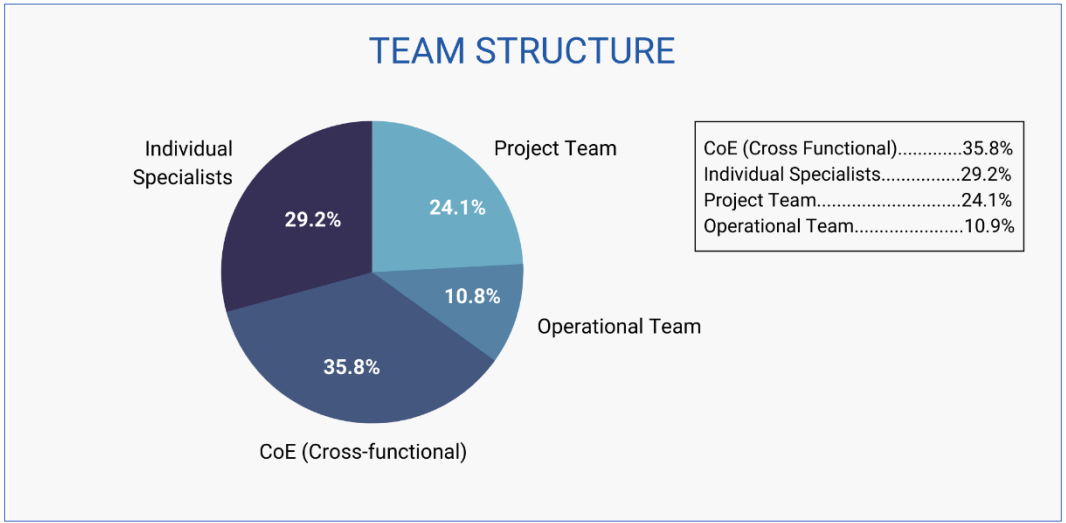

We observe a noticeable split between those that are still experimenting with knowledge graph (small teams) and those that have made the commitment to adopt knowledge graph to support mission-critical applications (centers of excellence).

KNOWLEDGE GRAPH SPONSORS

Leadership and executive 'air cover' are essential for the knowledge graph to operate on an enterprise level. Without a champion with organizational influence acting as advocate, the knowledge graph will live in jeopardy. Many of the existing initiatives for this study live inside of innovation groups and may struggle to become operational. And running the knowledge graph in production is essential.

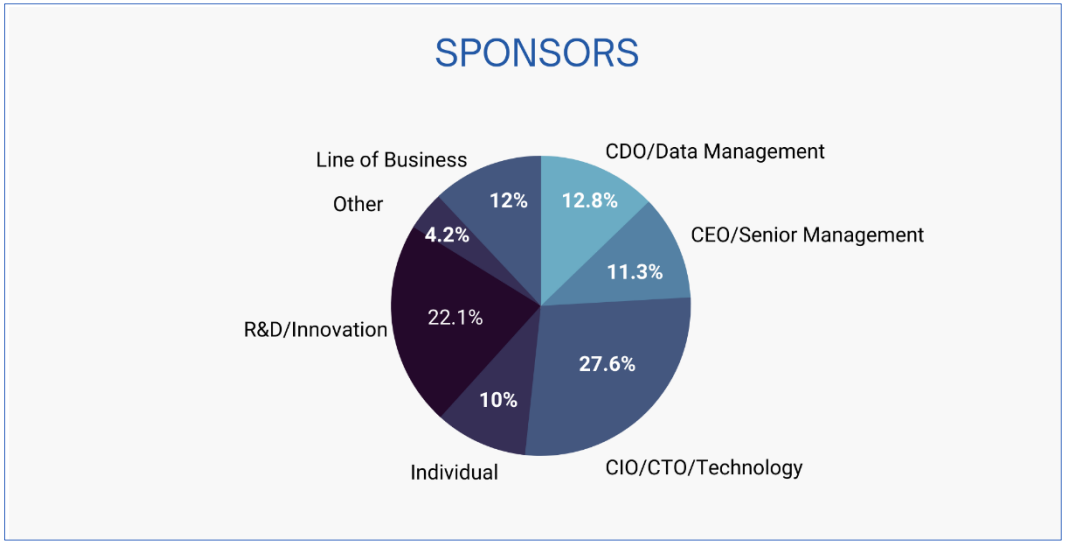

CIO/CTO/Technology.....27.6%
R&D/Innovation.....22.1%
CDO/Data Management.....12.8%
Line of Business.....12%
CEO/Senior Management.....11.3%
Individual.....10%
Other.....4.2%

Starting a knowledge graph initiative as an innovation project can be both wise and useful. It can give the organization time to build some of the foundational components that are needed to deliver for more urgent use cases. We view pilot projects and POCs as evidence of capability and consider them as 'table stakes' for internal marketing. There is no getting around the mandate to initially deliver something very practical. If you can't, there is often no runway for extension. The primary sponsors of the initiatives represented in this study are from research and development, innovation groups and technology. For some of these entities, knowledge graph has not yet penetrated the cognitive screen necessary to overcome organizational inertia and capture the attention of business stakeholders.

Not surprisingly, organizations that are running advanced knowledge graph initiatives – those that are both extensible and facilitating enterprise use cases – lean heavily toward sponsorship by executive management. This is reason why the objective of 'information literacy' is so important. If executive management makes knowledge graph and the adoption of semantic standards a priority, the rest of the organization is likely to follow.

MATURITY, DRIVERS, USE CASES, AND INHIBITORS

KNOWLEDGE GRAPH MATURITY

We measure maturity based on a capability model designed to document best practice. It covers standard evaluation criteria required for the design, implementation, and maintenance of an enterprise knowledge graph. Advancing through the levels of maturity requires leadership, commitment, and resources. The model (see page 11) is structured around four pillars:

1. *Business* – covering the mechanisms necessary to align business goals with knowledge graph objectives emphasizing use case prioritization, resource availability and maturity of business capabilities.
2. *Data* – covers operating models, policies and data architectures needed to execute the strategy. Emphasis is on the establishment of unique identification and unambiguous shared meaning.
3. *Technology* – covering capabilities related to managing the knowledge graph as the authoritative source for data including physical infrastructure, pipeline management and data integration.
4. *Organization* – examines the requirements for facilitating and governing change in complex organizational environments.

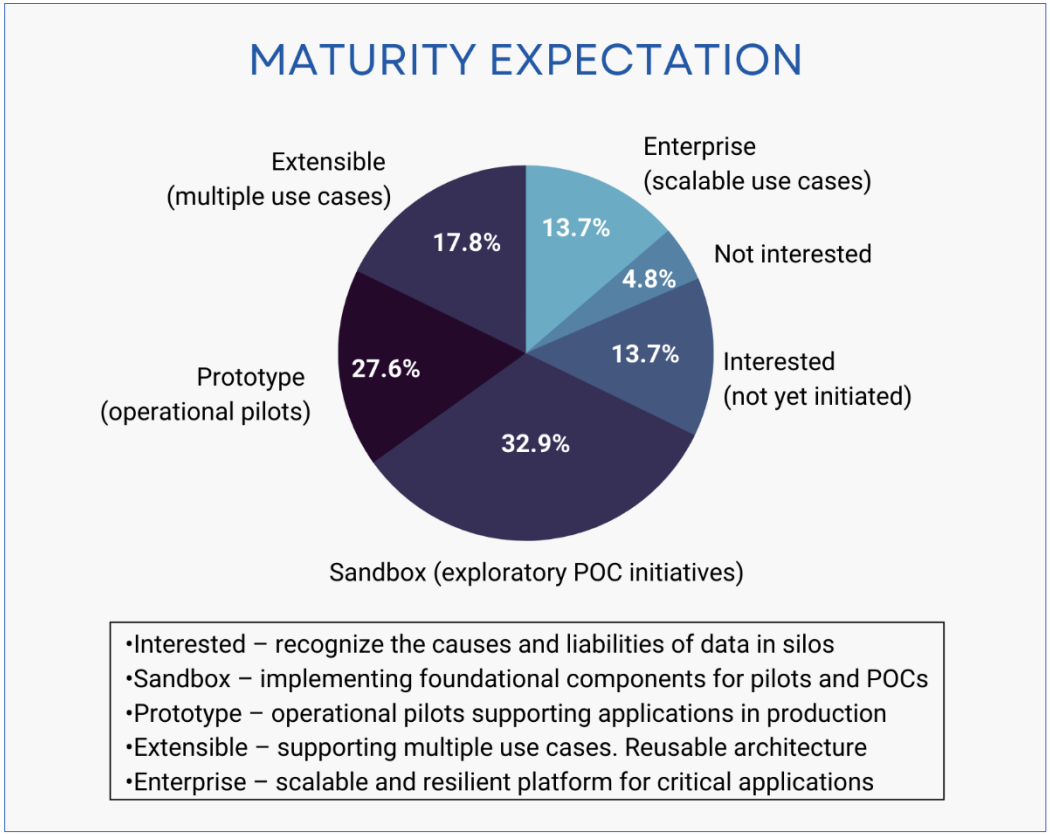

Most respondents (sandbox and prototype levels) are still implementing the foundational components needed for a minimum viable product including adopting the principles of data hygiene and defining the "isolated ontologies" required for specific use cases. Technology strategy is focused on experimentation and ETL processes are predominantly manual. This is the area of POCs and pilot initiatives.

The "extensible level" is the domain of parallel knowledge graph activities. This emphasizes reusable architecture based on expanded design principles. Organizations that have reached this level of maturity are defining their approaches to resolve identity and meaning. The knowledge graph is becoming the central point for integration across related use cases. Use cases are being extended incrementally with increasingly lower marginal costs.

Entities that have reached the "enterprise level" are building scalable and resilient platforms for business-critical applications. We expect to see inventories embedded into the graph and linked to governance. The "enterprise knowledge graph" is becoming the authoritative source for data. This represents a strategic commitment and recognition of the KG as a core service for the enterprise.

EKG Maturity Model Structure

BUSINESS

Business Strategy
Business Vision
Business Goals
Business Tactics

Business Architecture
Market Segmentation
Value Chain
Change Management

DATA

Data Strategy
Goals and Objectives
Business Case

Data Architecture
Identity Resolution
Ontologies/Data Models
Inventory Management
Business Glossary
Data Integration

Data Quality
Quality Framework
Business Rules
Data Quality Exceptions

Data Governance
Operating Model
Data Management Policy
Data Production Process
Entitlement Management
Classification Management
Risk & Control Framework

TECHNOLOGY

Technology Strategy
Technology Landscape
EKG Technology
Process Rationalization
Automation Strategy

Technology Execution
Data/Systems Integration
Deployment & Operations
Physical Architecture

User Experience
Search and Query
Components /Data Viewpoints
Change Management

Organizational Capabilities
Control Functions
Vendor Management
Talent Management
Knowledge Sharing

ORGANIZATION

Executive Leadership
Strategy and Policies
Funding and Resources
Technology Requirements
Processes and Governance

Product Ownership
Use Case Requirements
Funding and Budget
Measurement Criteria
Team Requirements Mgt.

EKG Delivery Management
ETL/Data Movement
Ontologies and Mapping
DataOps Process
User Interface & Access
EKG Platform Team

Organizational Culture
Implementation Approach
Agility and Innovation
Ecosystem Collaboration

EKG Maturity Model

This graphic is an overview of the components included in the knowledge graph maturity model as developed by the Enterprise Knowledge Graph Foundation.

Enterprise Knowledge Graph Foundation

KNOWLEDGE GRAPH PROGRESS

We are encouraged to see that organizations are moving their knowledge graph initiatives forward. Almost every respondent report progress. POCs are being initiated. Experiments are moving out of the innovation lab and running in production. Use case relationships are being modeled with the notion of capturing shared data relationships. Data is being expressed and onboarded as formal ontologies. Adoption is maturing and there is a projection of optimism among participants.

The companies that participated in our study represent information-intensive industry sectors where precision and nuance matter. 36% of these participants report to be in the 'experimental" phase where they are initiating a POC to demonstrate value to executive stakeholders. These entities are competing for mindshare. They are seeking evidence on the value of data connectivity. They are searching for acknowledgement by executives that the knowledge graph is transformational and therefore demands greater investment.

We note that well over half of the participants have found their way over the 'semantic hump' of awareness, buy-in and cognition. They are operationalizing their pilot projects as well as extending their ontology engineering efforts to cover additional use cases. This group of adopters are beginning to demonstrate the efficiency benefits associated with reusable data architecture. Those at the top tier of the progress chart are starting to adopt a 'data centric' (more strategic) mindset.

The comments provided by many participants do reveal some progress inhibitors. Companies are still swimming upstream due to the "lack of internal support" from executive management … competition for scarce "funding and resources" … skill set gaps and the "lack of experienced talent" … more "fluency in domain ontologies" … and enhanced "evidence of value" from successful implementation.

The knowledge graph industry is standing at somewhat of a pivot point. There is progress. Vendors can help by continuing to deliver a "better tool ecosystem" and more efficient ways to "extract knowledge to feed the graph." One participant put it succinctly – "we have a plan per our product roadmap to leverage knowledge graphs in innovative ways."

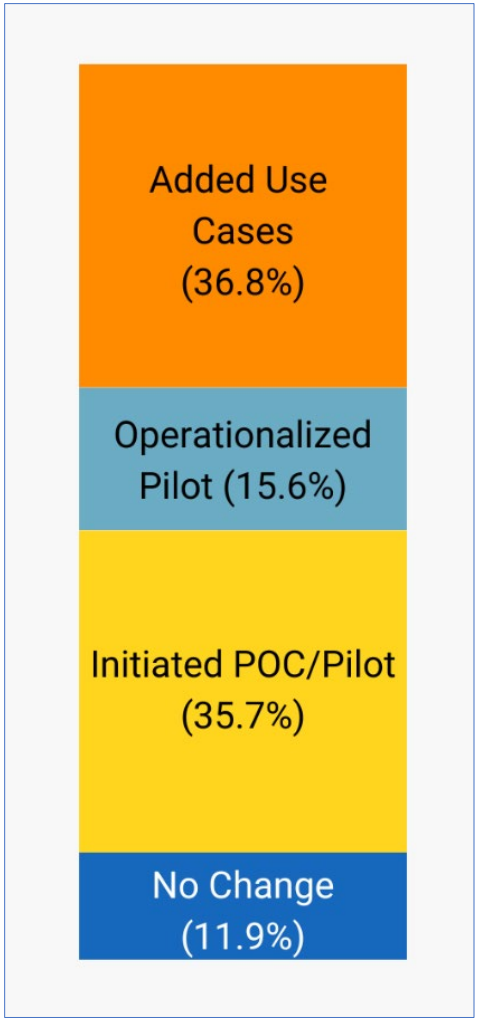

PROGRESS: IN THEIR OWN WORDS

Participant excerpts on requirements for moving their knowledge graph initiatives forward…

"Funding and resources"	"More and better use cases"	"Continuing to add conceptual definitions"	"Ontologies and taxonomies which are in short supply"	"Need buy-in from executive suite and engineering leadership"
"A hardened roadmap with clearly defined benefits for undertaking the effort"	"We need better mechanisms for cross-connecting different knowledge graphs"	"We have a plan per our product roadmap to leverage graphs in innovative ways"	"Technical skills about knowledge management"	"Data literacy, buy-in, sponsorship, talent"
"We are in full flight"	"Better tooling"	"Fluency in domain ontologies"	"Much more standardization and FAIR data principles"	"End-user interfaces. No users, no point"
"We are already moving forward"	"Data-centric mindset/culture"	"Structured curriculum on Basic KG"	"A defined reference architecture for knowledge graphs"	Continuing to show the value of KG to Executives"
"More management commitment"	"Executive level buy-in from Chief Data Office and Tech"	"Better user interfaces for data entry and visualization"	"Broader market adoption of the KG technology"	"Success stories shared by organizations to inspire stakeholders"
"We are working on several new initiatives with clients"	"Good skills in Machine Learning to extract knowledge and feed the graph"	"Acknowledgement by executive leadership that KG demands greater investment"	"Getting management buy-in that knowledge graphs fundamentally transform our business works is a hard sell"	"We recognize the need to incorporate the knowledge graph into our data catalogue "
"Awareness: Typically thought of as novel rather than useful in my organization"	"We convinced the organization with the POC and we are moving ahead"	"Rolling out products on multiple graph databases "	"Application of graph technologies is hindered by IT. Upskilling technologists will accelerate adoption."	"A C-suite that understands data (why are they there?)"

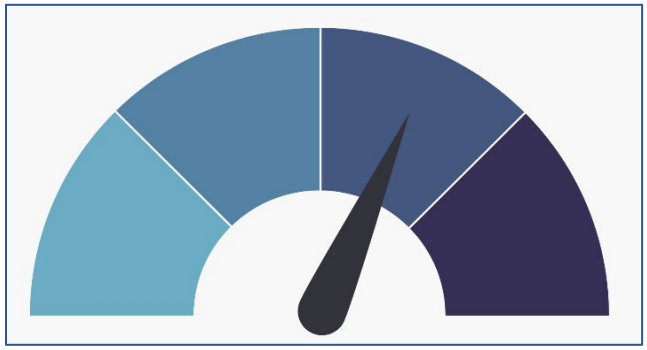

Weighted Average Scoring

We asked respondents to rank many of the questions in the following segments based on importance.

We used a basic weighting process (*three points for critical, two points for very important, one point for somewhat important and zero points for not important*) to arrive at a "weighted average" score. The weighted average score is expressed as a number in parentheses on the following graphs. The higher the number, the more important the category.

The weighted average has been applied to questions relating to adoption drivers, use cases, inhibitors to adoption, industry priorities and future priorities.

ADOPTION DRIVERS

The answer to the core question of why firms are deploying semantic standards and building knowledge graphs can be organized into three categories. The primary reason is for better insight and enhanced analytics. Organizations are seeking views into data that aren't evident at the surface. They want visibility into data connections.

The goal is "better business decision making" … "access to knowledge based on data relationships that are difficult to gather from relational systems" … and "the ability to formulate business questions and get business answers in return." Providing users with flexible analysis and tools to ask 'what if' questions is the prime motivation for knowledge graph adoption.

The implementation of a control environment to facilitate integration, track the flow of data, enhance quality assurance, and simplify governance is category number two. Almost every respondent cited the ability to "harmonize data for integration across repositories and silos" as an essential reason for adoption. The clear goal is to ensure a unified view of data because data with different structures, definitions and contextual meanings makes integration difficult. This is particularly true for firms that have dozens of systems of record all serving various operational processes and independent lines of business.

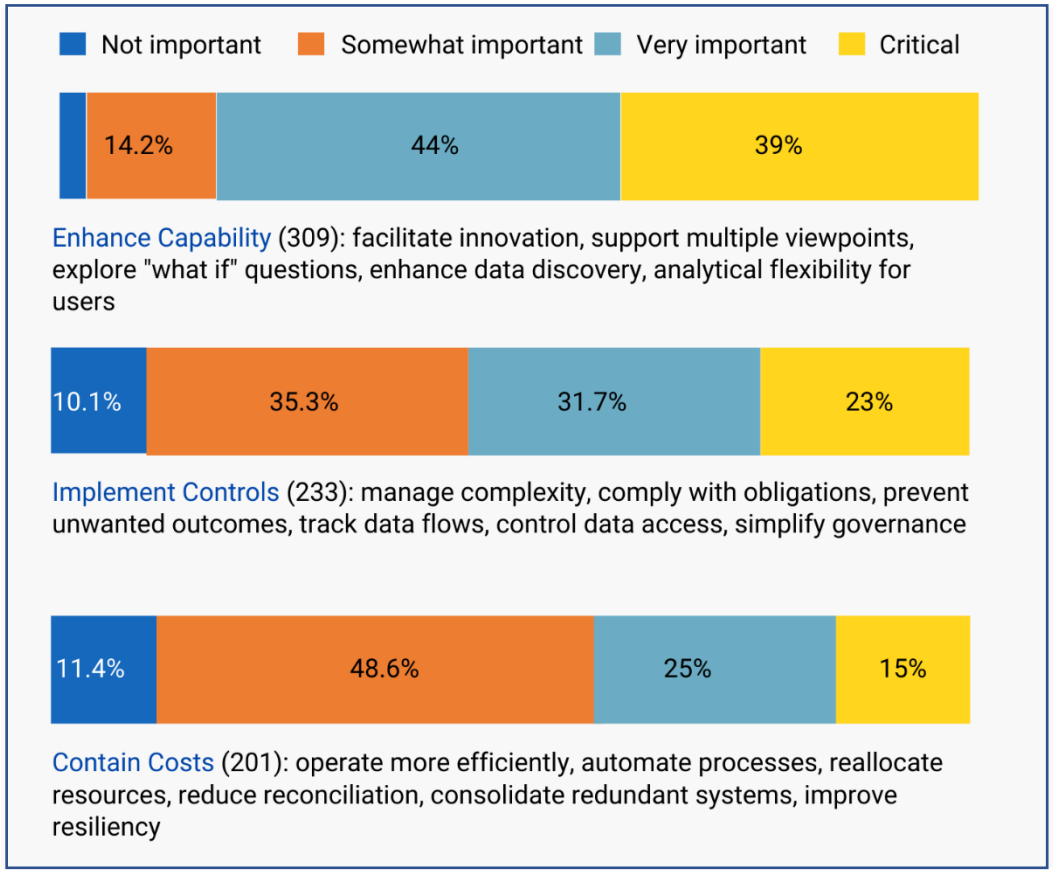

Cost containment is not the principal driver of knowledge graph technology. It is (however) a residual benefit enabling users to "offload repetitive tasks that require specific expertise" … "to replace legacy or complex processing pipelines" … "to achieve greater workflow efficiencies" … and to "reduce the total cost of ownership across the technology estate."

The composite case recognizes that the liabilities associated with technology fragmentation are significant. The inability to automate processes, explore 'what if' questions, aggregate data with confidence, secure sensitive data, respond to client needs and turn analytical ideas into action add up to competitive disadvantage in our complex and interdependent business environment.

DRIVERS IN THEIR OWN WORDS

Participant excerpts on the reasons for deploying semantic standards and building knowledge graphs…

"We want to digitize the vast traditional commerce experience of our company and make it available"	"Unlocking science trapped in silos"	"To document and enforce rules for data quality that tie directly to the business processes"
"We are creating a digital twin of parts of the world with conceptual inference and auto classification"	"Manage exploding complexity"	"The ability to get as close to a 360-degree view as possible"
It is replacing legacy or complex processing pipelines with a unified data fabric"	"Not using spreadsheets as production applications"	"Harmonize data integration in a distributed and siloed data landscape"
"Our Knowledge Graph supports our digital business transformation agenda for our holding company"	"Insights based on data relationships that are difficult to gather from relational systems"	"The centralized management of our core functional information accelerates many AI/ML capabilities"
"Find and uncover the richness of Enterprise Data"	"Uncover hidden insights in data and connections"	"Deliver data relevancy and knowledge insights to consumers in their terms (not technical ones)"
"Finally throwing away tech that has been legacy tech a decade ago"	"Finding the root cause of problems and fixing them quickly"	"Better modeling of reality and easier business analytics"
"Most of the problems are either not possible to solve efficiently or are limited by poor usage of relationships"	"Provide a trustable single point of truth for all data in the company"	"Connected data are a fundamental building block to getting new insights"

KNOWLEDGE GRAPH USE CASES

The use cases for knowledge graph can be broadly organized into these five major categories. Firms are taking advantage of semantic standards to both resolve entities and harmonize meaning. This is an essential part of their goal of data integration and mapping objectives. These are prerequisite activities for almost every "enhanced capability" use case from relationship analysis to aggregation for reporting.

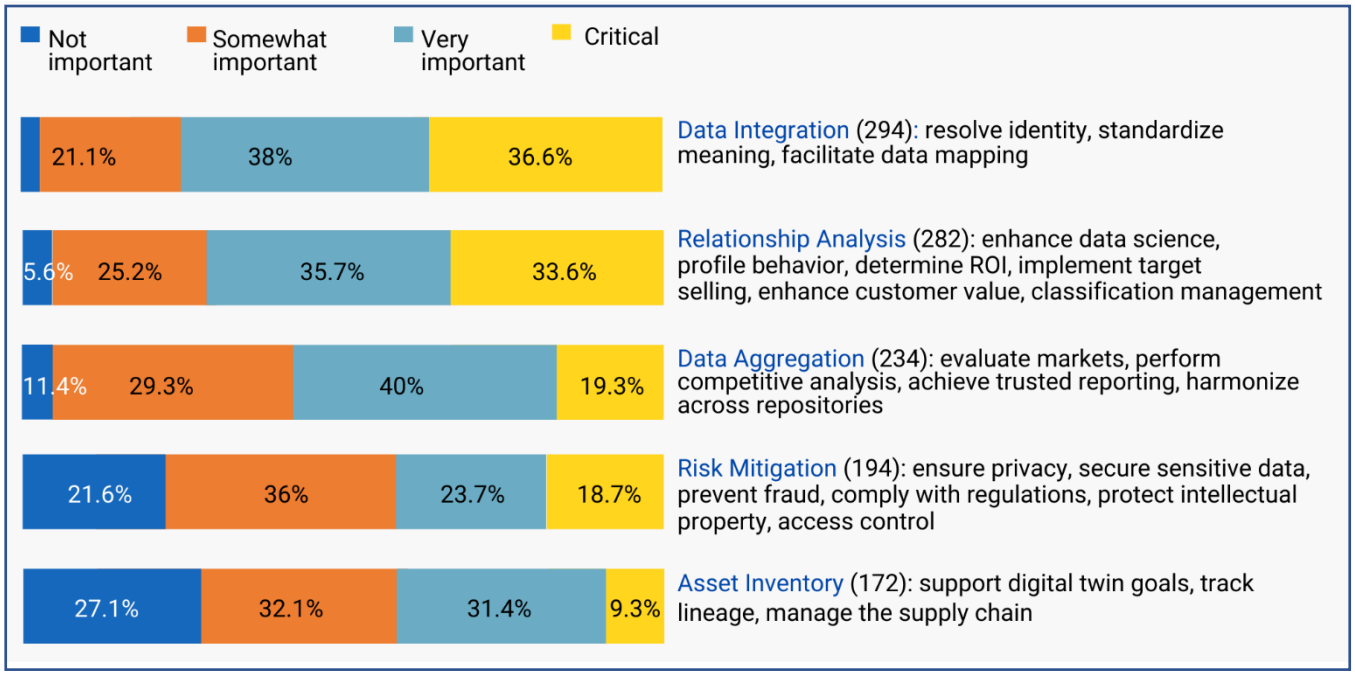

USE CASES	WEIGHTED	NOT IMPORTANT	SOMEWHAT IMPORTANT	VERY IMPORTANT	CRITICAL
Data Integration	294	4.20%	21.10%%	38.00%	38.00%
Relationship Analysis	282	5.60%	25.20%%	35.70%	33.60%
Data Aggregation	234	11.40%	29.30%	40.00%	19.30%
Risk Mitigation	194	21.60%	36.00%	23.70%	18.70%
Asset Inventory	172	27.10%	32.10%	31.40%	9.30%

Somewhere around 70% of cited use cases relate to this fundamental goal of data integration. When you unravel the specifics, the integration objective encompasses the construction of asset inventories for operational resilience and data quality analysis to identity criticality and manage KPIs across the data supply chain. We also find many using the capabilities for enterprise search to both build data catalogues and find hidden connections between entities.

Beyond these core initiatives of find, evaluate and harmonize, a growing percentage of firms are using semantic standards to turn ideas into action. These are more advanced decision-support types of use cases designed to perform scenario-based analysis, enhance visualization of data from multiple perspectives and traverse connections to unravel the intersection of clients, products and transactions.

USE CASES IN THEIR OWN WORDS

We asked participants to describe their own use case scenarios. These are excerpts of the responses.

"Fraud detection for the insurance and banking sector" *Chief Scientist, Business Automation*	**"Integrating multiple systems in enterprise digital ecosystems"** *Head of Data Architecture*	**"Manage the critical KPIs in the supply chain"** *Data Architect, Health Care Company*
"Funnel, normalize and aggregate data from diverse sources" *Enterprise Data Architect*	**"Machine-readable knowledge map of research topics"** *Research and Advisory Director*	**"A cross-repository data catalog"** *Knowledge Discovery Director*
"Curated regulatory reporting" *Enterprise Data Architect*	**"Identify anomalies, correct data, and enforce data quality rules"** *Data Governance Business Analyst*	**"Data alignment and remediation across global banking infrastructure"** *Director, Transformation and Agility*
"Entity resolution and associated relationships between those entities" *Data Manager, Pension Fund*	**"Risk assessment and mitigation on infrastructure"** *Manager, Emerging Technologies*	**"Better discovery - to find hidden connections between entities"** *CEO, Pharmaceutical Company*
"Drill down to the root cause without regard to the data source" *Senior Software Engineer*	**"Data mapping and visualization"** *CEO, Professional Services Firm*	**"Making sense of large amounts of unstructured data"** *CEO, Software Platform Company*
"Predicting material relationships between organizations, operations, and activities" *CTO, Financial Services*	**"Decision support for complex environments"** *Director, Data & Analytics*	**"Digital twin data model"** *Director, Business Analytics*

VALUE SUMMARY

We organized the value side of the knowledge graph equation into the three "Cs" of cost, capability, and control. These are standard KPIs that resonate with top executives (who think about growth and velocity), with technology executives (who think about resilience and scalability), with business executives (who think about use cases and time to market) and with regulatory executives (who think about transparency and traceability). The chart below is a further classification of responses on why firms are adopting semantic standards.

Cost Containment (~30% of operations)	
Support Data Integration	*standardize meaning, track sources, resolve identity, understand requirements*
Automate Business Processes	*reduce reconciliation, mitigate process failure, reallocate resources*
Consolidate Redundant Systems	*enhance scalability, simplify IT environment, improve resiliency, connected inventory*
Simplify Governance	*lock down meaning, track accountability, ensure quality, reduce transformation*

Capability Enhancement (facilitate innovation)	
Understand Relationships	*profile behavior, implement target selling, determine ROI, perform predictive modeling*
Flexible Query	*enterprise search, data discovery, ad hoc analysis, enhance personalization*
Support Multiple Viewpoints	*perform 'what if' analysis, gain insight without restructuring data*
Manage Complexity	*track flow, identify critical data elements, manage the chain of supply*

Control Environment (mitigate risk)	
Aggregate Data	*combine sources to understand links, relationships and interdependencies*
Comply with Obligations	*manage entitlements, support regulatory mandates, protect intellectual property*
Ensure Privacy	*prevent fraud, secure sensitive data, trace to authoritative sources*

INHIBITORS TO ADOPTION

For the second year in a row, skill set gaps top the list. These are the people that design the approach, build the use case trees, engineer the content, coordinate with DBAs, understand the data, manage the pipelines, and do the mapping. It is a significant gap that inhibits the capacity of the organization to operationalize their knowledge initiatives. The number of people citing skills as the critical challenge more than doubled over last year.

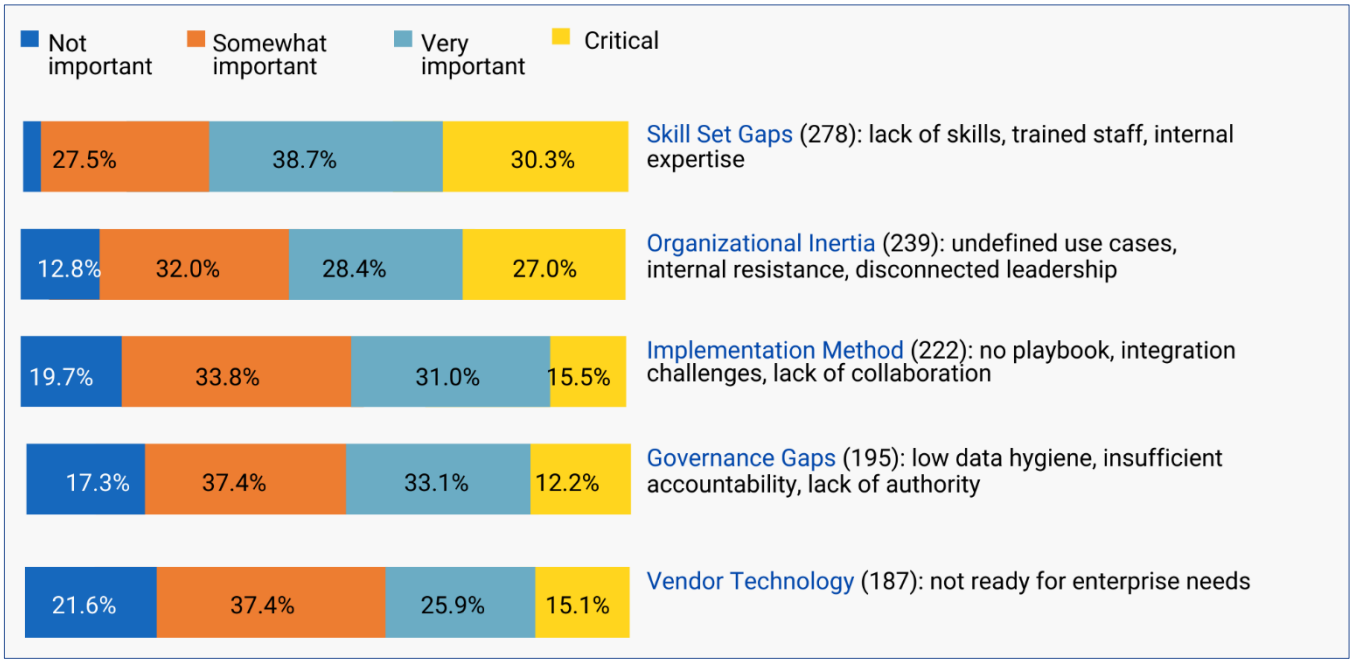

INHIBITORS	WEIGHTED	NOT IMPORTANT	SOMEWHAT IMPORTANT	VERY IMPORTANT	CRITICAL
Skill Set Gaps	278	3.50%	27.50%	38.70%	30.30%
Organizational Inertia	239	12.80%	32.0%	28.40%	27.00%
Method	222	19.70%	33.80%	31.00%	15.50%
Governance Gaps	195	17.30%	37.40%	33.10%	12.20%
Vendor Technology	187	21.60%	37.40%	25.90%	15.10%

Close behind skill sets are the challenges of overcoming organizational inertia. This means swimming upstream against layers of bureaucracy and dealing with both the hubris of systems owners and fear of rocking the boat. Overcoming inertia is the biggest part of the "cost equation." This is the domain of leadership and why it is so important to embrace the cause of information literacy.

The other inhibitors are less critical. There is a need for best practice. Vendor tools must continue to improve. There is a need for governance (governance in a semantic environment is different from traditional data governance). But the real obstacles are cultural resistance to change and talent.

INHIBITORS – IN THEIR OWN WORDS

We asked participants to identify their own inhibitors to adoption. These are excerpts of the responses.

"Technocrats serving as roadblocks who require proof of success before implementation - catch 22."	"Leadership highly risk averse and wedded to legacy methods."	"Entrenched data processing eco-systems. Culture."
"Lack of knowledge on the topic, c-levels are not involved in the proper way."	"Unstructured data is still a big black box as a source of knowledge."	"Lack of technical expertise to move beyond a proof of concept."
"Business units at varying degrees of sophistication with regards to data literacy."	"Corporate maturity among vendors. Can't get through the procurement gauntlet."	"Organization is too large, too complex, to siloed. Weight of politics and posturing."
"Delivery Managers who don't understand Information Architecture!"	"Reluctance to change, low sills, low accountability, zero jeopardy"	"Resources and capital have been the only barriers to our goals."
"Patience, leeway, mandate and skin-in-the-game of execs, strategic thinking."	"Takes a long time for people to understand the power of semantics and connected data."	"The benefits of the knowledge graph started to shine, and its success has become a headline."
"The technology stack is not understood by IT in general."	"Management is lost and very cautious about any decision."	"Short-term thinking. C-suite cost saving. No understanding of data."
"The organization to grasp the semantic EKG way of thinking."	"Technologies and approaches that address such goals have been ignored for years (e.g., linked data)."	"Lack of a reference architecture. We are making it up."

AWARENESS AND TIMEFRAMES

Level of Awareness Among Key Stakeholders

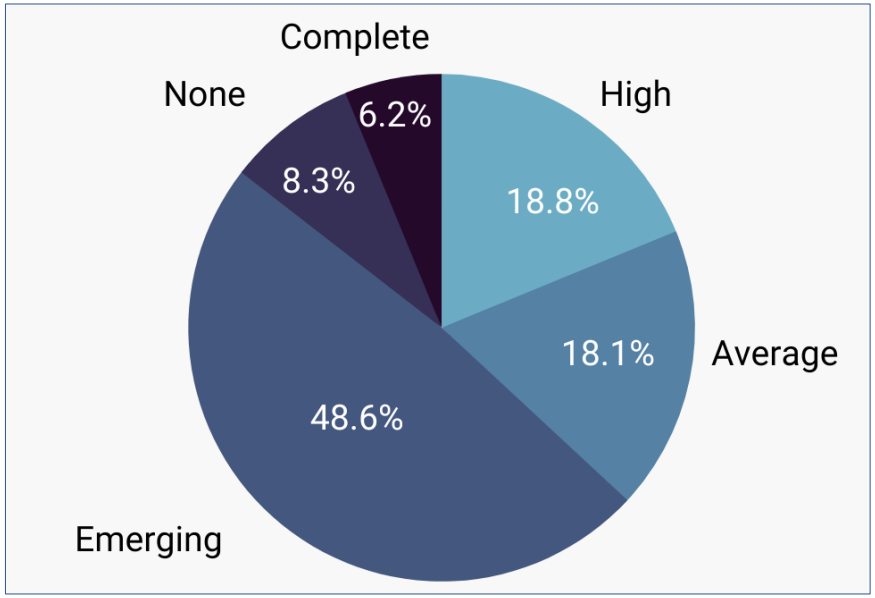

- Demystifying knowledge graph and presenting the value proposition remains a hurdle to address
- Practitioners should consider spending less time talking about how it works and more time talking about why this is smart policy for the organization
- Those that indicated “complete” levels of knowledge and awareness were all vendors

Timeframe for Next Stage of Maturity

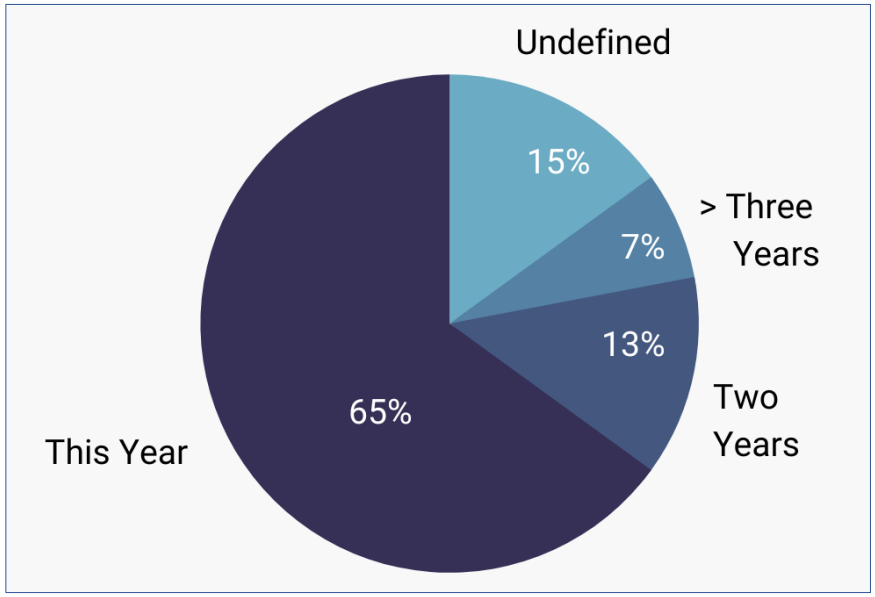

- Most respondents are planning next stage extensions to their knowledge graph initiatives (many said ‘immediately’ and ‘ASAP’)
- Moving activity to the next stage requires resources and commitment (and an internal client)
- Many internal clients view the capabilities of knowledge graph as ‘game changers’ after exposure

INDUSTRY CAPABILITES

INDUSTRY PRIORITIES

Use cases and ROI research top the list of priorities for those on the front lines of knowledge graph adoption. This is because knowledge graph is not well understood by executive stakeholders. This is also why respondents rank education (at both the executive and practitioner levels) as a high priority.

The correlation is clear as enhancing skill sets and overcoming organizational inertial were the number one and number two inhibitors to adoption. After that, the call is for standards interoperability. Users are frustrated at the reliance on proprietary approaches and with the 'holy wars' between labeled property and RDF graphs. We also see a fair amount of interest in engineering best practices, easier to use tools and guidance on implementation.

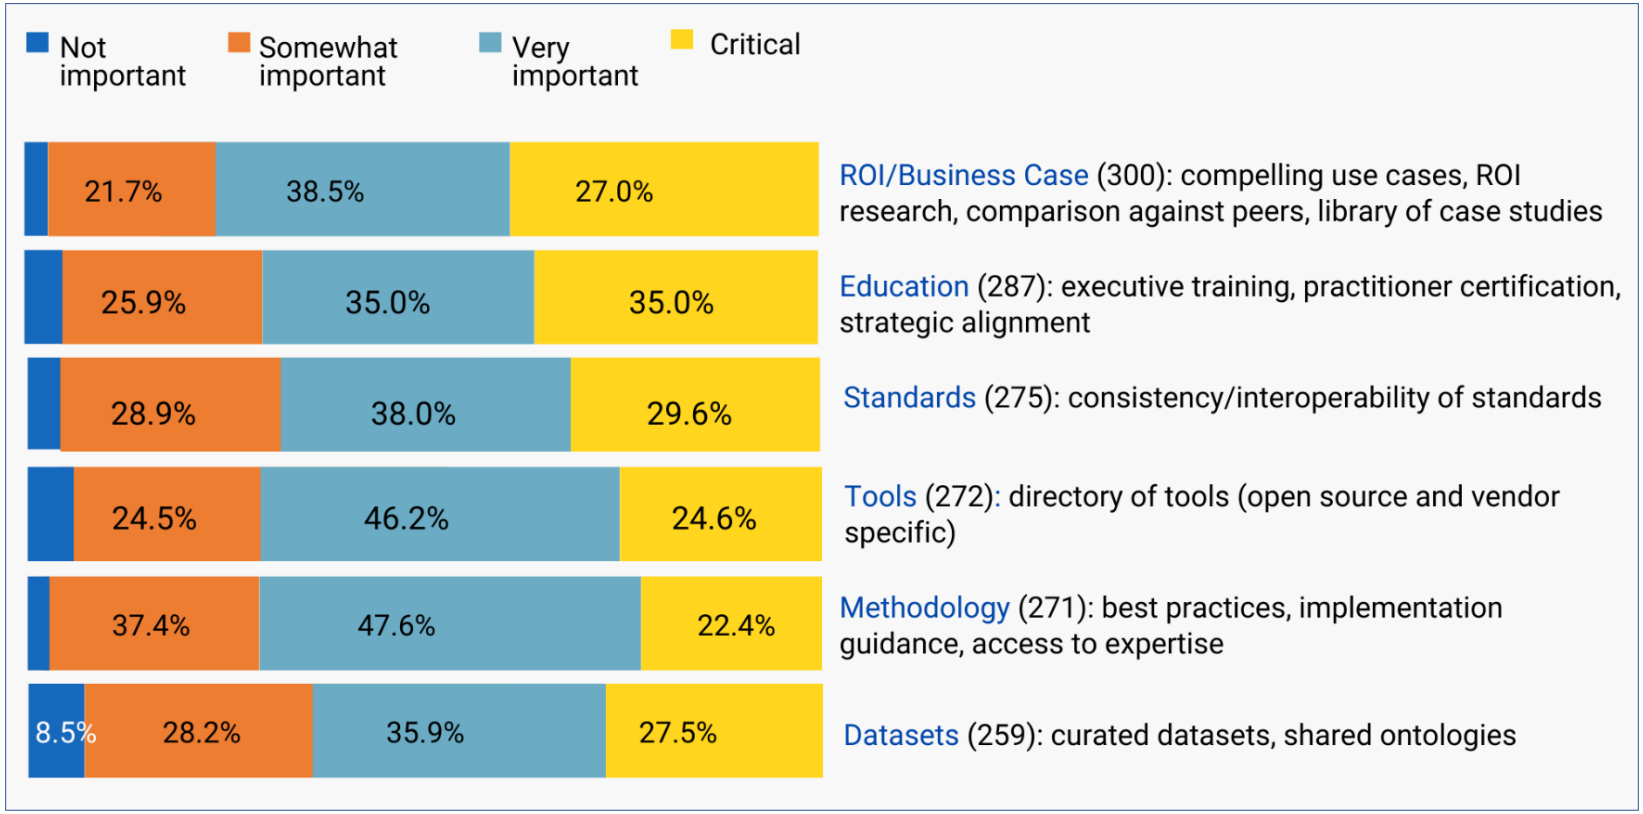

VENDOR ENHANCEMENTS

The message for vendors is about ease of adoption. Support for ontology development (particularly at the foundational level) mapping and better tools for data visualization are all at the top of the wish list from vendors.

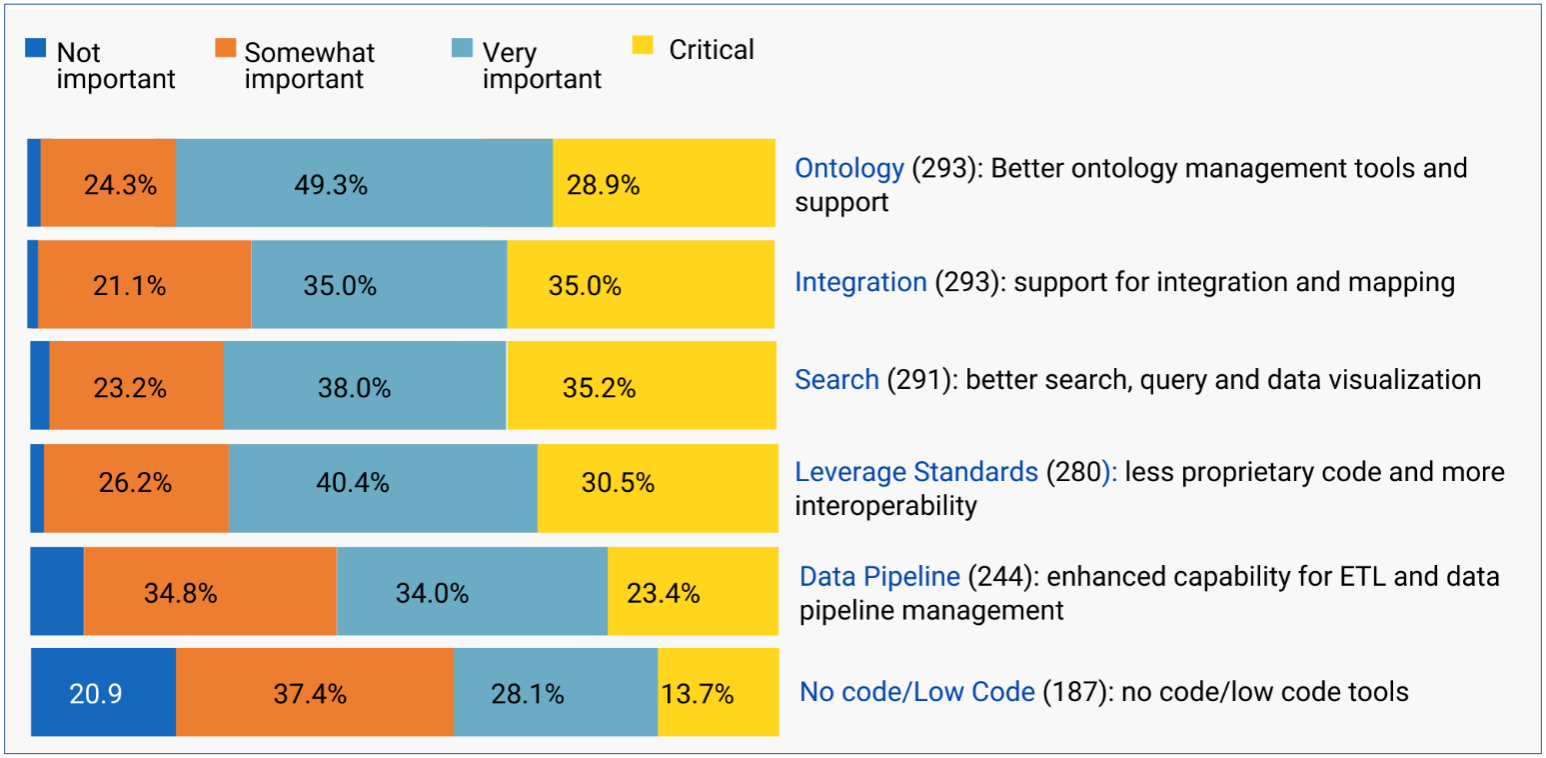

VENDOR PLATFORMS

We asked the respondents which knowledge graph vendors they were testing or using.

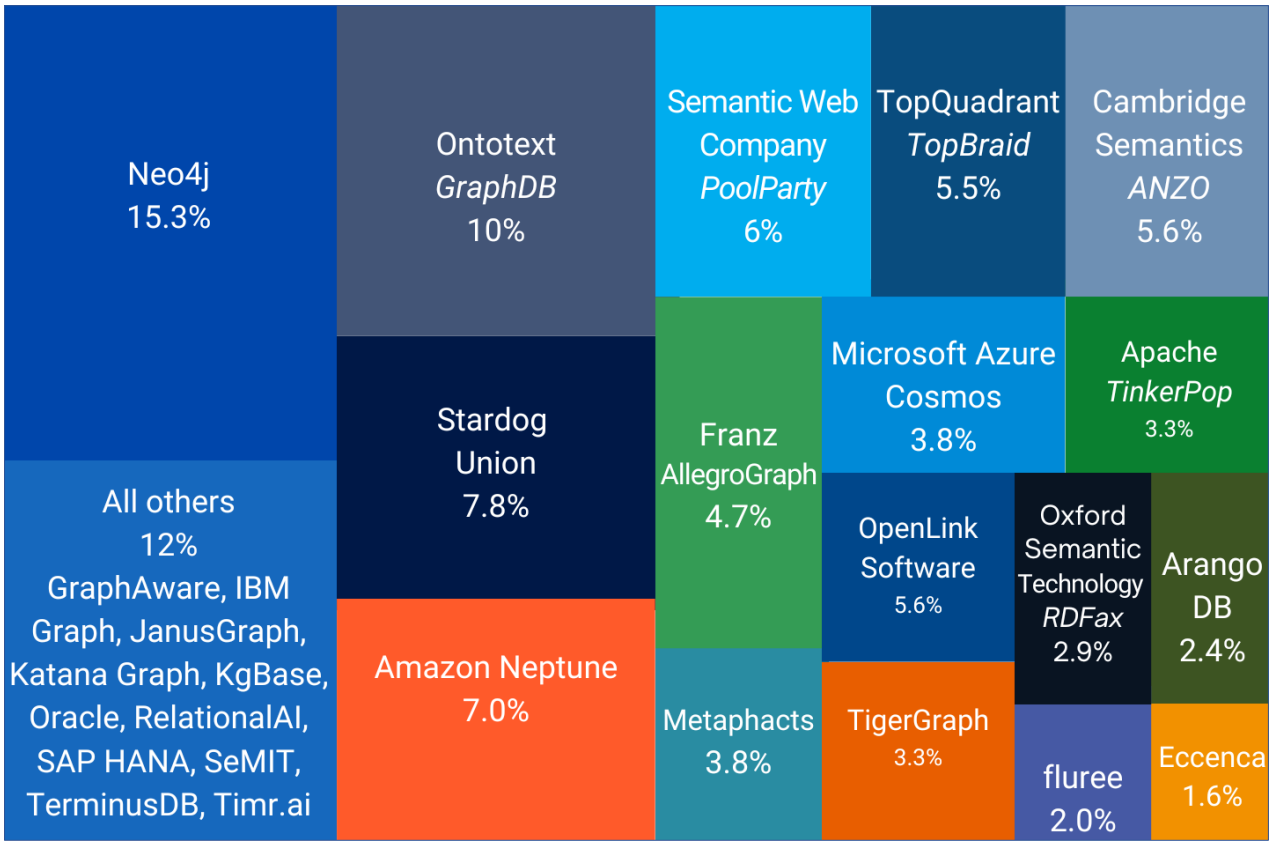

FUTURE PRIORITIES

For the future, respondents are seeking reusable components, better search capabilities, better interfaces, more value from inference and enhanced analytical capabilities. We again note the importance of (and interest in) shared ontologies. We know from our ongoing interactions that the long tail of investment necessary to get the foundation correct is still an inhibitor to adoption.

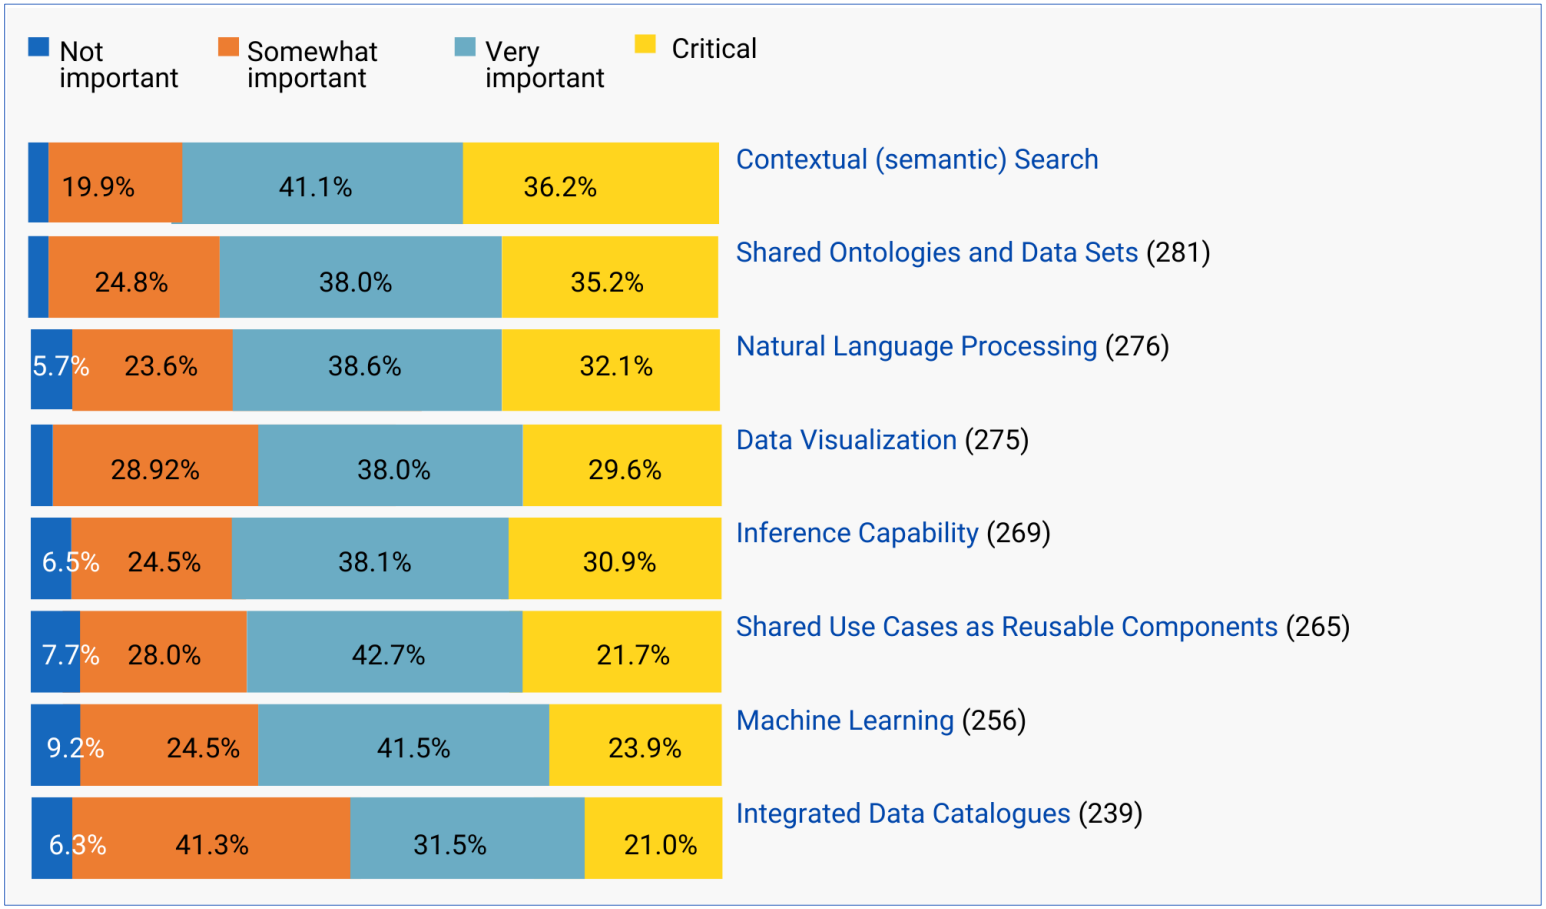

KEY TAKEAWAYS

1. The challenges associated with data incongruence (inconsistencies in meaning) and structural rigidity (limitations to analytical flexibility) are recognized as serious liabilities that must be addressed.

   The fundamental truth is that we have allowed data to become isolated, incongruent, and inflexible because of these core challenges. It diverts resources from business goals, extends time-to-value, leads to business frustration, and creates a mindset of mistrust across organizational boundaries. These liabilities are both real and significant.

2. Knowledge graph is at a transformation point. There is progress. POCs are being initiated. Experiments are running in production. There is a sense of optimism among participants.

   Knowledge graph adoption has crossed the Rubicon. Semantic standards are mature, and tooling from vendors continues to improve. The lesson from those that have been successful is to focus on practical value. Start small with a valuable use case. Model business goals before systems design. Get the engineering right as it is the foundation of the semantic house.

3. The business case for adopting semantic standards (knowledge graph) is strategic, not tactical.

   Executive stakeholders need to become more 'information literate' about the causes and implications of technology fragmentation. Without a champion with organizational influence acting as advocate, the knowledge graph will live in jeopardy. However, if top-of-the-house makes semantic data management a priority, the rest of the organization will follow.

4. A growing number of 'data centric' organizations are making knowledge graph the foundation of their data management operations.

   The root causes of data challenges can be permanently addressed without a significant investment in new technology and without the need to 'rip and replace' existing environments. Leading companies across many information-intensive industry sectors understand this is a solvable problem and are working on building the data infrastructure for the digital world.

COMPANY BIOS

Content Strategies LLC

Content Strategies is a strategic consulting firm focused on the principles, practices and operational realities of data management. The Managing Director is Michael Atkin who has been the analyst and advocate for data management since 1985. His experience spans from the foundations of the information industry to the adoption of semantic technology. Mike has served as an advisor to financial institutions, global regulators, publishers, consulting firms and technology companies. He is a frequent speaker at industry events and an active participant in the standards development process.

Enterprise Knowledge Graph Foundation

The Enterprise Knowledge Graph Foundation (EKGF) was incorporated in April 2020 to define best practices and mature the marketplace for enterprise knowledge graph adoption. The Foundation is a non-profit organization focused on the growth of semantic technology, the adoption of best practices, and the implementation of a shared infrastructure for evaluating data quality. As a consortium of data management and semantic technology advocates, the Foundation firmly establishes a collaborative community of vendors, academics, implementors and consultants as it seeks to build the marketplace for related services, products, and datasets. EKGF.org

Knowledge Graph Conference

The Knowledge Graph Conference is emerging as the premiere source of learning around knowledge graph technologies. We believe knowledge graphs are an underutilized yet essential force for solving complex societal challenges like climate change, democratizing access to knowledge and opportunity, and capturing business value made possible by advances in AI. We bring together leaders across industry sectors to cover the latest in innovation and adoption of knowledge technologies in finance, healthcare, drug discovery, privacy, cyber, media, education, supply chain, inventory management, e-commerce, personal knowledge graphs, visualization, recommender systems, law firms, real estate, and much more. We have organized hundreds of workshops, tutorials, presentations, keynotes, panel discussions, and demonstrations of knowledge technologies. knowledgegraph.tech

CONTACT INFORMATION

Michael Atkin
Managing Director
+1.240.602.8390
atkin@content-strategies.com *(principal analyst)*

Dennis Wisnosky
Chair,
Enterprise Knowledge Graph Foundation
+1.630.240.6910 dennis@ekgf.org

EKGF.org

François Scharffe
Co-Founder, Knowledge Graph Conference
+1.917.756.9535 francois@knowledgegraph.tech

Thomas Deely
Co-Founder, Knowledge Graph Conference
+1.917.385.4927 thomas@knowledgegraph.tech

KnowledgeGraph.tech