

CAPCO

KNOWLEDGE GRAPHS: BUILDING SMARTER FINANCIAL SERVICES



Outside financial services, a growing number of firms are mastering the ability to blend and exploit data, fueling decision-making with profound insights. But the benefits do not stop there: they are also able to offer personalization at an industrial scale, making deeper connections with their customer base; comply faster, compete harder and flourish. As a result, many are entering a period people are calling ‘the knowledge age’.

In contrast, in financial services, a lot of firms are still stuck in the past, using decades old tech to compete on a whole new data-driven playing field. They are unable to address and utilize legacy, siloed data; their clients and business users are frustrated by poor answers to simple questions, and those responsible for change are drowning in complexity. These firms face growing challenges from regulators, and as part-fossilized data dinosaurs, they are threatened with extinction.

SECRET SAUCE?

There is a potent technology that most of us use every day, often unknowingly. It uses data – just a bit differently – to create shared knowledge at incredible scale, from fragments of facts. This technology is at the heart (or brain) of tools we have come to depend upon, tools that help us to search (Google), navigate (SatNav), or turn on lights or music with our voices (Alexa/Siri).

This sounds high tech, but it was born all the way back in 1736.¹ It sounds complex, but its building blocks could not be simpler. It sounds expensive, but you could start using it today – with a

small, intuitive change to the way you work. You can create new value within minutes by capturing what you know, or reshaping information you already have.

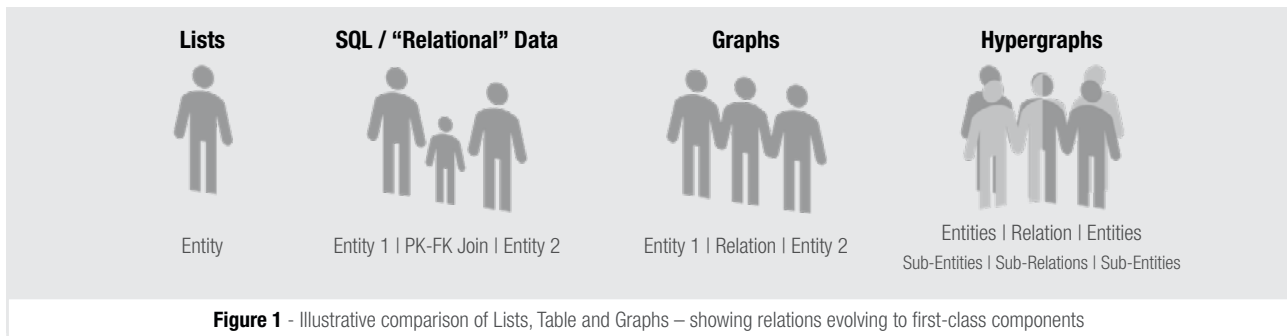
What is this mysteriously potent, common, hidden technology?

- Semantic and graph thinking and technology, specifically, **knowledge graphs**.

1. https://en.wikipedia.org/wiki/Seven_Bridges_of_K%C3%B6nigsberg

WHAT IS A KNOWLEDGE GRAPH?

The definition of a knowledge graph is a data structure (usually held in a graph database) that represents things and how they relate to each other. It is data that describes a network, where the relationships between things (or business objects) are at least as important as the things themselves.



Traditional databases (SQL / also confusingly called ‘relational databases’) are lists of things too, but awkwardly joined with keys (primary keys and foreign keys). Designing and using these joins requires data modelers and a good up-front understanding of questions and answers that the data must serve. SQL was designed in a time when cost of storage was a major factor. For the right problem, it remains the right tool and will have its place in the tech stack for the foreseeable future. But creating even simple networks using keys and joins, requires the rare combination of technical knowledge of the data and a broader awareness of the business context. Querying network data in SQL requires recursive queries that burn developer and processor time unnecessarily. Even then, some basic realities are simply impossible to represent in any useable way.

In knowledge graphs, the simple, atomic facts (known as triples) – state how things are related. These things (sometimes called entities, nodes or vertices) and their relationships (sometimes called edges) – self-assemble into chains, or paths that can be used to infer new, often unexpected, valuable connections. Inference and reasoning are just two examples of analytics capabilities that graphs enable. Graph data can be created

by extracting or virtualizing existing information, expert or crowdsourcing, and also machine learning – increasingly a combination of all three.

Capco expects that within the next two years, 80% of financial services firms will be building graphs. While some of these initiatives may stall – buried in tech labs, those driven by well understood business needs, will be game-changing.

“

*The application of graph processing and graph databases will grow at 100% annually*²

”

2. <https://www.gartner.com/smarterwithgartner/gartner-top-10-data-analytics-trends/>

USE CASES

Outside of financial services, graphs and specifically knowledge graphs are delivering transformational benefits.

- **In publishing** the BBC, the FT and Thomson Reuters all use graphs to curate and organize their content, improving journalism, search and personalized data feeds. The World Economic Forum has gone as far as publishing the articles and graph structure as a Transformation Map.
- **In pharmaceuticals** graphs are helping to match therapies to patients, and to find cures for many serious diseases.
- **Search and social:** Google's knowledge graph is powering more meaningful, intelligent search results. Social media platforms (e.g. LinkedIn) use knowledge graphs to curate posts and fuel their advertising algorithms. All use graphs for their ability to represent and analyze networks of users to identify influencers and clusters to drive their marketing power and products.

Meanwhile, **the financial services industry** has had a long-standing love affair with the spreadsheet. We all love lists and tables. We like data to fit in a box. Driven by overwhelming data volume, complexity and chaos, some firms are awakening to the potential of graph, helping their data to think and link 'outside the box'.

So, who is using graphs and knowledge graphs, and for which use cases?

CxO: for CEOs and other senior leaders

- **Strat-Nav:** Sat Nav for your business strategy? As leaders, it can feel as if the firm's great strategy fails to break

through the permafrost of middle management layers. Backing your strategy with a cascading knowledge graph allows for a clear and precise dialog with all levels of the business at any scale. Goals are cascaded and federated with great clarity and control, and metrics, issues, and reporting can be rolled up along the same structure.

- **Operational resilience:** UK financial regulators are now focusing on the largest firms' ability to continue providing key services during operational disruptions. FCA and PRA require compliance by the end of 2021. Operational resilience³ is the ability of a firm to respond to events and manage the resulting disruption to its services to avoid harm to clients and avoid threatening broader financial stability. A 'digital twin' for operational resilience builds on a knowledge graph to provide an efficient and effective tool to prepare for crises and respond when they occur.

Data: for chief data officers (CDOs)

- **Semantic / augmented data catalogs:** the first step in managing data is to name it and understand where it is coming from and where it is going to. Mapping this without some smart automation is expensive and can be very costly to maintain. Data catalogs as enhanced by knowledge graphs simplify this and allow data to be understood in context – relating data to people, systems, processes, changes, issues, rules and regulations.
- **Smart Data Strategy:** adding a semantic strategy to your data strategy is a great way to build capability and value. Ensuring that each step in taming your data results in assets that can form part of a connected data landscape will reduce the cost and increase the value of each forward step.

3. <https://www.capco.com/Intelligence/Capco-Intelligence/Managing-The-Inevitable-A-Primer-On-Operational-Resilience>

Technology: chief information officers (CIOs) and heads of data engineering

- **Relational + graph data:** at an application level, for decades, SQL data has been the only option or the default option. Now with graph emerging as one of the more potent NoSQL options, combinations of legacy and graph are a powerful cocktail – enabling the best of both worlds.
- **Semantic Data Fabric :** struggling with siloed, inconsistent data, that is hard to access, graphs offer the option to create a semantic data fabric – a layer of a modern architecture that ensures more consistent, more cohesive and more complete data.

Analytics: various data-driven business leaders

Data is no longer just for ‘data people’. Many business leaders want data-driven insights into clients, the market – all within the regulatory and ethical considerations of explainable AI (artificial intelligence). Graphs are not black box solutions like ML (machine learning models) but infer or reason over logical paths that can be exposed to explain how an answer was formed.

- **Deep text analytics - contract reviews, document sifters & next best action recommenders** – give sales and research teams graph-augmented superpowers, monitoring the firehose of news or libraries of contracts to identify what matters and spot nuggets of insight in the noise.
- **Compliance, KYC and fraud detection:** using graph to spot unusual identity or transaction patterns is one of the early point solutions where graphs are already acting as unblinking, expert detectives. Building on identity resolution, firms are turning KYC from a compliance-focused cost, to a revenue-driving, risk-mitigating ‘deep client insight’ capability.

Risk: risk & business managers

- **Enterprise digital twin:** digital twins are models of real-world systems. In this case a model of one or more areas or layers of the enterprise. This helps to better manage and transform the operating model – the complex interdependencies of people, processes, technology – as

well as data, policies, and change. Knowledge graphs offer a simple yet powerful way to model, manage and exploit this complexity. These can be built by virtualizing existing information with SMEs adding detail where necessary.

- **Graph-augmented policy, risk and control framework:** this relates the many policies and regulations to roles, providing a personalized, prioritized policy that shows staff what they need to know most.

Product innovation: banking and risk management product owners

- **Graph-enabled products:** build products that fit reality, rather than trying to make reality fit products. The need to fit products into tabular structures limits their ability to flex to real-world needs. Too often products cannot be offered because the data model cannot represent the demand.
- **Product interoperability:** even the most complex structured products can be more accurately defined and better managed using ontologies. Standards such as ISDA’s CDM (Common Domain Model) and FIBO (Financial Industry Business Ontology) – offer business concept model to describe instruments, parties and processes in the financial industry.

First steps: getting it, enabling agile knowledge curation

Building a knowledge graph may sound like a costly, significant commitment. However, the first step is to simply structure data differently – take existing tabular data, and just transform it to a set of entities and relationships. This can be used to explore and demonstrate the concept and evidence the potential value.

Graph data does not require an up-front schema or design and suits agile ways of working. Driven by business questions, it is not uncommon to reshape graphs on a daily cycle as the knowledge grows and the ability to answer more business questions develops.

Simple, low cost, high value first steps, and the ability to iteratively, or exponentially scale make building Knowledge Graphs an easy fit to most budgets or planning profiles.

A few key considerations:

- **Beware 'the hairball':** some clients (and some vendors) may believe that you choose a graph database, add some data, and visualize the output. A business-driven approach is essential to avoid creating a "hairball". Graph visualizations at any meaningful scale are tempting, and can be interesting, but are too complex for many use cases. Output lists or tables can be more powerful alongside enhanced search and browse capabilities.
- **Human-in-the-loop:** a graph is an opportunity to combine knowledge from different sources, and this includes experts and users. While ML can help with knowledge graphs, ensuring that the right people can contribute in the right way is key. Key roles to create shared value:
 - **Experts and ontologists:** within focused areas, your experts should curate and author knowledge, allowing its value to be exploited at scale.
 - **Data engineers:** automation should build much of your graph by extracting or virtualizing data you already have.
 - **Business users:** users offer a valuable chance to sense-check and crowd-source improvements to a knowledge graph by suggesting missing details or highlighting misinterpretation.
- **Full stack:** what are the components of a full graph technology stack?
 - **User interfaces:** familiar-looking interfaces with graph-powered search and browse features. Graph visualization, only where appropriate.
 - **Deep text analytics:** the ability to exploit semantics and AI to find meaning in masses of documents and data and then add that knowledge to the graph if appropriate. Can also extend to NLP (natural language processing).
 - **Taxonomy and ontology editor:** used by graph curators to define and refine ontologies and taxonomies. This shapes the blueprint for knowledge represented in the graph. Modern ontology editors are user-friendly enough to be used by SMEs directly.
 - **Virtualization / data integration:** allows information to be made available in the graph from existing data sources – both internal and external.
 - **Graph store / database:** stores the graph and supports efficient querying, reasoning and control.

CONCLUSION

In recent years, knowledge graphs have rocketed through the hype curve and are now ready-for-business. Capco's main focus is financial services. Across the many practices and sectors, we see complex business challenges that often have complex data challenges at their core. We continue to discover new ways that graphs and knowledge graphs can simplify and accelerate addressing these challenges.

Characteristics of opportunities where a knowledge graph may help include:

- **Need to build shared knowledge:** the information (in any form, about any topic) exists, but it is hard to find and bring together – possible due:
 - Variety of data sources – internal or external to a firm
 - Data with ambiguous identifiers or inconsistent levels of detail
 - When your data is rich in links, networks, or taxonomies with varying levels
- **Need to use shared knowledge:** explore or exploit large amounts of information to model complexity better and get meaningful, explainable answers
 - Using the power of inference and reasoning to join the dots at scale, especially where answers are not known in advance, and may be unexpected
 - Identifying items of interest in a static or streaming information including unstructured documents
 - Capco is building smarter financial services with knowledge graphs. We are proud to be partnering with the leading technology vendors in this fast-evolving space.

Knowledge graphs may be your golden ticket from 'the information age' of the late 20th century⁴ to the true capabilities of present day: the knowledge age. Pick a use case and partner with Capco and its network to see how knowledge graphs can unlock the value of your data - fast.

Webinar:

We are joined by PoolParty to discuss **Knowledge graphs: making smarter financial services.**

4. <https://www.youtube.com/watch?v=uLpK5Jz5l7Q>

CONTACTS

Chris Probert

Partner

chris.probert@capco.com

Mark Kitson

Author

mark.kitson@capco.com

ABOUT CAPCO

Capco is a global technology and management consultancy dedicated to the financial services industry. Our professionals combine innovative thinking with unrivalled industry knowledge to offer our clients consulting expertise, complex technology and package integration, transformation delivery, and managed services, to move their organizations forward.

Through our collaborative and efficient approach, we help our clients successfully innovate, increase revenue, manage risk and regulatory change, reduce costs, and enhance controls. We specialize primarily in banking, capital markets, wealth and asset management and insurance. We also have an energy consulting practice in the US. We serve our clients from offices in leading financial centers across the Americas, Europe, and Asia Pacific.

To learn more, visit our web site at www.capco.com, or follow us on Twitter, Facebook, YouTube, LinkedIn and Instagram.

WORLDWIDE OFFICES

APAC

Bangalore
Bangkok
Gurgaon
Hong Kong
Kuala Lumpur
Mumbai
Pune
Singapore

EUROPE

Bratislava
Brussels
Dusseldorf
Edinburgh
Frankfurt
Geneva
London
Paris
Vienna
Warsaw
Zurich

NORTH AMERICA

Charlotte
Chicago
Dallas
Houston
New York
Orlando
Toronto
Tysons Corner
Washington, DC

SOUTH AMERICA

São Paulo

WWW.CAPCO.COM



© 2020 The Capital Markets Company. All rights reserved.

CAPCO