ENGINE B

A guide to Knowledge Graphs

Creating the context for intelligent decision-making

“To know an object is to lead it through a context which the world provides.”

William James - The Meaning of Truth, 1909

Introduction

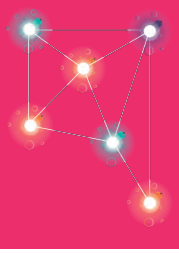

Corporate data can be a huge strategic asset. However, the sheer volumes of data, different types of data, including structured and unstructured, and the fact that in many cases data exists in silos can make analysing and using the data to make better operational decisions difficult. Data silos and data complexities impede everything from simple reporting to data science to machine learning. There have been many attempts over the years to resolve the siloed data challenge, for example, housing data in a data warehouse. However, a data warehouse contains widely varied types of data and datasets that are not easily linked, can contain dirty data and do not always contain unstructured data (data such as emails, PDFs, spreadsheets).

Unstructured data is a plentiful source in an organisation. IDG predicts that by 2022, 93 per cent of all data in the digital universe will be unstructured. Unstructured data generates immense business value, but most organisations have not been able to yield insights due to the challenges involved in analysing it - and **this is where Knowledge Graphs come into play.** In fact, for the past decade, knowledge graphs have been part of our daily lives. For example, Alexa, Siri or Google Assistant are all based on knowledge graph technology.

For enterprise use, Knowledge Graphs use artificial intelligence (AI) and graph technologies to compile organisational data (structured and unstructured), which is then served up visually, showing context and complex relationships between data entities. Knowledge Graphs work similarly to a human brain, by using context to create relationships between entities.

With more and more organisations today relying on AI for their decision-making processes, Knowledge Graphs will increasingly be used by corporates to create the context needed for more intelligent decision-making.

This guide discusses the data challenges of today and delves into the ways Knowledge Graphs are being used as well as exploring use cases for Knowledge Graphs, including auditing, fraud detection, anomaly detection and real estate examples.

ENGINE B

Contents

Today’s data challenges

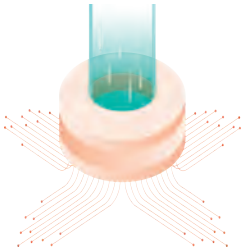

Siloed data

Today’s organisations are dealing with ever-increasing volumes of data and with increasing volumes comes growth in silos, which in turn leads to increasing costs of data integration. Data, that is spread across multiple disconnected data silos, for example in ERP and CRM systems and accounting systems, make efficient collection, aggregation, and analysis of data expensive and time-consuming. Organisations need to connect their disparate data so it can be analysed in a meaningful and more intelligent way.

Data types

IDG predicts that by 2022, 93 per cent of all data in the digital universe will be unstructured.

Not all data is created equal. Some data is structured, however, most data is unstructured, for example an invoice. The way data is collected, processed, and analysed all depends on its format. Structured data is highly-organised and formatted in a way so it's easily searchable in relational databases. Unstructured data has no pre-defined format or organisation, making it much more difficult to collect, process, and analyse. Unstructured data cannot be analysed with current databases because most data analytics databases are designed for structured data, and are not equipped for unstructured data. Therefore, organisations need to find new methods to analyse, locate, extract, organise data

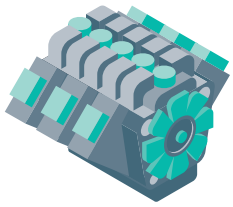

Data quality

It is unlikely an organisation's is going to be 100 per cent accurate. However, that does not mean that an organisation cannot control how reliable its data is. Not only can databases contain the wrong information, but data can also duplicate itself, as well as contain contradictions. It is doubtful that data of inferior quality can bring any useful insights or unearth opportunities to precision-demanding business analysis. However, that does not mean that an organisation cannot control how reliable its data is. Organisations need to ensure that any analysis is based on robust, good quality data.

What are Knowledge Graphs?

Knowledge Graphs combine and uncover connections across silos of information.

Searching for information across disparate systems is an ongoing challenge for most organisations. And, as a result, a plethora of search applications have come into force to try and solve the problem of accessing information for analytics purposes that will ultimately lead to improved operational decision-making.

> The application of graph processing and graph databases will grow at 100% annually through 2022 to continuously accelerate data preparation and enable more-complex and adaptive data science.
>
> Gartner Predicts

However, as analytics tools have become prominent and more sophisticated, user expectations have also increased. Organisations want the single correct answer without having to scroll through databases of results. They do not want to miss any critical information that could negatively impact on their organisation. It is the rise of AI that is changing how everyone wants to interact with information. This is where Knowledge Graphs come into play.

At a high level, Knowledge Graphs are interlinked sets of facts that describe real-world entities, facts or things and their interrelations in a human-understandable form. Unlike a simple knowledge base with flat structures and static content, a Knowledge Graph acquires and integrates adjacent information using data relationships to derive new knowledge, and it works very similarly to a human brain.

A relatable example

If you think of a colour, purple, for example, you may think of the flower Lavender, your favourite jumper or a sweet you liked as a child. Your brain is associating different things you know with the colour purple.

Your brain is behaving like a graph, a network of entities, where the entity in this example, is the colour purple, which is directly connected to flowers, sweets and a jumper that is purple.

If you brain worked in a linear, flat structure, like a database, it would have looked at every flower, shrub or tree in your garden and check the value of the attribute ‘colour’. This way of ‘thinking’ would be far more time consuming and not very practical at all!

What are Knowledge Graphs?

Context for decisions.

Some of the key characteristics of Knowledge Graphs:

- A Knowledge Graph needs to be connected around relevant attributes. It needs pertinent information that is contextually related.
- A Knowledge Graph is dynamic in that the graph itself understands what connects entities, eliminating the need to programme every new piece of information manually.
- A Knowledge Graph makes appropriate associations across attributes that are relevant to decision making due to them being programmed in.
- A Knowledge Graph needs to be understandable. Its use of intelligent metadata helps the graph to find answers to specific problems, even ones we do not know exactly how to ask for.
- Knowledge Graphs work from structured and unstructured data.

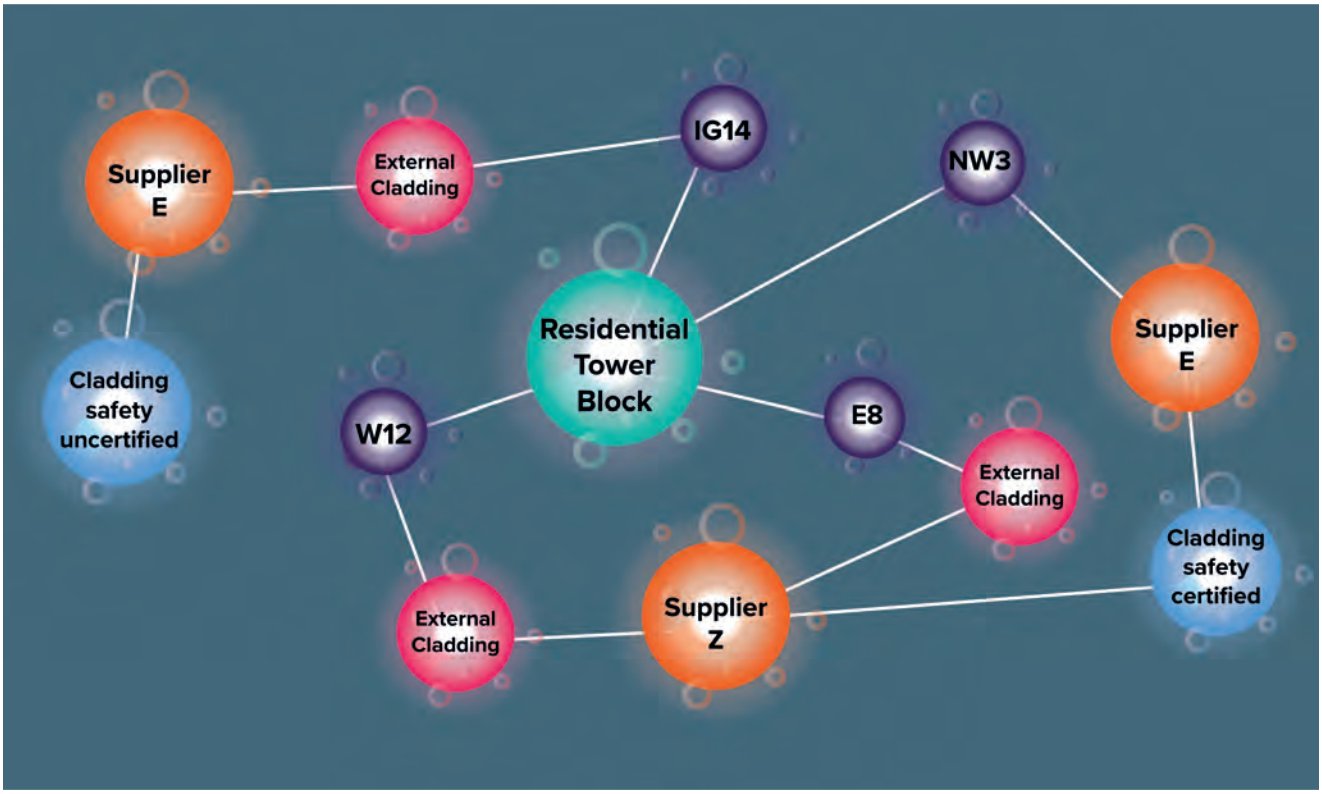

Benefits of Knowledge Graphs

One of the most valuable and fundamental uses of Knowledge Graphs is to express relationships between data. Knowledge Graphs visualise the nature of an underlying relationship and how they are interconnected to gain sufficient enough understanding of the landscape to inform more intelligent strategic business decisions.

At a high level, Knowledge Graphs provide the following main benefits:

Combine siloed data sources

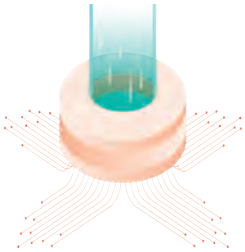

Knowledge Graphs help to combine disparate silos of data, giving you an overview of all of your knowledge – not only departmentally but also across departments and global organisations.

Combine structured and unstructured data

Accumulating data doesn't mean just assembling documents and excel sheets. Knowledge Graph technology means being able to connect different types of data in meaningful ways and supporting richer data services than most knowledge management systems. This information can then be used to extract and discover deeper and more subtle patterns

Help business leaders to make more intelligent decisions, quicker

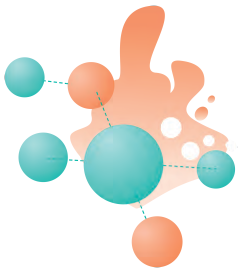

Using a Knowledge Graph means no more looking through piles of paper to find a particular document as Knowledge Graphs provide relevant facts and contextualised answers to your specific questions, rather than a broad search result with lots of irrelevant information.

Knowledge Graphs provide a model of how everything is related, having each subject or object represented only once with all its relationships, in the context of all of the other subjects and their relationships. This makes it possible to see how everything is related at a big picture level.

Benefits of Knowledge Graphs

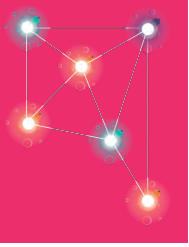

Summarise relationships

Knowledge Graphs summarise relationships so effectively and so efficiently, they can easily scale to accommodate more and more information. Knowledge Graphs are truly unparalleled in their capability to express interconnected relationships.

Insights from hierarchical data

Knowledge Graphs are also a great choice for gaining insights from hierarchical data. Hierarchically viewing data is particularly useful for representing sequences of decisions that lead to conclusions. Knowledge Graphs are also great at prioritising information gathering and provides an ideal, intuitive framework for portraying characteristics of and their relationships.

Revealing communities

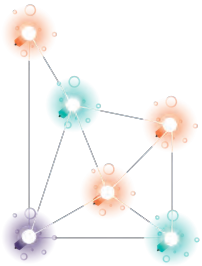

Knowledge Graphs are great for revealing communities, which are fundamental to understanding macro relationships and dynamics in business data. Communities in a Knowledge Graph visualisation are qualitatively reflected by clusters of related members in close proximity, distinguishable from the field of another graph member. Linking communities provides useful information about patterns of interests and influence.

Visualising a flow of information

Knowledge Graphs are also a great way to visualise a flow of information - showing information flows of a business process or any transactional activities that may reveal hidden risks or show patterns over time.

Network data

Knowledge Graphs can be used to show network data, how network routers and hubs are connected and their characteristics.

Difference between a Knowledge Graph and other analytics tools

Beyond the scope of traditional data analytics.

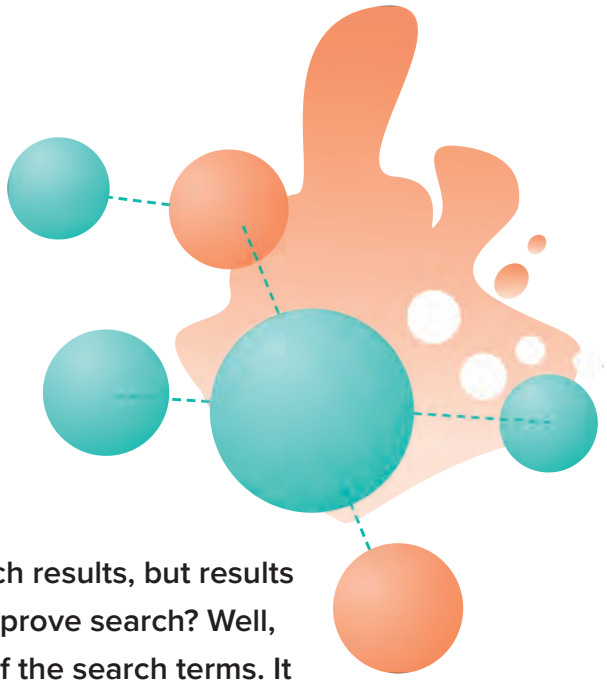

Knowledge Graph-powered search is often called “semantic search” — search enriched with meaning (logic and language). Non-graph related search tools tend to work from linear, static-relational schemes, with thousands, even millions of rows of data in datasets, making discoveries of patterns or anomalies very challenging and not very obvious.

A Knowledge Graph search typically returns fewer search results, but results are more relevant. So, how does a Knowledge Graph improve search? Well, typically a Knowledge Graph will capture the meaning of the search terms. It works by representing the layers of connections between various data sources.

A Knowledge Graph search typically returns fewer search results, but results are more relevant. So, how does a Knowledge Graph improve search? Well, typically a Knowledge Graph will capture the meaning of the search terms. It works by representing the layers of connections between various data sources.

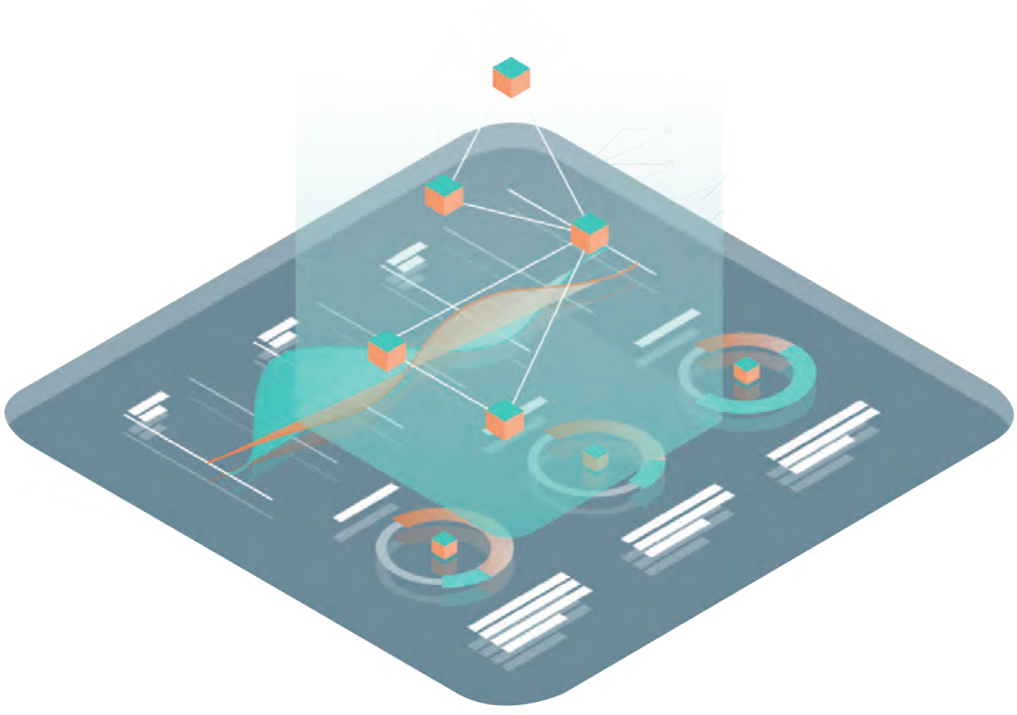

Difference between a Knowledge Graph and other analytics tools

Knowledge Graph characteristics:

- Knowledge Graphs show how tightly several trends or data points are related to each other.

- A Knowledge Graph will determine “connectedness” across data points and create clusters based on levels of influence, frequency of interaction and probability. Once highly complex models are developed and trained, the output is easier to store because of the expanded capabilities, computational power and adoption of graph databases. The user can interact directly with the graph elements to find insights, and the analytic results and output can also be stored for repeated use in a graph database.

- Generating a dynamic graph about how different entities of interest — people , places and things — are related, instead of more-static relational schemes, enables deeper insights that work similarly to the human brain.

- Graph technology underpins the creation of richer semantic graphs or knowledge bases that can enhance augmented analytics models, as well as the richness of conversational analytics.

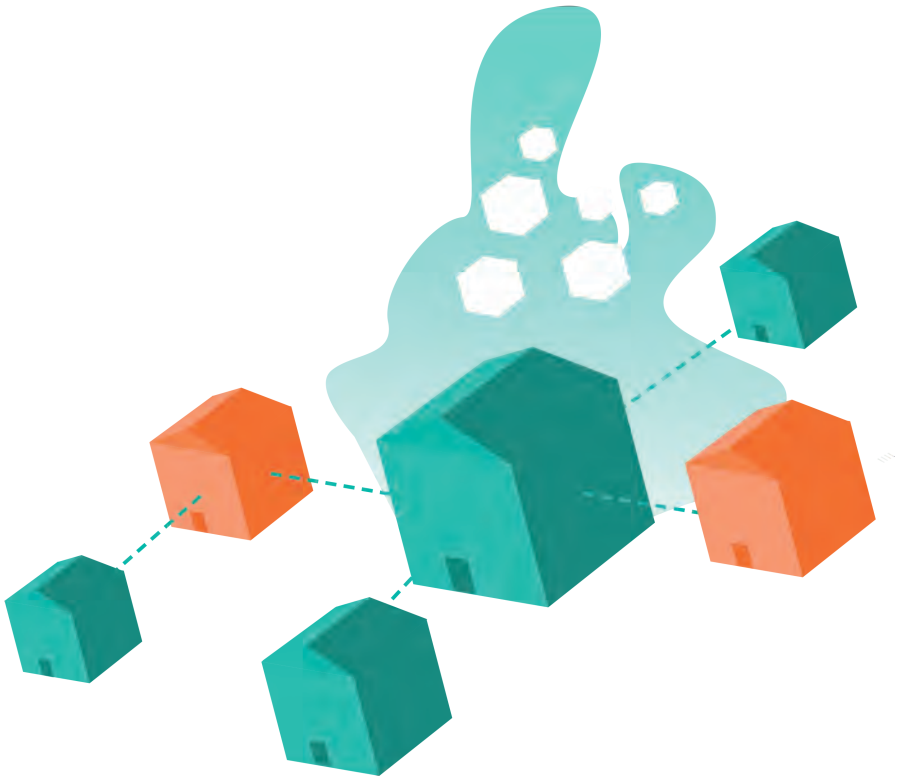

Knowledge Graph use cases

As discussed earlier in this guide, Knowledge Graphs are being used by Google and other voice-activated technologies. However, how are Knowledge Graphs uses in organisations? From an enterprise perspective, many organisations use Knowledge Graphs to give them a competitive edge, improve efficiencies, and to guide operational decisions.

Various types of analysis can be done with Knowledge Graphs in the business domain:

- Finding anomalies
- Managing networks and supply chains
- Identifying risk patterns
- Evidence-to-record confirmation
- Business process analysis and 4+ way matching
- Financial asset market value relationships
- Mapping hierarchies
- Detecting communities

Fraud Detection Knowledge Graph

Uncover and detect fraud quickly

Financial frauds are expensive and fraudsters are continually finding new ways to exploit monitoring systems. Many organisations try to tackle and identify fraud by analysing transactions or people in isolation. This method of fraud analyses has its limitations and prevents organisations from seeing the bigger picture.

Organisations need to adopt a context-driven view of information, one that uncovers connections across silos of information so data can be analysed in a meaningful and more intelligent way.

Fraud Detection using Knowledge Graphs

Engine B is focused on solving these challenges with our Fraud Detection Knowledge Graphs that allow organisations to interrogate and analyse their client's data in a powerful, context-driven way.

Our Fraud Detection Knowledge Graphs detect data abnormalities or problematic records. For example, audit firms, can identify hidden anomalies in a client's data, make sense of huge data volumes, and surface the most valuable connections to understand the context created. Our Fraud Detection Knowledge Graphs help organisations to:

- Identify more instances of fraud that would otherwise not been possible to see in traditional analyses tools.
- Reduce the time it takes to investigate fraud.
- Quickly eliminate non-fraudulent events.

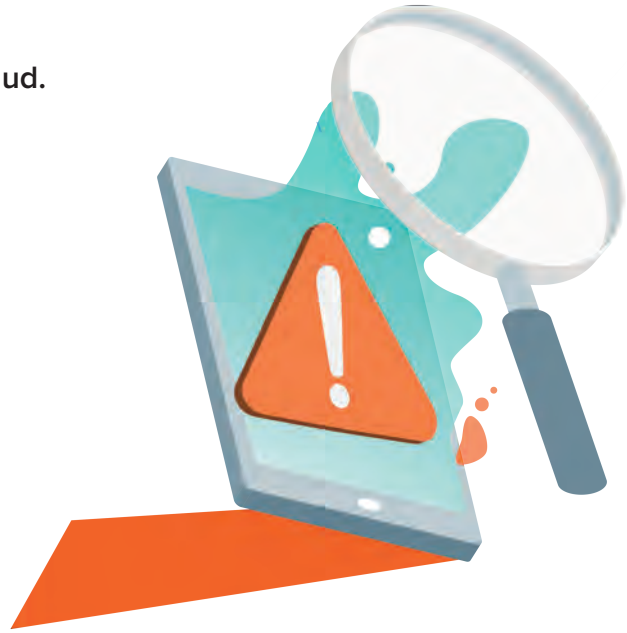

Anomaly Detection Knowledge Graph

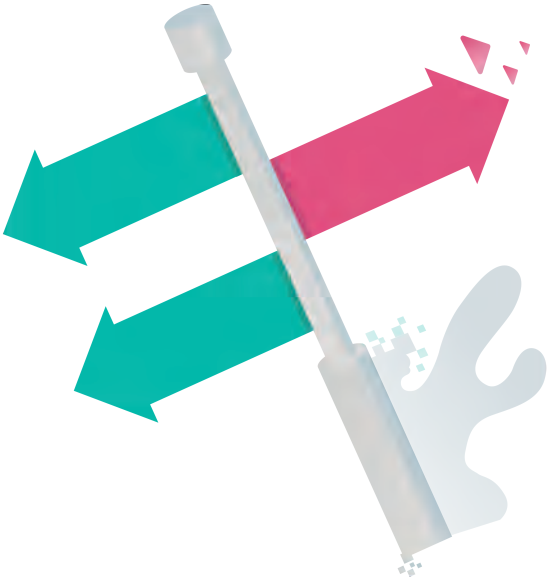

Today's organisations generate huge volumes of data daily. If data is used correctly, it can be a huge strategic asset to a company and help them to make faster, more intelligent decisions.

However, many organisations still rely on retrieving and analysing data from disparate, siloed systems. This method of detecting anomalies in data has its limitations and prevents organisations from zooming out and seeing the true scale of the picture.

Organisations need to adopt a context-driven view of data, one that uncovers connections between silos of data so the data can be used in a more meaningful way.

The difference between traditional linear models and Knowledge Graphs in detecting anomalies

Many analytics solutions designed to present anomalies are based on-linear models and simply present an enormous list of risks flagged as 'high, 'medium' or 'low', with no context behind why the data is risky. This means that each risk has to be investigated, which is very time-consuming, inefficient and costly.

Anomaly Detection using Knowledge Graphs

Engine B's Anomaly Detection Knowledge Graphs enable organisations to visualise unexpected changes or deviations from an expected pattern in a dataset and see reasons for the anomaly. This means there are far fewer anomalies to investigate and an audit, for example, can be completed much quicker with more intelligence behind decisions.

Our Anomaly Detection Knowledge Graphs also identify anomalies within different data types, such as structured, semi-structured or unstructured data.

There are many use cases for Knowledge Graphs, however, the principle for their use is the same - taking large amounts of data from various data sources and adding value to it, so that it can be used and re-used in a meaningful and more intelligent way.

Support resources

The Future of Audit event recording

Access Engine B's recent digital event to hear from Robert Hodgkinson of the ICAEW on the future of audit. See our Audit Common Data Model and Audit Knowledge Graphs working in tandem to provide any size of audit firm access to one common source of client data with intricate, context-driven decision-making capabilities.

Access the recording >>

Audit Common Data Model guide

If you would like learn more about Engine B's Audit Common Data Model, created in partnership with Microsoft, thirteen audit firms and the ICAEW, then you can download our guide to the Audit CDM which discusses its key benefits in tackling the data challenges of today.

Download >>

A guide to anomaly detection in audit

Access our guide to find out how Engine B's anomaly detection software is providing auditors with a laser-like focus on fraud risk using machine learning and knowledge graph technology.

Download >>

Product guides

For our full range Engine B product guides, please visit our website.

See more >>

Contact us

About Engine B

Engine B is a digital technology company specialising in AI and data analytics for the professional services industry. With our Audit Common Data Model and Knowledge Graphs, we are transforming the way that audit services are delivered and helping to create a level playing field, where any size of audit firm can compete.

We work with key industry partners, such as Microsoft, thirteen audit firms, ICAEW and educational institutions to vastly improve the quality of professional services practices, to empower organisations to make better operational decisions and to advance their digital growth.

Contact us for more information on our Audit Common Data Model and Knowledge Graphs.

Visit our website

ENGINE B

engineb.com