



IMPACT REPORT

Stardog pitches its knowledge graph as the data unification platform for the enterprise

FEBRUARY 22 2018

BY MATT ASLETT (/ANALYST-TEAM/ANALYST/MATT+ASLETT)

While perhaps best known for its graph database, Stardog has always had an eye on the bigger picture. By coupling the Stardog database with a Knowledge Toolkit, the company is offering a platform to provide a knowledge graph that unifies data from multiple data sources, across the enterprise, without it having to be moved into a central location.

The 451 Take

We have been aware of Stardog (the graph database) for some time, and have been interested to see Stardog (the company) elevating its marketing to talk about wider data management issues. The combination of Stardog with a Knowledge Toolkit to serve as an enterprise knowledge graph is an intriguing one, and we think that the data unification message will appeal to many enterprises as a potential solution to the data silo problem. Although interest in knowledge graphs is growing, getting that message in front of potential buyers is going to be an initial challenge, and one that should be aided by Stardog's message of data unification, rather than the potential advantages of knowledge graph.

Context

The company now known as Stardog was founded in 2006 as Clark & Parsia (it was briefly known as Complexible between 2015 and 2016 before settling on the name of its core product). The company can trace its origins to University of Maryland research lab, where its cofounders Kendall Clark (CEO) Evren Sirin (CTO) and Michael Grove (VP engineering) cut their teeth on semantic web standards.

Stardog made a start delivering semantic web projects for US government agencies (including NASA), as well as companies in the financial services, energy, health and retail sectors. The company has been focused throughout its lifespan on the development of the Stardog graph database, but has in recent years focused more explicitly on the combination of Stardog with its Knowledge Toolkit to serve as an enterprise knowledge graph.

Although the concept of the knowledge graph is relatively new to mainstream enterprises, it has become more widely acknowledged (if not yet adopted) thanks to the likes of Facebook, Google and LinkedIn, and their investment in graph-based maps of available data across the organization – and the relationships between that data – that can be used to combine data in

an agile manner in response to a wide range of use cases.

Stardog boasts close to 50 paying customers, most of which are Global 2000 enterprises. Named customers include BNY Mellon, FINRA, NASA, Cisco, Apple, eBay, Bosch, Ericsson, Nokia and Samsung. Key vertical markets include financial services, healthcare, government, high-tech and logistics/IoT.

Stardog employs more than 30 people today, up considerably from less than 10 a year ago. The company's growth has been fueled by a recent spate of VC funding. It announced in July 2016 that it had raised a \$2.3m seed round led by Core Capital and Boulder Ventures, which returned in late 2017 with Grotech Ventures to provide a \$6m series A round that was expanded to \$9m in January 2018 with the addition of Tenfore Holdings.

Products

As noted, at the heart of Stardog's enterprise knowledge graph platform is the graph database of the same name, but it would be wrong to think of the company as a database provider. Instead, Stardog couples the graph database with its knowledge toolkit and positions the combination as a platform for the unification of data from across the enterprise. Specifically, Stardog pitches the enterprise knowledge graph as a potential solution to the age-old problem of data silos, caused by data related to individual applications being stored in multiple databases.

Rather than proposing that potential customers move all their data into the Stardog database, the company positions it as an abstraction layer that is used to model data and map relationships between multiple data stores and federate queries across the underlying data estate. Stardog does this by providing a virtual graph engine for translating GraphQL queries into the appropriate queries (including SQL) for execution by a remote database – which could be relational, key value, document or graph databases – as well as unstructured content (PDFs, emails, text files, images, etc.).

By doing so, Stardog enables users to create enterprise graphs that span multiple data stores, while retaining existing data models, business logic and investments in data processing infrastructure. The company also notes that a prime advantage of this approach is its inherent flexibility, which enables users to quickly generate new applications based on the knowledge graph that utilize different data sources for different purposes without it needing to be copied, moved or remodeled.

In addition to the virtual graph engine, the knowledge toolkit also supports advanced analytics capabilities including logical reasoning (for inference, explanation and model checking), as well as machine learning (statistical inference and probabilistic reasoning), along with semantic search, geospatial semantics and query, and knowledge graph construction services. Stardog also provides support for the Apache Tinkerpop graph computing framework, as well as Linkurious's graph visualization and analysis software, along with Tableau for self-service visualization.

The company recently introduced Stardog Studio, a development environment including a SPARQL query notebook and tools for managing Stardog databases, users and roles. Besides automated functionality, there is some up-front work to be done to make all this happen. When engaging on a new project, the company's staff will sit down with a customer to identify the appropriate data sources, define the queries, and map the concepts to underlying databases and tables.

Competition

Stardog maintains that its biggest competition comes from the fact that potential customers are wedded to traditional approaches, and don't understand that there are potential ways of overcoming data silos without moving data into a single location (such as a data lake, for example). We would also observe that the understanding of semantic web and graph technologies is not widespread, and while it has improved in recent years, Stardog's virtual graph engine approach and data unification positioning is quite different from that of graph database specialists such as Neo4j, for example.

Stardog is more comparable with the approach taken in recent years by MarkLogic, which has pitched its enterprise NoSQL database as an operational data hub that integrates from multiple database silos to drive new application requirements. Similarly, Cambridge Semantics is positioning its Anzo Smart Data Lake technology as a unified platform for multi-structured data management (based on a massively parallel in-memory graph analytics engine). However, Stardog appears to be differentiated by its focus on the unification of data without requiring movement into a single platform.

Other potential competitors include Franz, which recently adjusted the messaging around its AllegroGraph RDF database to address knowledge graph use cases, and Ontotext, with its semantic knowledge graph platform driven by the GraphDB database. Another new player in this space is Siren, an Ireland-based company with a similar approach to federate queries enabled by semantic web technologies and virtual graph queries. The company recently raised a €3m (\$3.7m) seed round led by the Atlantic Bridge University Fund.

SWOT Analysis

Strengths

Stardog boasts undoubted expertise in graph technologies, as well as a differentiated offering that provides a potential solution to the age-old problem of data silos.

Weaknesses

Graph-based approaches have seen niche adoption in the past, and although understanding is growing, Stardog will likely have some evangelizing to do among mainstream enterprises. Reference customer stories will be invaluable.

Opportunities

Understanding of the concept of knowledge graphs is growing – thanks to the likes of Google, Facebook and LinkedIn – and Stardog's data unification positioning may strike a chord where arguments about the merits of graph may fall on deaf ears.

Threats

Solving the data silo problem is no simple task, and enterprises are bound to be dubious if Stardog overstates its claim to offer a potential answer. The company is aware of this, however, and is being careful not to over-promise and under-deliver.

Matt Aslett (/analyst-team/analyst/Matt+Aslett)

Research Director

M&A ACTIVITY BY SECTOR

Information management / Data management / Data integration (99) (https://makb.the451group.com/results?basic_selected_sectors=118)

Information management / Data management / Non-relational databases (21) (https://makb.the451group.com/results?basic_selected_sectors=432)

M&A ACTIVITY BY ACQUIRER

Apple/EMC/Ericsson/Microsoft/Research In Motion/Sony (70) (https://makb.the451group.com/results?basic_acquirers=Apple/EMC/Ericsson/Microsoft/Research+In+Motion/Sony)

Apple/EMC/Ericsson/Microsoft/Research In Motion/Sony (58) (https://makb.the451group.com/results?basic_acquirers=Apple/EMC/Ericsson/Microsoft/Research+In+Motion/Sony)

Bank of New York Mellon Corporation [aka BNY Mellon] (2) (https://makb.the451group.com/results?basic_acquirers=Bank+of+New+York+Mellon+Corporation+aka+BNY+Mellon)

Cisco Systems Inc. (135) (https://makb.the451group.com/results?basic_acquirers=Cisco+Systems+Inc.)

eBay Inc. (67) (https://makb.the451group.com/results?basic_acquirers=eBay+Inc.)

Elite [Thomson Reuters Corporation] (58) ([https://makb.the451group.com/results?basic_acquirers=Elite+\[Thomson+Reuters+Corporation\]](https://makb.the451group.com/results?basic_acquirers=Elite+[Thomson+Reuters+Corporation]))

Facebook (68) (https://makb.the451group.com/results?basic_acquirers=Facebook)

Google Inc. (190) (https://makb.the451group.com/results?basic_acquirers=Google+Inc.)

LinkedIn Corporation (19) (https://makb.the451group.com/results?basic_acquirers=LinkedIn+Corporation)

Lipper (11) (https://makb.the451group.com/results?basic_acquirers=Lipper)

NAVTEQ [Nokia] (50) ([https://makb.the451group.com/results?basic_acquirers=NAVTEQ+\[Nokia\]](https://makb.the451group.com/results?basic_acquirers=NAVTEQ+[Nokia]))

Robert Bosch GmbH (8) (https://makb.the451group.com/results?basic_acquirers=Robert+Bosch+GmbH)

Samsung Electronics Co. Ltd. (26) (https://makb.the451group.com/results?basic_acquirers=Samsung+Electronics+Co.+Ltd.)

Figures shown indicate number of transactions

COMPANY MENTIONS (PRIMARY)

Stardog (/search?company=Stardog)

COMPANY MENTIONS (OTHER)

Apple , The Bank of New York Mellon , Bosch , Boulder Ventures , Cambridge Semantics , Cisco , eBay , Ericsson , Facebook , Financial Industry Regulatory Authority , Franz , Google , Grotech Ventures , LinkedIn , Linkurious , MarkLogic , NASA , Neo4j , Nokia , Ontotext , Reuters Group , Samsung Electronics , Tableau , Thomson Reuters , University of Maryland (/search?company=University+of+Maryland)

CHANNELS

Data Platforms & Analytics (/dashboard?view=channel&channel=6)

SECTORS

All / Information management / Data management / Data integration (/search?sector=118)

All / Information management / Data management / Non-relational databases (/search?sector=432)

