

# Data-Centric Architecture Forum – DCAF 2021

Data and the subsequent knowledge derived from information are the most valuable strategic asset an organization possesses. Despite the abundance of sophisticated technology developments, most organizations don't have disciplines or a plan to enable data-centric principles.

## **DCAF 2021 will help provide clarity.**

Our overarching theme for this conference is to **make it REAL**. Real in the sense that others are becoming data-centric, it is achievable, and you are not alone in your efforts.

Join us in understanding how data as an open, centralized resource outlives any application. Once globally integrated by sharing a common meaning, internal and external data can be readily integrated, unlike the traditional “application-centric” mindset predominantly used in systems development.

The compounding problem is these application systems each have their own completely idiosyncratic data models. The net result is that after a few decades, hundreds or thousands of applications implemented have given origin to a segregated family of disparate data silos. Integration debt rises and unsustainable architectural complexity abounds with every application bought, developed, or rented (SaaS).

**Becoming data-centric will improve data characteristics** of findability, accessibility, interoperability, and re-usability (FAIR principles), thereby allowing data to be exported into any needed format with virtually free integration.\



**Dr. Jans Aasman to present –  
Franz's approach to Entity Event  
Data Modeling for Enterprise  
Knowledge Fabrics**

---

# **Data Fabrics and Knowledge Graphs – A Symbiotic Relationship**

Dr. Jans Aasman's recent article in Dzone.

The data fabric notion is gaining credence throughout the analyst community, in much the same way knowledge graphs have done so for years. Both technologies link all relevant data for a specific business purpose, which is why the most successful companies in the world employ them.



Amazon's knowledge graph retains metadata about its vast product array; Google's captures data about an exhaustive list of web entities of interest. Lesser-known organizations regularly deploy these mechanisms for everything from comprehensive customer views to manufacturing processes.

Data fabrics have a unique, symbiotic relationship with the knowledge graph movement because they substantially streamline the processes to extract data from the myriad sources that populate these platforms. In turn, knowledge graphs provide some of the fundamental capabilities enabling data fabrics to accomplish this objective.

Read the Full Article at Dzone.

---

## **Text Analytics Forum 2020 – KMWorld Connect**

**Join us November 17, 2020** – Text Analytics has the ability to add depth, meaning, and intelligence to any organization's most under-utilized resource – text. Through text analytics, enterprises can unlock a wealth of information that would not otherwise be available. Join us as we explore the power of text analytics to provide relevant, valuable, and actionable data for enterprises of all kinds.

Jans Aasman to present – Analyzing Spoken Conversations for Real-Time Decision Support in Mission-Critical Applications

November 17, 2020 at 2PM Eastern

---

## **Sharing Ontologies Globally**

# To Speed Science And Healthcare Solutions – OntoPortal

## International Ontology Sharing Is Becoming A Reality

A consortium of researchers recently formed an organization dedicated to standardizing how scientists define their ontologies, which are essential for retrieving datasets as well as understanding and reproducing research. The group called OntoPortal Alliance is creating a public repository of internationally shared domain-specific ontologies. All the repositories will be managed with a common OntoPortal appliance that has been tested with AllegroGraph Semantic Knowledge Graph software. This enables any OntoPortal adopter to get all the power, features, maintainability, and support benefits that come from using a widely adopted, state-of-the-art semantic knowledge graph database.

Read the full article at [HealthIT Outcomes](#) –

**As Dr. Jans Aasman, CEO of Franz Inc. explains,** “When building a Knowledge Graph as your enterprise’s single source of truth, it’s critical to include ontologies and taxonomies. AI applications and complex reasoning analytics require information from both databases and knowledge bases that contain domain information, taxonomies, and ontologies to solve complex questions. To make this possible, we developed a novel hybrid sharding technology called FedShard, which facilitates the combination of data and knowledge required by applications like Montefiore’s PALM. But this approach is not unique or specific to Healthcare, it is applicable in many other industries, which is why we are excited about OntoPortal’s plans to bring sharing of domain ontologies to a

broad audience.”



---

## **Knowledge Graphs: A Single Source of Truth for the Enterprise**



The notion of a “single source of truth” for the enterprise has been the proverbial moving goalpost for generations of CIOs, only to be waylaid by brittle technology and unending legacy systems. Truth-seeking visions rebuffed by technological trends have continuously confounded business

units trying to achieve growth and market penetration. But technology innovation has finally led us to a point where CIOs can now deliver that truth.

## **Graphing the Truth**

Knowledge graphs possess the power to deliver a single source of truth by linking together any assortment of data sources required, standardizing their diversity of data elements, and eliminating silos. They support the most advanced analytics options and decentralized transactions, which is why they’re now deployed as systems of records for some of the most significant, mission-critical use cases affecting our population.

Because they scale to include almost any number of applications – and link to other knowledge graphs as well – these repositories are the ideal solution for real-time information necessary to inform business users’ performances with concrete, data-supported facts. Most importantly, users can get an exhaustive array of touchpoints pertaining to any customer, product, or interaction with an organization from the knowledge graph, making it a single source of truth.

Read the full article at [Dataversity](#).

---

# Gartner Hype Cycle for AI – Knowledge Graphs

According to Gartner's 2020 Hype Cycle for Artificial Intelligence – Despite the global impact of COVID-19, 47% of artificial intelligence (AI) investments were unchanged since the start of the pandemic and 30% of organizations actually planned to increase such investments, according to a Gartner poll. Only 16% had temporarily suspended AI investments, and just 7% had decreased them.

“AI is starting to deliver on its potential and its benefits for businesses are becoming a reality”

Gartner's – AI Hype Cycle Article

The Hype Cycle growth is consistent with Franz's customer interest in our Enterprise Knowledge Graph Solutions – Read our recent White Paper.

# Hype Cycle for Artificial Intelligence, 2020



[gartner.com/SmarterWithGartner](https://gartner.com/SmarterWithGartner)

Source: Gartner  
© 2020 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner and Hype Cycle are registered trademarks of Gartner, Inc. and its affiliates in the U.S.

**Gartner®**

## Connected Data London – The



# Future of AI in the Enterprise

**The Future of AI in the Enterprise:**

**Entity-Event Knowledge Graphs for Data-Centric Organizations**

**Presented by: Dr. Jans Aasman**

**Register:**

**<https://enterprise-kg-cdl-online-meetup.heysummit.com/>**

Personalized medicine. Predictive call centers. Digital twins for IoT. Predictive supply chain management, and domain-specific Q&A applications. These are just a few AI-driven applications organizations across a broad range of industries are deploying.

Graph databases and Knowledge Graphs are now viewed as a must-have by Enterprises serious about leveraging AI and predictive analytics within their organization.

See how Franz Inc. is helping organizations deploy novel Entity-Event Knowledge Graph Solutions to gain a holistic view of customers, patients, students or other important entities, and the ability to discover deep connections, uncover new patterns and attain explainable results.

**Description:**

To support ubiquitous AI, a Knowledge Graph system will have to fuse and integrate data, not just in representation, but in context (ontologies, metadata, domain knowledge, terminology systems), and time (temporal relationships between components of data). Building from 'Entities' (e.g. Customers, Patients, Bill of Materials) requires a new data model approach that unifies typical enterprise data with knowledge bases such as industry terms and other domain knowledge.

Entity-Event Knowledge Graphs are about connecting the many dots, from different contexts and throughout time, to support and recommend industry-specific solutions that can take into account all the subtle differences and nuisances of entities and their relevant interactions to deliver insights and drive growth. The Entity-Event Data Model we present puts core entities of interest at the center and then collects several layers of knowledge related to the entity as 'Events'.

Franz Inc. is working with organizations across a broad range of industries to deploy large-scale, high-performance Entity-Event Knowledge Graphs that serve as the foundation for AI-driven applications for personalized medicine, predictive call centers, digital twins for IoT, predictive supply chain management and domain-specific Q&A applications—just to name a few.

During this presentation we will explain and demonstrate how Entity-Event Knowledge Graphs are the future of AI in the Enterprise.

---

## **Knowledge graphs enhance customer experience through speed and accuracy**

KMWorld's recent article covers AllegroGraph and Franz's customer N3 Solutions.

The Full Article – KMWorld



Knowledge graphs are a way to model enterprise knowledge and represent complex interrelationships in data. Information stored in a graph database can enable rapid retrieval of well-targeted results and provide insights into customers' interests and needs. Gartner predicts a 100% per-year growth in applications for graph analytics and databases for the next several years. Although knowledge graphs have been deployed by major companies such as Google, Amazon, and LinkedIn due to their ability to incorporate relationships in their analyses as well as their speed, only in the last 5 years has their use become more widespread.

N3 is an outsourced sales company for major organizations that sell complex B2B software, hardware, and tech solutions. It supports businesses in 92 countries, provides services in 25 languages, and holds thousands of hours of conversations every month with customers and prospects. "In today's world of complex products, it takes a well-educated team to tell the story about how this technology can help a company become more competitive," said Shannon Copeland, COO of N3. "The sales team needs to be able to instantly access the information they need to do their job."

### Faster insights

The company has been operating for 16 years, and in the last few years began an initiative to manage its knowledge in a more intentional way. "We generate a great deal of data," noted Copeland, "and we wanted to make more effective use of it to understand our customers. And because of the speed at which business is transacted now, we needed to get insights right away, not a month later in a report."

N3 built a data model to reflect the essential data elements

and the associations among them and decided that a knowledge graph was the best way to represent the information. After looking into partner options, N3 chose to work with Franz, Inc., which provides a semantic graph database called AllegroGraph. "We decided to work with Franz because of its extensive experience and the fact that it had worked with a variety of industries," Copeland said.

The system built by N3 allows sales teams to organize signals from the market in a way that allows them to better explain the products to prospective buyers. "We build relationships with tech buyers on behalf of our clients," continued Copeland. "Our employees are typically college graduates who would like to begin their careers in sales and marketing in tech solutions. They take ownership of their territory and we help them be as sophisticated as a future CMO would be." The resources supplied by the knowledge graph provide the support the sales team needs to tailor information to each prospective customer.

The specific expertise required by the team varies depending on the products being sold, the geographic region, and other factors, and the knowledge graph supports these needs. For example, if a team in southern Portugal needs to know the preferences of that market, the associations built into the graph database can provide the information that is essential for them. "The information we can access helps customers understand the answers to their questions very quickly," Copeland commented. "We believe the experience that the customers have helps them scope out what they need and what the road map might be."

### The strength of graph databases

A graph databases is a type of NoSQL database that stores data according to associations among data elements rather than in the rows and columns of a relational database. Because graph databases use a dynamic schema rather than a fixed, table-

based one, adding new data types and categories is much easier. And because they are semantics-based, graph databases have strengths in inferring intent, producing answers to questions, and making recommendations. They can also make inferences about possible associations from existing associations.

A graph database also provides much more context than a relational database and therefore can return more relevant results when a user is searching; they also integrate data from multiple sources. "At one telecom company we worked with, customer service reps might have [had] to open 15 databases to find out what went wrong and what the solution was," said Jans Aasman, CEO of Franz. "We took their core customer data, billing information, every CRM call, and every action and put them into AllegroGraph, and the customer service reps were finally able to respond in a meaningful way, whether that was to make an offer to the customer or provide appropriate technical support." The capability of graph databases to overcome silos and provide an integrated view of the customer is one of its strengths.

In order to create the graph database on which the knowledge graph is built, the relationships among entities need to be mapped. In the case of a hospital patient, the patient is the core entity, and the events are medical encounters or lab results, which may come out of different databases or a data warehouse. "The mapping is a major project, but it only needs to be done once," Aasman pointed out. "After that, the relationships do not need to be regenerated during the search because they are indexed in AllegroGraph, which makes retrieval very rapid."

---

# AllegroGraph Named to 100 Companies That Matter Most in Data

## *Franz Inc. Acknowledged as a Leader for Knowledge Graph Solutions*

**Lafayette, Calif., June 23, 2020** – Franz Inc., an early innovator in Artificial Intelligence (AI) and leading supplier of Semantic Graph Database technology for Knowledge Graph Solutions, today announced that it has been named to The 100 Companies That Matter in Data by Database Trends and Applications. The annual list reflects the urgency felt among many organizations to provide a timely flow of targeted information. Among the more prominent initiatives is the use of AI and cognitive computing, as well as related capabilities such as machine learning, natural language processing, and text analytics. This list recognizes companies based on their presence, execution, vision and innovation in delivering products and services to the marketplace.

“We’re excited to announce our eighth annual list, as the industry continues to grow and evolve,” remarked Thomas Hogan, Group Publisher at Database Trends and Applications. “Now, more than ever, businesses are looking for ways transform how they operate and deliver value to customers with greater agility, efficiency and innovation. This list seeks to highlight those companies that have been successful in establishing themselves as unique resources for data professionals and stakeholders.”

“We are honored to receive this acknowledgement for our efforts in delivering Enterprise Knowledge Graph Solutions,” said Dr. Jans Aasman, CEO, Franz Inc. “In the past year, we have seen demand for Enterprise Knowledge Graphs take off

across industries along with recognition from top technology analyst firms that Knowledge Graphs provide the critical foundation for artificial intelligence applications and predictive analytics.

Our recent launch of AllegroGraph 7 with FedShard, a breakthrough that allows infinite data integration to unify all data and siloed knowledge into an Entity-Event Knowledge Graph solution will catalyze Knowledge Graph deployments across the Enterprise.”

Gartner recently released a report “How to Build Knowledge Graphs That Enable AI-Driven Enterprise Applications” and have previously stated, “The application of graph processing and graph databases will grow at 100 percent annually through 2022 to continuously accelerate data preparation and enable more complex and adaptive data science.” To that end, Gartner named graph analytics as a “Top 10 Data and Analytics Trend” to solve critical business priorities. (*Source: Gartner, Top 10 Data and Analytics Trends, November 5, 2019*).

“Graph databases and knowledge graphs are now viewed as a must-have by enterprises serious about leveraging AI and predictive analytics within their organization,” said Dr. Aasman “We are working with organizations across a broad range of industries to deploy large-scale, high-performance Entity-Event Knowledge Graphs that serve as the foundation for AI-driven applications for personalized medicine, predictive call centers, digital twins for IoT, predictive supply chain management and domain-specific Q&A applications – just to name a few.”

### **Forrester Shortlists AllegroGraph**

AllegroGraph was shortlisted in the February 3, 2020 Forrester Now Tech: Graph Data Platforms, Q1 2020 report, which recommends that organizations “Use graph data platforms to accelerate connected-data initiatives.” Forrester states, “You

can use graph data platforms to become significantly more productive, deliver accurate customer recommendations, and quickly make connections to related data.”

### **Bloor Research covers AllegroGraph with FedShard**

Bloor Research Analyst, Daniel Howard noted “With the 7.0 release of AllegroGraph, arguably the most compelling new capability is its ability to create what Franz refers to as “Entity-Event Knowledge Graphs” (or EEKGs) via its patented FedShard technology.” Mr. Howard goes on to state “Franz clearly considers this a major release for AllegroGraph. Certainly, the introduction of an explicit entity-event graph is not something I’ve seen before. The newly introduced text to speech capabilities also seem highly promising.”

### **AllegroGraph Named to KMWorld’s 100 Companies That Matter in Knowledge Management**

AllegroGraph was also recently named to KMWorld’s 100 Companies That Matter in Knowledge Management. The KMWorld 100 showcases organizations that are advancing their products and capabilities to meet changing requirements in Knowledge Management.

### **Franz Knowledge Graph Technology and Services**

Franz’s Knowledge Graph Solution includes both technology and services for building industrial strength Entity-Event Knowledge Graphs based on best-of-class tools, products, knowledge, skills and experience. At the core of the solution is Franz’s graph database technology, AllegroGraph with FedShard, which is utilized by dozens of the top F500 companies worldwide and enables businesses to extract sophisticated decision insights and predictive analytics from highly complex, distributed data that cannot be uncovered with conventional databases.

Franz delivers the expertise for designing ontology and



taxonomy-based solutions by utilizing standards-based development processes and tools. Franz also offers data integration services from siloed data using W3C industry standard semantics, which can then be continually integrated with information that comes from other data sources. In addition, the Franz data science team provides expertise in custom algorithms to maximize data analytics and uncover hidden knowledge.

---

# Bloor Research Covers AllegroGraph with FedShard

**Written By:** Daniel Howard

**Published:** 3rd June 2020



[Link to Bloor Research Post on AllegroGraph 7](#)

Graph database and “*knowledge graph solution*” vendor Franz Inc. has recently released its latest version of its graph offering, AllegroGraph, which as of the new release is in version 7.0. Alongside this, the company also released version 8.0 of Gruff, its visualisation and discovery engine for AllegroGraph and SPARQL endpoints. As of version 8, Gruff is now available via a web client as well as on the desktop.

For those of you who are unfamiliar with Franz and AllegroGraph, here is a brief description of the latter taken from our 2019 deep dive:


*Franz AllegroGraph... is a semantic graph database focused on generating sophisticated semantic knowledge graphs, initially from your existing data. The graph database itself is an RDF-based quad store (in other words, a triple store where all the triples are named) with property graph support... The product's primary focus is on transactional processing; however, it is often used for analytics as well. Consequently, it is OLTP-enabled and fully ACID compliant, and additionally offers immediate consistency. The product is also highly secure, and supports the requirements for various government security standards, including HIPAA...It includes a wealth of features, including distributed deployment and querying, multi-modal ingestion, multi-master replication, AI and machine learning, and natural language processing (NLP)."*

The full text of this report can be found [here](#).

As for the version 7.0 release of AllegroGraph, arguably the most compelling new capability is its ability to create what Franz refers to as *"Entity-Event Knowledge Graphs"* (or EEKGs) via its patented FedShard technology, an example which can be seen below. These differ from regular graphs in that they are designed to capture a number of core entities (such as products, patients or customers) and any events (which may be time-stamped) relating to those entities within a hierarchical tree structure. Notice in the image below that each node – each event – *"branches"* into one or more additional events, and that these branches usually do not interact: this is characteristic of the tree model. Moreover, these events can terminate in a knowledge base, such as a taxonomy or an ontology. This allows you to bring a wealth of supporting information into your graph when that information may be relevant to your core entity (or entities). For example,

suppose you have a core entity that is a hospital patient. In this case, you might use a publicly available catalogue of drug interactions as one of your knowledge bases. This all creates an intuitive way to visually represent your entities and the events related to them while bringing in outside knowledge where it is useful. In this context it is also worth noting that AllegroGraph has probabilistic capabilities (this is not a new feature) so that, for example, a physician can assign a probability to a diagnosis.

Notably, your EEKGs can be built incrementally, starting with a simple model and extending gradually and as needed, but without requiring you to actually alter any existing parts of your model. EEKGs also store provenance information for your events, which captures where the initial data used to generate each of your events originally came from, and how its transformation into a graph object was achieved, thus providing data lineage.



The FedShard feature in the 7.0 release provides enhanced capabilities for horizontally distributed deployment, and EEKGs especially have been designed to take advantage of this. Other new features include improved JSON and JSON-LD document handling, Natural Language Processing (NLP) and speech recognition functionality. For the latter in particular, Franz

has been able to leverage voice to text capabilities to extract conceptual meaning from real speech, then store that meaning in a graph and subsequently run analytics on it. This is exciting, because in effect it allows you to analyse recorded conversations. For organisations that collect and store a lot of such conversations – call centres, for instance – this could prove very useful.

Franz clearly considers this a major release for AllegroGraph. Certainly, the introduction of an explicit entity-event graph is not something I've seen before. The newly introduced text to speech capabilities also seem highly promising.