Using AI and Semantic Data Lakes in Healthcare – FeibusTech Research Report

Artificial intelligence has the potential to make huge improvements in just about every aspect of healthcare. Learn how Montefiore Health Systems is using semantic data lakes, architectures, and triplestores to power AI patient-centered learning. With origins in post-9/11 municipal emergency projects, Montefiore Health Systems platform – called PALM, short for patient-centered Analytical Learning Machine – is beginning to prove itself out in the Intensive Care Unit, helping doctors save lives by flagging patients headed toward respiratory failure.

Intel and Montefiore in collaboration with FeibusTech have released a Research Brief covering Montefiore's PALM Platform (aka – The Semantic Data Lake) powered by AllegroGraph.

> *"Just atop all the databases is what's known as a triplestore, or triple, construct. That's a key piece of any semantic data architecture. A triple is a three-part data series with a common grammar structure: that is, subject-predicate-object. Like, for example, John Smith has hives. Or Jill Martin takes ibuprofen."*

> *"Triples are the heart and soul of graph databases, or graphs, a powerful, labor-saving approach that associates John and Jill to records of humans, hives to definitions of maladies and Ibuprofen*

to catalogues of drugs. And then it builds databases on the fly for the task at hand based on those associations."

Read the full article on Intel's website to learn more about healthcare solutions based on AllegroGraph.

The Cornerstone of Data Science: Progressive Data Modeling

From AI Business June 27, 2018

This article covers Single Schema, Universal Taxonomies, Time Series Analysis, Accelerating Data Science and features some thought leadership from Franz Inc.'s CEO, Jans Aasman:

'Contemporary data science and artificial intelligence requirements simply can't wait for this ongoing, dilatory process. According to Jans Aasman, CEO of Franz, they no longer have to. By deploying what Aasman called an "events-based approach to schema", companies can model datasets with any number of differences alongside one another for expedited enterprise value.'

'The resulting schema is simplified, uniform, and useful in multiple ways. "You achieve two goals," Aasman noted. "One is you define what data you trust to be in the main repository to have all the truth. The second thing is you make your data management a little more uniform. By doing those two things

your AI and your data science will become better, because the data that goes into them is better.”’

Dr. Aasman goes on to note:

‘The events-based schema methodology only works with enterprise taxonomies—or at least with taxonomies spanning the different sources included in a specific repository, such as a Master Data Management hub. Taxonomies are necessary so that “the type of event can be specified,” Aasman said.’

‘Moreover, taxonomies are indispensable for clarifying terms and their meaning across different data formats, which may represent similar concepts in distinct ways. Therefore, practically all objects in a database should be “taxonomy based” Aasman said, because these hierarchical classifications enable organizations to query their repositories via this uniform schema.’

Read the full article over at AI Business.

The Most Secure Graph Database Available

Triples offer a way of describing model elements and relationships between them. In come cases, however, it is also convenient to be

able to store data that is associated with a triple as a whole rather than with a particular element. For instance one might wish to record the source from which a triple has been imported or access level necessary to include it in query results. Traditional solutions of this problem include using graphs, RDF reification or triple IDs. All of these approaches suffer from various flexibility and performance issues. For this reason AllegroGraph offers an alternative: triple attributes.
Attributes are key-value pairs associated with a triple. Keys refer to attribute definitions that must be added to the store before they are used. Values are strings. The set of legal values of an attribute can be constrained by the definition of that attribute. It is possible to associate multiple values of a given attribute with a single triple.
Possible uses for triple attributes include:

- *Access control: It is possible to instruct AllegroGraph to prevent an user from accessing triples with certain attributes.*
- *Sharding: Attributes can be used to ensure that related triples are always placed in the same shard when AllegroGraph acts as a distributed triple store.*

Like all other triple components, attribute values are immutable. They must be provided when the triple is added to the store and cannot be changed or removed later.
To illustrate the use of triple attributes we will construct an artificial data set containing a log of information about contacts detected by a submarine at a single moment in time.

Managing attribute definitions

Before we can add triples with attributes to the store we must create appropriate attribute definitions.
First let’s open a connection

```
from franz.openrdf.connect import ag_connect

conn = ag_connect('python-tutorial', create=True, clear=True)
```

Attribute definitions are represented by **`AttributeDefinition`** objects. Each definition has a name, which must be unique, and a few optional properties (that can also be

passed as constructor arguments):

- *`allowed_values`: a list of strings. If this property is set then only the values from this list can be used for the defined attribute.*
- *`ordered`: a boolean. If true then attribute value comparisons will use the ordering defined by `allowed_values`. The default is false.*
- *`minimum_number`, `maximum_number`: integers that can be used to constrain the cardinality of an attribute. By default there are no limits.*

Let's define a few attributes that we will later use to demonstrate various attribute-related capabilities of AllegroGraph. To do this, we will use the **`setAttributeDefinition()`** method of the connection object.

```
from franz.openrdf.repository.attributes import AttributeDefinition

# A simple attribute with no constraints governing the set
# of legal values or the number of values that can be
# associated with a triple.
tag = AttributeDefinition(name='tag')

# An attribute with a limited set of legal values.
# Every bit of data can come from multiple sources.
# We encode this information in triple attributes,
# since it refers to the tripe as a whole. Another
# way of achieving this would be to use triple ids
# or RDF reification.
source = AttributeDefinition(
    name='source',
    allowed_values=['sonar', 'radar', 'esm', 'visual'])

# Security level - notice that the values are ordered
# and each triple *must* have exactly one value for
# this attribute. We will use this to prevent some
# users from accessing classified data.
level = AttributeDefinition(
    name='level',
    allowed_values=['low', 'medium', 'high'],
    ordered=True,
    minimum_number=1,
    maximum_number=1)

# An attribute like this could be used for sharding.
# That would ensure that data related to a particular
# contact is never partitioned across multiple shards.
```

```
# Note that this attribute is required, since without
# it an attribute-sharded triple store would not know
# what to do with a triple.
contact = AttributeDefinition(
    name='contact',
    minimum_number=1,
    maximum_number=1)

# So far we have created definition objects, but we
# have not yet sent those definitions to the server.
# Let's do this now.
conn.setAttributeDefinition(tag)
conn.setAttributeDefinition(source)
conn.setAttributeDefinition(level)
conn.setAttributeDefinition(contact)

# This line is not strictly necessary, because our
# connection operates in autocommit mode.
# However, it is important to note that attribute
# definitions have to be committed before they can
# be used by other sessions.
conn.commit()
```

It is possible to retrieve the list of attribute definitions from a repository by using the **`getAttributeDefinitions()`** method:

```
for attr in conn.getAttributeDefinitions():
    print('Name: {0}'.format(attr.name))
    if attr.allowed_values:
        print('Allowed values: {0}'.format(
            ', '.join(attr.allowed_values)))
        print('Ordered: {0}'.format(
            'Y' if attr.ordered else 'N'))
    print('Min count: {0}'.format(attr.minimum_number))
    print('Max count: {0}'.format(attr.maximum_number))
    print()
```

Notice that in cases where the maximum cardinality has not been explicitly defined, the server replaced it with a default value. In practice this value is high enough to be interpreted as ‘no limit’.

```
Name: tag
Min count: 0
Max count: 1152921504606846975

Name: source
Allowed values: sonar, radar, esm, visual
```

```
Min count: 0
Max count: 1152921504606846975
Ordered: N

Name: level
Allowed values: low, medium, high
Ordered: Y
Min count: 1
Max count: 1

Name: contact
Min count: 1
Max count: 1
```

Attribute definitions can be removed (provided that the attribute is not used by the static attribute filter, which will be discussed later) by calling **deleteAttributeDefinition()**:

```
conn.deleteAttributeDefinition('tag')
defs = conn.getAttributeDefinitions()
print(', '.join(sorted(a.name for a in defs)))
```

```
contact, level, source
```

Adding triples with attributes

Now that the attribute definitions have been established we can demonstrate the process of adding triples with attributes. This can be achieved using various methods. A common element of all these methods is the way in which triple attributes are represented. In all cases dictionaries with attribute names as keys and strings or lists of strings as values are used.
When **addTriple()** is used it is possible to pass attributes in a keyword parameter, as shown below:

```
ex = conn.namespace('ex://')
conn.addTriple(ex.S1, ex.cls, ex.Udaloy, attributes={
    'source': 'sonar',
    'level': 'low',
    'contact': 'S1'
})
```

The **addStatement()** method works in similar way. Note that it is not possible to include attributes in the **Statement** object itself.

```
from franz.openrdf.model import Statement

s = Statement(ex.M1, ex.cls, ex.Zumwalt)
conn.addStatement(s, attributes={
    'source': ['sonar', 'esm'],
    'level': 'medium',
    'contact': 'M1'
})
```

When adding multiple triples with **addTriples()** one can add a fifth element to each tuple to represent attributes. Let us illustrate this by adding an aircraft to our dataset.

```
conn.addTriples(
    [(ex.R1, ex.cls, ex['Ka-27'], None,
      {'source': 'radar',
       'level': 'low',
       'contact': 'R1'}),
     (ex.R1, ex.altitude, 200, None,
      {'source': 'radar',
       'level': 'medium',
       'contact': 'R1'})])
```

When all or most of the added triples share the same attribute set it might be convenient to use the `attributes` keyword parameter. This provides default values, but is completely ignored for all tuples that already contain attributes (the dictionaries are not merged). In the example below we add a triple representing an aircraft carrier and a few more triples that specify its position. Notice that the first triple has a lower security level and multiple sources. The common ‘contact’ attribute could be used to ensure that all this data will remain on a single shard.

```
conn.addTriples(
    [(ex.M2, ex.cls, ex.Kuznetsov, None, {
        'source': ['sonar', 'radar', 'visual'],
        'contact': 'M2',
        'level': 'low',
     }),
     (ex.M2, ex.position, ex.pos343),
     (ex.pos343, ex.x, 430.0),
     (ex.pos343, ex.y, 240.0)],
    attributes={
       'contact': 'M2',
       'source': 'radar',
       'level': 'medium'
    })
```

Another method of adding triples with attributes is to use the NQX file format. This works both with **addFile()** and **addData()** (illustrated below):

```
from franz.openrdf.rio.rdfformat import RDFFormat

conn.addData('''
    <ex://S2> <ex://cls> <ex://Alpha> \
    {"source": "sonar", "level": "medium", "contact": "S2"} .
    <ex://S2> <ex://depth> "300" \
    {"source": "sonar", "level": "medium", "contact": "S2"} .
    <ex://S2> <ex://speed_kn> "15.0" \
    {"source": "sonar", "level": "medium", "contact": "S2"} .
''', rdf_format=RDFFormat.NQX)
```

When importing from a format that does not support attributes, it is possible to provide a common set of attribute values with a keyword parameter:

```
from franz.openrdf.rio.rdfformat import RDFFormat

conn.addData('''
    <ex://V1> <ex://cls> <ex://Walrus> ;
              <ex://altitude> 100 ;
              <ex://speed_kn> 12.0e+8 .
    <ex://V2> <ex://cls> <ex://Walrus> ;
              <ex://altitude> 200 ;
              <ex://speed_kn> 12.0e+8 .
    <ex://V3> <ex://cls> <ex://Walrus> ;
              <ex://altitude> 300;
              <ex://speed_kn> 12.0e+8 .
    <ex://V4> <ex://cls> <ex://Walrus> ;
              <ex://altitude> 400 ;
              <ex://speed_kn> 12.0e+8 .
    <ex://V5> <ex://cls> <ex://Walrus> ;
              <ex://altitude> 500 ;
              <ex://speed_kn> 12.0e+8 .
    <ex://V6> <ex://cls> <ex://Walrus> ;
              <ex://altitude> 600 ;
              <ex://speed_kn> 12.0e+8 .
''', attributes={
    'source': 'visual',
    'level': 'high',
    'contact': 'a therapist'})
```

The data above represents six visually observed Walrus-class submarines, flying at different altitudes and well above the speed

of light. It has been highly classified to conceal the fact that someone has clearly been drinking while on duty — after all there are only four Walrus-class submarines currently in service, so the observation is obviously incorrect.

Retrieving attribute values

We will now print all the data we have added to the store, including attributes, to verify that everything worked as expected. The only way to do that is through a SPARQL query using the appropriate magic property to access the attributes. The query below binds a literal containing a JSON representation of triple attributes to the *?a* variable:

```
import json

r = conn.executeTupleQuery('''
   PREFIX attr: <http://franz.com/ns/allegrograph/6.2.0/>
   SELECT ?s ?p ?o ?a {
       ?s ?p ?o .
       ?a attr:attributes (?s ?p ?o) .
   } ORDER BY ?s ?p ?o''')
with r:
    for row in r:
        print(row['s'], row['p'], row['o'])
        print(json.dumps(json.loads(row['a'].label),
                         sort_keys=True,
                         indent=4))
```

The result contains all the expected triples with pretty-printed attributes.

```
<ex://M1> <ex://cls> <ex://Zumwalt>
{
    "contact": "M1",
    "level": "medium",
    "source": [
        "esm",
        "sonar"
    ]
}
<ex://M2> <ex://cls> <ex://Kuznetsov>
{
    "contact": "M2",
    "level": "low",
    "source": [
```

```
        "visual",
        "radar",
        "sonar"
    ]
}
<ex://M2> <ex://position> <ex://pos343>
{
    "contact": "M2",
    "level": "medium",
    "source": "radar"
}
<ex://R1> <ex://altitude> "200"^^...
{
    "contact": "R1",
    "level": "medium",
    "source": "radar"
}
<ex://R1> <ex://cls> <ex://Ka-27>
{
    "contact": "R1",
    "level": "low",
    "source": "radar"
}
<ex://S1> <ex://cls> <ex://Udaloy>
{
    "contact": "S1",
    "level": "low",
    "source": "sonar"
}
<ex://S2> <ex://cls> <ex://Alpha>
{
    "contact": "S2",
    "level": "medium",
    "source": "sonar"
}
<ex://S2> <ex://depth> "300"
{
    "contact": "S2",
    "level": "medium",
    "source": "sonar"
}
<ex://S2> <ex://speed_kn> "15.0"
{
    "contact": "S2",
    "level": "medium",
    "source": "sonar"
}
<ex://V1> <ex://altitude> "100"^^...
{
    "contact": "a therapist",
```

```
    "level": "high",
    "source": "visual"
}
<ex://V1> <ex://cls> <ex://Walrus>
{
    "contact": "a therapist",
    "level": "high",
    "source": "visual"
}
<ex://V1> <ex://speed_kn> "1.2E9"^^...
{
    "contact": "a therapist",
    "level": "high",
    "source": "visual"
}
...
<ex://pos343> <ex://x> "4.3E2"^^...
{
    "contact": "M2",
    "level": "medium",
    "source": "radar"
}
<ex://pos343> <ex://y> "2.4E2"^^...
{
    "contact": "M2",
    "level": "medium",
    "source": "radar"
}
```

Attribute filters

Triple attributes can be used to provide fine-grained access control. This can be achieved by using static attribute filters. Static attribute filters are simple expressions that control which triples are visible to a query based on triple attributes. Each repository has a single, global attribute filter that can be modified using **`setAttributeFilter()`**. The values passed to this method must be either strings (the syntax is described in the documentation of static attribute filters) or filter objects.
Filter objects are created by applying set operators to 'attribute sets'. These can then be combined using filter operators.
An attribute set can be one of the following:

- *a string or a list of strings: represents a constant set of values.*

- *TripleAttribute.name: represents the value of the name attribute associated with the currently inspected triple.*
- *UserAttribute.name: represents the value of the name attribute associated with current query. User attributes will be discussed in more detail later.*

Available set operators are shown in the table below. All classes and functions mentioned here can be imported from the `franz.openrdf.repository.attributes` package:

Syntax	**Meaning**
`Empty(x)`	True if the specified attribute set is empty.
`Overlap(x, y)`	True if there is at least one matching value between the two attribute sets.
`Subset(x, y)`, `x << y`	True if every element of *x* can be found in *y*
`Superset(x, y)`, `x >> y`	True if every element of *y* can be found in *x*
`Equal(x, y)`, `x == y`	True if *x* and *y* have exactly the same contents.
`Lt(x, y)`, `x < y`	True if both sets are singletons, at least one of the sets refers to a triple or user attribute, the attribute is ordered and the value of the single element of *x* occurs before the single value of *y* in the `lowed_values` list of the attribute.
`Le(x, y)`, `x <= y`	True if *y* < *x* is false.
`Eq(x, y)`	True if both *x* < *y* and *y* < *x* are false. Note that using the == Python operator translates to*Eqauls*, not *Eq*.
`Ge(x, y)`, `x >= y`	True if *x* < *y* is false.

Syntax	Meaning
`Gt(x, y)`, `x > y`	True if *y* < *x*.

Note that the overloaded operators only work if at least one of the attribute sets is a `UserAttribute` or `TripleAttribute` reference – if both arguments are strings or lists of strings the default Python semantics for each operator are used. The prefix syntax always produces filters.

Filters can be combined using the following operators:

Syntax	Meaning
`Not(x)`, `~x`	Negates the meaning of the filter.
`And(x, y, ...)`, `x & y`	True if all subfilters are true.
`Or(x, y, ...)`, `x | y`	True if at least one subfilter is true.

Filter operators also work with raw strings, but overloaded operators will only be recognized if at least one argument is a filter object.

Using filters and user attributes

The example below displays all classes of vessels from the dataset after establishing a static attribute filter which ensures that only sonar contacts are visible:

```
from franz.openrdf.repository.attributes import *

conn.setAttributeFilter(TripleAttribute.source >> 'sonar')
conn.executeTupleQuery(
    'select ?class { ?s <ex://cls> ?class } order by ?class',
    output=True)
```

The output contains neither the visually observed Walruses nor the radar detected ASW helicopter.

```
------------------
| class          |
==================
| ex://Alpha     |
| ex://Kuznetsov |
| ex://Udaloy    |
| ex://Zumwalt   |
```

```
------------------
```

To avoid having to set a static filter before each query (which would be inefficient and cause concurrency issues) we can employ user attributes. User attributes are specific to a particular connection and are sent to the server with each query. The static attribute filter can refer to these and compare them with triple attributes. Thus we can use code presented below to create a filter which ensures that a connection only accesses data at or below the chosen clearance level.

```
conn.setUserAttributes({'level': 'low'})
conn.setAttributeFilter(
    TripleAttribute.level <= UserAttribute.level)
conn.executeTupleQuery(
    'select ?class { ?s <ex://cls> ?class } order by ?class',
    output=True)
```

We can see that the output here contains only contacts with the access level of *low*. It omits the destroyer and Alpha submarine (these require *medium* level) as well as the top-secret Walruses.

```
------------------
| class          |
==================
| ex://Ka-27     |
| ex://Kuznetsov |
| ex://Udaloy    |
------------------
```

The main advantage of the code presented above is that the filter can be set globally during the application setup and access control can then be achieved by varying user attributes on connection objects.
Let us now remove the attribute filter to prevent it from interfering with other examples. We will use the **`clearAttributeFilter()`** method.

```
conn.clearAttributeFilter()
```

It might be useful to change connection's attributes temporarily for the duration of a single code block and restore prior attributes after that. This can be achieved using the **`temporaryUserAttributes()`** method, which returns a context manager. The example below illustrates its use. It also shows how

to use **getUserAttributes()** to inspect user attributes.

```
with conn.temporaryUserAttributes({'level': 'high'}):
    print('User attributes inside the block:')
    for k, v in conn.getUserAttributes().items():
        print('{0}: {1}'.format(k, v))
    print()
print('User attributes outside the block:')
for k, v in conn.getUserAttributes().items():
    print('{0}: {1}'.format(k, v))
```

```
User attributes inside the block:
level: high

User attributes outside the block:
level: low »
```

Making Big Data More Meaningful through Data Visualization

We've all heard the saying, "a picture says a thousand words." With today's millisecond attention spans, communicating a complex topic to any audience – business professional, consumer, doctor, investor, policy-maker, voter – has become more challenging than ever. Some industries are now taking this seriously and investing in new data visualization techniques.

Data visualization is a fundamental part of scientific research. In a scientific journal, pictures certainly do seem to be worth a thousand words, with graphs translating large amounts of data into insightful, visual representations.

Franz Inc. and The Wroclaw Institute of Spatial Information and Artificial Intelligence (The Wroclaw Institute) team up to deliver graph and A.I. solutions in Poland

A Wroclaw Institute News Release

OAKLAND, Calif. – March 15, 2016 – We are pleased to inform that Wroclaw Institute has been appointed as a partner by Franz Inc.– world's leading producer of semantic graph technologies. The agreement grants to Wroclaw Institute exclusive right to sell Franz's – AllegroGraph family of products for territory of Poland. AllegroGraph is best in class graph database, fully supporting W3C standards adopted by start-up's as well as vast number of Fortune 100 companies. AllegroGraph is a part of Big Data ecosystem as it could be integrated with Apache Hadoop and Amazon EC2.

The Wroclaw Institute CEO – Dr. Adam Iwaniak said "Partnership with Franz Inc. is a turning point for our company as semantic graph technology is gaining a lot of market attention in 'data tsunami' era. We are happy that we will be able to provide our customers with award winning solution to help them manage

their complex data resources. Moreover I'd like to emphasize that as a company we made a big progress in leveraging RDF graphs technologies also on our basic market – geoinformatics".

"We are excited about the opportunity to work with Dr. Iwaniak and the Wroclaw Institute team to demonstrate why Graph Databases deliver new, real time decision making capabilities for the Enterprise." said Dr. Jans Aasman, CEO, Franz Inc., "Organizations across Poland will benefit from AllegroGraph's ability to link highly complex data, generating new knowledge and insight for a significant competitive advantage."

AllegroGraph is a database technology that enables businesses to extract sophisticated decision insights and predictive analytics from their highly complex, distributed data that can't be answered with conventional databases. Unlike traditional relational databases, Franz's product AllegroGraph employs a combination of semantic, graph and spatial technologies that process data with contextual and conceptual intelligence. AllegroGraph is able to run queries of unprecedented complexity to support predictive analytics that help companies make better, real-time decisions.

AllegroGraph is commonly used in defense and intelligence, banking, and insurance, pharmaceutical, and healthcare, Linked Data publishing, as well as by organization dealing with complex, constantly changing knowledge bases.

About Franz Inc.

Franz Inc. is a leading vendor of semantic technology tools featuring AllegroGraph – high-performance, scalable, disk-based graph database, provides the solid storage layer for powerful GeoTemporal Reasoning, Social Network Analytics and Ontology Modeling. Based in Oakland, California, Franz Inc. is an American owned company that delivers leading-edge development products that enable software developers to build

flexible, scalable, semantic applications quickly and cost-effectively.

About The Wroclaw Institute

The Wroclaw Institute of Spatial Information and Artificial Intelligence is Wroclaw, Poland based technology company focused on knowledge engineering, data exploration and intelligent GIS providing products, services and solutions based on cutting-edge scientific and technological achievements.

Related Links

- WIZIPISI dystrybutorem oprogramowania AllegroGraph
- Oprogramowanie bazodanowe AllegroGraph dostepne w Polsce
- Wroclaw Institute of Spatial Information and Artificial Intelligence

---

Enriching Property Graphs with Relationship

Suppose we are creating a large graph database that contains information about payments between companies. A graph database analyst might start off modeling the payments as shown in Figure 1, which expresses who paid whom. (All graph figures in this article were produced using Gruff, a tool for visualizing graph databases, operating over the AllegroGraph graph database system.)

Figure 1: A Graph of a Payment

This seems straightforward enough. Now suppose that we want to record more information about payments, such as the amount of the payment, the means of payment (direct debit, e-check, wire, etc.), the date and time when the payment occurred, and so forth. A traditional property graph approach places these properties on the edge that connects the payor and payee nodes, as shown in Figure 2.

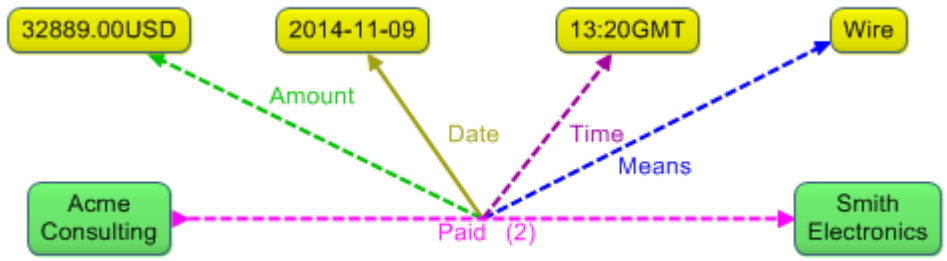

Figure 2: Attaching Properties to an Edge

Read the full blog post at DB-Engines

AllegroGraph Certification on Cloudera Enterprise Creates a Semantic Graph Data Platform for Hadoop

Leading Accountable Care Organization Gains Sophisticated Decision Insights from Complex, Distributed Big Data Using

AllegroGraph and Cloudera Enterprise

OAKLAND, Calif. – February 8, 2016 – Franz Inc., an early innovator in Artificial Intelligence (AI) and leading supplier of Semantic Graph Database technology, today announced the availability of AllegroGraph 6, the leading Semantic Graph Database with certification on the latest release of Cloudera Enterprise through the Cloudera Certified Technology Program (CCPT). AllegroGraph is the first Semantic Graph Database to be certified on Cloudera Enterprise, the global provider of the fastest, easiest, and most secure data management and analytics platform built on Apache Hadoop and the latest open source technologies.

The combination of AllegroGraph and Cloudera Enterprise eases the integration of highly complex Big Data, including large public datasets, and enables real-time analytics across distributed data, while leveraging the world's highest performance and most cost effective storage. Enterprises can run queries of unprecedented complexity to enable predictive analytics and real time decision-making within a myriad of industries including Healthcare, Life Sciences, Financial Services, Intelligence/National Security and Publishing. The hardened platform can run mission-critical applications that require uncompromised data resiliency using features like ACID compliance to ensure data is never lost.

"As the availability of large public datasets continues to rise, many organizations are looking to leverage these datasets to enrich enterprise analytics," said Tim Stevens, vice president of Business and Corporate Development at Cloudera. "A semantic graph approach to data sets provides a viable method to gleaning additional insights from data."

"Today's Big Data challenge is also a Cognitive Computing challenge," said Dr. Jans Aasman, CEO of Franz Inc. "We need to combine unstructured data with structured data to fuel real-time analysis, predictive analytics and deep learning.

But the ease of data integration largely depends on the type of database. With the Semantic flexibility of AllegroGraph, integrating databases is a virtually effortless, since the data can remain in its original databases and database designers do not have to create a schema up front. This capability is particularly valuable if organizations want to tap into the growing number of public datasets to enrich their analytics."

The powerful combination of AllegroGraph and Cloudera plays a critical role in the Semantic Data Lake for Healthcare, a collaboration between Montefiore Health System (The leading Accountable Care Organization in the U.S), Franz, Cloudera, Cisco and Intel to provide a scalable and extensible Big Data Analytic platform for Healthcare. The SDL for Healthcare is a shared vision between Montefiore Health System and Franz, for constructing longitudinally integrated, semantically enriched, scalable and secured analytics infrastructure necessary for next generation learning healthcare systems, and precision medicine. The SDL deploys Montefiore's innovative informatics solutions on Franz's AllegroGraph and Cloudera's Hadoop distribution, and enables modelers, data scientists and application developers to leverage complex information, biomedical knowledge-bases and ontologies, as well as the linked open data (LOD) for predictive modeling, care management, population and community health management, health systems research, and clinical and translational research.

"The Semantic Data Lake for Healthcare will help us to connect the dots to better understand the determinants of outcome, cost, and patient satisfaction in a complex ecosystem in which patients and clinicians interact with each other, with the delivery of care system, and with the research enterprise," said Dr. Parsa Mirhaji MD. PhD., Director of Clinical Research Informatics at Einstein College of Medicine and Montefiore Health System. "The problem is, there are billions of such

dots that needs to be connected meaningfully, and reproducibly. The complexity is that there are many different principles, pathways, and theoretical frameworks on how those dots would connect, and we just don't know which model or framework would yield the best answer. The SDL will enable us to address both the complexity, and scalability concerns efficiently, while maintaining a tight grasp on the semantic integrity and consistency of analysis over large, multi source, ever changing datasets."

Dr. Mirhaji added, "The SDL embodies Montefiore's incremental and measured approach towards Cognitive Computing in healthcare. Our ability to conduct real-time analysis over new combinations of data, to compare results across multiple analyses, and to engage patients, practitioners and researchers as equal partners in big-data analytics and decision support will fuel discoveries, significantly improve efficiencies, personalize care and ultimately save lives."

"Information has always existed everywhere but has often been isolated, incomplete, unavailable or unintelligible," according to Gartner. "Advances in semantic tools such as graph databases as well as other emerging data classification and information analysis techniques will bring meaning to the often chaotic deluge of information." (Source: Gartner Identifies the Top Strategic Technology Trends for 2016.)

AllegroGraph has been widely recognized and endorsed within the industry as the popularity of Graph databases has skyrocketed – growing nearly 500% in the past two years. In 2015 AllegroGraph was named a Leading Database Solution by CIOReview and awarded Best in Semantic Web Technology & Leader in Graph Database Products by Corporate America. This year, PharmaTech Outlook has named Franz a Top Ten Solution Provider.

"Franz has been leading the burgeoning Graph Database revolution with a highly sophisticated, yet elegant Semantic

Graph database solution," said Harvi Sachar, Publisher & Founder, CIOReview. "Franz's AllegroGraph continues to break new ground in predictive analytics and visual graph discovery capabilities- benefiting organizations around the globe within Healthcare, Intelligence/National Security, Life Sciences and Financial Services."

About AllegroGraph

Unlike traditional relational databases or Property Graph Databases, AllegroGraph employs semantic graph technologies that process data with contextual and conceptual intelligence. AllegroGraph is able run queries of unprecedented complexity to support predictive analytics that help organizations make more informed, real-time decisions. AllegroGraph is the first Graph Database to support analysis across N-dimensions – any conceivable measurement of an object, property or operation. AllegroGraph can analyze temporal (time) and geospatial (location) dimensions relative to any 'event,' such as a disease, drug interaction, genetic combination, biomarkers, observations, image or physical sensors. AllegroGraph is utilized by dozens of the top Fortune 500 companies worldwide.

About Franz Inc.

Franz Inc. is an early innovator in Artificial Intelligence (AI) and leading supplier of Semantic Graph Database technology with expert knowledge in developing and deploying complex Big Data analytics solutions. AllegroGraph, Franz's flagship, high-performance, transactional, and scalable Semantic Graph Database, provides the solid storage layer for Enterprise grade NoSQL solutions. AllegroGraph's Activity Recognition capabilities provides a powerful means to aggregate and analyze data about individual and organizational behaviors, preferences, relationships, plus spatial and temporal linkages between individuals and groups. For additional Franz Inc customer success stories please visit:

- AllegroGraph – http://franz.com/agraph/success/
- Allegro CL – http://franz.com/success/

Franz's Professional Service team is in the business of helping companies turn Data into Information and Information into Knowledge. We combine Data, Business Intelligence, and Analytics consulting services under one roof for our customers. Franz, an American owned company based in Oakland, California, is committed to market-driven product development, the highest levels of product quality and responsive customer support and service. Franz customers include dozens of Fortune 500 companies and span the healthcare, government, life sciences and telecommunications industries worldwide. Franz has demonstrated consistent growth and profitability since inception.

AllegroGraph Recognized as Best in Semantic Web Technology – USA & Leader in Graph Database Products

Franz's AllegroGraph Fueling Rapid Growth in Graph Database Category

OAKLAND, Calif. – February 3, 2016 – Franz Inc., an early innovator in Artificial Intelligence (AI) and leading supplier of Semantic Graph Database technology has been recognized As "Best in Semantic Web Technology – USA & Leader in Graph

Database Products” by Corporate America Software and Technology.

“At Corporate America, it’s our priority to showcase prominent professionals who are excelling in their industry and outperforming their competitors,” said Hannah Stevenson, Managing Group Editor, AI Global Media. “Franz Inc. have a reputation for innovation, utilizing their expert knowledge to create complex and exciting Graph Database solutions. Franz’s unique platforms offer highly scalable technologies for solving complex Big Data challenges.”

Corporate America is the definitive magazine for CEOs, top tier management and key decision makers across the US. Created to inform, influence, and shape the corporate conversation across the nation through high quality editorial, in-depth research and an experienced and dedicated network of advisers, Corporate America provides its readership with the most authoritative and current analysis of the major changes effecting the corporate landscape, and the latest deals and topical issues dominating the corporate universe. A multifaceted program, the awards are focused on rewarding excellence across all areas of the technology and software industries and all nominees are closely scrutinized to ensure that only the most deserving receive Corporate America’s prestigious awards.

“We are excited that Graph Databases, like AllegroGraph, have garnered the attention they deserve by Enterprise customers looking to innovate,” said Dr. Jans Aasman, CEO, Franz Inc. “In today’s data-driven environments, the ability to quickly analyze data from diverse sources is becoming critical. We are already seeing how Semantic Graph Databases with predictive analytics can help transform healthcare through Precision Medicine and make us safer through Insider Threat Detection.”

“Because it (AllegroGraph) is a Graph database, it can store pretty much any kind of data and query it, not just in the

time-worn relational fashion, but also in a graphical manner – carving out graphical maps of relationships. And on top of that, it can apply semantics to deduce as-yet-undiscovered knowledge from the data. Its capabilities are very broad, and they provide a glimpse of the shape of things to come," added Bloor. stated Robin Bloor, co-founder and Chief Analyst of The Bloor Group.

"Information has always existed everywhere but has often been isolated, incomplete, unavailable or unintelligible," according to Gartner. "Advances in semantic tools such as graph databases as well as other emerging data classification and information analysis techniques will bring meaning to the often chaotic deluge of information." (Source: Gartner Identifies the Top Strategic Technology Trends for 2016.)

A recent Forrester Research report stated, "Graph databases are a powerful optimized technology that link billions of pieces of connected data to help create new sources of value for customers and increase operational agility for customer service. Because graph databases track connections among entities and offer links to get more detailed information, they are well-suited for scenarios in which relationships are important, such as cybersecurity, social network analysis, eCommerce recommendations, dependence analysis, and predictive analytics." (Source: Forrester Research, Market Overview: Graph Databases, May 28, 2015)

Franz's recent announcement of the first Semantic Data Lake (SDL) for Healthcare, which was created in collaboration with Montefiore Medical Center (the eighth largest hospital group in the U.S.), Intel, Cloudera and Cisco. The SDL for Healthcare is a scalable and extensible Healthcare platform designed for Accountable Care and Personalized Medicine initiatives. AllegroGraph has played a critical role in the Semantic Data Lake for Healthcare, by facilitating integration of complex information for basic science, clinical, population, community, environmental, behavioral and

wellness research data to enable knowledge-based analytics, classification, pattern recognition, predictive modeling and simulations at scale.

About Corporate America

Corporate America is more than just a magazine. Alongside our quarterly publication, we also produce a website that is regularly updated with the latest news, features, opinion and comment, again in conjunction with a host of top-level advisers, experts and businesspeople, and throughout the year, you'll also get your chance to participate in our highly regarded awards programs, designed to pay tribute to the finest firms and individuals on the American business landscape.

About AllegroGraph

Unlike traditional relational databases or Property Graph Databases, AllegroGraph employs semantic graph technologies that process data with contextual and conceptual intelligence. AllegroGraph is able run queries of unprecedented complexity to support predictive analytics that help organizations make more informed, real-time decisions. AllegroGraph is the first Graph Database to support analysis across N-dimensions – any conceivable measurement of an object, property or operation. AllegroGraph can analyze temporal (time) and geospatial (location) dimensions relative to any 'event,' such as a disease, drug interaction, genetic combination, biomarkers, observations, image or physical sensors. AllegroGraph is utilized by dozens of the top Fortune 500 companies worldwide.

About Franz Inc.

Franz Inc. is an early innovator in Artificial Intelligence (AI) and leading supplier of Semantic Graph Database technology with expert knowledge in developing and deploying complex Big Data analytics solutions. AllegroGraph, Franz's

flagship, high-performance, transactional, and scalable Semantic Graph Database, provides the solid storage layer for Enterprise grade NoSQL solutions. AllegroGraph's Activity Recognition capabilities provides a powerful means to aggregate and analyze data about individual and organizational behaviors, preferences, relationships, plus spatial and temporal linkages between individuals and groups. For additional Franz Inc customer success stories please visit:

- AllegroGraph – http://franz.com/agraph/success/
- Allegro CL – http://franz.com/success/

Franz's Professional Service team is in the business of helping companies turn Data into Information and Information into Knowledge. We combine Data, Business Intelligence, and Analytics consulting services under one roof for our customers. Franz, an American owned company based in Oakland, California, is committed to market-driven product development, the highest levels of product quality and responsive customer support and service. Franz customers include dozens of Fortune 500 companies and span the healthcare, government, life sciences and telecommunications industries worldwide. Franz has demonstrated consistent growth and profitability since inception.

Franz's Vice President of Corporate Development to

Keynote at the 11th International Knowledge Management in Organizations Conference (KMO 2016)

OAKLAND, Calif. – February 2, 2016 – Franz Inc.'s Vice President of Corporate Development, Dr. Sheng-Chuan Wu, will be a keynote speaker at the 11th International Knowledge Management in Organizations Conference (KMO 2016) this July in Hagen, Germany. Knowledge Management is in the midst of a revolution. Traditional KM approaches have failed to meet the challenges posed by Big Data, mobility, social media, and customer demands. That's because the knowledge critical for customer service is everywhere, and new approaches are required to tap into its value.Continuing the success of the KMO conference series since 2005, the KMO 2016 conference will provide an international communication forum bringing together academia and industry for discussing the progress made and addressing the challenges faced by knowledge management.

Dr. Wu's talk, "Why Knowledge when Data Suffices"

According to the classic knowledge pyramid, we turn the data we collect into information by applying its context. We then interpret the information to derive knowledge from it. Our efforts on the management of knowledge stem from our belief that knowledge is what provides value to our endeavors. Is this paradigm still true with the explosive growth in Big Data? One of the most obvious examples is Google Translate. Despite employing machine learning on the massive multilingual text data instead of natural language understanding algorithms, Google Translate outperforms traditional natural language processing (NLP) methods when it comes to translation. Medical science is another potential example.

Since the sequencing of the human genome in 1996, we have dreamed about treating patients more effectively based on their genomic profile. Such a dream remains elusive due to the complexity of system biology. On the other hand, major progress can be made in “targeted medicine” with machine learning on the massive patient medical data accumulated. In essence, we can uncover ways to directly help patients from the data without precisely knowing how it works exactly. Using Big Data to derive value brings another set of management problems, namely the heterogeneous nature of data sources and taxonomies, the massive volume of data, and the analytic processing requirements. Dr. Wu will discuss all these issues and show some examples at this talk.

About Dr. Wu

Dr. Sheng-Chuan Wu received his Ph.D. in Scientific Computing and Computer Graphics from Cornell University in the US. He has, since graduation, involved in several software companies, including the founding of the first integrated CAD/CAM/CAE company. In the last 20 years, he worked as a senior corporate executive at the leading Artificial Intelligence and Semantic Technology company, Franz Inc in Silicon Valley, with responsibility in application development, marketing, consulting and new business development. Dr. Wu has also in many occasions collaborated with Bioinformatics experts from Harvard Medical School, Stanford University and Astra Zeneca, working with massive biological data.

Dr. Wu has been focusing on Semantic Technology over the last 8 years. He routinely lectured on AI and Semantic Technology at conferences. He has, since 2007, conducted more than 20 week-long workshops on Semantic Technology and Artificial Intelligence in Malaysia, China, Singapore, India and other Asian countries. Dr. Wu has also consulted on several Big Data and Semantic Technology projects in the US and Asia. Some of the projects include: Biodiversity Repository, Precision Agriculture for Citrus Plantation, Telecom Customer Relation

Management, Malaysia R&D Knowledgebase, Intelligence analytics, Meta Data Management, Smart City and E-Learning System.

About Franz Inc.

Franz Inc. is an early innovator in Artificial Intelligence (AI) and leading supplier of Semantic Graph Database technology with expert knowledge in developing and deploying complex Big Data analytics solutions. AllegroGraph, Franz's flagship, high-performance, transactional, and scalable Semantic Graph Database, provides the solid storage layer for Enterprise grade NoSQL solutions. AllegroGraph's Activity Recognition capabilities provides a powerful means to aggregate and analyze data about individual and organizational behaviors, preferences, relationships, plus spatial and temporal linkages between individuals and groups. For additional Franz Inc customer success stories please visit:

- AllegroGraph – http://franz.com/agraph/success/
- Allegro CL – http://franz.com/success/

Franz's Professional Service team is in the business of helping companies turn Data into Information and Information into Knowledge. We combine Data, Business Intelligence, and Analytics consulting services under one roof for our customers. Franz, an American owned company based in Oakland, California, is committed to market-driven product development, the highest levels of product quality and responsive customer support and service. Franz customers include dozens of Fortune 500 companies and span the healthcare, government, life sciences and telecommunications industries worldwide. Franz has demonstrated consistent growth and profitability since inception.

AllegroGraph Recognized Among Top 10 Analytics Solution Providers by Pharma Tech Outlook

Franz's AllegroGraph powers Pharma Analytics for Sophisticated Decision Insights from Complex, Distributed Big Data

OAKLAND, Calif. – January 28, 2016 – Franz Inc., an early innovator in Artificial Intelligence (AI) and leading supplier of Semantic Graph Database technology has been named to Pharma Tech Outlook's Top 10 Analytics Solutions Providers for 2016.

"Franz Inc. has been selected as a Top 10 Analytics Solution Provider after careful evaluation across a dozen quantitative and qualitative elements," said Stacey Smith, Editor of Pharma Tech Outlook. "Our selection process takes into consideration a company's experience, industry recognition, technical certifications, market presence and positive client reviews. Franz Inc. and their Semantic Graph Database, AllegroGraph, are clear market leaders for Analytics in the Pharmaceutical Industry."

Pharma Tech Outlook covers the latest developments in the pharmaceutical industry. They provide valuable updates – news, views and trends, expert opinions, studies, discoveries, R&D and clinical trials – essential for decision-makers in the industry. Covering all the novel outcomes, Pharma Tech Outlook aims at contributing to the transformation of innovations into services as well as creating a healthy and productive society.

Pharma Tech Outlook's "Top 10 Analytics Solution Providers"

are selected annually by a panel of experts and members of Pharma Tech Outlook's editorial board to recognize and promote technology entrepreneurship.

"Using AllegroGraph, Enterprises can run queries of unprecedented complexity to enable predictive analytics and real time decision-making within a myriad of industries including Healthcare, Life Sciences, Financial Services, and Publishing," said Jans Aasman, CEO of Franz Inc. "Integrating databases is a virtually effortless which is particularly valuable if organizations want to tap into the growing number of public datasets to enrich their analytics."

"Information has always existed everywhere but has often been isolated, incomplete, unavailable or unintelligible," according to Gartner. "Advances in semantic tools such as graph databases as well as other emerging data classification and information analysis techniques will bring meaning to the often chaotic deluge of information." (Source: Gartner Identifies the Top Strategic Technology Trends for 2016.)

A recent Forrester Research report stated, "Graph databases are a powerful optimized technology that link billions of pieces of connected data to help create new sources of value for customers and increase operational agility for customer service. Because graph databases track connections among entities and offer links to get more detailed information, they are well-suited for scenarios in which relationships are important, such as cybersecurity, social network analysis, eCommerce recommendations, dependence analysis, and predictive analytics." (Source: Forrester Research, Market Overview: Graph Databases, May 28, 2015)

Franz's recent announcement of the first Semantic Data Lake (SDL) for Healthcare, which was created in collaboration with Montefiore Medical Center (the eighth largest hospital group in the U.S.), Intel, Cloudera and Cisco. The SDL for Healthcare is a scalable and extensible Healthcare platform

designed for Accountable Care and Personalized Medicine initiatives. AllegroGraph has played a critical role in the Semantic Data Lake for Healthcare, by facilitating integration of complex information for basic science, clinical, population, community, environmental, behavioral and wellness research data to enable knowledge-based analytics, classification, pattern recognition, predictive modeling and simulations at scale.

About Pharma Tech Outlook

Pharma Tech Outlook is an online and a monthly magazine which covers most important and latest developments in the pharmaceutical industry. Through nominations and consultations with industry leaders, its editors choose the best in Pharma domains. Pharma Tech Outlook's December-January Edition is an annual listing of Top 10 Analytics Solution Providers. For more information, visit the website at: http://www.pharmatechoutlook.com/

About AllegroGraph

Unlike traditional relational databases or Property Graph Databases, AllegroGraph employs semantic graph technologies that process data with contextual and conceptual intelligence. AllegroGraph is able run queries of unprecedented complexity to support predictive analytics that help organizations make more informed, real-time decisions. AllegroGraph is the first Graph Database to support analysis across N-dimensions – any conceivable measurement of an object, property or operation. AllegroGraph can analyze temporal (time) and geospatial (location) dimensions relative to any 'event,' such as a disease, drug interaction, genetic combination, biomarkers, observations, image or physical sensors. AllegroGraph is utilized by dozens of the top Fortune 500 companies worldwide.

About Franz Inc.

Franz Inc. is an early innovator in Artificial Intelligence (AI) and leading supplier of Semantic Graph Database technology with expert knowledge in developing and deploying complex Big Data analytics solutions. AllegroGraph, Franz's flagship, high-performance, transactional, and scalable Semantic Graph Database, provides the solid storage layer for Enterprise grade NoSQL solutions. AllegroGraph's Activity Recognition capabilities provides a powerful means to aggregate and analyze data about individual and organizational behaviors, preferences, relationships, plus spatial and temporal linkages between individuals and groups. For additional Franz Inc customer success stories please visit:

- AllegroGraph – http://franz.com/agraph/success/
- Allegro CL – http://franz.com/success/

Franz's Professional Service team is in the business of helping companies turn Data into Information and Information into Knowledge. We combine Data, Business Intelligence, and Analytics consulting services under one roof for our customers. Franz, an American owned company based in Oakland, California, is committed to market-driven product development, the highest levels of product quality and responsive customer support and service. Franz customers include dozens of Fortune 500 companies and span the healthcare, government, life sciences and telecommunications industries worldwide. Franz has demonstrated consistent growth and profitability since inception.