



# Combining Semantic, Relational and Document NoSQL worlds, and doing something useful with it too

Emil Zegers, Systems Engineer, MarkLogic Corporation

# We Are The New Generation Database

## Hierarchical Era

*For your application data!*

- Application- and hardware-specific



## Relational Era

*"For all your structured data!"*

- Normalized, tabular model
- Application-independent query
- User control



## Any Structure Era

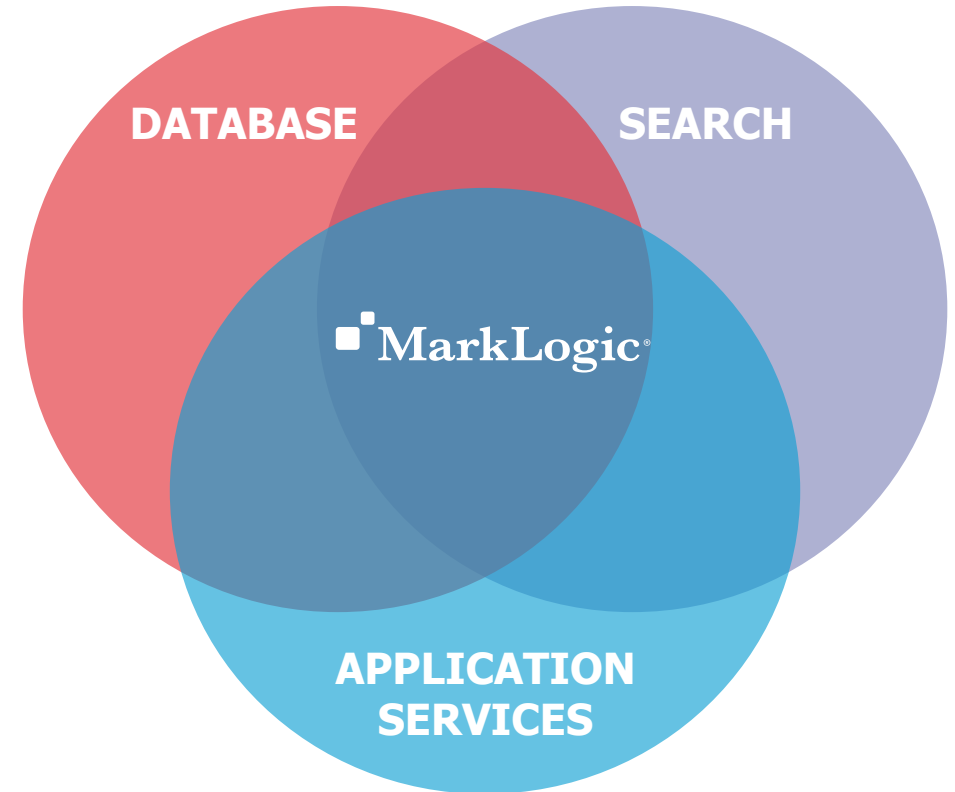
*"For all your data!"*

- Schema-agnostic
- Massive scale
- Query *and* search
- Analytics
- Application services
- Faster time-to-results



# The Only Enterprise NoSQL Database

- Search & Query
- ACID Transactions
- High Availability / Disaster Recovery
- Replication
- Government-grade Security
- Elasticity
- Cloud Deployment
- Hadoop for Storage & Compute
- Semantics



# How Does it Work?



Schema-agnostic design



Real-time indexing for search and query



*Reverse queries for event processing and alerting*



*Scale-out shared-nothing cluster topology*



*Analytics and Visualization*

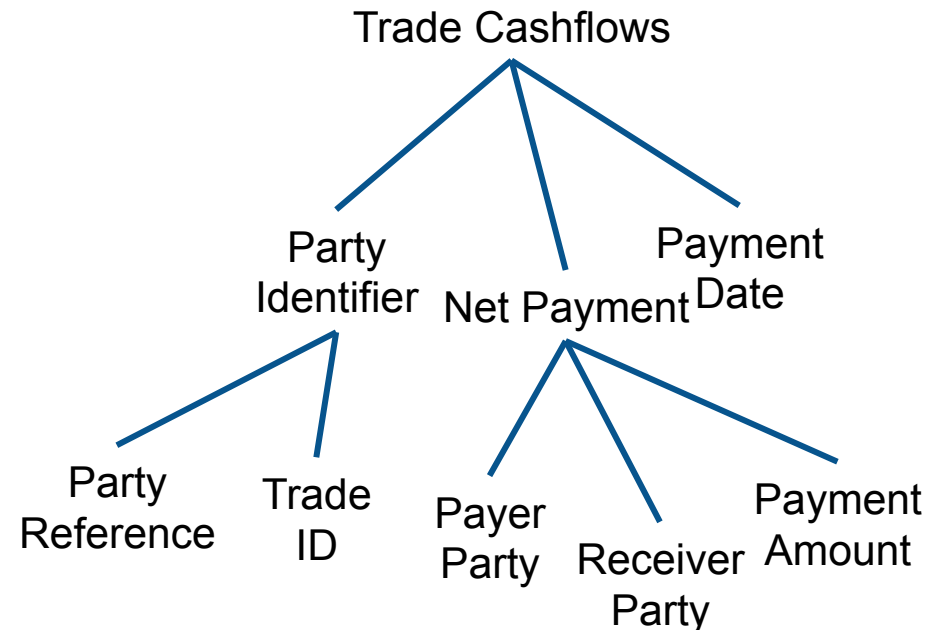
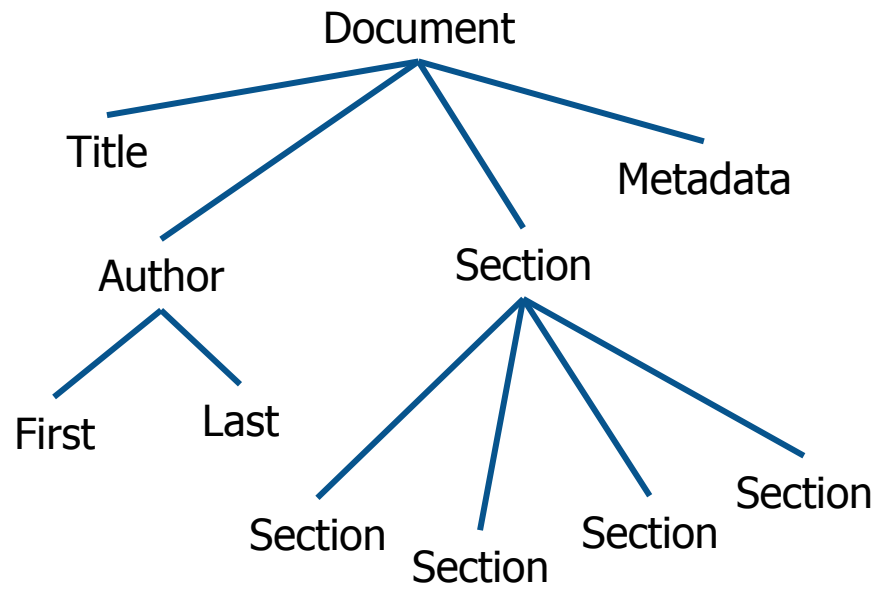


*Transactions*



# Hierarchical Data Model

- MarkLogic Server is a document-centric database
  - Supports any-structured data via hierarchical data model
  - Stores compressed binary trees

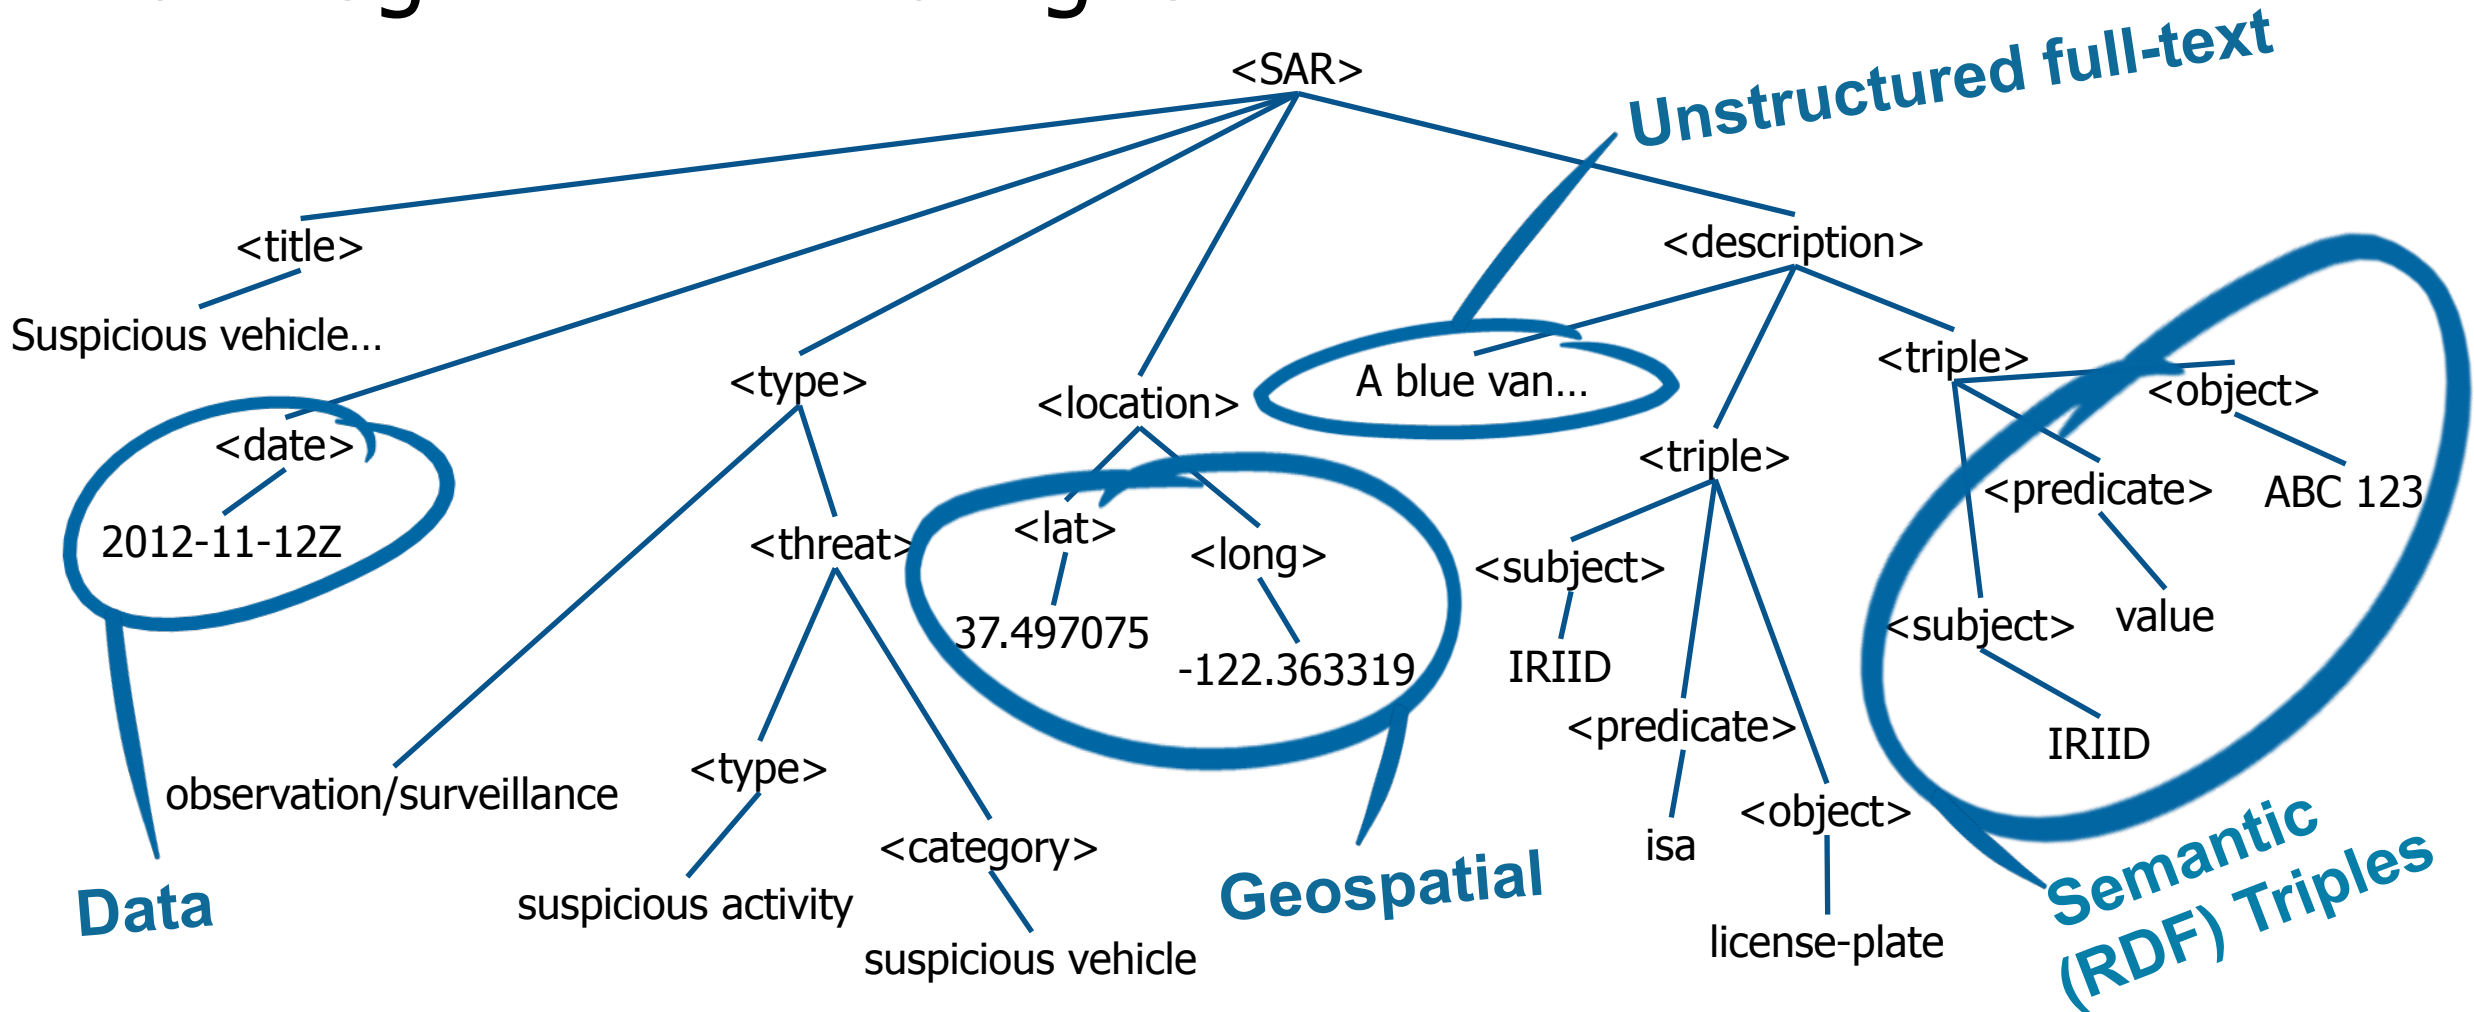


# MarkLogic is Schema-Agnostic

JSON and XML are self-describing:

```
<SAR>
  <title>Suspicious vehicle near airport</title>
  <date> 2012-11-12Z</date>
  <type> observation/surveillance</type>
  <threat>
    <type> suspicious activity</type>
    <category> suspicious vehicle</category>
  </threat>
  <location>
    <lat>37.497075</lat>
    <long> -122.363319</long>
  </location>
  <description>A blue van with license plate ABC 123 was observed parked behind the airport sign...
    <triple>IRI</subject>is</predicate>license-plate</object></triple>
    <triple>IRI</subject>value</predicate>ABC 123</object></triple>
  </description>
</SAR>
```


# MarkLogic is Schema-Agnostic



# Universal Index

*Which vetted SAR reports contain the phrase **blue van**?*

Term	Term List
"blue"	123, 125, 129, 152, 344, 491, ...
"van"	123, 125, 126, 129, 130, 152, ...
"observed"	125, 152, 516, 522, 765, 890, ...
"blue van"	123, 125, 126, 129, 130, 152, ...
STEM "observe"	125, 152, 516, 522, 765, 890, ...
<SAR>	...
<SAR>/<location>	...
<threat>/<category>	...
<type>suspicious activity</type>	...
<date>2012-11-12Z</date>	...
Collection:Vetted	...
Role:Analyst + Action:Read	...
...	...
...	...
...	...




## MarkLogic indexes...

- Words
- Phrases
- Stemmed words and phrases
- Structure
- Data
- Collections
- Security Permissions

# Range Index

*Which vetted SAR reports containing the phrase **blue van** were submitted before 2013?*

Term	Term List
"blue"	123, 125, 129, 152, 344, 491, ...
"van"	123, 125, 126, 129, 130, 152, ...
"observed"	125, 152, 516, 522, 765, 890, ...
"blue van"	123, 125, 126, 129, 130, 152, ...
STEM "observe"	125, 152, 516, 522, 765, 890, ...
<SAR>	...
<SAR>/<location>	...
<threat>/<category>	...
<type>suspicious activity</type>	...
<date>2012-11-12Z</date>	...
Collection:vetted	...
Role:Analyst + Action:Read	...
...	...
...	...
...	...



DocID	<date>
1	<date>
3	1999-12-31Z
4	2003-02-16Z
5	2007-06-01Z
8	2008-01-01Z
10	2009-04-28Z
11	2011-05-12Z
17	2012-11-12Z
...	2013-09-30Z
...	...
...	...
...	...




Range indexes map document IDs to values, and vice-versa in a compact in-memory representation.

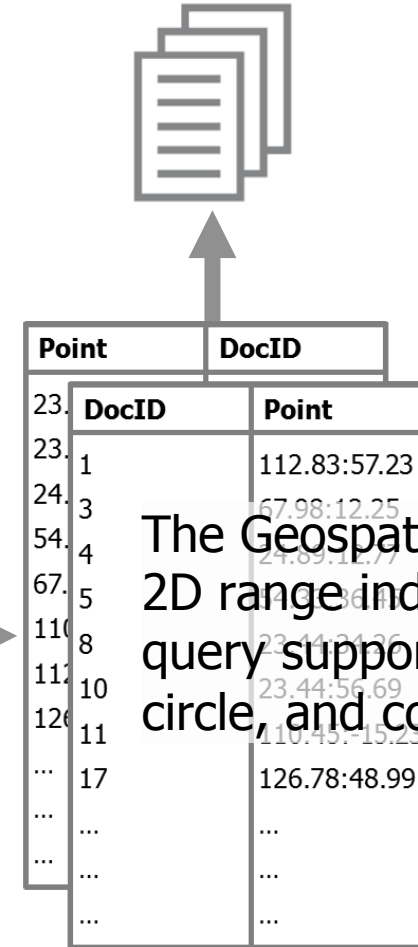
# Geospatial Index

*Which vetted SAR reports about a blue van from before 2013 refer to a location near the airport?*

Term	Term List
"blue"	123, 125, 129, 152, 344, 491, ...
"van"	123, 125, 126, 129, 130, 152, ...
"observed"	125, 152, 516, 522, 765, 890, ...
"blue van"	123, 125, 126, 129, 130, 152, ...
STEM "observe"	125, 152, 516, 522, 765, 890, ...
<SAR>	...
<SAR>/<location>	...
<threat>/<category>	...
<type>suspicious activity</type>	...
<date>2012-11-12Z</date>	...
Collection:Vetted	...
Role:Analyst + Action:Read	...
...	...
...	...
...	...



DocID	<date>
1	<date>
3	1999-12-31Z
4	2003-02-16Z
5	2007-06-01Z
8	2008-01-01Z
10	2009-04-28Z
11	2011-05-12Z
17	2012-11-12Z
...	2013-09-30Z
...	...
...	...
...	...



The Geospatial index is like a 2D range index, with built-in query support for point, box, circle, and complex polygons.

# Triple Index

*Which vetted SARs about a blue van from before 2013 with this location refer to partial plate **ABC**?*

Term List
123, 125, 129, 152, 344, 491, ...
123, 125, 126, 129, 130, 152, ...
125, 152, 516, 522, 765, 890, ...
123, 125, 126, 129, 130, 152, ...
125, 152, 516, 522, 765, 890, ...
...
...
...
...
...
...
...
...
...
...

Document References
125, 516, 890, ...

DocID	<date>	DocID
1	<date>	DocID
3	1999-12-31Z	8
4	2003-02-16Z	10
5	2007-06-01Z	4
8	2008-01-01Z	5
10	2009-04-28Z	3
11	2011-05-12Z	11
17	2012-11-12Z	1
...	2013-09-30Z	17
...	...	...
...	...	...
...	...	...
...	...	...
...	...	...
...	...	...

*All indexes fully composable for sub-second response to very complex queries*

Point	DocID
23	DocID
23	1 112.83:57.23
24	3 67.98:12.25
54	4 24.89:12.77
67	5 54.33:36.45
110	8 23.44:34.26
112	10 23.44:56.69
120	11 110.45:-15.23
...	17 126.78:48.99
...	...
...	...
...	...
...	...
...	...

DocID	Subject	Predicate	Object
1	Object	Subject	Predicate
3	Object	Subject	DocID
4	Predicate	Subject	Object
5	Subject	Object	Predicate
8	Entity1	Entity1	Entity1
10	Entity1	Entity1	Entity1
11	Entity1	Entity1	Entity1
17	Entity2	Entity2	Entity2
...	Entity2	Entity2	Entity2
...	Entity2	Entity2	Entity2
...	Entity3	Entity3	Entity3
...	Entity3	Entity3	Entity3
...	Entity3	Entity3	Entity3
...	Entity3	Entity3	Entity3
...	Entity3	Entity3	Entity3
...	Entity3	Entity3	Entity3
...	Entity3	Entity3	Entity3
...	Entity3	Entity3	Entity3
...	Entity3	Entity3	Entity3
...	Entity3	Entity3	Entity3

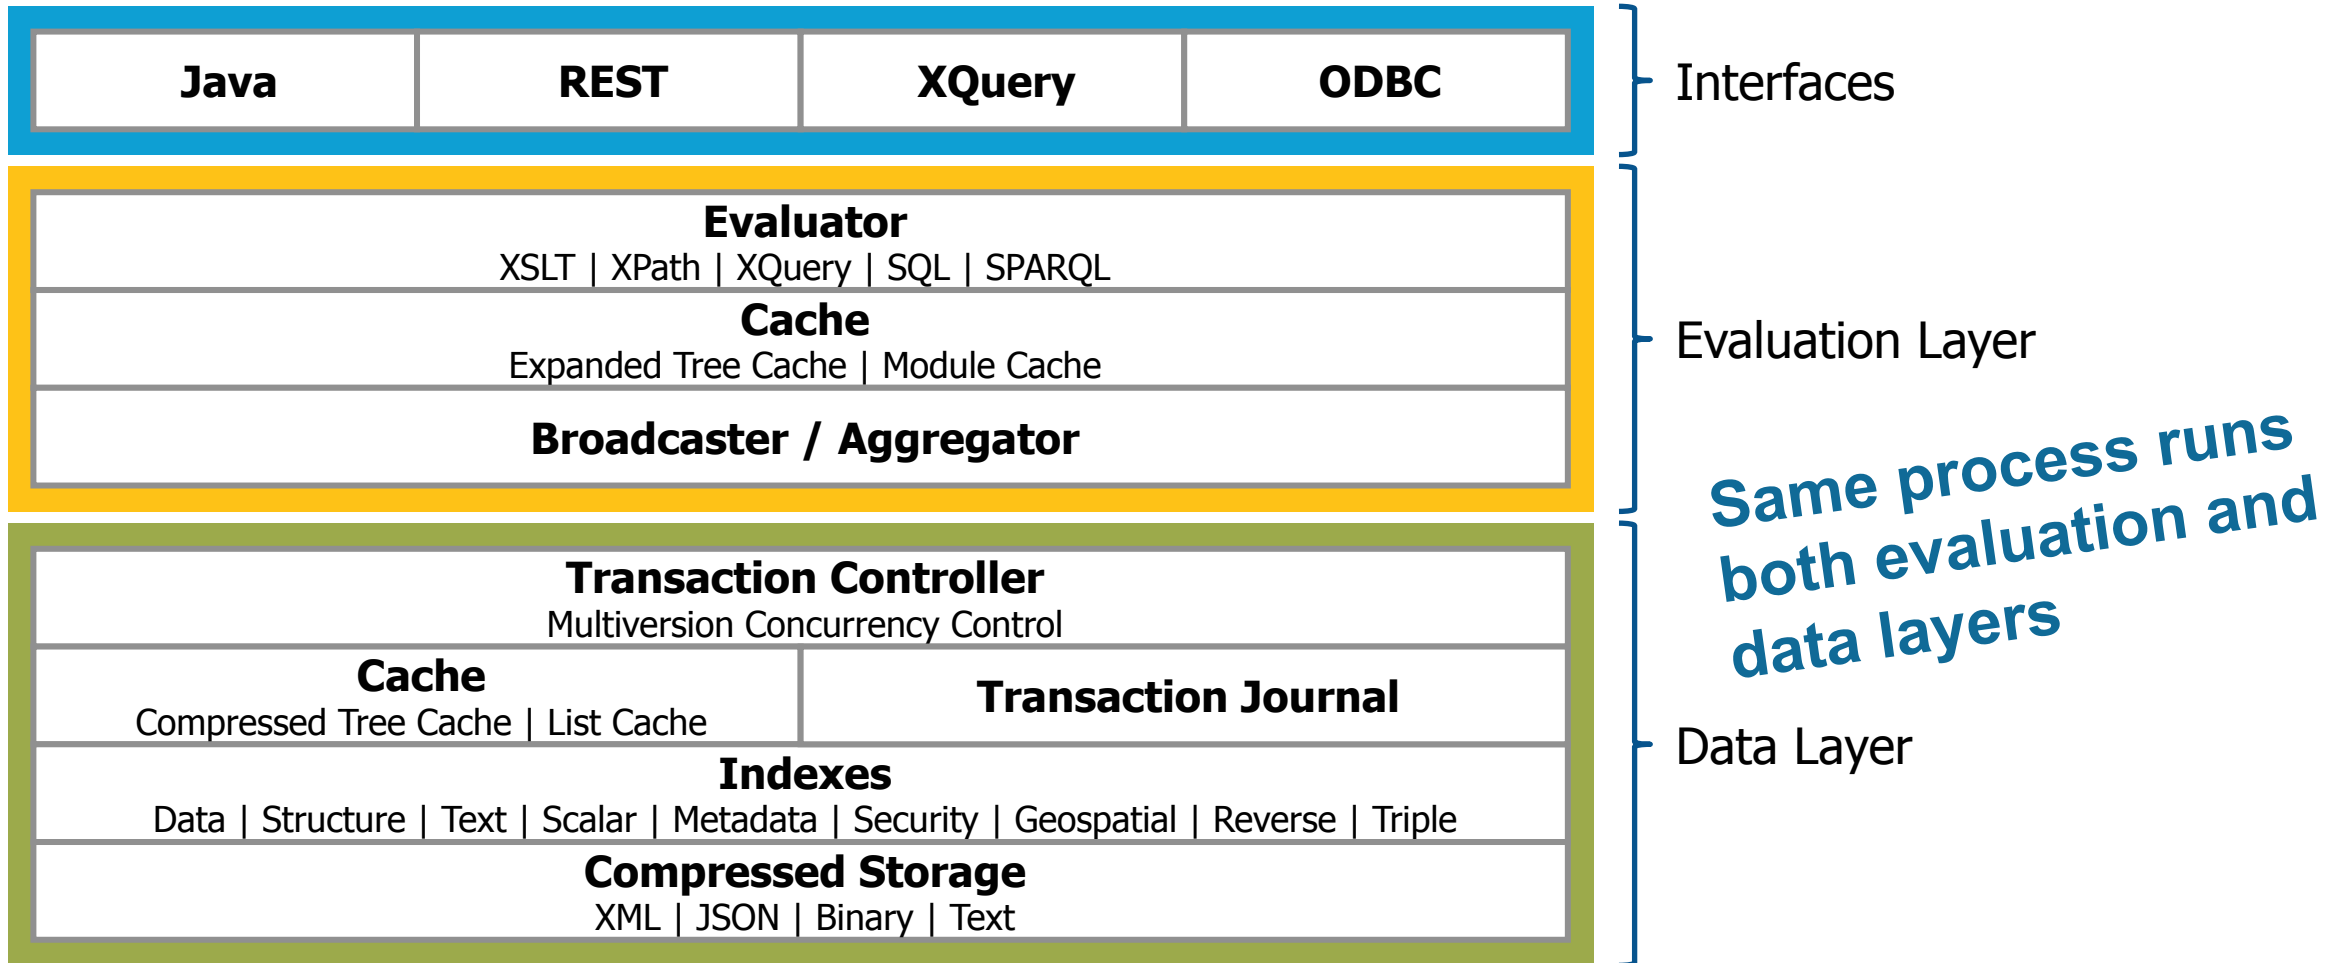
The Triple index is an index of "facts" expressed as Semantic triples. It can efficiently query and join billions of "linked data" triples.

# Reverse Indexes

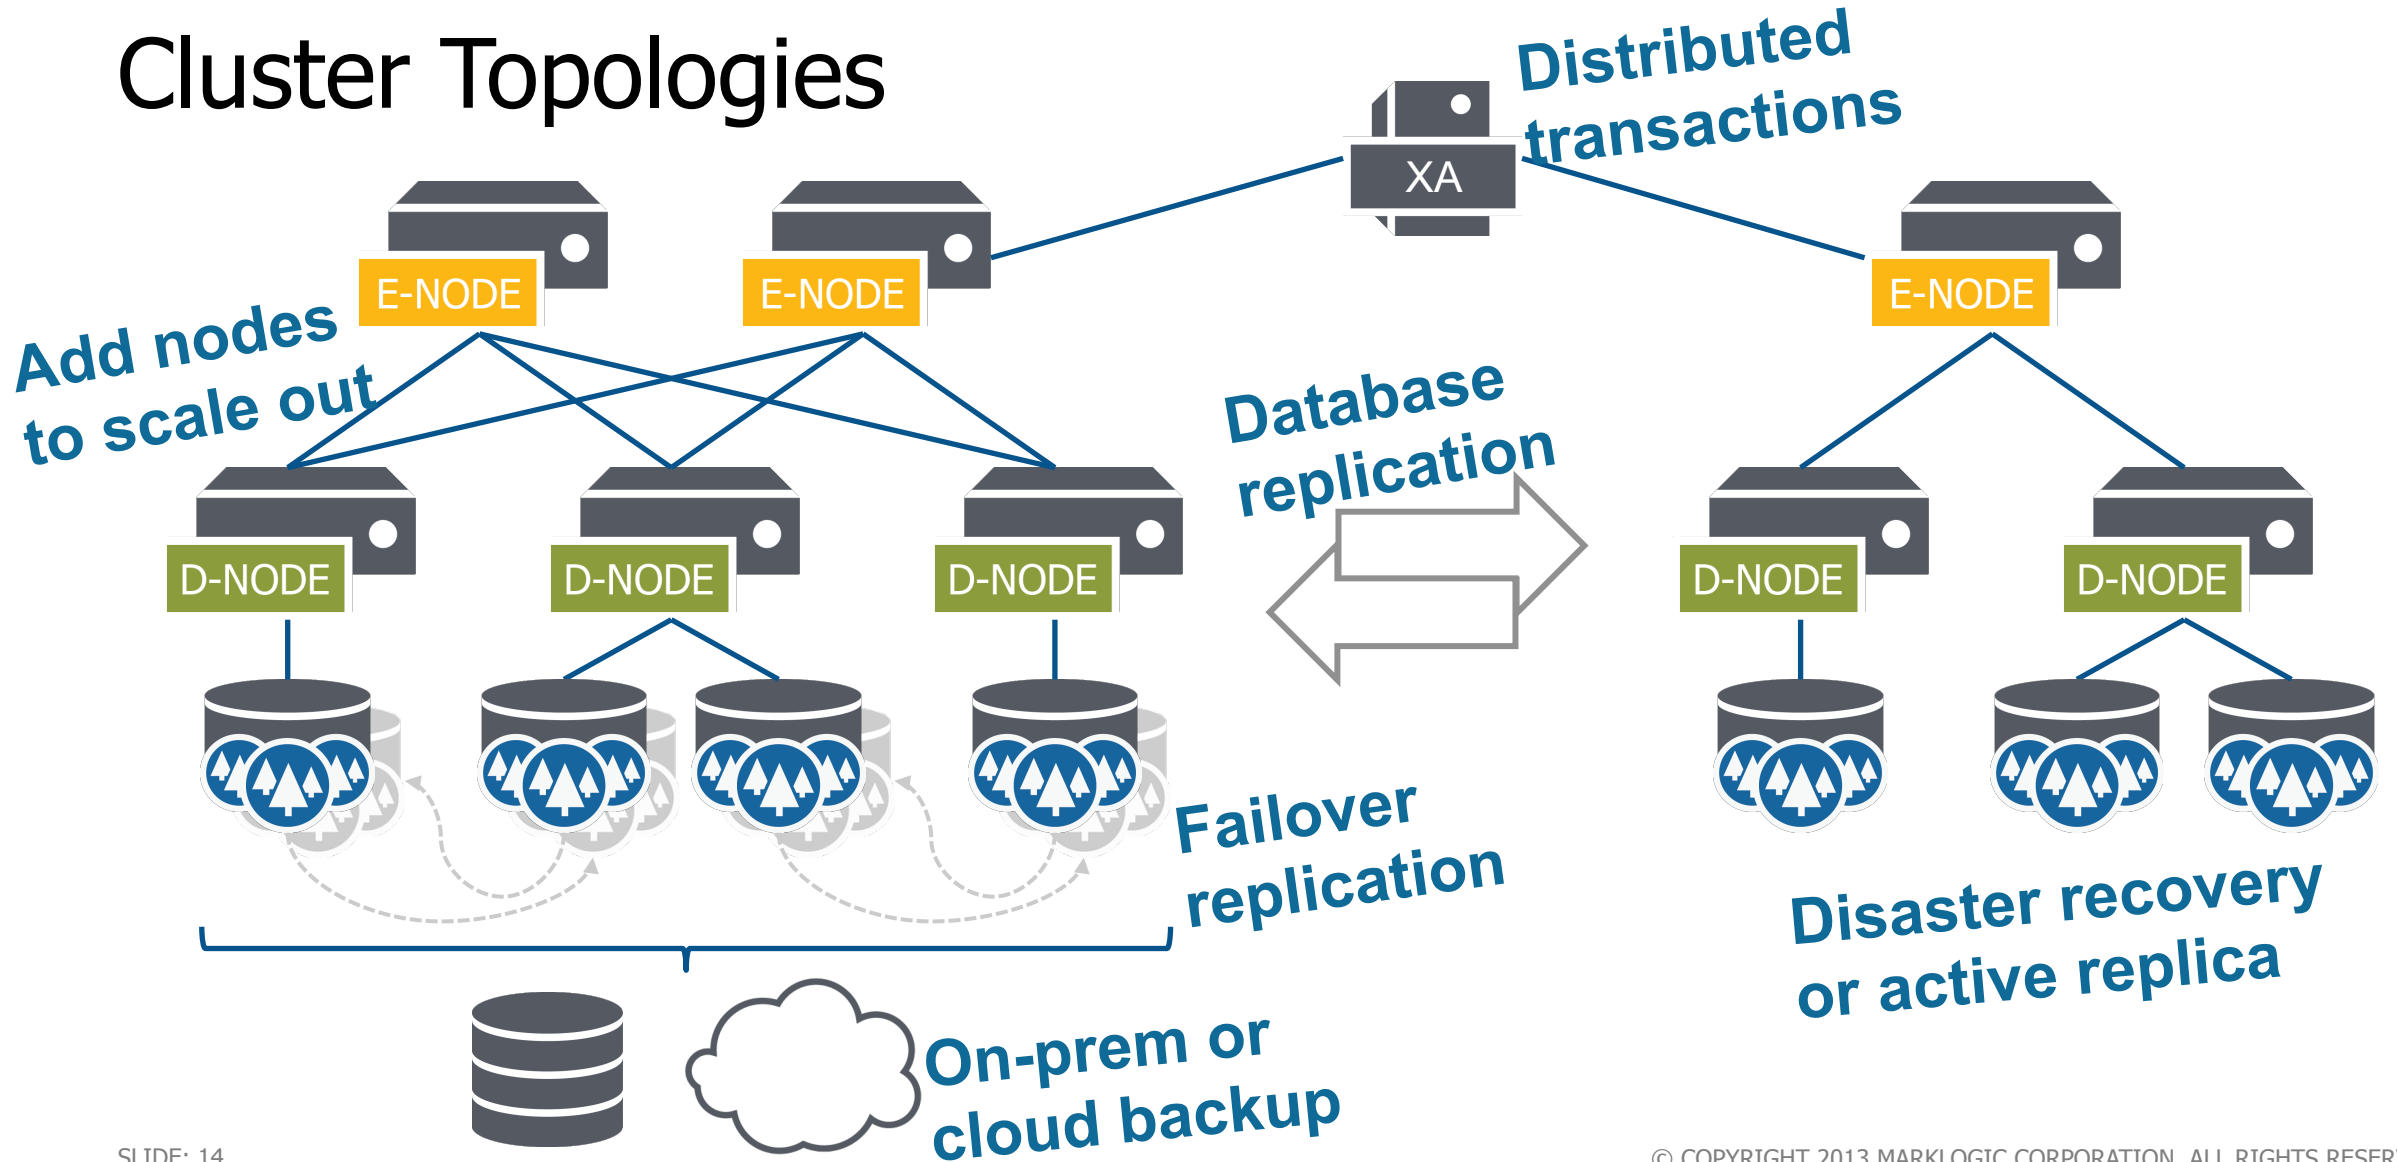
- Queries can be serialized as XML and indexed, just like documents
  - `word-query("blue van")`  
→ `<word-query><text>blue van</text></word-query>`
- MarkLogic's Alerting feature can
  - Index these queries
  - For a given data document, quickly find all possible matching queries
  - Narrow down these candidates to exact matches using a unified expression tree



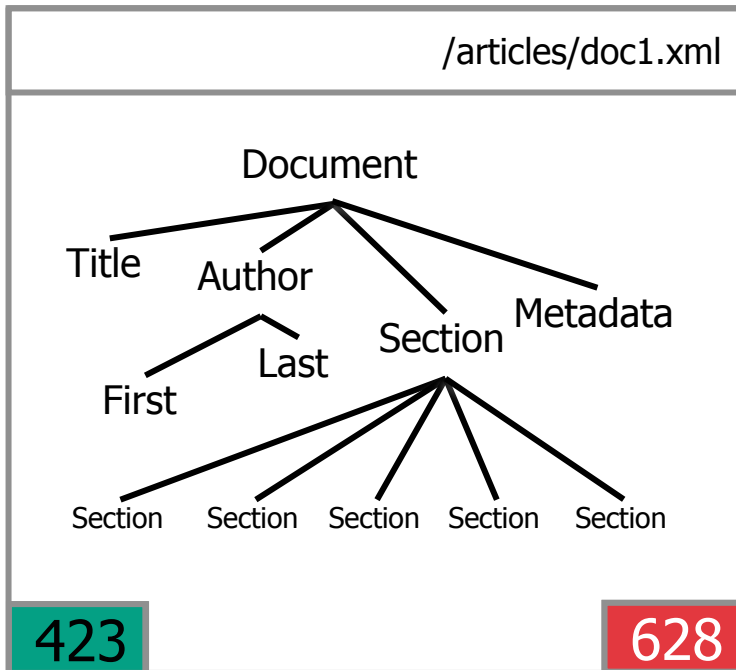
# MarkLogic Architecture



# Cluster Topologies

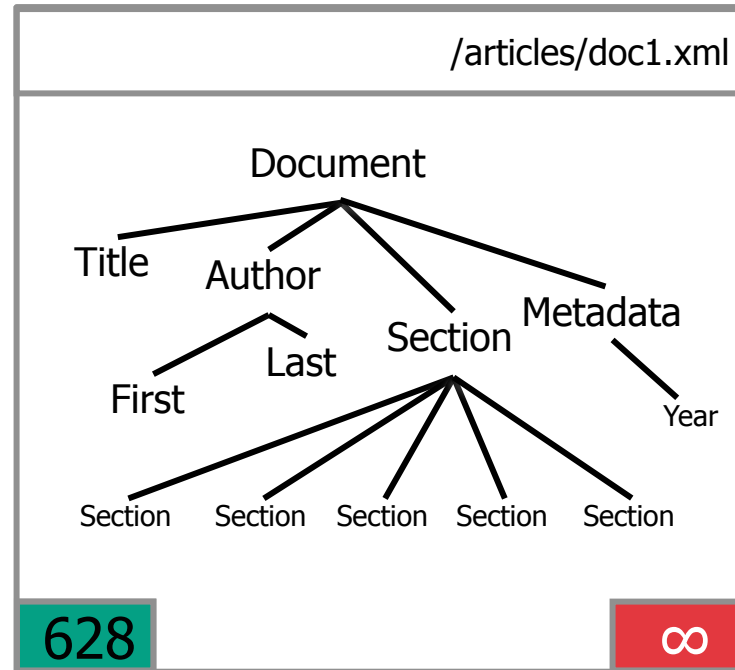


# MVCC



**C** Creation Timestamp

**D** Deletion Timestamp



## MVCC Benefits:

- ACID transactions
- Zero-latency search indexing
- High throughput
- Lock-free reads
- Point-in-time query
- Fast database rollback
- Enables serial writes






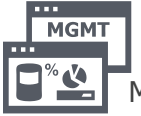






# Enterprise NoSQL

Same trusted,  
enterprise-class  
features you find in  
a RDBMS



## Trusted

Enterprise-ready and secure for mission-critical apps

 <p>ACID Transactions</p>	 <p>XA Distributed Transactions</p>
 <p>Common Criteria Security Certification</p>	 <p>Role-based Security &amp; LDAP Support</p>
 <p>Configuration Management</p>	 <p>Monitoring &amp; Management</p>
 <p>Automated Failover</p>	 <p>Backup/Restore</p>
 <p>Point-in-time Recovery</p>	 <p>Replication</p>
 <p>Journal Archiving</p>	 <p>Database Rollback</p>

# DATA

**DATA ≠  
INFORMATION**

**DATA +  
CONTEXT =  
INFORMATION**

# CONTEXT

DOCUMENTS

DOMAIN

WORLD AT LARGE



# Context from the World at Large

## Linked Open Data


- Facts that are *freely available*
- In a form that's *easily consumed*

## DBpedia (wikipedia as structured information)

- Einstein was born in Germany
- Ireland's currency is the Euro

## GeoNames

- Doha is the capital of Qatar
- Doha has these lat/long coordinates



"Linking Open Data cloud diagram, by Richard Cyganiak and Anja Jentzsch. <http://lod-cloud.net/>"

# Context from Domain

Like Open Data, but domain specific

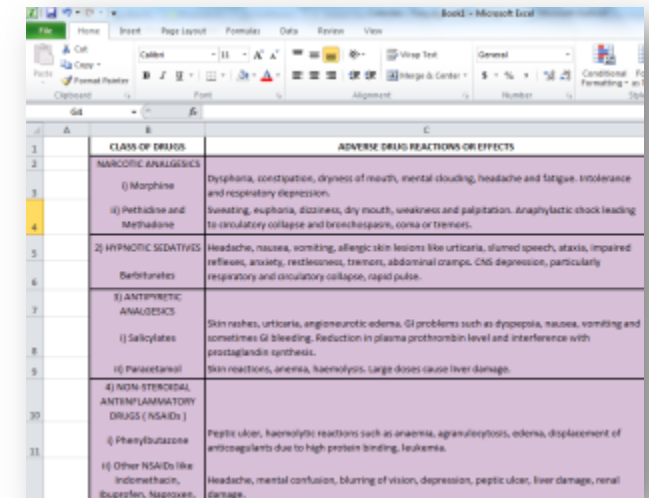
- Might be proprietary within a company
- Or shared across an industry
- Includes data and ontologies

## Some Examples

- A bank's proprietary reference data
- A pharmaceutical company's drug ontology
- An industry-wide ontology such as FIBO

## Proprietary Semantic Facts

(Facts and Taxonomies in your organization or industry)



	CLASS OF DRUGS	ADVERSE DRUG REACTIONS OR EFFECTS
1	<b>1) NARCOTIC ANALGESICS</b>	
2		
3	i) Morphine	Dysphoria, constipation, dryness of mouth, mental clouding, headache and fatigue, intolerance and respiratory depression.
4	ii) Pethidine and Methadone	Sweating, euphoria, dizziness, dry mouth, weakness and palpitation. Anaphylactic shock leading to circulatory collapse and bronchospasm, coma or tremors.
5	<b>2) HYPNOTIC SEDATIVES</b>	
6	Benzotriazoles	Headache, nausea, vomiting, allergic skin lesions like urticaria, slurred speech, ataxia, impaired reflexes, anxiety, restlessness, tremor, abdominal cramps. CNS depression, particularly respiratory and circulatory collapse, rapid pulse.
7	<b>3) ANTIPIRETIC ANALGESICS</b>	
8	i) Salicylates	Skin rashes, urticaria, angioneurotic edema, GI problems such as dyspepsia, nausea, vomiting and sometimes GI bleeding. Reduction in plasma prothrombin level and interference with prostaglandin synthesis.
9	ii) Paracetamol	Skin reactions, anemia, haemolysis. Large doses cause liver damage.
10	<b>4) NON-STEROIDAL ANTIINFLAMMATORY DRUGS (NSAIDs)</b>	
11	i) Phenylbutazone	Peptic ulcer, haemolytic reactions such as anaemia, agranulocytosis, edema, displacement of anticoagulants due to high protein binding, leukaemia.
	ii) Other NSAIDs like Indomethacin, Ibuprofen, Naproxen.	Headache, mental confusion, blurring of vision, depression, peptic ulcer, liver damage, renal damage.

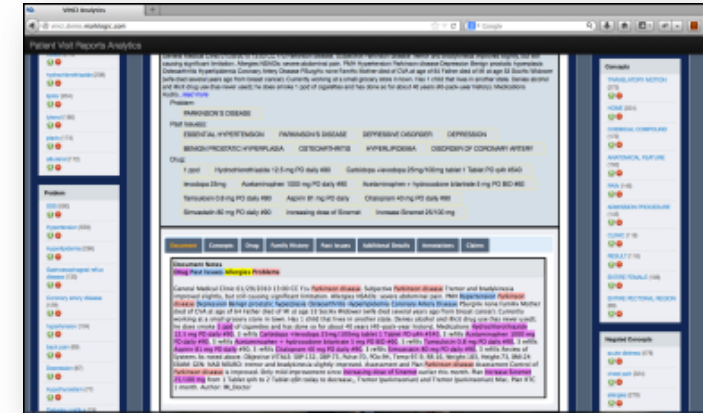
# Context from Documents

## Document metadata

- Ex: Categories, author, publish date, source

## Facts in free-flowing text

- Entities: this document mentions the person Richard Nixon, the product Advil, the company IBM
- Events: this document says that Nixon *went to* China, John Smith *met* Jane Doe, Barclays *acquired* Lehman Brothers
- Found automatically or provided at authoring time



```
Documents can contain triples
<article>
<meta>
<title>Man bites dog</title>
<sem:triple>
  <sem:subject>http://example.org/news/42</sem:subject>
  <sem:predicate>http://example.org/published</sem:predicate>
  <sem:object>2013-09-10</sem:object>
</sem:triple>
...
```

# What's behind this?

The image shows a Google search for "new york city". The search results include several links: "Welcome to NYC.gov | City of New York", "New York City - Wikipedia, the free encyclopedia", "The Official New York City Guide to NYC Attractions, Dining, Hotels ...", "News for new york city" (with a sub-link for "Security Overhaul Planned For New York City Marathon"), "Small Drone Crashes Into New York City Sidewalk (VIDEO)", and "New York City Tourism and Vacations: 2,073 Things to Do in New ...".

On the right side of the search results, there is a knowledge panel for "New York City". It features a map of the city and the following information:

- New York City**
- 8,506 followers on Google+
- New York is the most populous city in the United States and the center of the New York Metropolitan Area, one of the most populous urban agglomerations in the world. *Wikipedia*
- Area:** 468 sq miles (1,213 km<sup>2</sup>)
- Weather:** 72°F (22°C), Wind S at 4 mph (6 km/h), 62% Humidity
- Local time:** Thursday 8:54 PM
- Population:** 8.337 million (2012)
- Unemployment rate:** 8.7% (Aug 2013)
- Awards:** Top Intelligent Community of the Year
- Upcoming events:**
  - StubHub!
  - Audra McDonald Oct 5
  - Kaskade Oct 5
  - Bollywood Bash Oct 5
  - Torment & Girl in the back - Do... Oct 5
  - MoMA Studio: Sound in... Oct 3

# SEMANTICS



## Semantics: A New Way to Organize Data

**Data** is stored in Triples, expressed as: **Subject** : **Predicate** : **Object**  
John Smith : livesIn : London  
London : isIn : England




**Rules** tell us something about the triples:

If (**A livesIn X**) AND (**X isIn Y**) then (**A livesIn Y**)



# Semantics: A New Way to Search Data

- **Document search:**  
Search full-text and structure
  - ➔ **Database query:**  
Query data (dates, numbers, strings)  
and ranges
  - ➔ **Semantics:**  
Query facts and relationships from  
many sources
- SPARQL** is a language designed to query  
triples. It looks a bit like SQL.


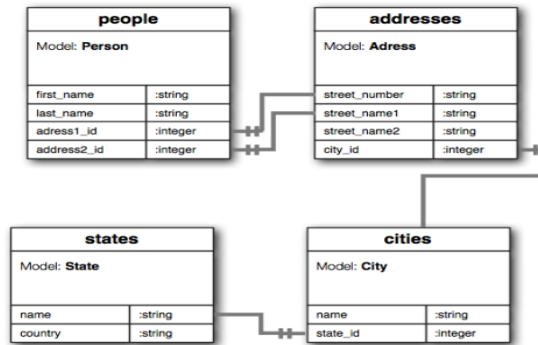
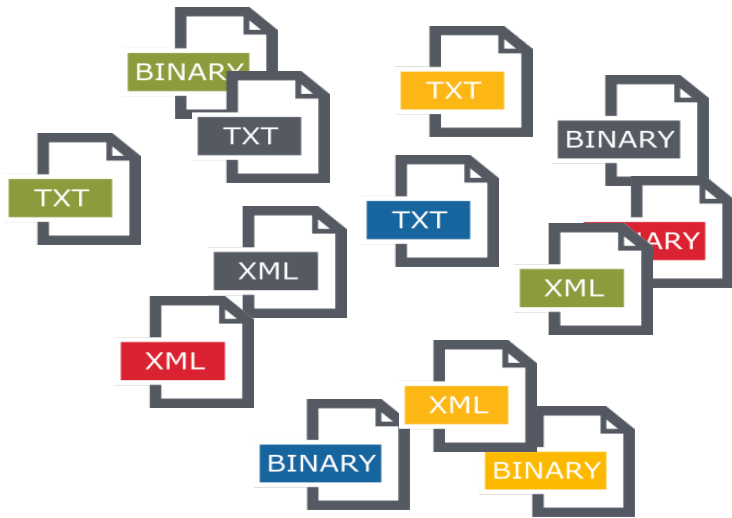


# Executing SPARQL

```
sem:sparql(  
  prefix : <http://example.org/kennedy/>  
  select * {  
    ?person :first-name ?first;  
            :last-name ?last;  
            :alma-mater [:ivy-league :true]  
  }",  
  map:entry("first", "John"),  
  (),  
  cts:collection-query("mycollection")  
)
```



# MarkLogic + MarkLogic Semantics: A New Way to Understand and Discover Information



Document Store + Data Store + Triple Store



# DYNAMIC SEMANTIC PUBLISHING

# Dynamic Semantic Publishing

## BBC Sports

### The Challenge

- Size and Complexity:
  - # of athletes
  - # of teams
  - # of assets (match reports, statistics, etc.)
  - # of relations (facts)

### Goals

- Rich user experience
  - See information in context
  - Personalize content
  - Easy navigation
  - Intelligently serve ads (outside of UK)
- Manageable
  - Static pages? Too many, changing too fast
- Limited number of journalists
  - Automate as much as possible

# Dynamic Semantic Publishing

## A Solution

### **XML Database**

- Store, manage documents
  - Stories
  - Blogs
  - Feeds
  - Profiles
- Store, manage values
  - Statistics
- Full-Text search
- Performance, scalability
- Robustness

### **Triple Store**

- Metadata about documents
  - Tagged by journalists
  - Added (semi-)automatically
  - Inferred
- Facts reported by journalists
- Linked Open Data for real-world facts

The image shows a screenshot of the BBC Sport Football website for Chelsea. Several elements are circled in red and annotated with black arrows and text:

- Chelsea**: Circled in red at the top left of the main content area.
- Lampard**: Circled in red in the main article header.
- plays for**: An arrow points from the word "Lampard" to the word "Chelsea".
- played in**: An arrow points from the word "Chelsea" to the match listing "Southampton v Chelsea".
- plays in**: An arrow points from the word "Lampard" to the match listing "Southampton v Chelsea".
- Southampton v Chelsea**: Circled in red in the "Latest Football" section.
- Premier League**: Circled in red in the "League Table" section.

The page content includes a navigation bar with "SPORT FOOTBALL", a sub-navigation bar with "Home Football Formula 1 Cricket Rugby U Rugby L Tennis Golf", and a main navigation bar with "All Teams Chelsea Results Fixtures Live Scores All Teams Leagues & Competitions". There is an advertisement for Doctors Without Borders, a main article about Bobby Tambling support, a match schedule table, a video and audio section, and a league table.

Last Match	Next Match
Sun 17 March 2013 PREMIER LEAGUE vs West Ham United (Home) Won 2 - 0 View match report	Sat 30 March 2013 PREMIER LEAGUE vs Southampton (Away) 15:00

Pos	Team	P	GD	Pts
1	Man Utd	29	38	74
2	Man City	29	25	59

## Facts into actions:

Include a story about Lampard

Include in the list of recent matches

Include in the League Table

# MARKLOGIC SEMANTICS



## Triple Store

Enterprise ready & horizontally-scalable. Store RDF triples with documents and data



## Triple Index

In addition to data, structure, text, scalar, metadata, security, and geospatial




## SPARQL

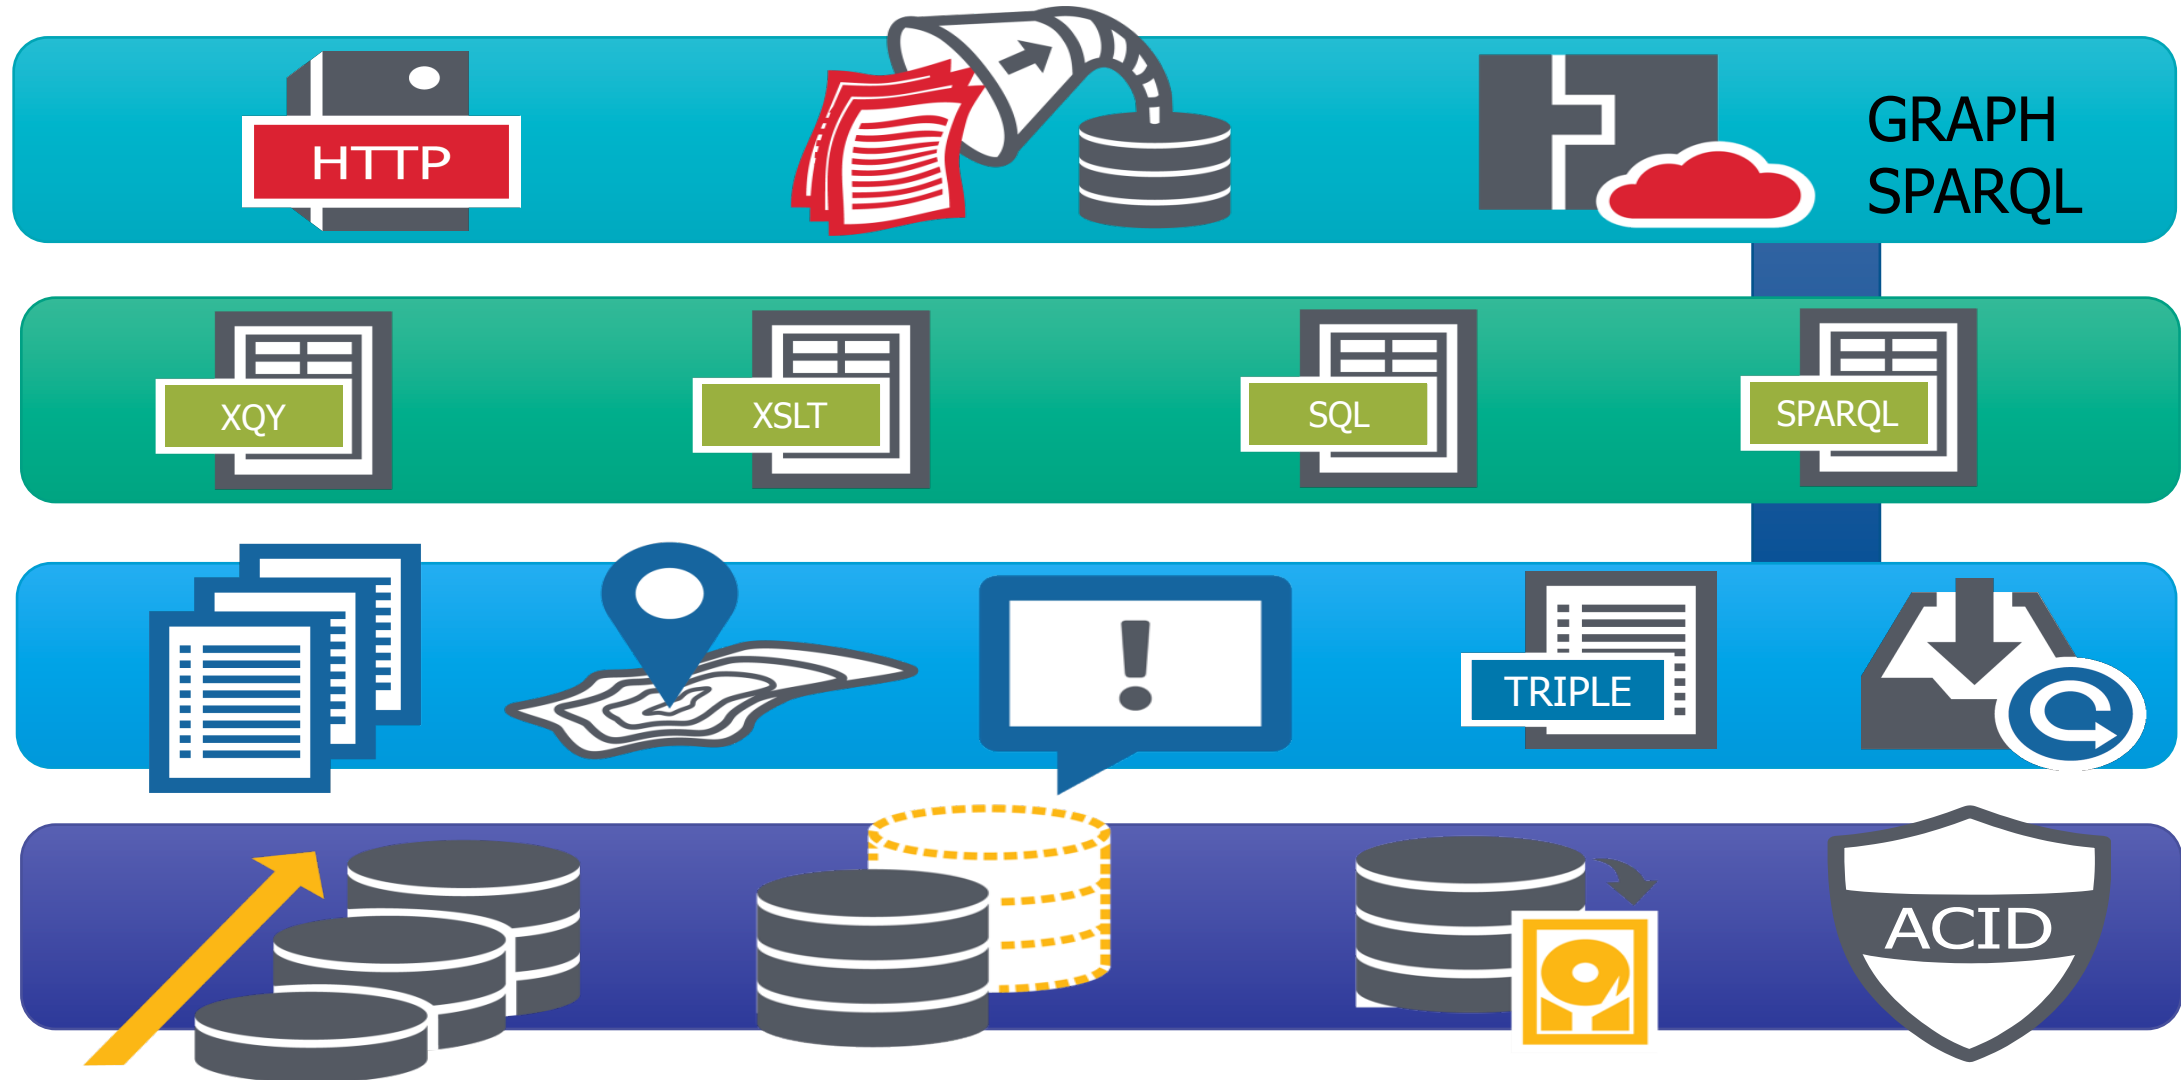
Industry-standard query language: compose with data, text, geo & other indexes

# Triple Index

- 3 triple orders
- Cached for performance
- Works seamlessly with other indexes
- Security
- 150 bytes per triple on disk
- Billions of triples per host
- Scaling out horizontally



# Semantics Architecture

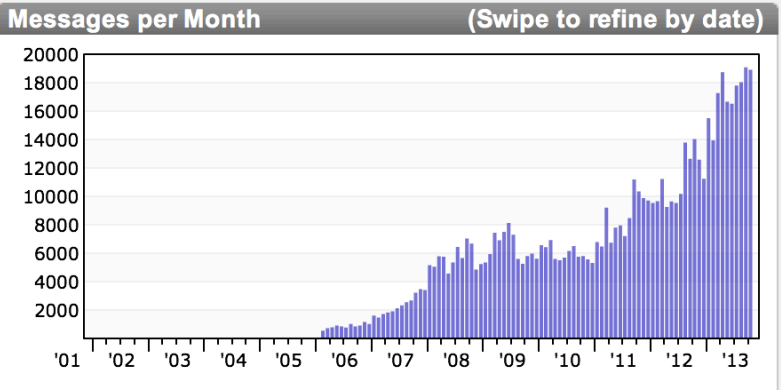
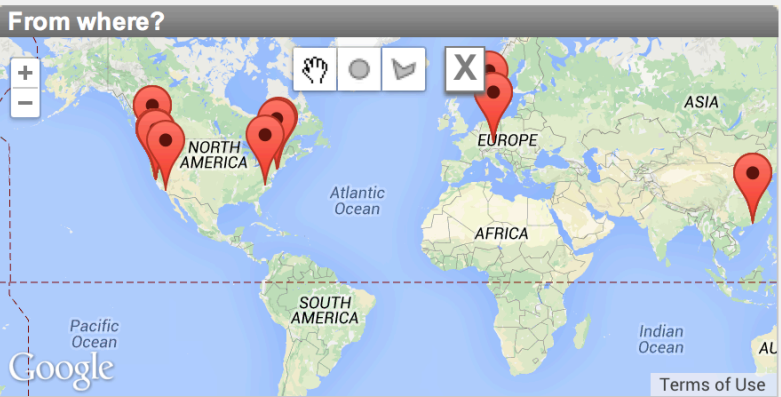




# SEMANTICS IN ACTION

# Demo time

- MarkMail  
<http://markmail.org>
- Newz  
<http://newz.nl/?lang=en>  
<http://www.youtube.com/watch?v=Mf8Lb1Cmhus>
- Gospel Topic Explore  
<http://cdev.lds.org/semantic>
- Combining Semantic with Document search and migrating RDBMS data to the MarkLogic NoSQL DB  
<https://github.com/mustard57/SYTYCD>  
<http://www.youtube.com/watch?v=4cUqMzsu0F4>



What List? [View more](#)

org.apache.hadoop.core-dev	83,180
org.apache.hadoop.hive-dev	67,340
org.apache.hadoop.core-user	46,813

Who Sent It?	View more	Any Attachments?	View more
Apache Jenkins Server	20,431	png	851
Hadoop QA (JIRA)	11,574	jpg	352
Hudson (JIRA)	10,345	txt	259

Sort by Relevance 1 to 10 of about 661576

[Re: \[ANN\] HBase 0.20.0-alpha available for download](#)  
 Is there a requirement for **hadoop** 0.20 for HBase 0.20?  
 Jun 17, 2009 - jason hadoop - org.apache.hadoop.hbase-user

[re about translating Hadoop:The Definitive Guide into chines...](#)  
 So am I, although i'm just a novice to **hadoop**. I worked on a project based on the Zookeeper the... .. when I go deep into **hadoop**. Maybe I can help with chapter 11 Pig. I am now pig's committer.  
 Jan 9, 2010 - hadoop - org.apache.hadoop.general

[Re: Running hadoop in standalone mode using Eclipse](#)  
 ...JAVA System Property **hadoop**.log.dir to something  
 Jun 17, 2009 - jason hadoop - org.apache.hadoop.core-dev

[Re: Hadoop 20.2 release compability](#)  
 required to execute the same on any current **hadoop** 0.20.20x releases. Both the Subject: **Hadoop** 20.2 release compability  
 Nov 3, 2011 - bejo...@gmail.com - org.apache.hadoop.core-user

[Re: Re: Re: Regarding "Hadoop multi cluster" set-up](#)  
 ...0, depending on your **hadoop** version working in this case. I am getting the following error in the **hadoop**-root-datanode-slave log file while running the co...  
 Feb 7, 2009 - jason hadoop - org.apache.hadoop.core-user

[Re: Hadoop Installation Mappers setting](#)  
 Other services include non **hadoop** services as well as **hadoop** daemons. Subject: **Hadoop** Installation Mappers setting  
 May 23, 2013 - bejo...@gmail.com - org.apache.hadoop.mapreduce-user

[Re: Is there any performance issue with Jrockit JVM for Hado...](#)  
 ... lock with profiling **Hadoop** with SUN JVM? Is there any tricky or special setting for Jrockit vm on **Hadoop**?  
 May 9, 2009 - JQ Hadoop - org.apache.hadoop.core-user

[Re: incremental loads into hadoop](#)  
 Sam Your understanding is right, **hadoop** definitely works great with largeGiga,Tera or Peta bytes. Mostly when said **hadoop** process tera bytes of data. It Subject: Re: incremental loads into **hadoop**

# Apache Hadoop



Apache Hadoop is a software framework that supports data-intensive distributed applications under a free license. It enables applications to work with thousands of computational independent computers and petabytes of data. Hadoop was derived from Google's MapReduce and Google File System (GFS) papers. Hadoop is a top-level Apache project being built and used by a global community of contributors, written in the Java programming language.

## Releases

- Stable: 1.0.4 [download](#)
- Unstable: 2.0.2-alpha [download](#)

## Key Personnel

- [Doug Cutting](#)
- [Raghu Angadi](#)
- [Owen O'Malley](#)

## Related Projects

- [Apache Zookeeper](#)
- [Apache Pig](#)

Sort by Relevance 1 to 10 of about 661576

Re: [ANN] HBase 0.20.0-alpha available for download
Is there a requirement for hadoop 0.20 for HBase 0.20?
Jun 17, 2009 - jason hadoop - org.apache.hadoop.hbase-user

re about translating Hadoop:The Definitive Guide into chines...
So am I, although i'm just a novice to hadoop. I worked on a project
based on the Zookeeper the... when I go deep into hadoop. Maybe I
can help with chapter 11 Pig. I am now pig's committer.
Jan 9, 2010 - hadoop - org.apache.hadoop.general

Re: Running hadoop in standalone mode using Eclipse
...JAVA System Property hadoop.log.dir to something
Jun 17, 2009 - jason hadoop - org.apache.hadoop.core-dev

Re: Hadoop 20.2 release compability
required to execute the same on any current hadoop 0.20.20x releases.
Both the Subject: Hadoop 20.2 release compability
Nov 3, 2011 - bejo...@gmail.com - org.apache.hadoop.core-user


Re: Re: Re: Re: Regarding "Hadoop multi cluster" set-up
...0, depending on your hadoop version working in this case. I am
getting the following error in the hadoop-root-datanode-slave log file
while running the co...
Feb 7, 2009 - jason hadoop - org.apache.hadoop.core-user

Re: Hadoop Installation Mappers setting
Other services include non hadoop services as well as hadoop
daemons. Subject: Hadoop Installation Mappers setting
May 23, 2013 - bejo...@gmail.com - org.apache.hadoop.mapreduce-
user

Re: Is there any performance issue with Jrockit JVM for Hado...
... lock with profiling Hadoop with SUN JVM? Is there any tricky or
special setting for Jrockit vm on Hadoop?
May 9, 2009 - JQ Hadoop - org.apache.hadoop.core-user

Re: incremental loads into hadoop
Sam Your understanding is right, hadoop definitely works great with
large Giga, Tera or Peta bytes. Mostly when said hadoop process tera
bytes of data. It Subject: Re: incremental loads into hadoop

Jason Venner



Public profile

San Francisco

Known for: Pro Hadoop, Apress



On Wed, Jun 17, 2009 at 1:44 AM, Andrew Purtell <apur...@apache.org> wrote:

Minor correction/addition: "Stargate" is undergoing shared development in
two github trees:
http://github.com/macdiesel/stargate/tree/master
http://github.com/apurtell/stargate/tree/master
The github network view can show which branch is leading at any point in
time:
http://github.com/apurtell/stargate/network

- Andy

From: stack <sta...@duboce.net>
To: hbas...@hadoop.apache.org; core...@hadoop.apache.org
Sent: Tuesday, June 16, 2009 10:23:20 PM
Subject: [ANN] HBase 0.20.0-alpha available for download

An alpha version of HBase 0.20.0 is available for download at:

Apache Hadoop



Apache Hadoop is a software framework that supports
data-intensive distributed applications under a free
license. It enables applications to work with thousands
of computational independent computers and petabytes
of data. Hadoop was derived from Google's MapReduce
and Google File System (GFS) papers. Hadoop is a top-
level Apache project being built and used by a global
community of contributors, written in the Java
programming language.

Releases

- Stable: 1.0.4 download
Unstable: 2.0.2-alpha download


Key Personnel

- Doug Cutting
Raghu Angadi
Owen O'Malley

Related Projects

- Apache Zookeeper
Apache Pig








## Gerelateerde concepten

## Personen

▼8




-  Helima Obama
-  Kezia Obama
-  Fernando Obama Galindo

## Gebeurtenissen


-  Thomas More Law Center v. Obama

## Organisaties

▼8

-  Republicans for Obama
-  Obama campaign
-  Obama Works

## Locaties

-  Ohama

geen resultaten

## Verfijnen


## Datumfilter

Start: Eind:

# GTE beta

The Gospel Topics Explorer is a concordance of various LDS-related study helps, including the Bible Dictionary, Topical Guide, Guide to the Scriptures, Teaching by Topic and many others.

**Entity enrichment has been added!** Now you can see the important entities in your content and link them into the GTE graph.

[Learn more »](#)[Get started »](#)

## Explore

Navigate the Gospel Ontology by zooming in anywhere among the 6,600+ connected topics, or click a node for more details. Focus on a given cluster of nodes or scroll back for the bigger picture.

[View details »](#)

## Search

Search more than 11,000 topics, 23,000 relationships and 67,000 facts by keyword, phrase or topic. Results are ordered by the number of relationships each topic contains and includes direct links to examples of the topic.

[View details »](#)

## Find relationships

Ever wonder what the relationship between sacrament and revelation is? Or how sacrifice and the atonement are connected?

[View details »](#)

## Enrich

Find the topics, people, organizations and roles embedded in content. Just copy and paste some text into the form and the GTE will identify entities of interest.

[View details »](#)

Search inspiration,worthiness

Inspiration (12)

Worthiness (13)

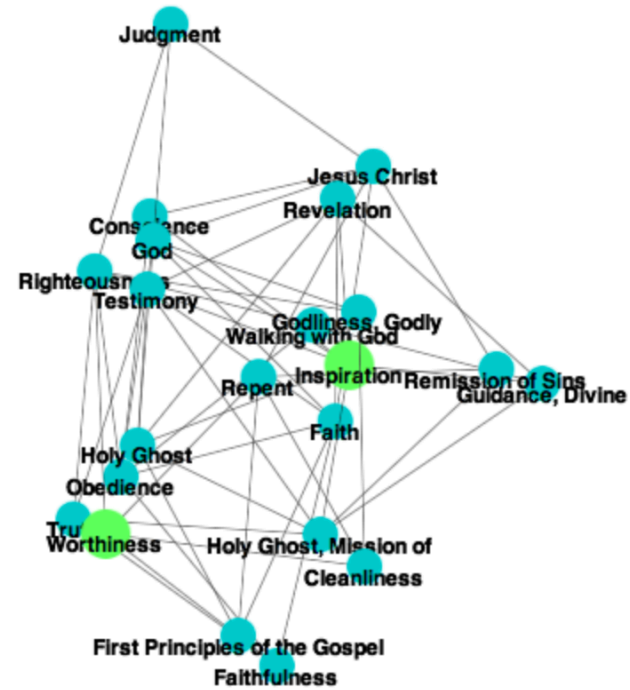
Relate topics

3895 nodes examined

Getting started: In the Search box above, enter two terms or phrases separated by a comma, then click Search. Select other topics from one or both dropdowns to refine your search.

Examples

- atonement, joseph smith
wealth, misery
inspiration, worthiness
sacrament, atonement
worthiness, mission
cleanliness, godliness
tea, coffee





Search for:  which [+ Child]

Are related to  AS  relationship [-] [+ Child]

Are related to  AS  relationship [-] [+ Child]

Have property  less than  [-]

Are related to  AS  relationship [-] [+ Child]

Are related to  AS  relationship [-] [+ Child]

Search

## Subject Search Results

[Load related content](#)

1. [Ms. Theodora Broussard \(Joint Customer\)](#)
2. [Mr. Jeffrey Blow \(Joint Customer\)](#)
3. [Mr. John Castillo \(Joint Customer\)](#)
4. [Mr. James Bowers \(Joint Customer\)](#)
5. [Mr. Cory Jackson \(Joint Customer\)](#)
6. [Mrs. Yvone Crawford \(Joint Customer\)](#)
7. [Mr. Robert Nelson \(Joint Customer\)](#)
8. [Dr. Carole Smoot \(Joint Customer\)](#)
9. [Mrs. Joanne Miller \(Joint Customer\)](#)

## Entity Facts

NKB Customer: Broussard

[Load related content](#)

Account: [96302787 \(NKB Account\)](#)

Address: [/docs/2.html](#)

Branch: [Broussard \(NKB Customer\)](#)

Home Phone Number: 01422 518885

Mobile Phone Number: 07838 147787

Address ID: 11

Customer ID: 11

Date Of Birth: 1964-12-19

Branch ID: 9

Initial: A

## Related Documents

Search:

## Results

No Results

# Enterprise NoSQL: Powerful, Agile & Trusted

## Powerful

Deliver more value, build more powerful applications



Semantics:  
RDF &  
SPARQL



Flexible  
Indexes



Triple  
Index



JSON  
Storage



Full Text  
Search



Alerting  
& Event  
Processing



Scalable



Geospatial  
Query



REST &  
Java APIs



In-database  
MapReduce



Analytic  
Functions



Visualization  
Widgets

## Agile

Prepare for and respond to change quickly



Schema-  
Agnostic



Elastic



Tiered  
Storage



Cloud  
Ready



mlcp  
Content  
Pump



Programmatic  
Controls &  
Metering



HDFS &  
Amazon S3  
Storage



Application  
Builder



Hadoop  
Connector



SQL  
Support



Information  
Studio



BI  
Integration

## Trusted

Enterprise-ready and secure for mission-critical apps



ACID  
Transactions



XA  
Distributed  
Transactions




Common  
Criteria  
Security  
Certification



Role-based  
Security &  
LDAP  
Support



Configuration  
Management



Monitoring  
&  
Management



Automated  
Failover



Backup/  
Restore



Point-in-  
time  
Recovery



Replication



Journal  
Archiving



Database  
Rollback



ANY QUESTIONS?

# Contact Information

Emil Zegers

Pre-Sales Consultant Benelux & Nordics

[emil.zegers@marklogic.com](mailto:emil.zegers@marklogic.com)

+31 619929703

Graadt van Roggenweg 328-334, Blok D

3531AH Utrecht, Netherlands

<http://marklogic.com>

## Useful links [1]

MarkLogic Developer Community

<http://developer.marklogic.com>

Inside MarkLogic Server

<http://developer.marklogic.com/pubs/architecture/inside-marklogic-server-r7.pdf>

MarkLogic Application Development

<http://docs.marklogic.com/guide/app-builder/intro>

## Useful links [2]

What's New in MarkLogic 7

<http://www.marklogic.com/what-is-marklogic/marklogic>

MarkLogic MLJS

<https://github.com/adamfowleruk/mljs>

MarkLogic Blog

<http://developer.marklogic.com/blog>