

Neo4j : Konzepte, Anwendungsfälle, Live-Demo

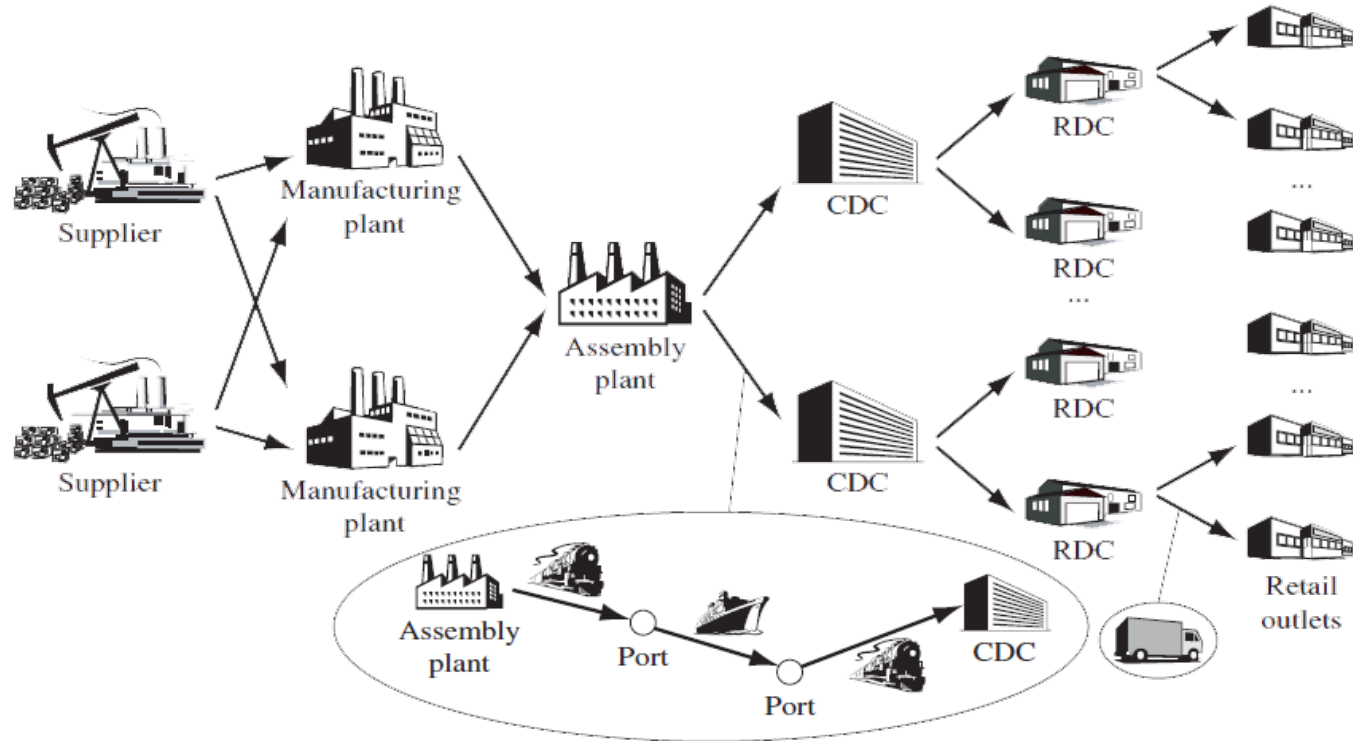
Herzlich Willkommen!

Juni 2016

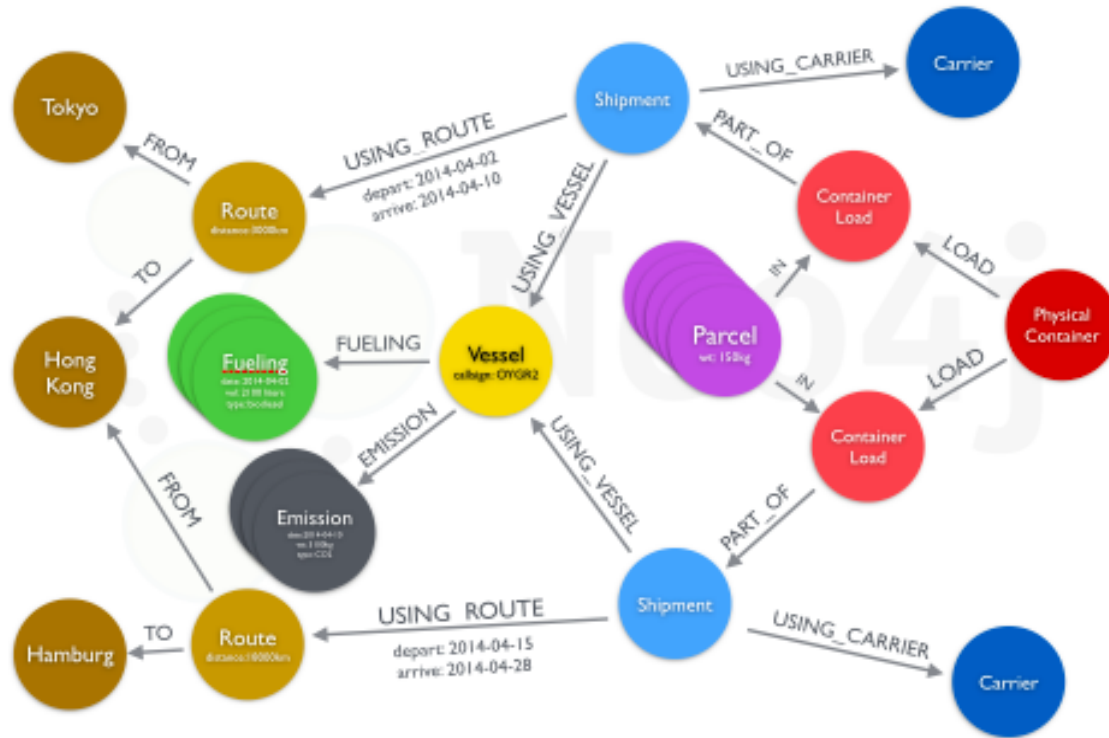
stefan.armbruster@neotechnology.com



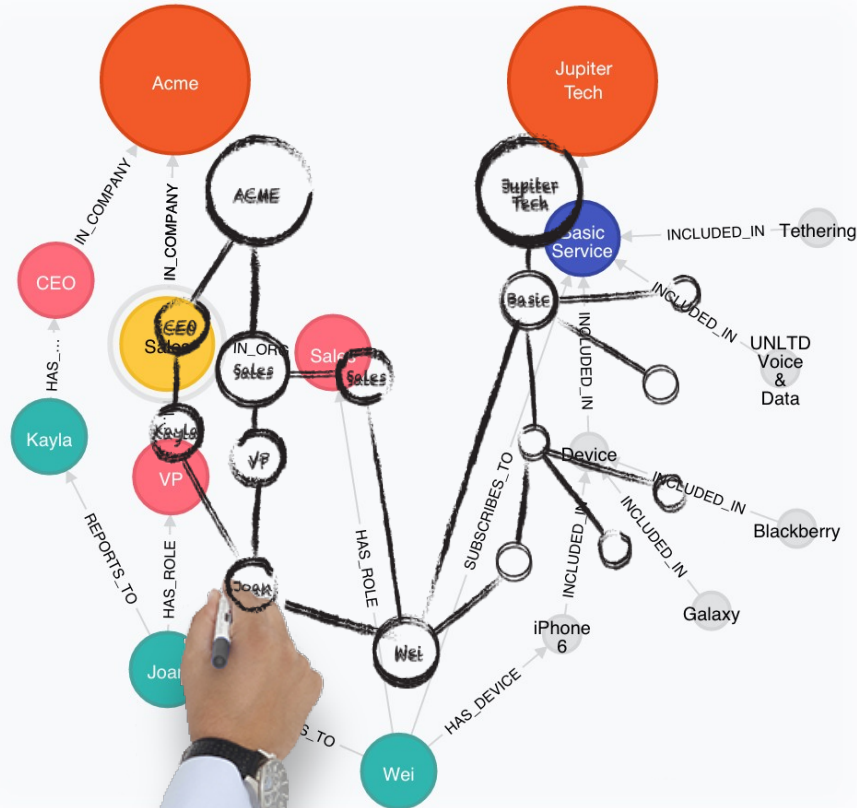
Beispiel: Logisches Modell Logistikprozess



Graphmodell: Knoten und Kanten, kein Schema



Intuitiv, “white board friendly”




performant, flexibel, agil



Graphs are the most efficient and natural way of working with data.

They are deeply intuitive, and mimic the interconnectedness of concepts and ideas in the human mind.



Flexibility

Model, develop, and visualize the world as you experience it: slice away logical-to-physical friction and the need to wrestle reality into predefined tables. It's simply nodes and relationships, and the world becomes flexible: fast development and agile evolution.

Performance

Neo4j's native graph engine is engineered to let you navigate hyper-connectivity at speed. Built from the bottom up to support property graphs, Neo4j lets you connect the "dots" easily, and with unparalleled performance and reliability.

Scalability

Neo4j scales up and out, supporting tens of billions of nodes and relationships, and hundreds of thousands of ACID transactions per second. We believe your data is one of your most valuable assets, and you need be able to trust your database with it. [Learn more →](#)

Use the Right Database for the Right Job




Discrete Data

*Minimally
connected data*

Connected Data

*Focused on
Data Relationships*



Neo4j is designed for data relationships

Development Benefits

Easy model maintenance
Easy query

Deployment Benefits

Ultra high performance
Minimal resource usage

High Business Value in Data Relationships



Data is increasing in volume...

- New digital processes
- More online transactions
- New social networks
- More devices

... and is getting more connected

Customers, products, processes, devices interact and relate to each other

Using Data Relationships unlocks value

- Real-time recommendations
- Network and IT operations
- Identity and access management
- Fraud detection
- Graph-based search
- Meta data management

Early adopters became industry leaders

LinkedIn **Google** **facebook**

Neo4j Leads the Graph Database Revolution



“Graph analysis is possibly the **single most effective competitive differentiator** for organizations pursuing data-driven operations and decisions after the design of data capture.”



“Forrester estimates that **over 25% of enterprises** will be using graph databases by 2017”



“Neo4j is the current market **leader in graph databases.**”

IT Market Clock for Database Management Systems, 2014

<https://www.gartner.com/doc/2852717/it-market-clock-database-management>

TechRadar™: Enterprise DBMS, Q1 2014

<http://www.forrester.com/TechRadar+Enterprise+DBMS+Q1+2014/fulltext/-/E-RES106801>

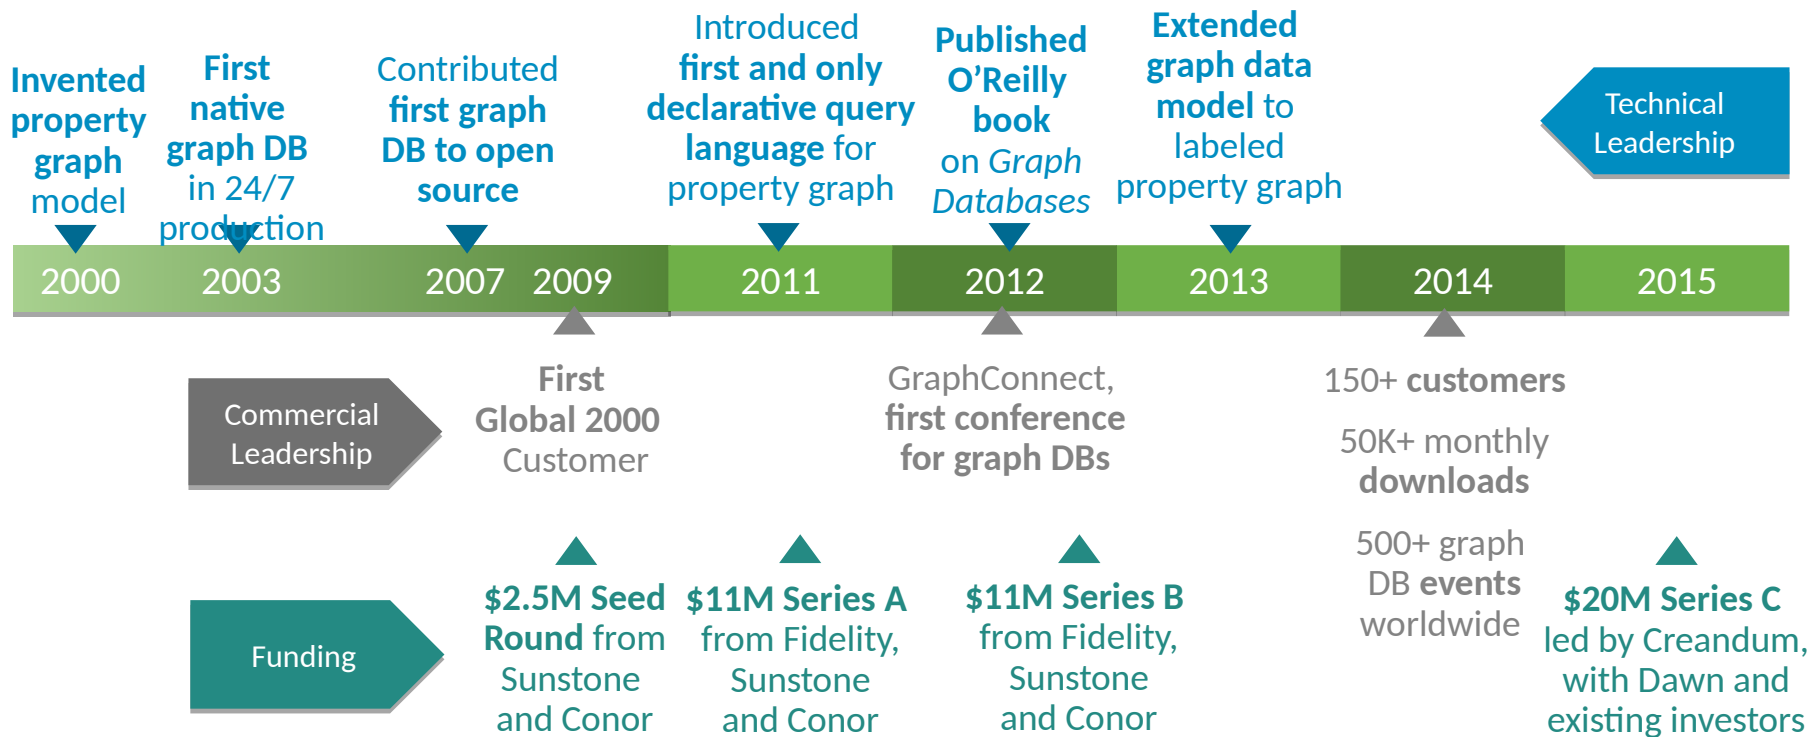
Graph Databases - and Their Potential to Transform How We Capture Interdependencies (Enterprise Management Associates)

<http://blogs.enterprisemanagement.com/dennisdrogseth/2013/11/06/graph-databasesand-potential-transform-capture-interdependencies/>

2012 ☾ 2015



Neo4j: The Graph Database Leader



Neo4j Adoption by Selected Verticals



Financial Services	Communications	Health & Life Sciences	HR & Recruiting	Media & Publishing	Social Web	Industry & Logistics
--------------------	----------------	------------------------	-----------------	--------------------	------------	----------------------



Entertainment	Consumer Retail	Business Services	Information Services
---------------	-----------------	-------------------	----------------------



Background

- Founded in 1962 and based in Arkansas
- 11,000+ stores in 27 countries with *walmart.com* online store
- 2M+ employees and \$470 billion in annual revenues



Business Problem

- Optimize *walmart.com* user experience
- Connect complex buyer and product data to gain super-fast insight into customer needs and product trends
- RDBMS couldn't handle complex queries

Solution and Benefits

- Replaced complex batch process real-time online recommendations
- Built simple, real-time recommendation system with low-latency queries
- Serve better and faster recommendations by combining historical and session data

Background

- One of the world's largest logistics carriers
- Projected to outgrow capacity of old system
- New parcel routing system

*Single source of truth for entire network
B2C and B2B parcel tracking
Real-time routing: up to 7M parcels per day*



Business Problem

- Needed 365x24x7 availability
- Peak loads of 3000+ parcels per second
- Complex and diverse software stack
- Need predictable performance, linear scalability
- Daily changes to logistics network: route from any point to any point

Solution and Benefits

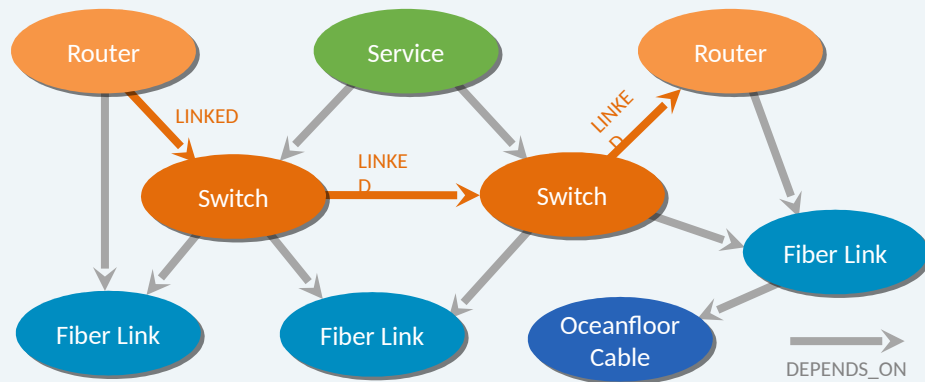
- Ideal domain fit: a logistics network is a graph
- Extreme availability, performance via clustering
- Greatly simplified routing queries vs. relational
- Flexible data model reflect real-world data variance much better than relational
- Whiteboard-friendly model easy to understand

Background

- Second largest communications company in France
- Based in Paris, part of Vivendi Group, partnering with Vodafone

Business Problem

- Infrastructure maintenance took week to plan due to need to model network impacts
- Needed what-if to model unplanned outages
- Identify network weaknesses to uncover need for additional redundancy
- Info lived on 30+ systems, with daily changes



Solution and Benefits

- Flexible inventory management supports modeling, aggregation, troubleshooting
- Single source of truth for entire network
- New apps model network via near-1:1 mapping between graph and real world
- Schema adapts to changing needs

Background

- Top investment bank with \$1+ trillion in assets
- Using a relational database and Gemfire to manage employee permissions to research document and application-service resources
- Permissions for new investment managers and traders provisioned manually



Business Problem

- Lost an average of 5 days per new hire while they waited to be granted access to hundreds of resources, each with its own permissions
- Replace an unsuccessful onboarding process implemented by a competitor
- Regulations left no room for error

Solution and Benefits

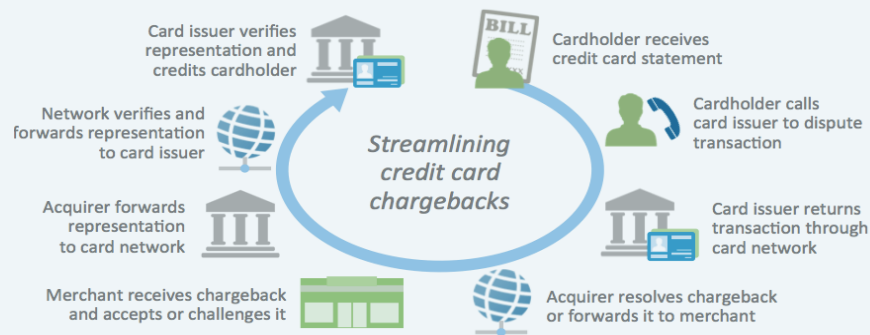
- Store models, groups and entitlements in Neo4j
- Exceeded performance requirements
- Major productivity advantage due to domain fit
- Graph visualization ease permissioning process
- Fewer compromises than with relational
- Expanded Neo4j solution to online brokerage

Background

- Global financial services firm with trillions of dollars in assets
- Varying compliance and governance considerations
- Incredibly complex transaction systems, with ever-growing opportunities for fraud

Business Problem

- Needed to spot and prevent fraud detection in real time, especially in payments that fall within “normal” behavior metrics
- Needed more accurate and faster credit risk analysis for payment transactions
- Needed to dramatically reduce chargebacks



Solution and Benefits

- Lowered TCO by simplifying credit risk analysis and fraud detection processes
- Identify entities and connections uniquely
- Saved billions by reducing chargebacks and fraud
- Enabled building real-time apps with non-uniform data and no sparse tables or schema changes

INTERNATIONAL CONSORTIUM

ICIJ

OF INVESTIGATIVE JOURNALISTS

A GLOBAL INVESTIGATION

THE PANAMA PAPERS

Politicians, Criminals and the Rogue Industry That Hides Their Cash

#PanamaPapers




Amount of data compared to previous leaks

1,7 GB Cablegate/Wikileaks (2010)



260 GB Offshore-Leaks/ICIJ (2013)




4 GB Luxemburg-Leaks/ICIJ (2014)




3,3 GB Swiss-Leaks/ICIJ (2015)



≈2,6 TB Panama Papers/ICIJ (2016)



Context is King




name: "John"
last: „Miller“
role: „Negotiator“

name: "Jose"
last: "Pereia"
position: "Governor"



name: "Alice"
last: „Smith“
role: „Advisor“

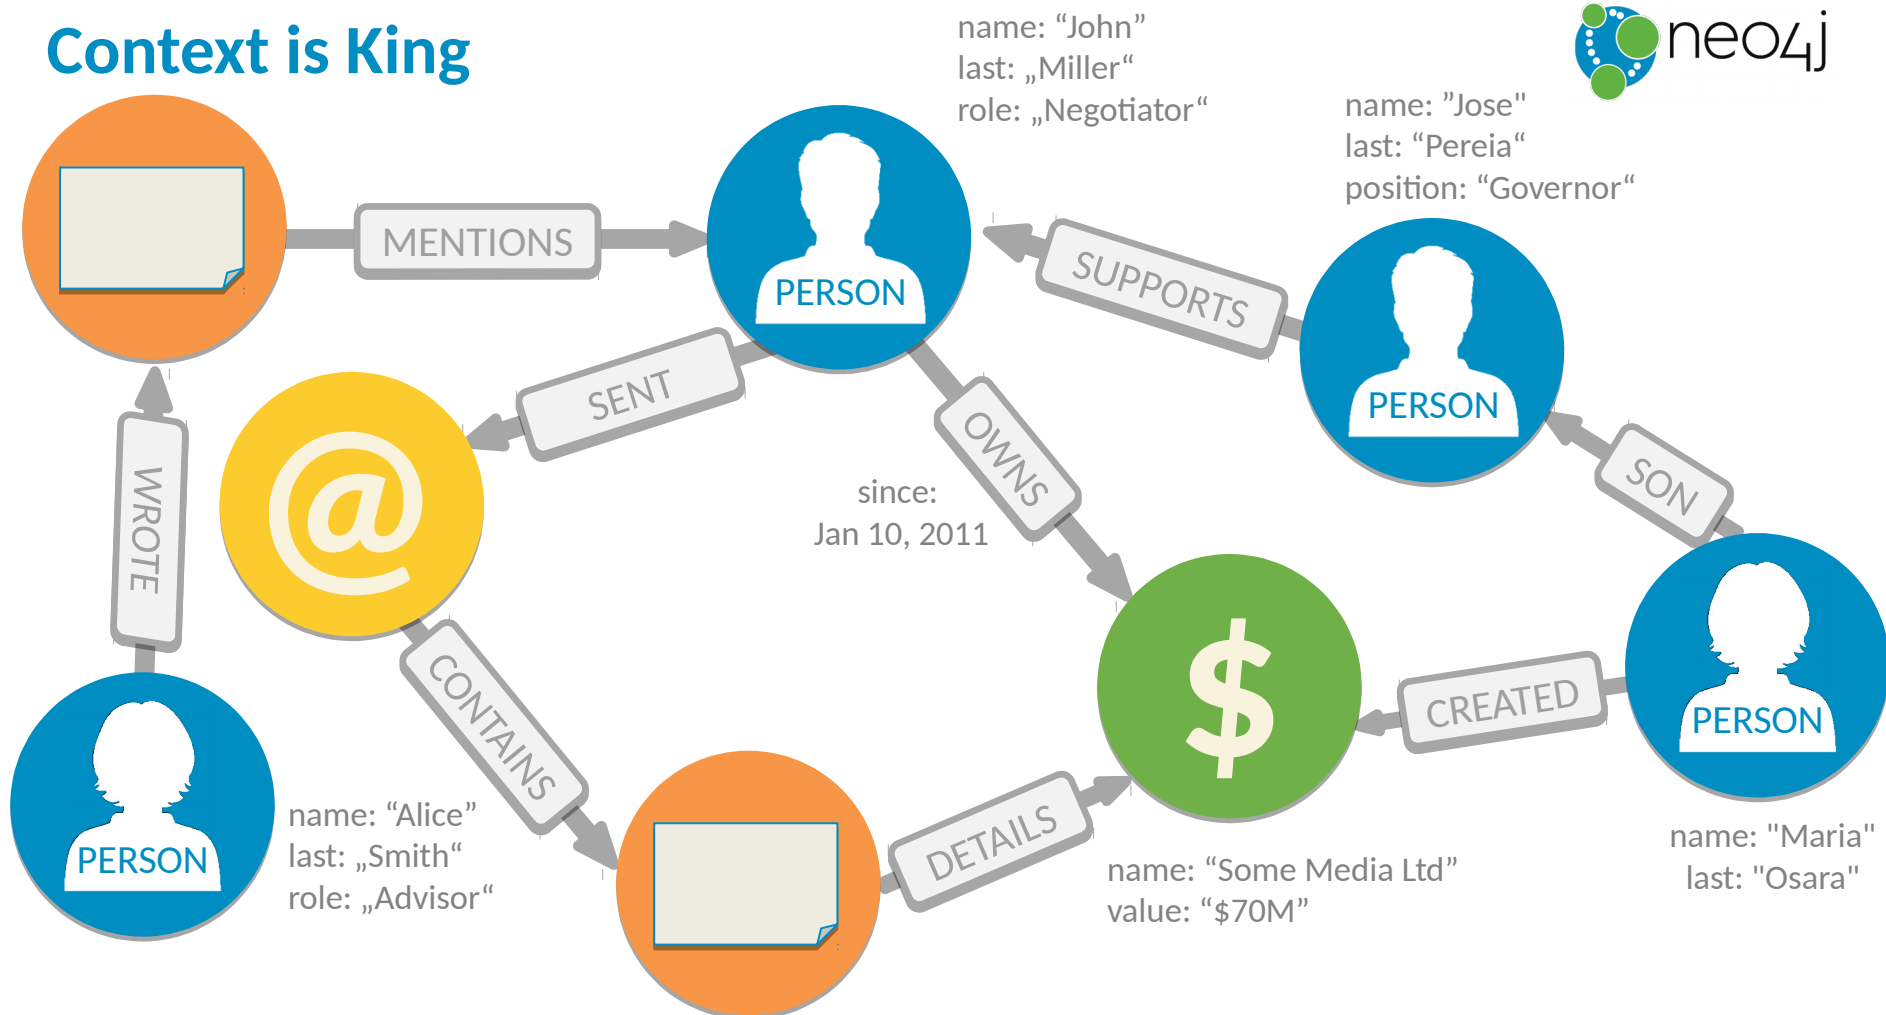


name: "Some Media Ltd"
value: "\$70M"



name: "Maria"
last: "Osara"

Context is King



ICIJ editor Mar Cabra presenting at GraphConnect



Mar Cabra is the Editor of the Data and Research Unit at [the International Consortium of Investigative Journalists](#) (ICIJ), the organization responsible for breaking [the Panama Papers](#) story.

Mar has over 11 years of experience working in data journalism, including the BBC, CNN and the Miami Herald.

At GraphConnect, Mar will be presenting on “**How the ICIJ Used Neo4j to Unravel the Panama Papers.**”



Background

- Global leader in sporting goods industry services firm footwear, apparel, hardware, 14.5 bln sales, 53,000 people
- Multitude of products, markets, media, assets and audiences

Business Problem

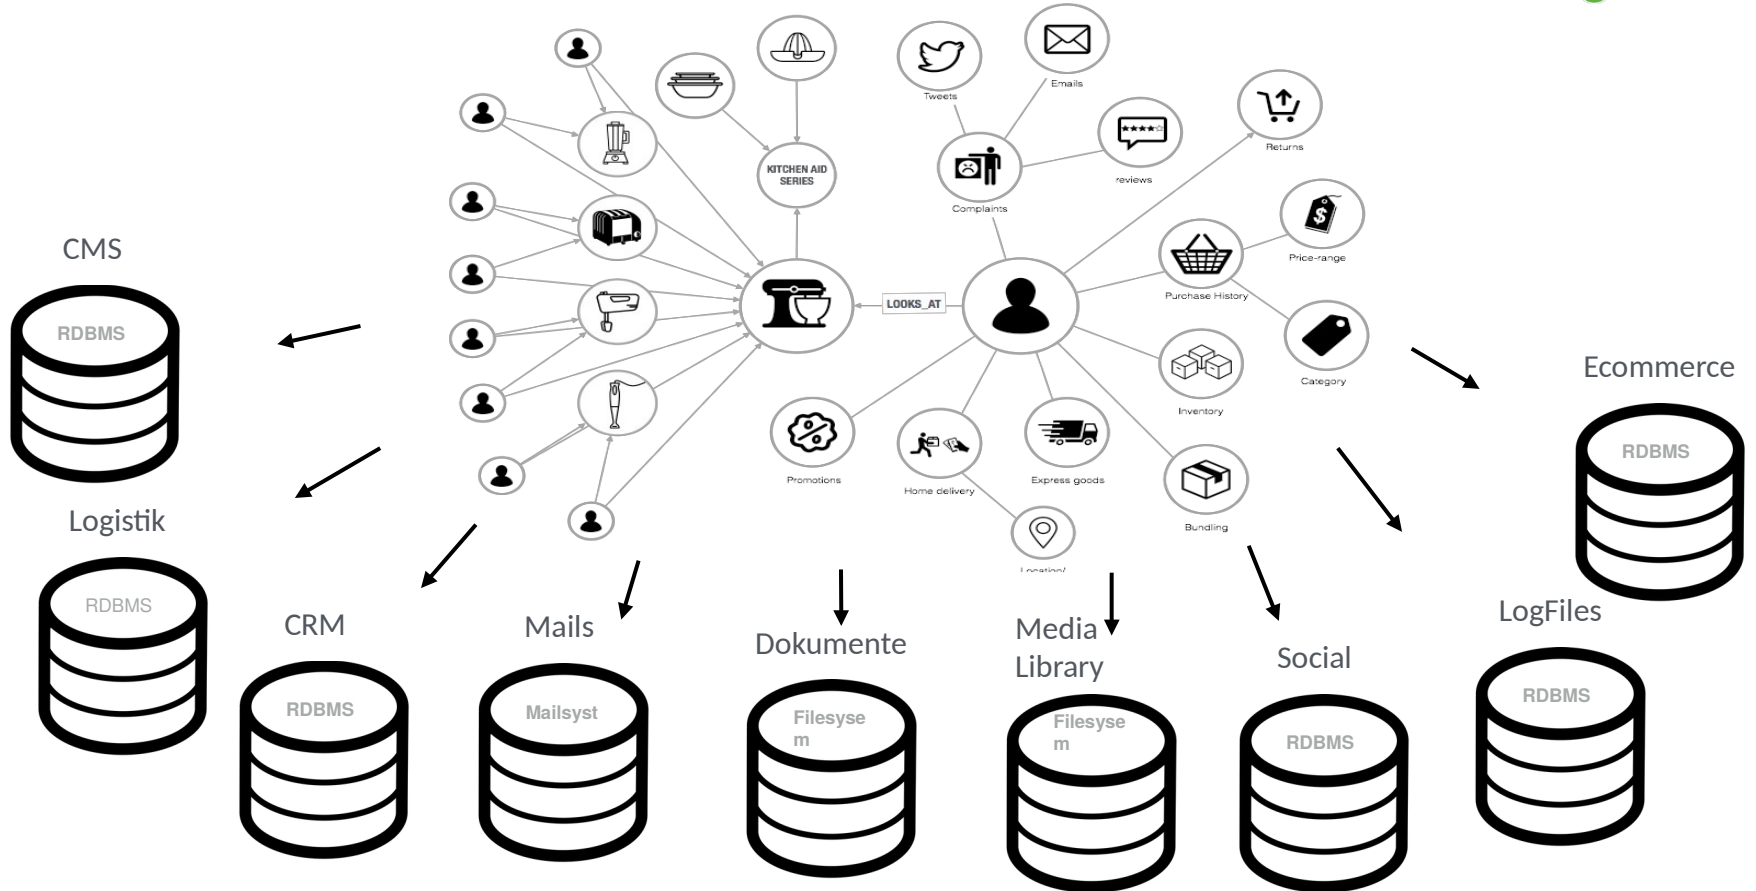
- Beset by a wide array of information silos including data about products, markets, social media, master data, digital assets, brand content and more
- Provide the most compelling and relevant content to consumers
- Offering enhanced recommendations to drive revenue



Solution and Benefits

- Save time and cost through standardized access to content sharing-system with internal teams, partners, IT units, fast, reliable, searchable avoiding redundancy
- Improve customer experience and increase revenue by providing relevant content and recommendations

Metadata-Management ..



[Sammlung Use Cases](#)

[Präsentationen Videos...](#)

[Beispiel-Modelle](#)

Fragen?

stefan.armbruster@neotechnology.com

Adidas: Shared Metadata Service



Delivering the Most Compelling Customer Experience with the Most Relevant Content

INDUSTRY

Retail

GOAL

Increase revenue and customer loyalty by delivering a better online customer experience

CHALLENGE

Data required to provide personalized experience

The adidas Group has a multitude of products, markets, media, assets and audiences to track and target on a daily basis. In order to deliver the most compelling customer experience across these various audiences and markets, the adidas Group uses Neo4j to serve up the most relevant content for each precisely targeted audience – and all in real time.

The Company

The adidas Group is a global leader in the sporting goods industry and one of the world's most valuable brands. With annual sales of €14.5 billion, the adidas Group offers a broad

Background

- One of the world's oldest and largest banks
- 100+ year-old bank with more than 1000 predecessor institutions
- 500,000 employees and contractors
- Needed to manage and visualize ~50,000 Unix servers in its network

Business Problem

- Original RDBMS solution could handle only 5,000 servers
- Improve net performance company-wide
- Leverage M&A legacy systems with no room for error

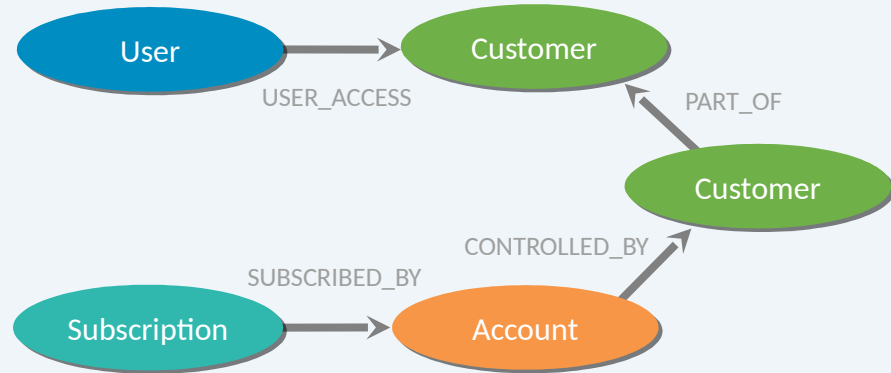
Solution and Benefits

- Store UNIX server and network config in Neo4j
- Combine Splunk log data into an application that visualizes events on the network
- Neo4j vastly improved app performance
- New apps built much faster with Neo4j than SQL



Background

- Oslo-based telcom provider is #1 in Nordic countries and #10 in world
- Online, mission-critical, self-serve system lets users manage subscriptions and plans
- availability and responsiveness is critical to customer satisfaction



Business Problem

- Logins took minutes to retrieve relational access rights
- Massive joins across millions of plans, customers, admins, groups
- Nightly batch production required 9 hours and produced stale data

Solution and Benefits

- Shifted authentication from Sybase to Neo4j
- Moved resource graph to Neo4j
- Replaced batch process with real-time login response measured in milliseconds that delivers real-time data, not yesterday's snapshot
- Mitigated customer retention risks

Background

- Mid-size German insurer founded in 1858
- Project executed by Delvin, a subsidiary of die Bayerische Versicherung and an IT insurance specialist



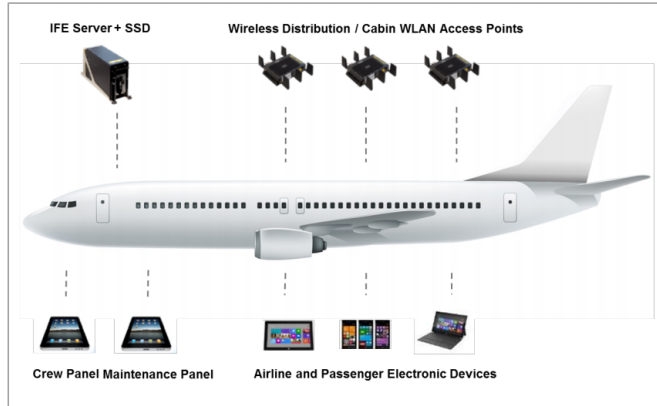
Business Problem

- Field sales needed easy, dynamic, 24/7 access to policies and customer data
- Existing DB2 system unable to meet performance and scaling demands

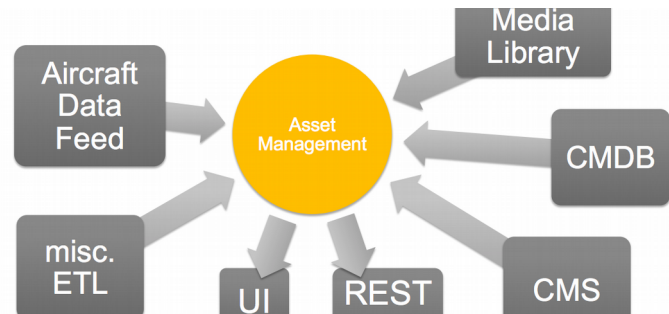
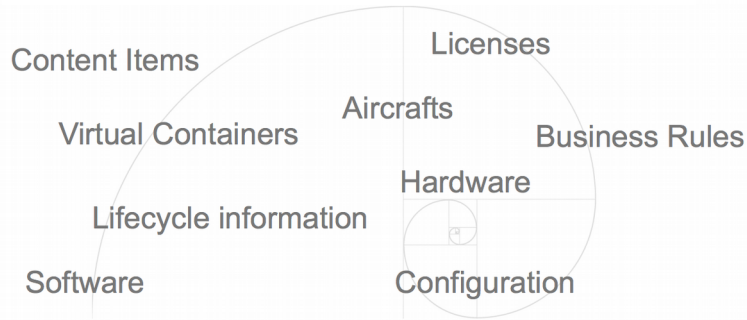
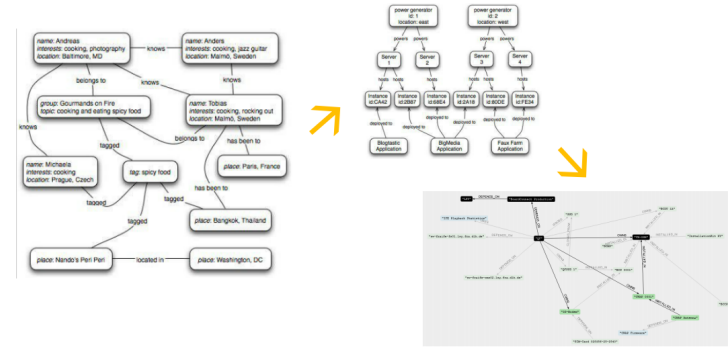
Solution and Benefits

- Enabled flexible searching of policies and associated personal data
- Raised the bar on industry practices
- Delivered high performance and scalability
- Ported existing metadata easily

Lufthansa: Content/Digital Asset Management



Does the problem nature really differ from social networks?

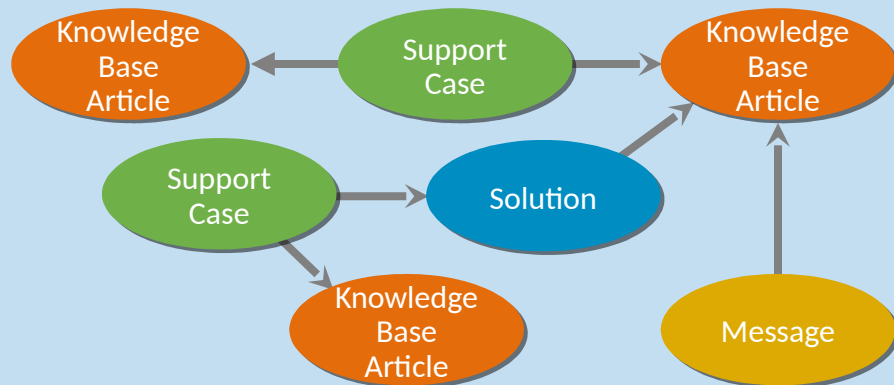


Background

- San Jose-based communications equipment giant ranks #91 in the Global 2000 with \$44B in annual sales
- Needed real-time recommendations to encourage knowledge base use on company's support portal

Business Problem

- Reduce call-center volumes and costs via improved online self-service quality
- Leverage large amounts of knowledge stored in service cases, solutions, articles, forums, etc.
- Reduce resolution times and support costs



Solution and Benefits

- Faster problem resolution for customers and decreased reliance on support teams
- Scrape cases, solutions, articles et al continuously for cross-reference links
- Provide real-time reading recommendations
- Uses Neo4j Enterprise HA cluster

Background

- eBay acquired London-based Shutl bring same-day delivery to London to counter Amazon Prime and to expand its global retail presence
- Founded in 2009, Shutl was the UK leader in same-day delivery with 70% of the market



Business Problem

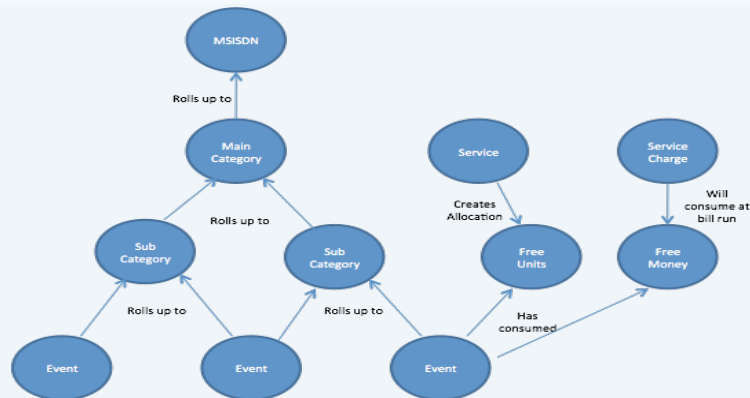
- Enable delivery in London within 90 minutes
- Manage network of routes, carriers and couriers
- Calculate delivery options and times in real time across all possible routes
- Scale to enable a variety of services, including same-day and consumer-to-consumer shipping

Solution and Benefits

- Calculates all possible routes in real time
- Thousands of times faster than MySQL solution
- Queries require up to 100 times less code, improving time-to-market and code quality
- Adding new functionality that was previously impossible

Background

- Part of Hutchison Whampoa, one of the world's largest telecom conglomerates
- Operates in the Nordics and UK
- Moving toward real-time customer profiling and analytics



Business Problem

- New business requirement to give customers more insight into their own usage patterns
- Changing data model was slow and painful
- New queries were difficult to write
- Very large RDBMS data sets creating serious connected query (>L2) performance issues

Solution and Benefits

- Customer-facing apps access Neo4j cluster containing a billing-information graph
- Graph model gives services reps timely and insightful customers profiles
- Much faster query performance
- Faster app and feature development

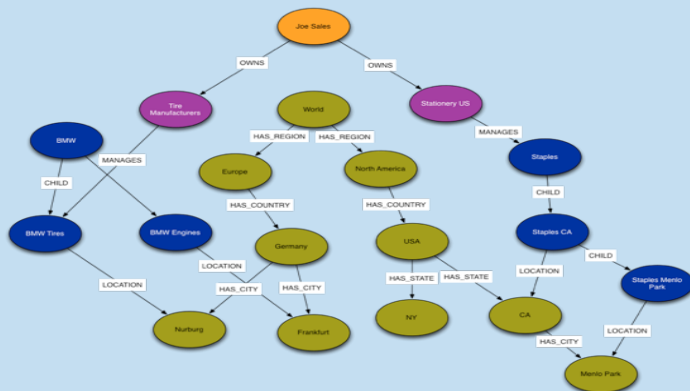


Background

- San Jose-based communications equipment giant ranks #91 in the Global 2000 with \$44B in annual sales
- Needed high-performance system that could provide master-data access services 24x7 to applications company-wide

Business Problem

- Sales compensation system didn't meet needs
- Oracle RAC system had reached its limits
- Inflexible handling of complex organizational hierarchies and mappings
- "Real-time" queries ran for more than a minute
- P1 system must have zero downtime



Solution and Benefits

- New *Hierarchy Management Platform* (HMP) manages master data, rules and access
- Cut access times from minutes to milliseconds
- Graphs provided flexibility for business rules
- Expanded master-data services to include product hierarchies