Building and Sustaining an Enterprise-Wide Semantic Web Ecosystem

Ramanathan Srikumar

Senior Vice President and Head of Industry Solutions

18th September, 2013

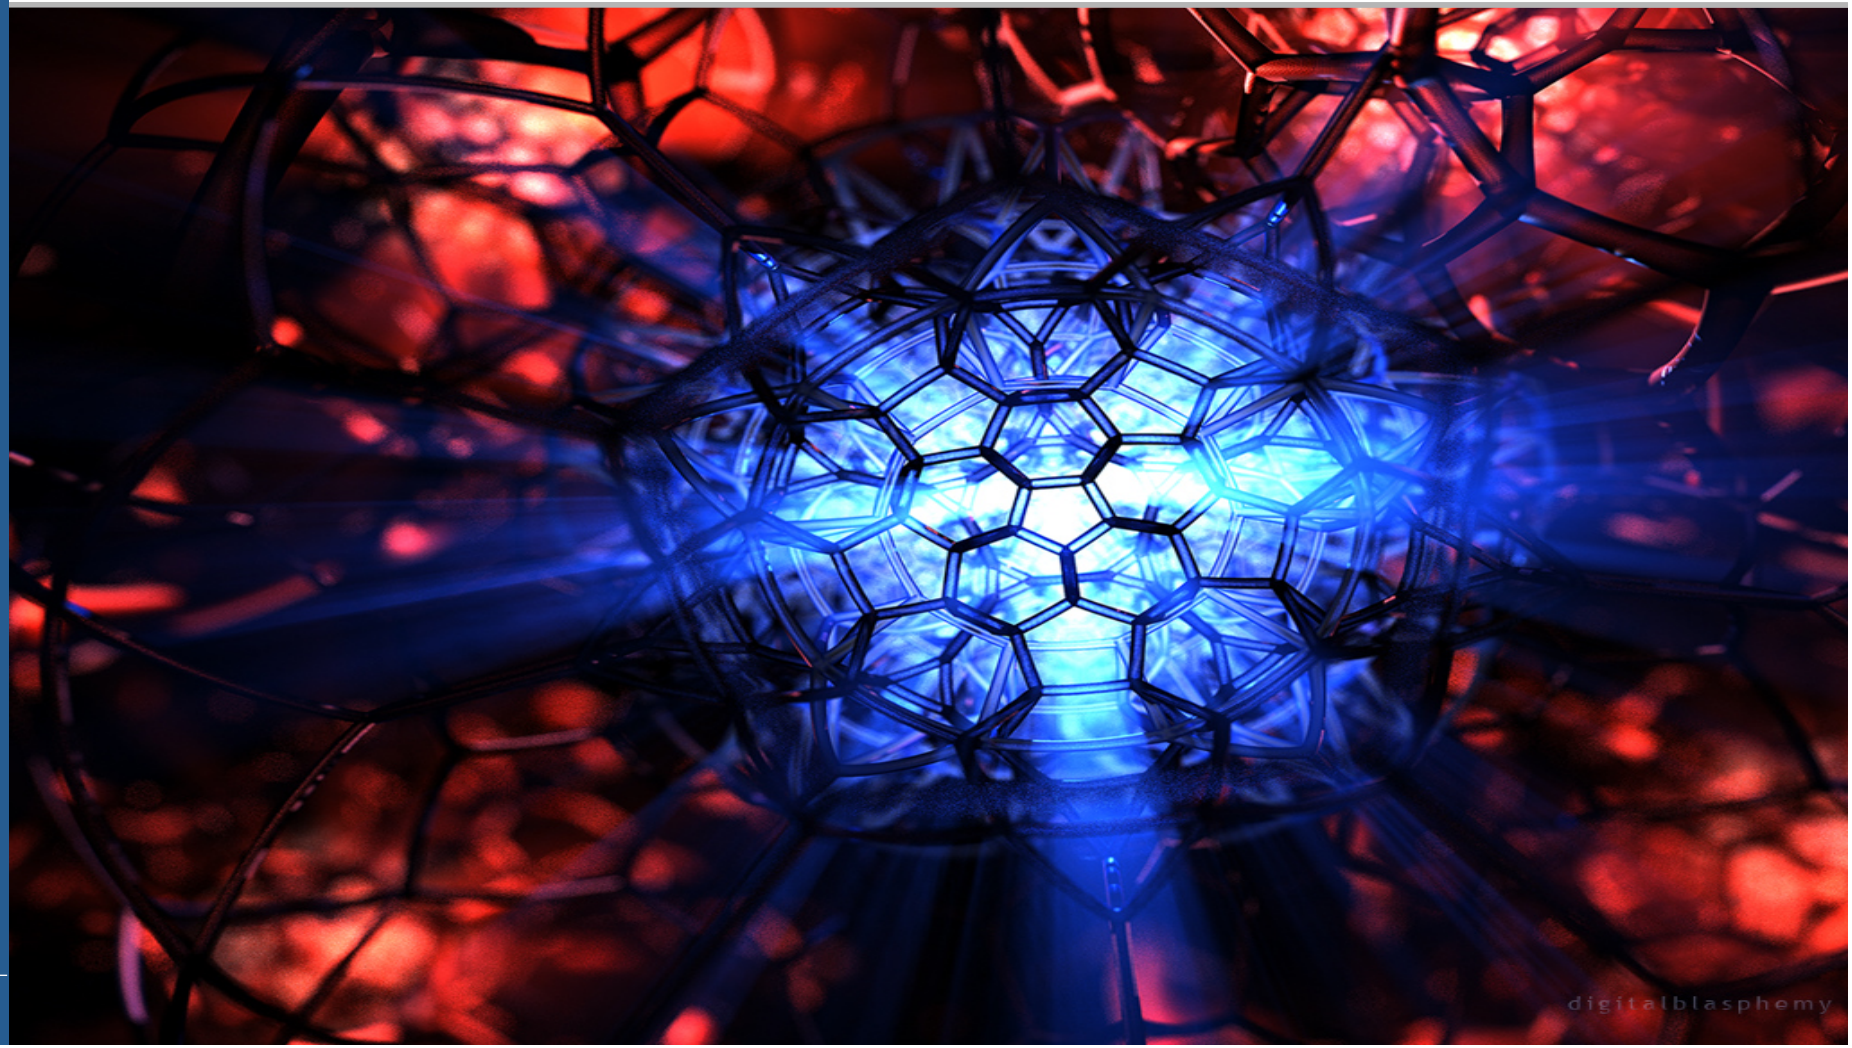

Semantic Model: A data model linked to the real world through a conceptual model

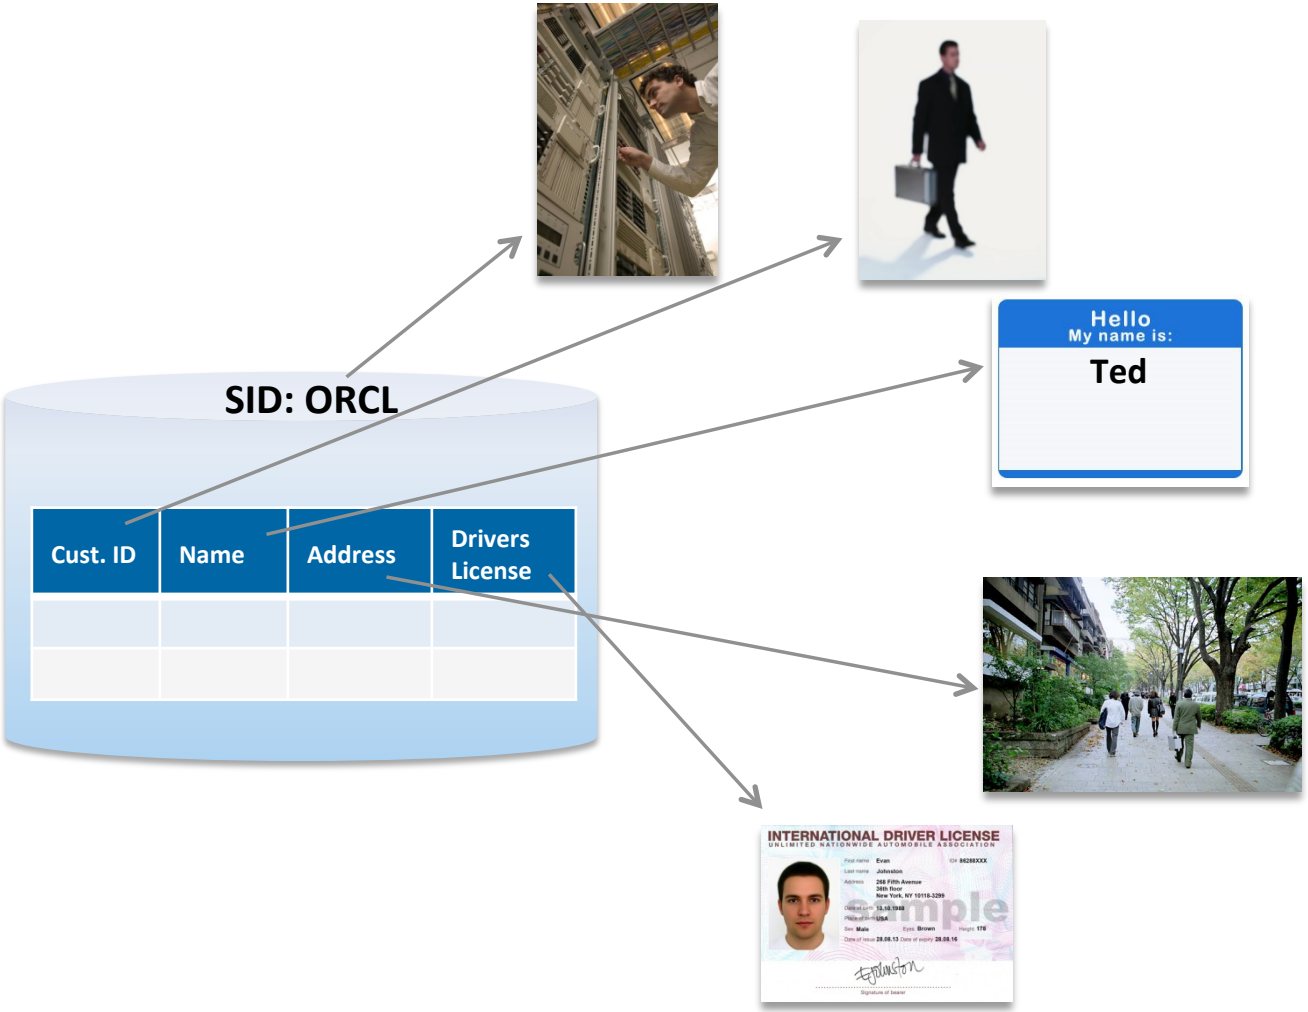

Semantic Model

MPHASIS an HP company

Discreet Fact Sets

John Smith is a customer of the bank with a large investment portfolio. He is a cautious investor with a preference for energy stocks and commodity futures. PB banker is James M

John Smith holds positions in Venture Solar., an energy company, currently headquartered in California, USA

John Smith currently works in Germany and is subject to German tax laws

Venture Solar is corporate customer. Frank K from corporate advisory services is currently helping them put together a reverse merger with Benthik Petroleum .

Benthik Petroleum is authorized to operate in Syria

Unified Semantic Model

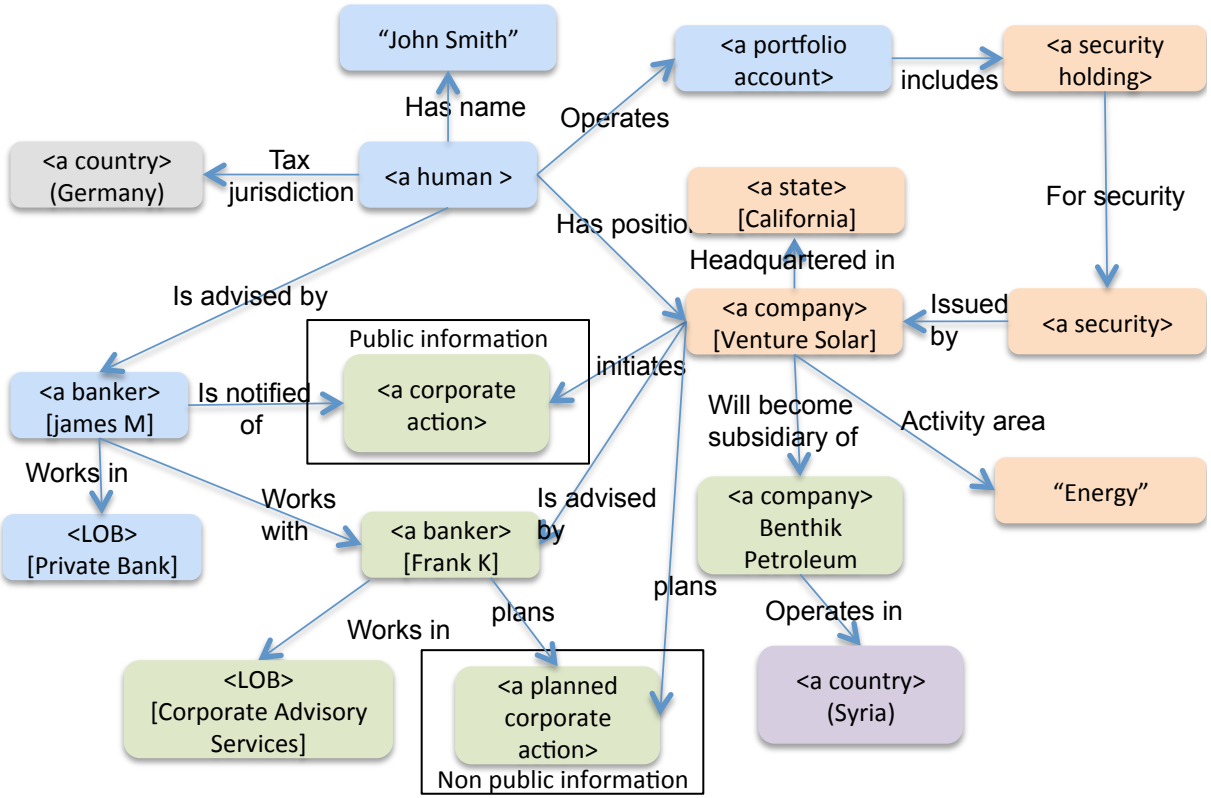

A model of the real world that both humans and computers can understand. Ontologies are formalizations of concepts, using "graphs" to model the highly interconnected nature of facts in the real world

What do we mean by an Enterprise Semantic Web Ecosystem

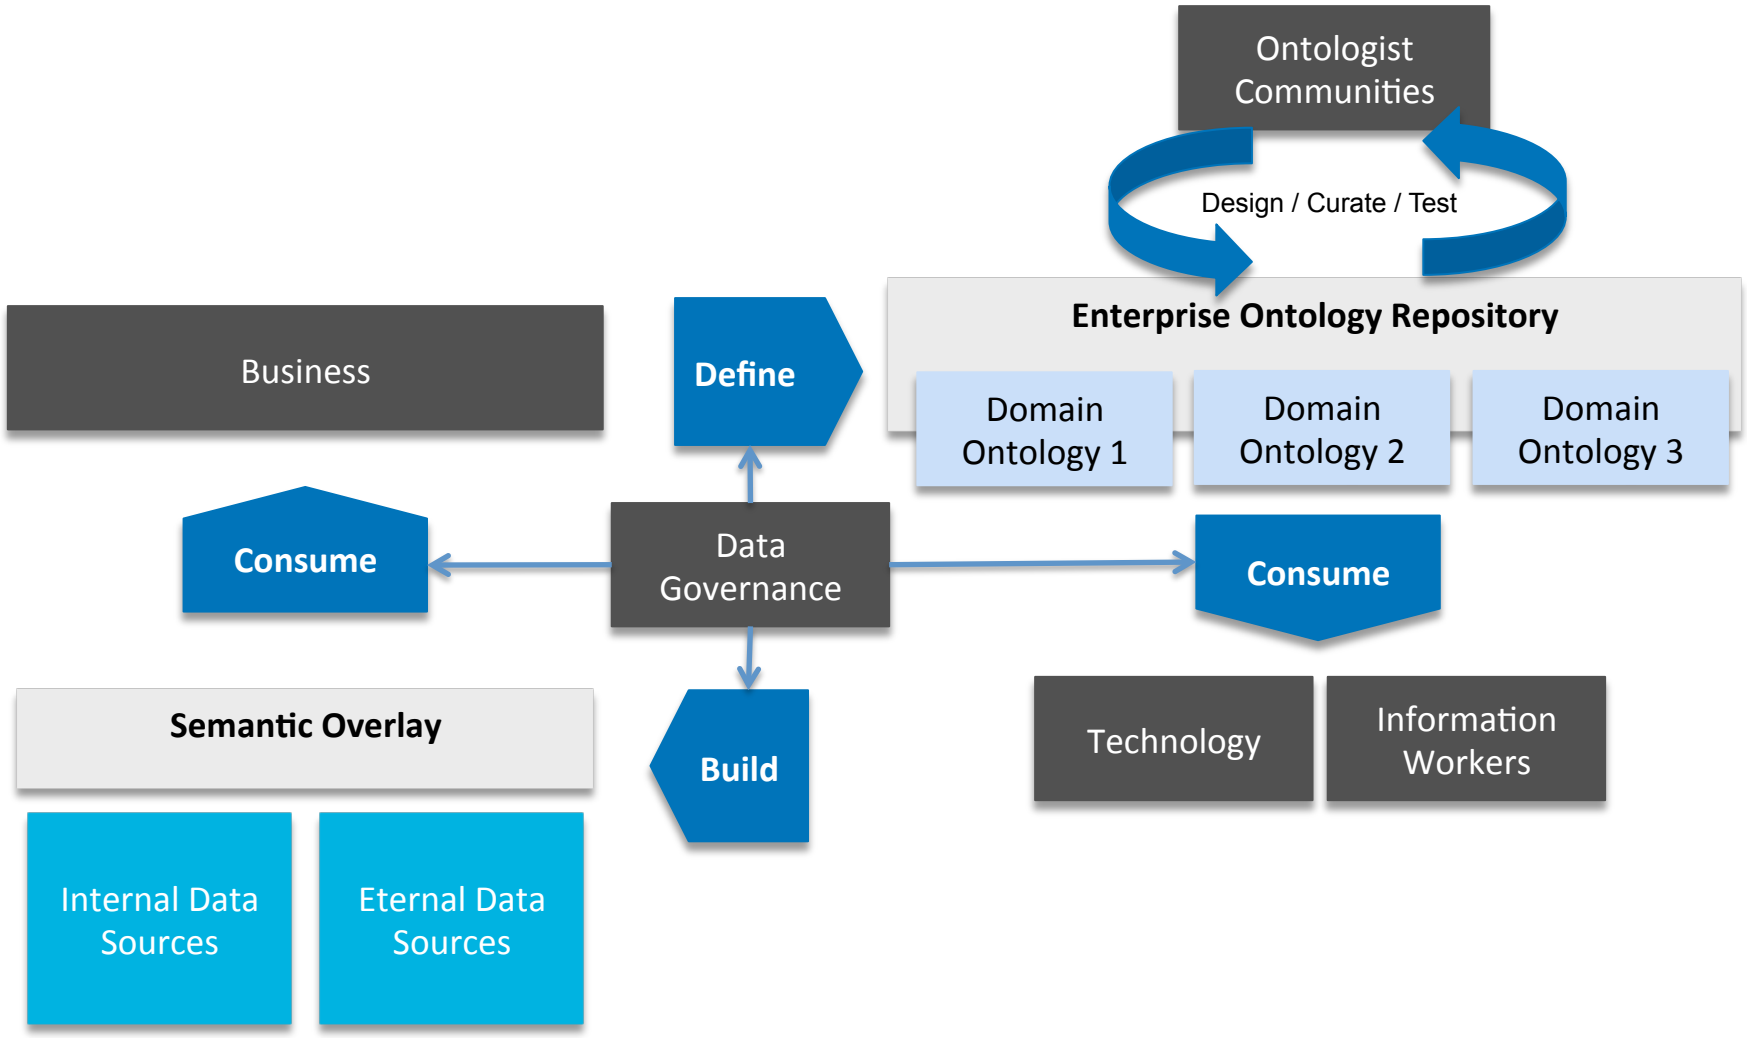

Value of Semantic Models to the Enterprise

	Reduce Build Costs	Reduce operating costs	Reduces risk	Improves innovation	Knowledge Retention & data quality
Precise, shared and enforced vocabulary across the enterprise	✔		✔	✔	✔
Reduced cost to change processes as data models built to the real world do not change often	✔		✔		✔
Eliminate "discovery" in projects, since existing corporate knowledge is encoded	✔			✔	
Breaks down data silos and defeats "data hoarders"		✔	✔	✔	
Gives people the ability to repurpose their own data in completely different contexts		✔		✔	✔
Eliminates the root causes for unofficial data ecosystems such as inflexibility, mismatched definitions or missing relationships across domains		✔	✔	✔	✔
Improve the efficiency of knowledge workers by bringing data from multiple sources together based on the specific context of a situation		✔	✔		✔
Allows for rules, process and data to coexist in one model, making analytics much easier		✔		✔	✔
Exposes what you don't know that you don't know			✔		✔

Industry figures:
60% to 80% reduction in cost of change
10% to 40% increase in productivity of knowledge workers
80% to 90% reduction in "rediscovery" due to loss of corporate data memory

Major Challenges

The People Problem:	Requires the consistent application of common sense … and occasionally very deep philosophy
	Why do software programmers use up shampoo so quickly – because it says "apply, lather, rinse, repeat"
The Open World Assumption:	Three possible answer to a question: Yes, No, I don't know
	Requires a major shift in thinking on the part of designers
The Visibility Problem:	Everyone can see the data (and meta data) Inherently different from the mature models already in place
	Industry solutions are not as good as they need to be
Semantic Modelers vs. Enterprise KPI's:	Purists in semantic modeling have a very different set of objectives, methods and tools from the folks focused on project milestones and deliverables
Incentive Problem	Primary incentives today: Job security, personal productivity and reuse of existing investments
	None of these are healthy for a semantic web initiative
What is old is new again	We were here in the 1970's with graph data stores – and we switched to RDBMS because of performance

Advice and Lesson Learnt

Use Public Ontologies, but don't wait for them	The reason semantic web might succeed: the availability of public ontologies
	The reason enterprise projects might fail: overreliance on public ontologies
Stay with the standards	Stay current with the standards, because that gives you the maximum possibility for interoperability
	Temporarily going off standard becomes going permanently going off standard very quickly
Solve the knowledge problem, not the data problem	If you only want to solve the data problem, other tools may be better
	Solving the knowledge problems solves your data problem as well, but gets you so much more
Do not over think the ontologies	Too many projects fail because of time spent on abstract concepts
	Operational ontologies are more useful to the enterprise, and are sufficient for enterprise wide linked data
	When public ontologies become available, link to them to get the benefits of linked open data

The Rest of the Afternoon

CASE STUDY - Mastering Enterprise Metadata with Semantic Modelling	While there is a lot of work being done on using semantic modelling on enterprise data, modelling the enterprise itself model can yield a much higher value. This meta data model of the enterprise can revolutionize the way applications are designed and built. This case study will explain what an enterprise metadata model means, how it has been used in projects, and some of the key lessons learnt
PANEL DISCUSSION - The Business Case for Financial Data Semantics	The use of data and knowledge inherent in the data is rapidly evolving as the key differentiator between peers in the market. Factors for success include: Managing complexity / inference-based processing / formal ontologies/ knowledge representation / incremental implementation over existing infrastructure This panel will discuss the benefits of semantics in managing data and what it takes to achieve these
PANEL DISCUSSION: Semantic Business Analytics - An Analytics Ecosystem	Today's enterprise data ecosystem is more complex than ever. Challenges include increased regulatory reporting requirements, increased risk & control needs, new political realities and rapidly changing market environments Opportunities include the advent of big data, access to previously unavailable sources of information and an increasingly tech savvy work force. Compares traditional and semantic analytics, and their impact on enterprise ecosystems