Enterprise Semantic Web Solutions

What some commercial buyers of Semantic Web technologies are building

Jeffrey T. Pollock
Vice President Technology, Network Inference
Jeff.Pollock@networkinference.com

Abstract

This paper examines two implementations that highlight the use of OWL and RDF in innovative ways to solve significant business problems. First, a Fortune 500 electronics manufacturer uses an inference-driven system to resolve chart-of-accounts financial reporting problems. Second, a major NATO force uses an inference-driven system to get better visibility into battlefield situations. Taken together, these enterprise production deployments prove the commercial and military viability of OWL, RDF, inference, and, more broadly, the concept of the enterprise semantic web.

1 Enterprise Semantic Web

The enterprise semantic web is a well-scoped, more tightly focused implementation of the Semantic Web vision. It does not attempt to create a virtual database of web content, nor does it rely only on HTTP accessible data. Instead, the enterprise semantic web takes advantage of the inherent power of RDF and OWL to enable more flexible instance data, evolvable schemas, federated ontologies instead of common data models, and the power of inference to generate non-intuitive, yet correct answers to queries.

The enterprise semantic web enables corporations to gain efficiency and reduce operating costs by becoming model-driven. Both OWL and RDF provide for machine-actionable data, logic, and process models that reduce the amount of code that must be written, rewritten, and compiled.

The enterprise semantic web must be scalable to enterprise needs. It should support high transaction volume, be fault tolerant, fail-over ready, clusterable, and support data volumes ranging from millions to hundreds of millions of records. Above all, it should provide for a platform that is highly reliable and predictable.

Network Inference has found that organizations who are adopting the enterprise semantic web seek to reduce their vendor dependence, preserve legacy investments in object-oriented and XML systems, while creating a truly adaptive organization.

2 Adaptive Organizations

The notion of an adaptive organization, sometimes referred to in the industry as an adaptive enterprise, is actually much broader than software systems, or IT in general. The adaptive enterprise is a commitment to a broad range of flexible organizational structures, human resource practices that encourage innovation, business processes designed for agility, as well as corporate information systems that are change resilient and dynamic.

The following scenarios, drawn from actual Network Inference customer projects with a Fortune 500 company and an influential government agency, should serve to highlight how the path to adaptive organizations naturally follows from implementing an enterprise semantic web.

2.1 Chart of Accounts Solution

A Fortune 500 electronics manufacturer had trouble reporting quarterly results to Wall Street. The trouble was that their products were always changing and converging, necessitating different algorithms for determining financial results.

Like most businesses in the electronics domain, they are faced with a perpetual blurring of lines about what their products really are. Consider the home audio domain, today we can find devices that are DVD players, DVD recorders, VHS players, digital video recorders (DVR), and TIVO services all in one box. How would a large company with tens of thousands of products, discrete and converged, report the sales of the converged devices? Wall Street does not change their predefined market segments quickly enough to account for new electronics devices, so a

labor-intensive accounting process usually has to occur.

The old way of doing things involved a complicated series of web applications written in J2EE that required quarterly code updates, testing, and redeployment. Even with this application in place, the data was usually exported to spreadsheets for further analysis by dozens of analysts. Overall, the process took about six weeks with more than 20 people to generate quarterly statements. Because the process was so human-intensive, reporting errors did eventually find their way in quarterly statements to Wall Street.

The new way of doing things is to connect sales databases, operational databases, and financial systems to an RDF data store. Simple conversions take raw ERP data from a variety of places and reduce it to collections of triples. Then, the RDF data is connected to master ontologies written in OWL. The OWL ontologies represent three fundamental concept spaces: (1) product line hierarchies, (2) financial market segments, and (3) business rules for how products are classified. The resulting solution allows business users to "tweak" business rules and/or product line hierarchies and watch how the actual business data sorts out. At runtime, inference engines are used to query data and (re)classify how sales results are bucketed into various market segments.

The resulting solution has taken the overall process from six weeks to one day. Errors have been reduced because humans are no longer required to "fat-finger" data into spreadsheets. The value of automated classification and total visibility truly becomes apparent when financial analysts have the ability to test different reporting scenarios and instantly see impacts to the bottom line.

2.2 Battlespace Awareness

A major NATO military force chose to prototype an Enterprise Semantic Web architecture as part of a major initiative under the umbrella of the Network Centric Warfare doctrine. Net-Centric Warfare means that military operations, in battle and for supply-chain logistics, are deeply connected in ways that improve force projection, flexibility, and coordinated cross-military operations.

The focus of this ESW prototype was on battlespace awareness, also referred to as a common operating picture, which can dynamically show commanders exactly what is happening in battle. The ESW solution focuses on a built-in reasoning capability to discern the important information from the static. In battle, data can be received in realtime from a number of different kinds of sources such as: radar units, GPS devices, intelligence reports (sensor, electronic, and imagery), and other command and control systems. Making sense of this data is crucially important in the field.

The resulting system used RDF as a way to store and manipulate instance data information. Standard RDF query interfaces provided realtime access to live data represented as triples. In many cases, but not all, these RDF triples were linked to OWL ontologies. The OWL information layer provided the ability to make inferences about data without prior knowledge about what kind of data would appear in the system.

Perhaps the single largest advantage that the Enterprise Semantic Web approach offered in this capacity was the ability to get reliable and flexible visibility into a wide range of data – without needing to effect a common data model. Typically, with other information visibility tools and techniques – such as EAI or EII – a common data model, or a series of "view" models are created for application access. The trouble with the common data model approach is two-fold: (1) the politics of creating common models is a significant barrier, and (2) common data models cannot adapt to changing physical schemas.

For Network-Centric Warfare, the ESW offers a way to enable the information supply chain to adapt, reconfigure, and reason about new information in realtime. This prototype used a robust Service-Oriented Architecture (SOA) for the message-level interfaces, but leveraged the Enterprise Semantic Web for its information architecture, thus enabling systems to intelligently fuse information without requiring troublesome common data models.

3 Summary Conclusions

OWL and RDF are viable IT solutions today. There is no need to envisage a holy-grail in order to reap benefits from the technology. We've shown how mission-critical applications can make use of, indeed require, the value provided by RDF, OWL, inference, and the enterprise semantic web concept. The twin values of machine-actionable models of logic and federated ontologies drive tangible business savings, new capabilities, and offer a roadmap towards a truly adaptive organization.