
ACHIEVING SEMANTIC INTEROPERABILITY

EXPLOITING THE FUSION OF SOA AND SEMANTIC WEB CAPABILITIES IN SUPPORT OF THE NETWORK-CENTRIC ENTERPRISE

Executive Summary

Radical improvements for information sharing are nearing readiness for widespread adoption. The fusion of service-oriented architectures (SOA) with semantic Web data specifications means that organizations will soon be able to decouple services from systems and information from software. These new IT capabilities have the potential to save billions of wasted dollars and enable simpler compliance with modern regulatory and security restrictions. If adoption trends and technical advances continue at current rates, 2005 will be a break-out year for the fusion of these two technology bases. Much of the innovation, adoption, and momentum for this technology fusion has been driven by federal efforts with Network-Centric Enterprise Service goals, and commercial organizations stand to benefit greatly from the information capabilities generated by these and other programs. A new era of meaningful service-oriented network computing is taking shape – be prepared.

Motivation for Change

Information outlives software. Since the birth of the Internet in the mid-1990's our ability to create, publish, and disseminate information has exploded. Our information economy operates on no fewer than 200 billion lines of legacy code and our world population has produced over 5 exabytes of content in 2002 alone. We now generate digital information at a rate exceeding half a million new libraries of Congress each year. In short, our ability to create data has finally outstripped our ability to understand it.

Problem

It has been said that people will only change when the fear of not changing is greater than their fear of change itself. Today, with decades of global investment in silos of data with hardwired software embedded to that data's true meaning, people are starting to fear the systemic chaos our IT malpractice has created. With few exceptions, there have been no capabilities for planned data interoperability in future IT systems – nor has there been planned obsolescence. Lack of planning for either data retirement or data longevity means that silos of corporate, federal, community, and societal information cannot be reused in unforeseen ways. Simply put, information cannot evolve.

Pain

Unfortunately this is a hidden pain. It is a systemic problem, without an epicenter. Decision-makers, typically lacking hands-on IT implementation experience, cannot fathom the billions of lines of code, the billions of person-hours of effort, nor the wasted expenditures that go towards making IT systems interoperability possible. NIST, the National Institute of Standards and Technology, estimates that, in the Automotive sector alone, \$5.9 Billion is wasted each year due to imperfect interoperability of manufacturing systems. Extrapolate this to other industries, to

federal government IT systems, to lost opportunity in drug research or biotechnology – we are faced with a problem of immense scale.

Hundreds of billions of wasted dollars and unknown lost opportunities are the reality of data interoperability today. Unfortunately, it is indeed a hidden pain – because everybody shares a small part of the burden. There is no Chief Interoperability Officer with organizational mandate, there are no line items in the CFO's budget tagged to identify data interoperability costs. Instead, the costs are buried in organizational, or departmental IT budgets. They are expressed as 10, 20, or 30% inflation costs on software upgrades, deployments, or maintenance costs. They are the labor an organization needs to support software, or in the budgetary overruns so common for IT projects.

False Hopes

Compounding the problem is the false hope sold to buyers of IT consulting services, enterprise software, and middleware platforms. The tradition of silver-bullet sales techniques is alive and well across the globe. Again and again the purveyors and pundits of technology have over-stated capabilities and under-delivered interoperability. From CORBA to Object-Oriented software; monolithic ERP packages to single customer-view CRM systems; XML document formats to federated data warehouses – the IT buyer has heard it all. Reusability, maintainability, low-cost, future-proofing, and self-healing software are just some of the promises made, and promises broken. It is no wonder that Gartner accepts the “trough of disillusionment” as a natural part of software adoption.

Semantics and Loose-Coupling

The most recent entrant to the mix of fad IT technology is something called the Service-Oriented Architecture. Introduced with fanfare in 2000-2001, Web services promised “dynamic discovery” – the ability to have automated discovery of network application services and corresponding connectivity, collaboration, and business functionality. It was predicated on the notion of loosely-coupled service interfaces – where an intermediary software service identifies the appropriate connections between services based on need and context. This loose-coupling of service connections was not revolutionary or break-through in and of itself, rather, the fact that it was intended to be done in a standardized manner generated much buzz. That excitement quickly turned to doubt, and expectations were quickly reset as reality set in.

Reality dictated that capturing what software services did, and what their data meant, was impossible with the technology at hand for Web services. XML does not capture data meaning, nor can simple taxonomies and registries tell software what services are supposed to do. Web services need semantics to fulfill their original vision, hopes, and promises.

Service-Oriented Architectures

It would be unwise to be dismissive of Web services inherent value, even without semantics. The capabilities provided by the SOAP (Simple Object Access Protocol) and WSDL (Web Services Description Language) are helping to enforce good software practices such as: coarse-grained interfaces, contract-based development, standard syntactic and protocol definitions, and common payload structures. These technical patterns are already gaining wide adoption and fostering business value in

terms of cost savings and IT flexibility. Service-Oriented Architectures (SOA) – as a pattern of deployment for Web services – are not a passing fad, they will be the foundation upon which future IT infrastructures are built to.

However, among the plethora of small problems being addressed by the Web services communities, one major problem looms, which is not being addressed by the mainstream – SOA, Web services has no inherent data strategy.

Simply put, creating a Web service is just like creating a new data silo. Behaviorally, each Web service is a silo that requires manual efforts, programming in some language (.NET, Java, XSL/T), or deployment of a proprietary tool, to interpret, manipulate, or otherwise make use of it. Some would point to XML canonical exchange formats as a Web services data strategy, but the argument doesn't hold – canonical formats are community agreed upon terminologies for data meaning and process sequencing. At best this is akin to an EDI style technical architecture, at worst it is simply a data dictionary for XML. Neither metaphor is particularly attractive since the issues with EDI (misuse of “agreed” formats) and data dictionaries (usually wrong) are widely known to practitioners.

Semantic Web Technologies

Seemingly orthogonal to SOA's, a class of technology has emerged from the World-Wide-Web Consortium (W3C) that is squarely focused on data – some would say, information. In fact, Sir Tim Berners-Lee, Director of the W3C and universally acknowledged father of the Web, underwent extreme scrutiny in 2001 when he asked the W3C to give equal attention to the Semantic Web as it did for Web services and Service-Oriented Architectures. But he saw then, what many now see: that a more dynamic, flexible, and reconfigurable format for IT data/information was required. Without something new, global IT practitioners would be forced to rely upon existing data structures such as XML, relational database formats, and object-oriented systems. What is the problem with that, you ask?

Rigidity and ambiguity are the defining characteristics of XML, relational models, and OO systems. For IT systems to scale, for digital information to outlive the software it was created in, for mere mortals to find needles in the global digital haystack, rigidity and ambiguity must go.

In February 2004 the W3C officially recommended two new specifications: the Resource Description Framework (RDF) and the Web Ontology Language (OWL). RDF is a graph-based data format that enables a schema-less, evolvable structure for persisting data. OWL is a logic-based data format for describing the meaning of terms in relation to other terms. In practice, OWL acts as a highly evolved schema for describing the meaning of IT information, while RDF provides a flexible representation system to manipulate it and associate it with legacy data. Neither RDF or OWL are intended to supplant relational systems, XML, or OO systems – they enable a separate metadata layer than can describe, and connect, the multifaceted meanings of business information, process, and rules to more rigid and ambiguous existing legacy data formats.

Fear, Uncertainty, and Doubt

With all the false hope proliferated in past decades, there is more than a little skepticism about the ability of semantic technology to really add substantial value.

People fear the new OWL and RDF data languages because they are new and because they have unfamiliar new attributes. People are uncertain about how the Semantic Web architecture fits together, and more importantly, how it fits with all the legacy data already in place. People doubt that this new technology base is really different and that it can actually provide benefits that they cannot readily get elsewhere using tried and true technology.

This fear, uncertainty, and doubt (FUD) will persist for some time. Existing large vendors will attempt to spread FUD to IT buyers because if they accept the premise of the data semantics argument – that existing technology bases are insufficient for substantial new progress – they must essentially admit that their current product lines have minimal impact on the broader systemic issues at hand. Likewise, conservative IT buyers will shy away from adoption because they must accept the premise that their existing information strategies are potentially flawed. Taken together, these are classic signs of a disruptive new technology force, with early adopters driving innovation and developing “chasm-crossing” improvements that ease the pain of change and simplify adoption for late-movers.

The Semantic Wave

Despite the many hurdles of FUD and early adoption, the industry is moving forward at a rapid pace. The proliferation of these new semantic technologies, like OWL and RDF, has been gaining momentum. Indeed, a semantic wave is building that will fundamentally alter how digital information is captured, manipulated and disseminated over the lifetime of the software that makes use of it. The following section takes a snapshot at how industry is advancing with semantic technologies during late 2004.

Industry Adoption of Semantic Technologies

By all signs, the infusion of machine-interpretable semantics into traditional IT systems is an inevitability which is occurring at a rapid pace. A recent analyst report identified more than 50 software vendors actively adopting semantic-rich metadata strategies for their product offerings. Many recent books on the Semantic Web have identified dozens of commercial use cases where Global 2000 organizations are adopting semantic technologies. Many US Federal agencies and user groups are driving adoption within the government sector for semantic technologies. Perhaps most telling for this trend, every significant IT standards body has announced plans to align data models with emerging specifications for semantically rich metadata and ontology schema formats.

Commercial Adoption

The Global 2000 are researching and adopting the semantic technology base in droves. Organizations such as Audi, General Motors, Cisco Systems, Fed-Ex, France Telecom, Nokia, Fujitsu, Hitachi, NEC, Unisys, Pfizer and others have all pursued research or adoption strategies for a semantics-aware technology base. The applicability of these commercial use cases has proven immense value in areas such as financial reporting, drug discovery, product lifecycle management, logistics and supply chain, as well as various consumer-oriented applications. TopQuadrant Research estimates that, based on known use cases, companies can expect to reduce labor hours by as much as 90%, reduce overall cost of ownership by up to 80%, and realize up to a 300x positive ROI over 3 years.

Federal Adoption

Government agencies within the US have created, funded, and motivated adoption of semantic technologies for many years. Historically, research agencies such as DARPA have funded the creation of semantic standards including OWL, RDF and evolving standards OWL-S, SWRL. Today, there is a community within the US Government called the Semantic Interoperability Community of Practice (SICoP) which is an active promoter of semantics-aware solutions. Agencies such as DISA, Navy, Army, DLA, EPA, and NCI have all publicly announced programs that have leveraged semantic approaches toward data fusion, integration, logistics management, and registry/repositories. Later in this whitepaper, the Department of Defense's ambitious Net-Centric Enterprise Services plan will be evaluated in the context of SOA and semantic technology fusion.

Vendor and System Integrator Adoption

Over 50 software vendors world-wide and nearly all major systems integrators have publicly acknowledged efforts – research or product support – in the areas of data semantics. Major software vendors such as Oracle, IBM, Microsoft, Adobe, and Hewlett Packard have all publicly announced product improvements and/or significant public research efforts in the semantics space – most of them using W3C Semantic Web specifications. System integrators such as IBM, CGEY, BearingPoint, SAIC, and Lockheed Martin have all leveraged semantic technologies in various programs on research or customer deployment projects. There is clearly a growing base of COTS support and trained engineers who know how to leverage semantics to produce new capabilities and cost savings.

Standards Body Adoption

Every significant standards organization has aligned themselves around semantics as the path toward interoperability. First, and most importantly, the W3C – who also manages most of the relevant internet and SOA specifications – has formally recommended the Web Ontology Language (OWL) and the Resource Description Framework (RDF) as global standards. These represent the de-facto standards for ontologies and metadata. Likewise, the Object Management Group (OMG) have begun to align their Model-Driven Architecture (MDA) vision with W3C's OWL and further plan to leverage the UML as a visual modeling interface for OWL. OASIS, the Organization for the Advancement of Structured Information Standards, has begun infusing OWL into a number of their standards including ebXML, UDDI and DCML (Data Center Markup Language), they also proliferate some of their own semantic standards such as Topic Maps. ISO, the International Standards Organization, is also working with various semantic standards including KIF (Knowledge Interchange Format) and SCL (Simple Common Logic) with plans to branch out in more directions as well (UDEF). Clearly the trends are for all IT standards bodies to begin the work to embed formal semantics in everything they provide/ratify and also to leverage the W3C OWL/RDF specifications to the greatest extent possible.

In summary, with the vast amount of work being put into semantics-based technology solutions across commercial organizations, federal agencies, vendors, system integrators, and global standards bodies – it should be apparent that overall maturity has reached critical mass for widespread exploitation of the technology base.

Capability and Value Drivers

The global investment to date in semantic technologies, about one billion dollars since 2000, is not just because academics, logicians, and researchers need something to do. Rather, the promise of new information capabilities and business value have prompted this investment. With new ROI and TCO studies in place, the early signs are supportive that the technology will pay off. Simple calculations on the maintenance of middleware environments, costs of data federation, and development of reusable business rules show that there is tremendous potential to redefine value metrics by orders of magnitude.

While the broader Semantic Web vision will likely contain broad societal impacts derived from information transparency at a global scale, here at Network Inference we are focusing on narrower goals at the corporate level. This “Enterprise” Semantic Web promises new information capabilities and return on investment through improved efficiencies. Some of the more important capabilities Network Inference customers are achieving include the following value drivers.

Regulatory Policy Application

Today’s IT emphasis on regulatory controls is so pervasive as to be cliché. Every software vendor in the market has a Sarbanes Oxley proposition, a HIPPA story, and sophisticated audit capabilities. Unfortunately, the ability of existing technologies to provide these regulatory features is limited to their own sphere of control. Said another way, most regulatory approaches are silos unto themselves.

The fundamental problem with implementing effective regulatory controls is that they must take into account all pertinent organizational information with exceptional reliability and deterministic results. OWL and RDF provide the most viable infrastructure to solve this broader systemic issue. These technologies enable (a) the normalization of any data into a flexible format and (b) the application of deterministic business policies to that very same data. Using OWL/RDF for applied regulatory policies allows IT workers to describe domain knowledge, business rules, and information constraints in the same metadata layer – while linking it directly to existing legacy data.

Security Policy Application

As with the regulatory environment, security, permissions, and views on corporate data and identity have taken on increased importance in recent years. OWL/RDF enable IT architects to gain control over fine-grained information views and permissions on data that is sourced across many business silos. This capability is especially important in the SOA and Web services world because current security approaches are usually limited to interface, application, and protocol security – leaving the fine-grained data security issue untouched for the most part. In tools where data-level security is offered, it is managed during the transformation process – which is unwieldy and exceedingly complex to manage when the scope of an SOA surpasses just a few services. OWL/RDF provide a way to manage secure data views from within a standardized metadata framework that relies less on transformation code and more on query/inference capabilities.

Decentralized, but Unified Visibility

CIO’s have been under immense pressure for years. Paradoxically, their internal customers want more local control over systems and information, yet also want more

efficient and effective global visibility into enterprise resources. This push-pull scenario is a driving force for much of the infrastructure spending going on in IT today. As a horizontal and internal service provider to an organization, the office of the CIO needs to be responsive and keep costs under control. OWL/RDF provide the ideal solution to this paradox because of their defining inherent characteristic: a webbed network of data meaning expressed independently of physical location. Used to augment traditional infrastructure services (middleware, portals, data warehouses, and application servers), OWL/RDF provides a mechanism to enable local control over information – yet ensure global understandability and visibility for information consumers across different organizations.

Reactive IT Infrastructures

Another CIO concern for years has been in the area of increasing infrastructure maintenance costs. More and more of the CIO budget is being consumed by the fixed costs of change management and maintenance so that existing systems can keep pace with business events. As a consequence, more and more software is geared towards providing agile and adaptive capabilities to minimize the impact of change. On a point basis, large gains have been made. But for the overall system, information, data, process, and business rules are no more interoperable than they were five or ten years ago. OWL/RDF have been noted by many to be a key enabler of semantic interoperability – which speaks directly towards an ability to manage change more easily. If future IT systems have built-in interoperability of data meaning, the pain and costs associated with effecting change to organizational information will significantly drop over time. This can make what is today an exponential problem of costs and effort more linear over time.

Automated Service Integration

A significant capability that is targeted, but not yet realized, is for the automatic discovery, negotiation, and collaboration of enterprise services. Still more of a vision than reality, many industry groups are focused on improving the directory and interface descriptions of Web services so that service brokers may perform complete transactions without any prior coding or specialized programming efforts. This overall vision is part of the industry movement towards composite applications. A composite application makes use of information and business rules that may have originated from legacy applications. Today, the vision of building composite applications is accomplished with a great deal of manual effort. Tomorrow, with the advent of “Semantic Web Services,” composite applications and portals may be capable of recovering disparate information in a highly automated manner. Important work with OWL-S, WSML, and other business process execution languages (like BPEL) indicates that this could be a reality as soon as 2006.

Network Centric Enterprise Services

In a dramatic turn of events from the 1990’s, the federal and military IT sectors are a driving force in the commercialization of innovative technologies. At least three key forces are at play here: (1) the military has a mandate to modernize, (2) the military has significant budget to drive innovation, and (3) the military has incentives to buy COTS – commercial off-the-shelf software. The combination of these three forces means that more commercial research and development is being funded to fulfill federal IT needs. Sometimes forgotten, the governments and military are among the largest organizations in the world – many times larger than some of the largest Fortune 500 companies. Fundamentally, the issues of regulatory, security, visibility,

flexibility, and automation for IT information are as large and complex as any in the commercial world.

One response to this from the US Department of Defense (DoD) and the Defense Information Services Agency (DISA) has been to envision a Service-Oriented Architecture of reconfigurable IT systems that can be repurposed dynamically and flexibly to meet ever-changing mission objectives. The similarity to commercial requirements is not hard to grasp. DISA's Net-Centric Enterprise Services (NCES) program was launched in support of the Global Information Grid (GIG) initiative. The vision is that various communities of interest can access secure, reliable, and relevant information for support of various federal and DoD missions. The foundation is grounded upon an SOA and adds key service elements to support network operations.

Revolution in Military Affairs

Historically, the emergence of NCES was spawned from DoD-wide Revolution in Military Affairs (RMA) strategy undertaken by the US Pentagon planners in the 1990's and still underway today. Andrew Krepinevich, in "Calvary to Computer, The Pattern of Military Revolutions" describes RMA as, "what occurs when the application of new technologies into a significant number of military systems combines with innovative operational concepts and organizational adaptation in a way that fundamentally alters the characteristics and conduct of conflict. It does so by producing a dramatic increase--often an order of magnitude or greater--in the combat potential and military effectiveness of armed forces." The effects of RMA can be felt in the military in areas ranging from stealth technology to supply-chain logistics planning to wearable computers for the warfighter. NCES is only one part of a many faceted RMA initiative that will likely continue for decades to come.

Ideas Seeking Delivery

Net-Centric Enterprise Services are built around the notion that an SOA can provide the pluggable backbone for IT systems (battlefield, intelligence, command and control, etc) to interchange information in a flexible and dynamic way. The NCES vision is that 9 key "service areas" – or grid functionality – will be embedded within the network for use by any agency or command that is connected. The following areas are considered key service areas by most:

- **Collaboration Service** – a service that can manage complex and long-lived collaborations among other services to accomplish some objective output.
- **Security Service** – a service that can authorize and authenticate users of NCES services.
- **Mediation Service** – a service that can interpret and repurpose information such that information can be utilized in services that it did not originate.
- **Directory Service** – a service that will enable other services to lookup and discover available services that may fulfill some requirement for processing.
- **Repository Service** – a service that will store/persist metadata that describes information resources on the network in a way that is easily consumable and interpretable.
- **Messaging Service** – a service that provides queuing capability for guaranteed delivery of messages.

- **Management Service** – a service that provides a range of management features for other NCES services such as quality of service, including timeliness, bandwidth, robustness, reliability, etc.
- **User Assistant Service** – a portal-like service that allows human users to interact with available network services for information visibility on the network, there may be a provision for Google-like learned capabilities that can evolve usage scenarios over time.
- **Application Services** – application services are generally meant to include any system or subsystem that is a participant of the GIG/NCES infrastructure, typically these would be legacy environments that have had Web services wrappers put in place.

As of late 2004 the NCES vision is still just that, a vision. Agency's ranging from DISA to the Navy, the Army to the Air Force, have all implemented IT systems that claim some compliance to the NCES vision. However, there is no "gold standard" to date and there is no widespread agreement as to the low-level technical capabilities or specifications of NCES service instantiations. Today, NCES is still an idea seeking delivery.

At the Tipping Point

The advent of semantic technologies has caused a rethinking, by many, of the NCES service areas. In particular, the arrival of commercially available tools that make use of RDF/OWL for intelligent information possessing has opened up new avenues to NCES planners. Inferencing and business rule engines can dramatically improve the information capabilities of embedded NCES services like mediation, directory, repository, and collaboration. In fact, much of the DARPA funded work from the late 1990's and 2000 has directly translated into agent-based information capabilities targeted by the NCES vision.

Some would say that the NCES vision has languished for a few years, as technologists struggled to determine how such advanced capabilities could be deployed using only XML-based technologies. Today, with the OWL/RDF standards in place, widening vendor support for the standard, and new tools (like inference engines) available for use within NCES services, we may be at a tipping point of practicability for the original NCES vision.

But just how practical are these new semantic technologies for NCES use?

Enterprise Scalability of W3C Semantic Web Approaches

Initially a source of concern for many, work during 2004 on several fronts (academic, commercial and federal) have proven the ability of W3C Semantic Web approaches to scale to enterprise needs. However, it is readily apparent that more work has to be done – end-to-end scalability is quite good, but still trails that of more hard-wired IT solutions. Recent gains make the technology viable for enterprise demands in both horizontal and vertical scaling scenarios.

Horizontal Infrastructure Scalability

Horizontal scalability is the ability of an infrastructure to handle large numbers of sub-systems and system-to-system transactions. The first battle line of horizontal scalability is the tried-and-true local-director, round-robin, style distribution of transactions. Since the semantic services are just SOA services in the traditional

sense, local directors can load-balance based on demand with the same simplicity that they do so with application servers in traditional IT environment. At a sub-system level, several information broker vendors (Network Inference and Racer Systems) support the instantiation of multiple inference engines per broker – enabling a cluster-style management of multi-threaded transactions as load increases to any given broker service. Taken together, these are known scaling techniques that can scale virtually unimpeded to massive numbers of transactions.

Vertical Infrastructure Scalability

Vertical scalability is the ability of a given component or service to handle large quantities of data or ontology concepts in a single transaction. Today's commercial-grade inference engines have proven the ability to handle ontologies of several tens of thousands of concepts in an inference engine, while only adding an additional 30-60 milliseconds to the overall transaction. With regard to physical data handling, as opposed to models, commercial-grade RDF databases have proven the ability to handle billions of instances while maintaining sub-second query retrieval times. Areas of need for further work in the area of vertical scaling include high-entailment inferencing on data instances, optimizations for certain ontology modeling styles, and write/delete operations on large numbers of RDF instances.

Taken together, the ability of COTS semantic services to scale to large numbers of transactions, sub-systems, ontology concepts, and data instances in commercial-grade systems with audit, logging, management, and security features – the scalability of W3C Semantic Web approaches is no longer a barrier to adoption.

Integration of SOA and Semantic Web Specifications

The integration of SOA and Semantic Web should be evaluated upon two axis; (1) the information or logical architecture of the data that is leveraged to fulfill mission goals and, (2) the messaging or exchange architecture that orchestrates and discovers various services embedded within the system of systems.

This distinction is important from a technical standpoint because IT integrators have often tightly-coupled enterprise information with the software that processes it, typically resulting in brittle, inflexible, and expensive to maintain IT systems. In contrast, separating the information architecture from the messaging architecture in loosely-coupled approach produces a highly-evolvable framework that can withstand change in a highly resilient manner.

The distinction is also important from an exploitation standpoint because the technology base to support the information architecture within an enterprise is far more mature and ready for immediate exploitation. Leveraging a well-scoped Semantic Web implementation – the Enterprise Semantic Web – is achievable today using approved standards (OWL, RDF) and commercial off-the-shelf (COTS) software.

In contrast, the fusion of Semantic Web specifications with the messaging architectures of an SOA is well underway, but un-standardized and un-implemented by commercial software vendors. Details regarding the maturity of both the information and messaging architectures are discussed in the following two sections.

Information/Logical Architecture

The primary question to be answered with the information architecture of an SOA is: how shall the content, or payload, of an SOA exchange environment be represented, transmitted, transformed, queried, and secured?

From its inception, SOA has relied on the use of XML specifications for expressing content and payload in the messaging framework. In fact, the earliest implementations, and many of the modern ones, use point-to-point techniques for handling the information architectures. These rudimentary P2P approaches are successful in SOA environments with few services and data structures, but quickly become exponentially more difficult to control as complexity emerges. Technically, this approach relies on end-point transformation routines expressed in Java, .NET, or XSL/T programs. These transformation routines end up containing some of the most important business semantics – logic, rules, and domain knowledge – in compiled and brittle code. However, P2P translators are a well-understood and mature approach for a SOA information architecture.

A canonical information architecture for the SOA is a more highly-evolved approach. Essentially, it is a hub-and-spoke style architecture for data – relying on the use of a common data dictionary/thesaurus to act as a logical pivot point for mapping disparate schema into. This is also a well-understood and mature technique for a SOA information architecture.

However, like the P2P approach, the canonical approach has obvious and severe limitations. First, and most limiting is the necessity for service owners on an SOA to agree to a common data vocabulary. This is often not achievable and when implemented is usually abused in the same way that EDI exchanges were/are abused – people will use “standard fields” in unintended and arbitrary ways on a point-to-point basis. Further, the technical architecture remains quite brittle due to the fact that transformation programs (Java, .NET, and XSL/T) are still leveraged to convert local data to the canonical form. If you change the canonical form, everything breaks. In reality, the canonical form is not a model at all. Models should contain rich semantics about relationships, logic and rules of the information set. With XML alone, this level of specificity is not achievable and typically requires proprietary metadata and executable code to move beyond the limiting semantics of the taxonomic style that is inherent with XML schemas.

With the canonical approach, reusability suffers due to the fact that the most useful business semantics are still persisted in executable code/algorithms instead of metadata. Finally, there is absolutely no capacity for intelligent querying of the internal SOA data set. Because the information architecture is essentially document-centric, any need to query or secure the information contained within a document (canonical or otherwise) necessitates a separate XML database that can execute XPath style XQueries and implement an entity-level security paradigm. Implementing these capabilities is difficult, and maintaining it on a large scale is prohibitive.

In contrast, the use of ontologies (rather than canonical XML or P2P translators) enables users of an SOA to access information (data equipped with semantics) at the entity-level – while simultaneously allowing for a coarse-grained and document-centric message exchange. As has been demonstrated by several NATO member nation COP (Common Operational Picture) programs, OWL ontologies can be

mapped to XML documents, which in turn may be queried and secured from the business information models rather than the technical data models. Inferences, or automated joins, can be effected with a small number of manually expressed links between disparate ontologies – making queries far simpler and much more change resilient. Because business models/ontologies are connected via URIs at the entity level, the schemas themselves can evolve with minimal impact to the overall ecosystem. Most importantly, the whole business ontology domain is “virtual” in the sense that the data will never be persisted in the structure of the ontologies – instead, the information broker lets consumers query the ontology and it handles the retrieval, inference, transformation, and content-level security directly from XML-aware SOA services in realtime.

The intent of these COP programs was specifically to prove the maturity and viability of this approach to SOA information architecture. The programs set out to demonstrate the value of ontologies in accomplishing information fusion without manually persisting or moving data into warehouse style central repositories. While it is clear that the canonical and P2P approaches are more mature, the ontology-driven approach to SOA information architecture is indeed viable and proven.

To date, this information-broker technique has been leveraged for request/reply style operations to SOA aware services. In the future, the addition of publish/subscribe operations to the information broker will make this technical component a full-fledged Mediation Service. A Mediation Service vision, which is unproven to-date, would enable subscriber/alert type operations, driven from joined business ontologies, to post XML documents to underlying SOA services that have set triggers within the Mediation Service. Capabilities for write operations, into the target legacy application would be handled within the application Web services that are tightly-coupled and managed by the legacy application team. This broader publish/subscribe capability, and the realization of a true Mediation Service capability, is expected to be fully online sometime during 2005.

In conclusion, to answer the questions posed with regard to the SOA information architecture, the following has been recommended:

- The payload of SOA exchange data shall be XML documents complying with DTD or XSD schemas
- The representation of SOA information shall be OWL models containing logic, rules, and domain metadata expressed with OWL axioms and class constructors
- The transformation of SOA data to information shall be accomplished with mappings from XML schema to OWL ontologies
- The transformation of SOA information from one context to another shall be accomplished with an inference engine's ability to dynamically reclassify information in OWL format
- The SOA consumer shall query SOA business models (OWL ontology) using XQuery statements issued to an inference engine
- The SOA consumer shall be capable of specifying the XML schema format to which the results of the query should comply – making it simpler for consumption of query results

- The security of SOA content will be ensured with separate security models (OWL ontology) that identify access privileges at the information entity level – enabling secure and auditable reassembly of SOA data into query result sets issued to the SOA consumer

Most of these capabilities, in a request/reply capacity, have already been proven in the previously mentioned COP programs work to date. Further work on publish/subscribe operations – using the ontology-centric paradigm – will result in a powerful and robust Mediation Service capability that can fulfill an advanced NCES capability.

Messaging/Exchange Architecture

The primary question to be answered with respect to the messaging or exchange architecture of an SOA is: how shall services be discovered, coordinated, managed, registered, and made compliant with business processes and workflow?

The aforementioned COP programs did not attempt to address questions at this architectural level, and thus neither the Systems Integrators nor the defense agencies in question have any direct validation of success toward fusing the Semantic Web technology base with an SOA in these terms. However, there is awareness of successful work in these areas from a number of prominent technology leaders in the industry including: IBM, Booz Allen Hamilton, MITRE, France Telecom, Hewlett Packard, and others.

Technically speaking, the goal of research work in this area is to understand how existing SOA specifications such as WSDL and UDDI can be augmented with OWL and inferencing capability to further automate the discovery, coordination and management of SOA services.

To be clear, there are no standards yet in place that exploits the fusion of the two technology bases for messaging. However, there is significant work occurring in the space. OWL-S is a DARPA funded program intended to augment WSDL specifications with enough metadata to enable a transaction or process broker to automatically coordinate and assemble services in a long-lived, multi-step transaction environment. UDDI 3.0 work, in part driven by Carnegie Mellon University researchers, has encompassed more than three years on UDDI upgrades to fuse OWL semantics with the registry specification so that automated service discovery becomes a reality. Similarly, ebXML Semantic Content Registry/Repository work, primarily sponsored by Sun Microsystems, has been focused on providing a semantically rich interface to the repository (using OWL/RDF) in support of dynamic discovery of SOA services.

In Europe, work in Ireland's DERI program has focused on the use of Frame-Logic approaches for the expression of SOA workflow "conversations" such that automated orchestration of long-lived transactions can be machine-managed in a dynamic environment. These efforts are primarily supported by DERI's staff researchers and Germany's University of Karlsruhe, who have significant research efforts supporting DERI's WSML (Web Services Modeling Language) and other F-Logic based approaches to SOA semantics.

None of the research efforts have produced any global standards for SOA semantics as of yet – with the singular exception of the UDDI 3.0 work, which is already embodied in as an OASIS recommendation.

In conclusion, the fusion of the SOA and Semantic Web technology bases for the purposes of service discovery, management, registration, or process orchestration are still immature. Important indicators of the technical requirements for such fusion are apparent from the work at IBM Research and Booz Allen Hamilton in support of the C4 NCES vision advocated by the US DoD. However there is a strong probability that the core specifications of the W3C's OWL-S and DERI's WSML will converge to produce a new specification for the orchestration of services. Work with UDDI and ebXML registries are somewhat mature and can certainly be implemented in support of a dynamic service discovery capability for the SOA.

The SOA software services impacted by these developments would be the "Collaboration Service," "Discovery Service," and "Repository Service." The Collaboration Engine is more than likely a vision that cannot be realized using adopted standards until sometime in 2007. The Discovery Service can already make use of OWL with UDDI 3.0 and will likely find production uses during 2005. Finally, the Repository Service is likely to evolve towards an RDF-based data model, also during 2005.

Technology Base Maturity Conclusions

The maturity of fusion between SOA technology and Semantic Web technology varies by architectural level and service capability. The following matrix summarizes the maturity of various capabilities organized by architectural level. The projected maturity dates are best-case scenarios assuming sufficient funding for development and prototype phasing. These projections are independent of any DISA or DoD planning elements pertaining to official NCES timelines or objectives. The intention with the following table is simply to give the reader a better feel for what level of effort may be required to fuse Semantic Web technology with specific SOA and/or NCES service areas in support of improved information capabilities.

[Capability matrix is on the following page]

Capability Maturity Matrix

Information/Logical Architecture			
	Maturity Level	Maturity Date	Notes
Mediation Service	Medium	Q3 2005	Full publish/subscribe capabilities via broker
<ul style="list-style-type: none"> Semantic Query Broker 	High	Today	Query and result set capabilities will be augmented during 2005
<ul style="list-style-type: none"> Content Security 	High	Today	Manual development of security ontology is still required at present
Messaging/Exchange Architecture			
	Maturity Level	Maturity Date	Notes
Collaboration Service	Low	Q3 2007	No standards, competing visions (US v. EU)
Discovery Service	Medium	Q2 2005	UDDI 3.0 with OWL and inference
Repository Service	Medium	Q3 2005	ebXML/JAXR repository with RDF and OWL/DL inference
Other SOA Service Capabilities			
	Maturity Level	Maturity Date	Notes
Application Services	High	Today	Proprietary COTS “any-to-XML” mappings for legacy support
User Assistant Services	Low	Unknown	Highly dependent on evolved technologies at other service tiers
Messaging Services	High	Today	Proprietary queuing environment for guaranteed delivery
Security Services	High	Today	Network level security only, at document level and higher
Management Services	High	Today*	*for WSDL 2.0 support only (there are implications when collaboration services fuse with Semantic Technologies)

Summary Conclusions on SOA and Semantic Web Fusion

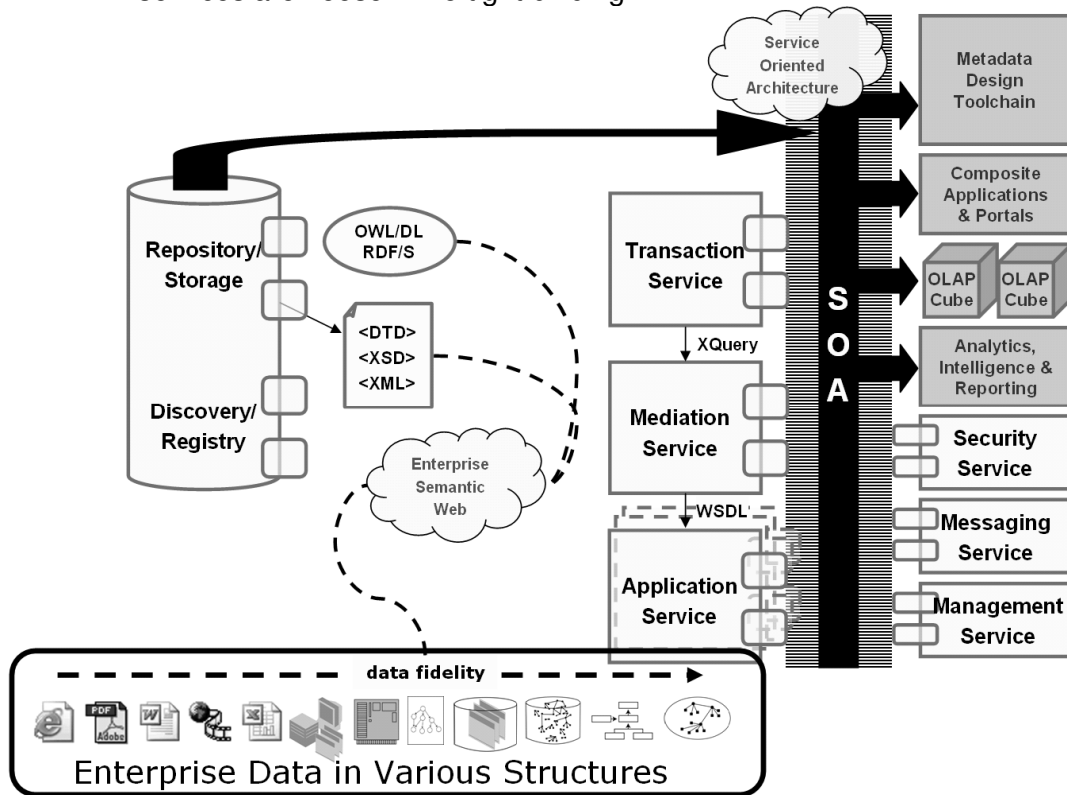
1. Semantic Technologies, in general, are quite mature, with a long history of development and adoption by commerce, federal, software vendors, and standards bodies
2. Scalability of the W3C standards base (both SOA and Semantic Web) is mature enough for exploitation, and have been proven to scale both horizontally and vertically at commercial-grade levels – however more work in this area is sure to follow

3. Fusion of the two technologies bases is (a) mature for exploitation along the information/logical architecture layer, but (b) still immature for many of the messaging/exchange architectural services

Future-State Conceptual Architectures: SOA Services View

The following simplistic conceptual architecture denotes several key characteristics:

- The paradigm is “services” – not servers
- Any service may address any other service
- The SOA is a transport mechanism only
- The ESW is a logical data/information architecture
- Metadata inside repositories describes physical legacy sources
- Physical legacy sources are “wrapped”, and become application services
- End service consumers include portals and composite applications
- OLAP and BI intelligence may be generated from services
- All services are “loose” – no tight binding

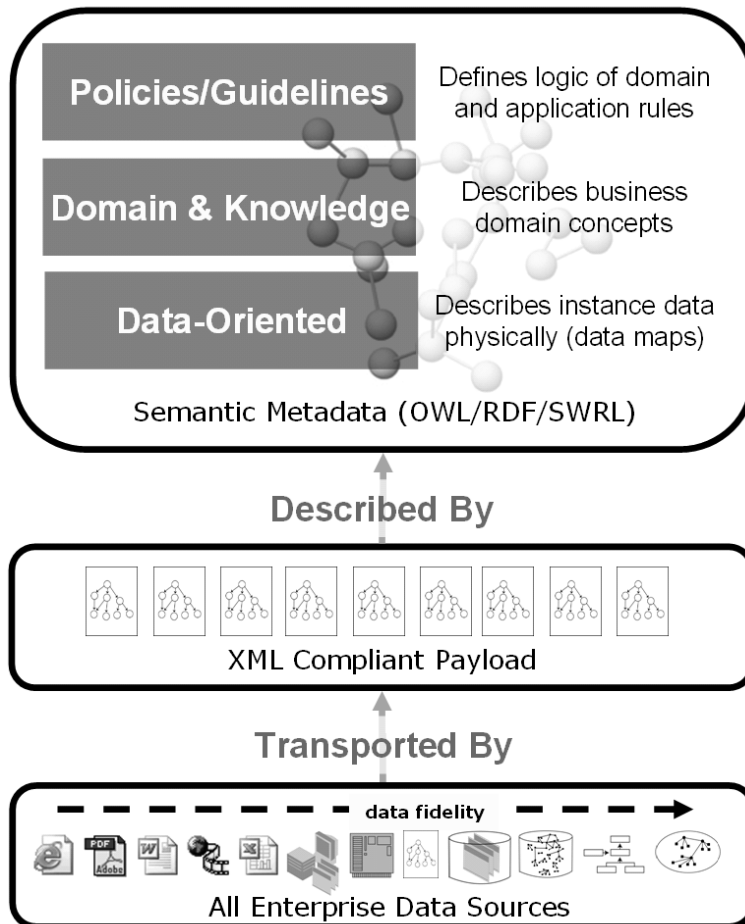


Future-State Conceptual Architectures: Logical Information View

The following simplistic logical information architecture denotes several key characteristics:

- Legacy data sources expose their data as simple XML compliant payload
- A particular XML data architecture is not assumed (here)
 - Eg: XML schema architecture is independent of semantic metadata
- Semantic metadata is used to describe XML document contents
- Three conceptual layers of semantic metadata are typical
 - Data-oriented metadata provides mappings from XML content to OWL
 - Domain metadata is typical domain ontology – as OWL

- Policies may be business rules, security, or regulatory constraints
- Connections between metadata layers are URI-based (per OWL/RDF standards)
- Connections between metadata layers may be axiomatic
 - Making use of standard OWL semantics



Call to Action

Radical improvements for information sharing are nearing readiness for widespread adoption. The fusion of service-oriented architectures (SOA) with semantic Web data specifications means that organizations will soon be able to decouple services from systems and information from software. These new information capabilities have the potential to save billions of wasted dollars and enable simpler compliance with modern regulatory and security restrictions. If current adoption trends and technical advances continue at current rates, 2005 will be a break-out year for the fusion of these two technology bases. Much of the innovation, adoption, and momentum for this technology fusion has been driven by federal efforts with Network-Centric Enterprise Service goals, and commercial organizations stand to greatly benefit from the information capabilities generated by these and other programs. A new era of meaningful service-oriented network computing is already taking shape – be prepared.

About the Author

Jeffrey T. Pollock

Mr. Pollock is a technology leader and author of "Adaptive Information: Improving Business Through Semantic Interoperability, Grid Computing, and Enterprise Integration." (John Wiley & Sons, 2004) As VP of Technology at Network Inference, Mr. Pollock is responsible for product strategy and envisioning next-generation adaptive enterprise software. Previously, as Chief Technology Officer of Modulant, he delivered one of the industry's first market-ready semantic integration middleware platforms in 2001. Throughout his career, he has architected, designed, and built application server/middleware solutions for Fortune 500 and US Government clients. Prior to Modulant, Mr. Pollock was a Principal Engineer with Modem Media and Senior Architect with Ernst & Young's Center for Technology Enablement. He is also a frequent speaker at industry conferences, author for industry journals, active member of W3C and OASIS, and engineering instructor with University of California at Berkeley's Extension on the subjects of Object-Oriented systems, software development process and software systems architecture.