A Platform for Cooperative SOA Research
in
High Performance Semantic Interoperability

Working Paper - November 20, 2007

Abstract

Major architectural and operational information-system changes are coming to the World-Wide-Web, enterprise infrastructure on a global scale, and military defense strategy everywhere. These changes come as Service Oriented Architecture (SOA), Web Services, and System Interoperability. All share a fundamental Achilles heel of semantic interoperability, the glue that enables and limits the capabilities of the systems-of-systems that these concepts propose to enable. Technologies and standards for semantic interoperability are being worked on by various industry groups and product vendors, but their focus is on enablement without concern for operational performance. Deployment scaled for real-world net-centric and World-Wide-Web applications will need to provide semantic mediation services at speeds, capacities, and increasing complexities ill suited to traditional computing approaches. Current intentions limit performance potential by design (tough problems ignored in favor of manageable solutions), and ignore real-time world-scale operational needs (tough problems in denial). Yet multi-billion dollar military programs are rolling forward expecting already-late just-in-time solutions, commercial products are beginning to flow into the market, and expensive infrastructure and standardization commitments are being made.

This initiative proposes to address the research and facilitate the cross-discipline collaboration needed to achieve high performance semantic interoperability. This document proposes a laboratory platform for hosting an open-community collaborative applied research effort, one that enables academic, commercial, and government entities to team on identifying and solving urgent and tough problems in this area. Additionally, this document proposes an initial applied research project that will both shape and demonstrate the nature of the platform, with a focus on general-purpose, infrastructure-embedded, hardware acceleration of complex semantic mediation at network speeds.

Introduction

The glue of a system of systems (SoS) is the meaningful interchange of information among its constituent systems. Without it there is no SoS. The quality of this glue inhibits and enables SoS cohesiveness, synergy, and effectiveness. Semantic interoperability is both the name and the measure of this glue.

When information flows between two systems it has a syntactic element of form and a semantic element of meaning. Among artificial systems, such as data bases and software services, the semantic transformations range from equating different data labels, to rationalizing different data representations, to inferring ontological equivalencies among application needs and service capabilities. Among human activity systems, such as extended enterprises or war fighter coalitions, the range of transformations is similar, but exacerbated greatly by ontological complexity and discontinuity natural to the irrationalities and dynamics of the human milieu.

Performance and failure are in-your-face issues for today's high profile SoS: service oriented architecture, semantic web, extended enterprise, health care data fusion, emergency response, intelligence community practice, net-centric war fighting and Future Combat System programs. These and similar SoS are the application targets of this research initiative.

Semantic interoperability is the ability for two entities to communicate effectively in notification of events, in request for action, in query and response, in ascertaining capability, in negotiating responsibility, in broadening a common body of knowledge, in creating new knowledge. Human beings are the model. Yet frequently, no matter how closely related, two thinking individuals have what they believe is congruent conversation that is eventually realized as conceptually disjoint. Often, that realization never occurs. Computer and system science now knocks at this door for systems of unconscious entities.

The Software Engineering Institute's report on Systems of Systems Interoperability [MORR 04] defines interoperability as: The ability of a set of communicating entities to (1) exchange specified state data and (2) operate on that state data according to specified, agreed-upon, operational semantics. Commenting on the state of technology support for interoperability, they begin (just) to set the stage for our research focus:

> Experts claimed that much of the technology needed to support current interoperability needs already exists. Two exceptions involved real-time applications and multilevel security. First, "Internet Protocol Version 6 fixes many network interoperability problems, but it does not sufficiently address high-speed issues." … The market has converged on data transmission protocols. While this basic technology is present there's been no convergence on common semantics for messages or data. "Systems simply don't operate the same way. They use different time concepts, different tagging of information, and different expectations regarding the order of information, and so on."...Experts were more uncertain about the technologies needed for the future. Organizations are struggling with Information Exchange Requirements (IERs) in a netcentric context. They claim that "IERs make best sense point to point but if the system is providing a service to a general audience, it is not clear the unique needs of systems are specified. New IER approaches must build in flexibility to expect the unexpected."

In the preparatory background discussion we look at government needs for semantic interoperability first, principally military – as they have open and definitive funding solicitations and clear paths of critical urgency. Next we look at the status of SOA and Semantic Web convergence on semantic interoperability with a general commercial tenor, but with recognition that military interoperability expects to build upon this foundation as well. Then we focus on performance needs and issues, and show candidate performance-acceleration opportunities. Finally, we propose three initial research projects that together challenge, elicit, demonstrate, and enable open community application-driven collaboration on high performance SOA and SoS semantic interoperability. Each stands on its own rights independently, while offering and drawing synergy in relationship with the others.

Background

Two humans in effective conversation do not share perfectly identical knowledge models, so each employs a model of the other, which interprets the interaction as well as the interchange. Providing directions to a tourist from another country with little common language makes great allowances for interpretation, while a tourist from a neighboring city is also filtered, but differently. The way each frames and states their question is questioned, and translated into a dynamic succession of convergent interpretations. We cannot expect the success of system-to-system and system-to-human interactions to avoid these same realities – we will settle for no less than equivalent capability, and strive for more. That is our natural benchmark.

Semantic interoperability may be mitigated by a translation service. This works for certain types of human information interchange: document translation, polite introductions, formal UN speech rituals – where time is not a factor. It also works for information systems that batch process, do periodic updates, asynchronously connect, service non-critical or relaxed completion needs. But for both humans and information systems the interesting and critical applications require real-time interoperability: war fighters covering each others back, human-in-the-loop information systems, customer service interaction, electricity grid anomaly response, critical medical diagnosis, sensor network data fusion, C4I tactical decision making. Even web services have a real-time requirement – consider our decided preference for Google's instantaneous reply.

In the end semantic interoperability must facilitate real-time interactive discourse among entities – where real-time may be relative in the abstract sense, but means immediately by whatever waits for closure. For humans in collaborative conversation real-time is measured in fractions of seconds. The very purpose of information systems is to provide improvement on human capability by many orders of magnitude.

Humans exhibit an illusive superiority in the semantic arena. The pattern complexity so easily managed by the human mind defies current computational engines. Machine language translation is literally a joke, as is speech to text within a single language, as is search query quality within bounded data environments.

A typical information system approach is to simplify the nature of the problem: establish a universal common data interchange language that all entities share. It is manageable. It is based on controlling the environment. It is a failure for numerous reasons well cataloged in "Interoperability" [COMM 99], a chapter from *Realizing the Potential of C4I, Fundamental Challenges*. It suffices to say that we live in a fast changing technological ecology of increasingly interconnected and complex systems. No strategy that demands semantic commonality among all entities can survive. Entities based on such a strategy become isolated in sub-domains, which become increasingly inconsequential and eventually liabilities to the ecology as a whole.

Semantic interoperability has front row attention in the Department of Defense: the Future Combat System program [PUCK 06] and net-centric warfare infrastructures cannot deliver without it:

> "Network Centric Warfare is based upon the ability of a force to develop shared situational awareness in the cognitive domain. Technical interoperability will get us to the point where the information is correctly represented in distributed systems, but does not ensure that the individuals in different locations, in different organizations, at different echelons have a similar understanding even though they "see" the same thing. With the added complexity of coalition operations that involve different cultures, the problem is greatly compounded. Semantic interoperability is the capability to routinely translate the same information into the same understanding. This is, of course, necessary to develop the shared situational awareness upon which mature forms of Network Centric Warfare are based [ARMY 07a]."

Pending and outstanding (in May 2007) semantic interoperability proposal solicitations illuminate some of the recognized front-burner issues:

- AFRLS BAA 06-08- IFKA, "Semantic Interoperability, Initial Announcement," FY06: Semantic Interoperability Studies. FY07: Ontological Interoperability….seeks innovative approaches and technologies for achieving ontological mediation between different domains, vocabularies, information sources, and services. FY08: Cross-domain Information Sharing and Systems Interoperability…seeks innovative approaches for the use of semantic technologies for achieving cross-domain information sharing and systems interoperability. FY09: Situational Awareness…seeks innovative approaches for the use of semantic technologies for achieving situational awareness. Our inability to autonomously manage and network these assets must be overcome to achieve and maintain on-demand situational awareness. Semantic service descriptions for the discovery, mediation, orchestration, and composition of networks of ad-hoc information assets could play an important role in servicing the need to have these networks put in place quickly to deliver the right information to those who need to be aware of it. [AFRL 07]

- DARPA solicitation Number SN07-28 posted Apr 02, 2007, "Request for Information (RFI): Strategic Collaboration," DARPA/IXO, Kendra Moore. This RFI seeks information on relevant technologies, capabilities, and concepts in [areas including] Semantic Glue: Provide individuals from diverse organizations and cultures, speaking different languages and having different goals, the ability to collaborate in support of SSTR/HADR operations. "Semantic Glue" refers to the collection of technologies and capabilities that will enable strategic collaboration by providing lightweight, adaptive and interactive representation schemes (language, icons, and geospatial graphics) for exchanging information, across a wide array of devices with different form factors (ranging from cell phones to desk top computers), in environments with varying network bandwidth and reliability, while providing semantically meaningful information appropriate to the individual, organization, form factor, and network capacity. [DARP 07a]
- Army SBIR 07.2: "A07-143: Healthcare Interface Engine To Support Health Level Seven (HL-7), Version 3.0 Data Standard," formal Solicitation period (14 May – 13 June 2007). PHASE I: …research alternative methods to achieving semantic interoperability which might employ terminology or ontology mediation services…Phase I work will also develop a Concept of Operations, functional and technical requirements, and a system design document for the prototype. PHASE II: …the vendor will build a prototype to demonstrate semantic data interoperability of lab, pharmacy, radiology, and clinical encounter notes between DOD and the VHA. PHASE III:…The ability of this engine to deal with semantic interoperability issues is the greatest challenge faced by civilian physician practices and hospitals, healthcare payers, and State Regional Health Information Offices, as they try to comply with the President's April 2004 Executive Order. [ARMY 07b]
- DARPA pre-solicitation preparation for potential BAA 07-Q4, "The Challenge of Data Interoperability from an Operational Perspective," Todd Hughes, DARPA/IXO, Data interoperability is rightly regarded as a pervasive, longstanding, and costly problem. Data management systems do work well at the enterprise level. Certain aspects make them not a good fit for agile, dynamic organizations . . . such as the military in times of conflict. In times of conflict, operational timelines are to far too short for the enterprise acquisition model. The Semantic Web offers a better value proposition by enabling interoperability on an open scale. Price of admission is still high: Semantic Web services, ontology engineering, client applications, service oriented architectures, etc. Future joint and multinational military operations will need to integrate in days, not months. Data interoperability technology must support ad hoc communities of interest with their respective legacy data sources. What technology framework would make this possible? Proposal: A Data Translation Appliance. [DARP 07b]

Semantic interoperability is the most recent focus in information systems research as seen in Table 1, a partial reproduction from "Changing Focus on Interoperability in Information Systems: From System, Syntax, Structure to Semantics [SHET 99]". The time periods refer to the emphasis in R&D, not to the state-of-practice.

	Generation I to 1985	**Generation II 1986-1995**	**Generation III 1996 +**
Level of interoperation concern (new emphasis underlined)	system, data	system, data, information	System, data, information, knowledge (incl. social), process
Types of interoperability emphasized	system (computer system and communication); limited aspects of syntax and structure (data model); transparency of location, distribution, replication, data models	syntax (data types and formats), structure (schematic, query languages and interfaces)	semantic (increasingly domain-specific)
Dominant interoperability architecture	multidatabases or federated databases	federated information systems, mediator	mediator, information brokering
Software and information system architecture	terminal access, point to-point; also mainframes and minicomputers with remote access, clientserver (two-tier)	client-server (three tier);	network, distributed, and mobile
Data/ information interoperability approaches	structural and data model, data representation	understanding of a variety of metadata, comprehensive understanding of schematic heterogeneity	comprehensive use of metadata, increasing emphasis on semantics and ontology supported approaches
Interoperability techniques (representative samples)	data-level relationships, common/canonical data models, mappings, database exchanges, remote database interfaces, query transformations, schema translation, schema integration	schematic and metadatalevel relationships, wrappers, extractors, single ontology, metadatabase, schematic heterogeneity, multidatabase consistency, mediators	multiple ontologies, information or semantic level relationships, context, mediaindependent information correlation, interontological relationships, metadata consistency

Table 1: From "An overview of three generations of interoperability R&D [SHET 99]"

Semantic technology has made great strides in recent years, principally due to two related and driving application goals: Service Oriented Architecture and the Semantic Web [BERN 01]. Nevertheless, a recent article in the IBM Systems Journal sums up the current state of SOA:

> "On the way to information systems integration in an ever growing distributed world, developers come face to face with one of the earliest and most venerable disciplines: semantics. SOAs, armed with suitable descriptive languages and powerful reasoning algorithms, provide a solid and standardized basis that facilitates the design and implementation of semantics-enabled IT infrastructures. Nevertheless, at the current state of the art, although service oriented infrastructures are rapidly becoming a reality, the adoption of specific frameworks to address the semantic layer is still at an early stage. Many kinds of difficulties hinder the acceptance of semantic-oriented artifacts, technologies, methodologies, practices, and standards [VETE 05]."

According to the World Wide Web Consortium (W3C): "*The Semantic Web* provides a common framework that allows data to be shared and reused across application, enterprise, and community boundaries. It is a collaborative effort led by W3C with participation from a large number of researchers and industrial partners [WWWC 07].” Though W3C has been effective in developing technology and standards that will help enable the semantic web, their web intro speaks ahead of reality: “The emergence of Web services and service oriented architectures is leading to new innovative enterprise solutions based on composition of Web services to realize business and scientific processes. …[but] … One of the biggest stumbling blocks in the grand vision proposed by SOA is data heterogeneity between interoperating services [NAGA 06].”

Figure 1 relates interoperability, ontology, and recent standards from the World Wide Web Consortium. The workhorse tool for semantic interoperability is the ontology: “a description (like a formal specification of a program) of the concepts and relationships that can exist for an agent or a community of agents [GRUB 07] … Pragmatically, a *common ontology* defines the vocabulary with which queries and assertions are exchanged among agents [GRUBB 93a, GRUBB 93b].”

A *common ontology* … neither scalable nor sustainable in the accelerating dynamics of this new millennium. Reality is the need to mediate among heterogeneous and constantly changing ontologies. Pragmatically, a system’s ontology can neither remain static nor share congruence in lock step with the ontologies in other systems. To underscore: In *Software Agents as Facilitators of Coherent Coalition Operations* the authors note that “heterogeneity should be accepted and embraced as it is *seen as being inevitable* and can actually be beneficial in a number of cases - especially in security terms [ALLS 01].”

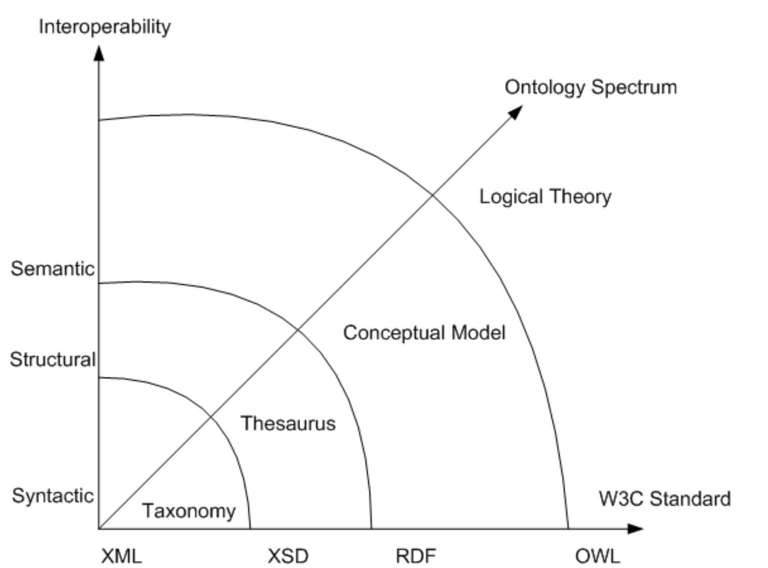

Figure 1: Three Dimensions of Semantic Computing

From Daconta, Obrst, Smith 2003 [DACO 03];
Adapted by Richard Murphy [SICoP 05]

Each system in a system of systems has an ontology, however individual or inexplicit it may be. The trick is to mediate, align, match, translate (roughly interchangeable terms) among the ontologies of interacting agents – not too different from interacting with the lost tourist who speaks little English.

“The recognized benefits of SOA can be enhanced by *semantically enabling* the software so that flexible and dynamic solutions for organizing, managing and handling business interaction can be developed. On its own, SOA offers the advantages of loose coupling between components, well defined interfaces, peer-to-peer interactions, etc. but does not solve the problems of semantic interoperability…the data to be exchanged between the business partners can be semantically described in an unambiguous manner in terms of ontologies, enabling a high degree of automation when solving the heterogeneity problems [MOCA 06].”

What we have is a complex network-speed pattern-processing problem in this mediation of ontologies. Software solutions on traditional computing engines have already been shown inadequate in applications of lesser pattern-processing complexity: special purpose processors and hardware accelerators are standard practice in network routing and intrusion detection [KAU 05], and even then, as Figure 2 depicts, many cannot keep up with network speeds [PETE 05]. The best source for comprehensive and current state-of-the-art network-speed pattern detection can be found in a November 2006 UC Berkeley Ph.D. dissertation: High Speed Deep Packet Inspection with Hardware Support [YU 06].

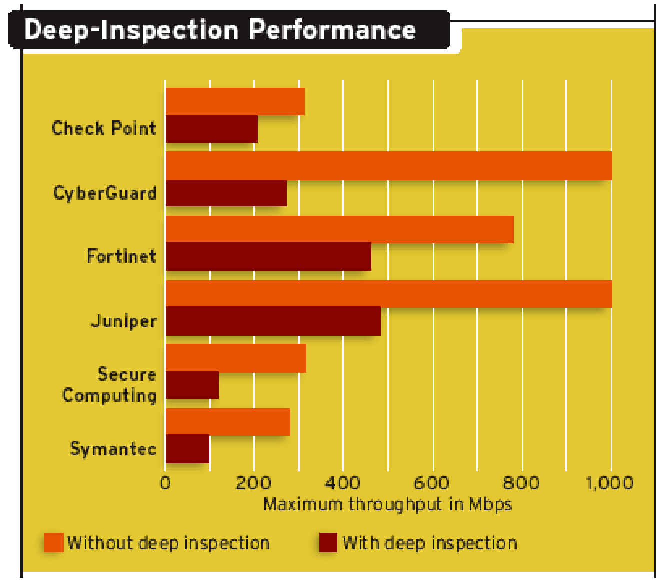

Figure 2: Intrusion Pattern Detection at Network Speeds Degrades Throughput [PETE 05]

Collectively those elements of the networked communications infrastructure that facilitate a Service Oriented Architecture are called an Enterprise Service Bus, typically a distributed software construct. “Essential characteristics of an Enterprise Service Bus (ESB): the meta-data that describes service requestors and providers, mediations and their operations on the information that flows between requestors and providers, and the discovery, routing, and matchmaking that realize a dynamic and autonomic SOA [SCMM 05].” In SOA, the ESB is the home of ontological mediation, or more correctly, *will* be the home. Solely as a software construct, the ESB will come up wanting when significant ontological mediation becomes the norm.

Trends: Forces advancing the Semantic Web recognize synergy with Service Oriented Architecture concepts, and the Military expects both to provide the foundation for SoS semantic interoperability.

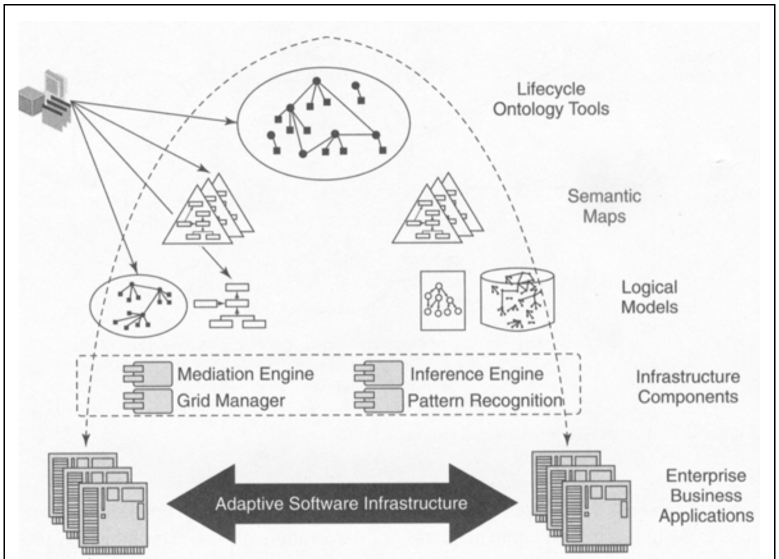

Figure 3: from *Adaptive Information – Improving Business Through Semantic Interoperability, Grid Computing, and Enterprise Integration* [POLL 04],

But semantic interoperability on any grand scale is a tough problem when left to un-aided software solutions running on traditional computing engines.

We leave this background section with evidence of urgency – DoD’s Joint Tactical Radio System (JTRS) program initiated in 1997 is budgeted at $37 billion, expects to put anyone in touch with anyone, relies on interoperability among different network domains, is a core component of the Future Combat System strategy, and, without even considering deep semantic mediation, has failed to solve the critical interoperability need, according to a 2006 GAO report [GAO 06]:

> “The proposed interim technical solutions enabling network interoperability have yet to be developed. To achieve DOD’s desired networking capabilities, waveforms must be able to

communicate and interoperate with each other. However, technologies and radio designs are not mature enough at this point to develop an interoperability capability that would function inside individual JTRS radios. As a result, the program plans to meet network interoperability requirements for the initial increment through the use of gateways. A gateway is a separate node within a network equipped to interoperate with another network that uses different protocols. As such, key functions facilitating interoperability between waveforms may be performed outside of the JTRS radio rather than inside.

"At this point, the JPEO is assessing different options to achieve the gateway function and anticipates that development will start in 2007. The JPEO expects that the development of the gateway will result in a separate acquisition decision but is uncertain as to whether it will be acquired through the forthcoming Airborne, Maritime, Fixed Site system development contract or through a separate contract. In addition, the JPEO is uncertain as to whether the gateway will be employed as a separate piece of hardware or whether it will leverage an existing radio in the network. According to JPEO officials, employing the gateway as a separate piece of hardware could result in additional size, weight, and power risks for some platforms. JPEO officials also noted that without a fully functioning gateway capability, users operating in separate networks will not be able to communicate directly with one another. For example, a ground soldier operating on a Soldier Radio Waveform with a handheld radio would not be able to call directly for fire support from an aircraft operating on the Joint Airborne Network—Tactical Edge Waveform with a Multifunctional Information Distribution System-JTRS radio.

"Conclusions - U.S. military forces' communications and networking systems currently lack the interoperability and capacity DOD believes are needed to access and share real-time information, identify and react quickly to threats, and operate effectively as a joint force. JTRS is critical to providing the capabilities to support DOD's future vision of net-centric warfighting. Yet, since its inception, the JTRS development effort has struggled due to unrealistic cost, schedule, and performance expectations. As a consequence, DOD and the military services have had to make adjustments and acquire interim communications solutions to meet their near-term communications requirements.

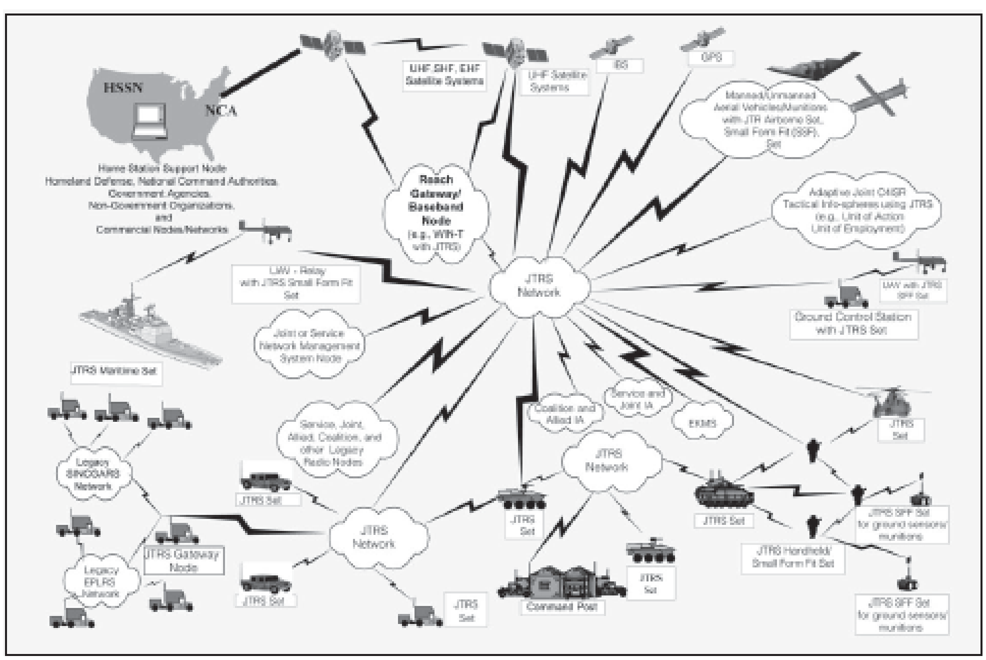

Source: JTRS Operational Requirements Document, April 2003.

Figure 3: JTRS Operational Overview [GAO 06]

Research Initiative

Semantic interoperability is not without serious attention elsewhere. This initiative is differentiated by its focus on high performance hardware-assisted real-time interoperability, its scope of application directed research, and its architecture for open-community collaborative engagement.

High-Performance Hardware-Assisted Real-Time Semantic Interoperability

In a pure sense the metrics of semantic interoperability include both accuracy (preservation of meaning) and scope (one for one concept mapping). In a pragmatic sense the speed of interoperability mediation will make or break its usefulness. This research initiative will initially investigate and deploy a breakthrough pattern-recognition processor as its benchmark for network-speed mediation without compromise in either accuracy or scope.

The processor architecture is patent-pending IP owned by Kennen Technologies, LLC. It is being developed first as an FPGA instantiation on a PCI board for security application development with Phase 2 SBIR prototyping funds from DHS. It is also in early stages of VLSI design by a major memory-device company. As a VLSI RAM-like device, a single chip at RAM-like production costs, using today's memory technology, should have a capacity of 1 million pattern features and a data-stream processing throughput of 1 GigaByte per second. The architecture accommodates any number of patterns in any configuration without affecting throughput. Processors can be ganged serially to increase pattern capacity, and in parallel to increase throughput. Integrated combinatorial Boolean logic allows the processor to perform high-level logical aggregations of sub-patterns appropriate to semantic processing needs. Notably, the architecture does not employ latency-incurring pipeline techniques, and it eliminates the typical exponential relationship between pattern complexity and throughput. It is a breakthrough that could launch an era of semantic technologies. SDK and emulation software also exists, and will be made available to this research initiative. Documentation currently is limited to patent documents, the DHS SBIR Phase 1 final report [KENN06a], and the DHS Phase 2 project proposal (underway) [KENN06b]. The DHS project will produce processor and SDK documentation by Q3 2007 [KENN 06b].

Application Directed Research

Applied research fits with the tenor and historical origins of Stevens Institute of Technology, which traces its origins to John Stevens, pioneer of steamboat development and designer of the first American-built steam locomotive. The School of Systems and Enterprises in particular has placed a heavy and successful emphasis on both professionally engaged graduate students and professionally experienced faculty. Research interests run decidedly toward problems with immediate and significant application.

Semantic interoperability is both the frontier and the fulcrum of major high-stakes systems-of-systems efforts underway, which are expected to change the fundamental nature of global networks – a result expected to have social effects at least as far reaching as the locomotive and steamship in their days.

Open Community Collaborative Engagement

The new School of Systems and Enterprises is the result of phenomenal growth from its origins as the Systems Engineering Department in the School of Engineering. This growth in academics is largely the result of an open community collaborative approach, which will be leveraged, grown-further, and harvested in this and other research initiatives. Systems engineering is the discipline that coordinates, rationalizes, and integrates across the many other disciplines that collectively work together to realize a complex system. Semantic interoperability is a problem that requires rational collaboration among many different disciplines to deliver on its promise.

High Performance Interoperability Laboratory

This initiative will create a research laboratory for high performance mediation research.

Semantic interoperability will achieve necessary network-speed performance only through hardware assisted semantic mediation. This project will develop a general purpose reusable hardware-assisted prototype of ontology mediation capability, based initially on the Kennen Pattern Recognition Coprocessor (KPRC). The architecture of the system and its individual components will be compatible with SOA and web-service concepts. The resultant system will be accessible as a suite of SOA web-services suitable for employment in SOA and web-service research projects, which may want to use the suite as a composite service or explore alternate and improved methods for some or all of its components.

Some immediate clarification on the 3 projects are in order. The first-draft project description shown below is a bit of a mash-up of what will be three separate but related projects. Some of the aspects of the project described in the current draft need to be separated out and moved into the other two projects.

The three projects I will have detailed (two pages each) in the next version of the white paper are strongly related. In the end their results will be fully integrated, but each can proceed independent of the others (though Project C will want to incorporate the results of the others). The three projects are as follows:

----------------- What Follows is an unfinished Work-in Process ----------------

Technology Demonstrator – Hardware Mediation – Research Platform

Technical Objectives – (A mashup of three individual projects at the moment – to be separated later). 1) Show that KPRCs operating at 1G-Hertz cycle times can provide ontological mediation at1GB throughput, 2) provide actual ontological mediation at throughput limited only by the cycle time of the particular KPRC available at project time, 3) provide general purpose mediator development and deployment capability to translates a data stream between two typical SOA service ontologies, 4) make mediation development and testing accessible from the Internet, 5) provide a total-system architecture that delivers the service agility expected of SOA, 6) enable remote researchers and developers to access the laboratory hardware and software components, 7) provide a demonstration of hardware assisted technology and a algorithm challenge for conceptually addressing ontology mediation.

Tasks – This proposal is preliminary and general as it expects to gain specific shape after an academic/commercial team and a target application are identified. Thus, there are two categories of tasks shown here: A) proposal finalization and B) strawman project shape.

A) Proposal finalization:

   1. Identify and engage commercial partner(s) – appropriate profile would include a large-business with a strong existing commitment to the SOA infrastructure market, augmented as appropriate with small-business semantic technology expertise.
   2. Identify specific application to shape the proposal details – based upon needs of target funding sources (government solicitations and/or commercial partner.
   3. Identify appropriate research team members from among the Stevens schools of Systems and Enterprises and Computer Engineering, commercial partner(s) and potential academic research collaborators from other institutions.

4. Develop funding proposal(s) with the team identified in 3.commercial partner as necessary.

B) Strawman project shape – Tasks are divided among four principle areas:

1. Mediator Developer Kit – this area falls into the computer science and software engineering realms, and entails three primary tasks:
    i. Develop algorithms for enabling ontological transformation that leverage the architecture of the KPRC, and specify the additional control and memory resources needed on the PCI-board that employs the KPRC. Approaches will be inspired by (for instances) first order logic methods for reasoning [HAAS06], surveyed reasoners in WSML Reasoning Implementation [BRUI 05], various Datalog Engines, RDF Gateway by Intellidimension (http://www.intellidimension.com/), and WonderWeb: Ontology Infrastructure for the Semantic Web [HORR 03].
    ii. Develop a suitable compiler for translating ontologies into algorithm-compatible KPRC loads.
    iii. SDK Other – which encompasses typical SDK-like routines for loading, controlling, emulating, and monitoring the KPRC device specifically, and the Ontology-Mediator Engine more generally.
2. System Access – A software engineering task. Develop Web-Service and SOA compatible interfaces to the system and its individual components to facilitate remote access by researches, developers, and testers. This task will also enable substitution of any of the components so that researches and developers may mix and match alternate component instances.
3. Ontology-mediation Engine – Hardware computer engineering that encompasses two primary tasks:
    i. Instantiation (FPGA coding) and integration (FPGA or ASIC) of a KPRC device for use on the Ontology-Mediation Engine PCI board [KENN 06a, KENN 06b].
    ii. Design, fabrication (or COTS FPGA-board acquisition) of an appropriate PCI board that will contain the KPRC and all ancillary resources needed to service it, such as a power-PC processor, on-board memory, network interface, whatever is deemed necessary.
4. Direction and Oversight – Top level project management, and system engineering with three primary tasks:
    i. Conceptual architectures for software and hardware engineering, principally oriented toward ensuring an agile-systems result.
    ii. Top-level requirements for the other three task areas.
    iii. Research and development methods that will enable collaborative open-community participation.
5. Demonstration and Challenge Environment -

Costs – To be determined during proposal finalization.

What follows is a breakdown of three separate projects that by themselves stand alone but together provide and offer synergy to each other. In the end they will appear to be an integrated suite of platform services.

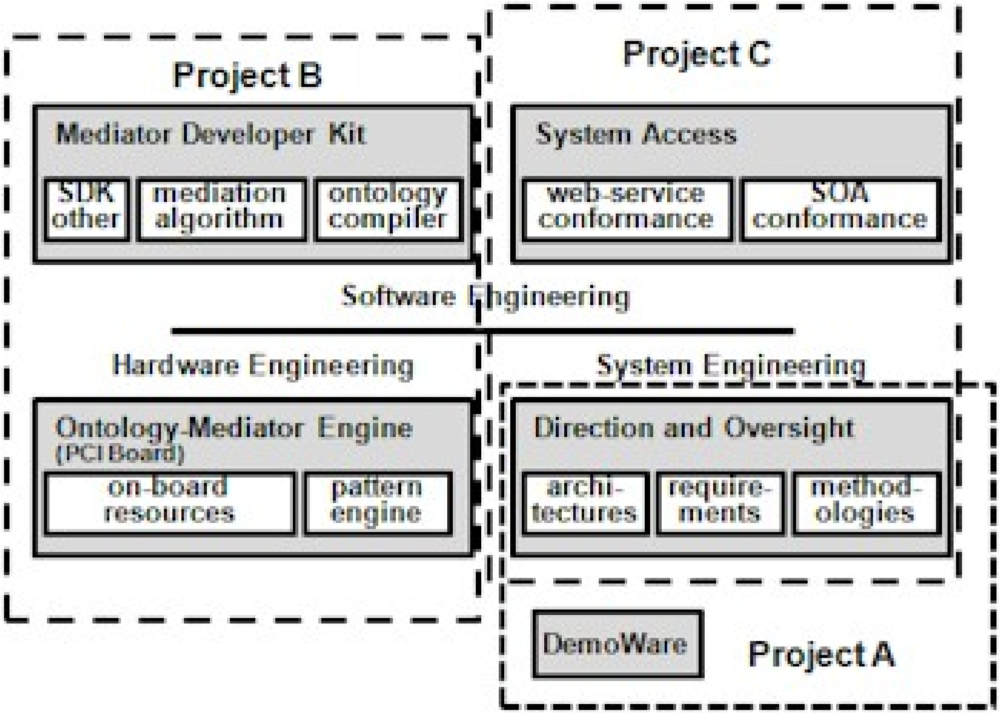

Project A: Technology and Platform Demonstration and Challenge Environment

Abstract – To be completed. This is a laboratory demonstration, educational platform, and test bed that to engage researchers and developers interested in learning about the capabilities of the underlying technology employed in Project 1.

With Company-X (enjoined from naming at the moment) expressing major likelihood as a natural partner/sponsor, this project is a laboratory demonstration of the underlying pattern-recognition technology - it is the smallest of the three projects, quite likely to begin by September 2007, and should be finished 6-9 months after initiation. It will be a $100k project total, funded half by a "sponsor" and matched by funds earmarked for this purpose already available. It will make use of the Kennen software "emulator" and pattern compiler currently available and can make use of the DHS-funded SDK work-in-process at Kennen. Purpose of this project is to: a) provide a first rate demonstration available on the web of the kinds of things that this technology can enable in semantic interoperability (as a minimum), b) act as a recruitment tool for semantic and other SoS interoperability research opportunities, c) to entice and obtain open-community research collaborators, and d) provide the initial Systems Engineering architecture and requirements development in anticipation of the other two projects. [lower right quad of the master-project diagram in the white paper and below]

Project B: Hardware Ontology-Mediation Engine and Infrastructure

Abstract – (To be completed)

With a company like Lockheed in mind as a natural partner/sponsor, this is the initial applied research project (with government partner involved with real problem), with a focus on general-purpose, infrastructure-embedded, hardware acceleration of complex semantic mediation at network speeds. This project can occur completely independently of Project-C (the research platform project), but will utilize the SE results available from Project-A, and provide an exemplary demonstration of the research platform (Project-C) configuration. Preferably this project will involve an application needed by a Government partner.

Technical Objectives – 1) Show that KPRCs operating at 1GB cycle times can provide ontological mediation at1GB throughput, 2) provide actual ontological mediation at throughput limited only by the cycle time of the particular KPRC available at project time, 3) provide general purpose mediator development and deployment capability to translates a data stream between two typical SOA service ontologies, 4) make mediation development and testing accessible from the Internet.

Project C: High Performance Mediation Collaborative Research Open Platform

Abstract – (To be completed) This will be an open community resource that will engage researchers interested in hardware-assisted high performance interoperability.

With a company like IBM in mind as a natural partner/sponsor, this project creates a laboratory platform for hosting an open-community collaborative applied research effort, one that enables academic, commercial, and government entities to team on identifying and solving urgent and tough problems in the area of high performance semantic interoperability in an SOA/Web-Services context. This project inherits and takes over where project A left off. It also integrates the results of, and provides a "home" for, the results of Project B.

Management and Resources

The designated Stevens faculty who will work in support of this proposal and form part of the management structure giving oversight to research effort include: Professor Rick Dove (PI), Dr. John Boardman, Dr. Victor Lawrence, Dr. David Nowicki, Dr. Art Pyster, Dr. Brian Sauser, and Dr. Roshanak Nilchiani. PhD students, especially those in the program that accommodates employed professionals, will support and drive research topics in this area of semantic interoperability in systems-of-systems. Members of TopQuadrant, a commercial small business, recognized internationally for leadership and pioneering development in semantic technologies, semantic web, and SOA, are part of the research team.

Summary

Semantic Interoperability is a major and urgent problem. Major because large efforts (JTRS, SOA, FCS Semantic Web) rely on a solution. Urgent because major programs-in-process assume a solution will arrive in time.

Semantic Interoperability in-the-large does not have a demonstrable <u>effective</u> solution. Current technical approaches impose restrictive and unrealistic constraints and controls on the scope of the problem (common "language", customization expertise).

Interoperability in general is insufficiently respected/addressed by systems that must collaborate in a System of Systems (DoD Joint Programs).

Stevens sits at the nexus of advanced Systems Engineering education, Systems of Systems research, agile systems and systems in-the-large thinking, professional practicing system engineers, the interoperability-frontier community, the open-community research and education concept, and a breakthrough semantic processor technology.

The nature of the problem and the nature of these assets is a unique and synergistic combination. New thinking, new technology, and new process will naturally converge on a reusable, scalable, effective solution.

Benefits of this research:

- Hardware accelerated semantic interoperability performance
- Scalable solutions
- Reusable technologies
- Affordable deployment
- Effective application

Enabled by…

- Systems engineering leadership, discipline, architecture
- Breakthrough-capability semantic processor technology

- Research and development grounded by practicing professionals
- Multidisciplinary collaborative team
- Application deployment-need as driver

It is natural for the School of Systems and Enterprises to take this approach. Its focus is systems engineering and its core is systems thinking. By charter the interest of system engineering and systems thinking is the coherency of systems that span knowledge disciplines. Here we enable the cooperative integration of disciplines with a collaborative process architecture, and we enable the engineering of a solution with an open colaboratory system. And…we build the first instance of a functioning collaborative laboratory resource in the area of high-performance semantic interoperability.

References

[AFRL 07] "Semantic Interoperability, Initial Announcement," AFRLS BAA 06-08- IFKA: http://www.fbo.gov/spg/USAF/AFMC/AFRLRRS/Reference-Number-BAA-06-08-IFKA/Modification%2005.html

[ALLS 01] "SOFTWARE AGENTS AS FACILITATORS OF COHERENT COALITION OPERATIONS," David Allsopp, Patrick Beautement, Jeffrey Bradshaw, John Carson, Michael Kirton, Niranjan Suri and Austin Tate, Informatics Research Report EDI-INF-RR-0113, University of Edinburgh, Centre for Intelligent Systems and their Applications, June 2001, http://www.inf.ed.ac.uk/publications/online/0113.pdf

[ARMY 07a] "FCS Network," harvested May6, 2007 from http://www.army.mil/fcs/network.html

[ARMY 07b] "A07-143: Healthcare Interface Engine To Support Health Level Seven (HL-7), Version 3.0 Data Standard," Dr. James Phillips, Army SBIR 07.2, http://www.dodsbir.net/solicitation/sbir072/army072.doc

[BERN 01] "The Semantic Web,", Tim Berners-Lee, James Hendler and Ora Lassila, Scientific American, May 2001

[BRUI05] "WSML Reasoning Implementation," Jos de Bruijn, Cristina Feier, Uwe Keller, Rub¶en Lara, Axel Polleres and Livia Predoiu, WSML Working Draft October 3, 2005, http://www.wsmo.org/TR/d16/d16.2/v0.2/20051003/d16.2v0.2_20051003.pdf

[COMM 99] "What is Interoperability and Why is it Important?" Chapter 2 from *Realizing the Potential of C4I, Fundamental Challenges,* Committee to Review DOD C4I Plans and Programs, National Academies Press, 1999 http://books.nap.edu/html/C4I/ch2.html

[DARP 07a] "Request for Information (RFI): Strategic Collaboration," DARPA/IXO, Kendra Moore, Solicitation Number SN07-28 Posted Apr 02, 2007 http://www.fbo.gov/spg/ODA/DARPA/CMO/SN07-28/Synopsis.html

[DARP 07b] "The Challenge of Data Interoperability from an Operational Perspective," Todd Hughes, DARPA/IXO, http://iswc2006.semanticweb.org/program/panel-todd-huges.ppt

[GAO 06] "Restructured JTRS Program Reduces Risk, but Significant Challenges Remain," GAO Report to Congressional Committees, GAO-06-955, September 2006, www.gao.gov/cgi-bin/getrpt?GAO-06-955

[GRUB 07] "What is an Ontology?", Harvested May 3, 2007 from http://www-ksl.stanford.edu/kst/what-is-an-ontology.html

[GRUB 93a] "A Translation Approach to Portable Ontology Specifications," Thomas R. Gruber, *Knowledge Acquisition*, 5(2):199-220, 1993, and Technical Report KSL 92-71, Knowledge Systems Laboratory, Stanford University, 1993, http://tomgruber.org/writing/ontolingua-kaj-1993.pdf

[GRUB 93b] "Toward Principles for the Design of Ontologies Used for Knowledge Sharing," Thomas R. Gruber, International Journal Human-Computer Studies 43, p.907-928. Substantial revision of paper presented at the International Workshop on Formal Ontology, March, 1993, Padova, Italy. Technical Report KSL 93-04, Knowledge Systems Laboratory, Stanford University, 1993, http://tomgruber.org/writing/onto-design.pdf

[HAAS 06] "Practical Reasoning with OWL and DL-Safe Rules," Peter Haase, Pascal Hitzler, Markus Krötzsch, Jürgen Angele and Rudi Studer, http://km.aifb.uni-karlsruhe.de/ws/prowl2006/prowl06_4on1.pdf

[HORR 03] "Rule Language, WonderWeb: Ontology Infrastructure for the Semantic Web," Ian Horrocks, Raphael Volz, Stefan Decker, Benjamin Grosof and Boris Motik, IST Project 2001-33052, http://wonderweb.semanticweb.org/deliverables/documents/D2.pdf

[KAU 05] PERPOS Information Assurance, Jason Kau and Son Nguyen, PERPOS Technical Report ITTL/CSITD 05-1, Georgia Tech Research Institute, Georgia Tech Information Security Center, May 31, 2005, http://perpos.gtri.gatech.edu/publications/PERPOS%20TR%202005-01.pdf

[KENN 06a] DHS HASSM SBIR Phase 1 Final Report, Kennen Technologies, June 2006

[KENN 06b] DHS HASP SBIR Phase 2 Project Proposal, Kennen Technologies, July 2006

[MOCA 06] "Filling the Gap – Extending Service Oriented Architectures with Semantics,", Adrian Mocan, Matthew Moran, Emilia Cimpian and Michal Zaremba, IEEE International Conference on e-Business Engineering (ICEBE'06), 2006

[MORR 04] "Systems of Systems Interoperability (SOSI)," Edwin Morris, Linda Levine, Craig Meyers, Pat Place and Dan Plakosh, Technical Report CMU/SEI-2004-TR-004, http://www.sei.cmu.edu/pub/documents/04.reports/pdf/04tr004.pdf

[NAGA 06] "Semantic Interoperability of Web Services – Challenges and Experiences," Meenakshi Nagarajan, Kunal Verma, Amit P. Sheth, John Miller and Jon Lathem, 2006 IEEE International Conference on Web Services, http://lsdis.cs.uga.edu/library/download/techRep2-15-06.pdf

[PETE 05] "Clash of the Titans," Adrian Peters and Michael Jones, *Network Computing*, April 28, 2005, pps 55-67, http://www.networkcomputing.com/showitem.jhtml?docid=1608f2

[POLL 04] *Adaptive Information – Improving Business Through Semantic Interoperability, Grid Computing, and Enterprise Integration*, Jeffrey T. Pollock and Ralph Hodgson, John Wiley & Sons, 2004

[PUCK 06] "FCS Battle Command," Tim Puckett, FCS Update, August 2006, http://www.ausa.org/pdfdocs/ARMYMag/Puckett_Aug05.pdf .

[SCHM 05] "The Enterprise Service Bus: Making service-oriented architecture real," M.-T. Schmidt, B. Hutchison, P. Lambros and R. Phippen, IBM SYSTEMS JOURNAL, VOL 44, NO 4, 2005, http://www.research.ibm.com/journal/sj/444/schmidt.pdf

[SHET 99] "Changing Focus on Interoperability in Information Systems: From System, Syntax, Structure to Semantics," Amit P. Sheth, in Interoperating Geographic Information Systems. M. F. Goodchild, M. J. Egenhofer, R. Fegeas, and C. A. Kottman (eds.), Kluwer Academic Publishers, 1999, pp. 5-30 http://lsdis.cs.uga.edu/lib/download/S98-changing.pdf

[SICoP 05] "Introducing Semantic Technologies and the Vision of the Semantic Web," Semantic Interoperability Community of Practice (SICoP), 02/16/05, http://colab.cim3.net/file/work/SICoP/WhitePaper/SICoP.WhitePaper.Module1.v5.4.kf.021605.doc

[VETE 05] "Models for semantic interoperability in service oriented architectures," G. Vetere and M. Lenzerini, IBM SYSTEMS JOURNAL, VOL 44, NO 4, 2005, http://www.research.ibm.com/journal/sj/444/vetere.pdf

[WWWC 07] World Wide Web Consortium, Semantic Activity Page, Harvested May 3, 2007, http://www.w3.org/2001/sw/

[YU 06] *High Speed Deep Packet Inspection with Hardware Support*, *Fang Yu, PhD Disertation,* Electrical Engineering and Computer Sciences, University of California at Berkeley, November 22, 2006, http://www.eecs.berkeley.edu/Pubs/TechRpts/2006/EECS-2006-156.pdf

Other Relevant Resources

[DIMA 06] "Systems of Systems Interoperability Types and Characteristics in Joint Command and Control," Michael J. DiMario, 2006 IEEE/SMC International Conference on Systems of Systems Engineering, Apr 2006

[EUZE 04] "D2.2.3: State of the art on ontology alignment", Coordinator: Jérôme Euzenat - Thanh Le Bach, Jesús Barrasa, Paolo Bouquet, Jan De Bo, Rose Dieng, Marc Ehrig, Manfred Hauswirth, Mustafa Jarrar, Ruben Lara, Diana Maynard, Amedeo Napoli, Giorgos Stamou, Heiner Stuckenschmidt, Pavel Shvaiko, Sergio Tessaris, Sven Van Acker and Ilya Zaihrayeu, IST Programme of the Commission of the European Communities project number IST-2004-507482, http://www.aifb.uni-karlsruhe.de/WBS/meh/publications/euzenat04state.pdf .

[GANE 03] "The dawning of the autonomic computing era," A. G. Ganek and T. A. Corbi, IBM SYSTEMS JOURNAL, VOL 42, NO 1, 2003, http://researchweb.watson.ibm.com/journal/sj/421/ganek.pdf

[KEPH 03] "The Vision of Autonomic Computing," Jeffrey O. Kephart and David M. Chess, IEEE Computer Society, Jan 2003, http://researchweb.watson.ibm.com/autonomic/research/papers/AC_Vision_Computer_Jan_2003.pdf

[PEDR 06] "Towards a Semantic Event-Based Service-Oriented Architecture," Carlos Pedrinaci, Matthew Moran, and Barry Norton, In Proceedings of 2nd International Workshop on Semantic Web Enabled Software Engineering (SWESE 2006), Workshop at 5th International Semantic Web Conference (ISWC 2006) , Athens, Georgia, USA, 2006, http://kmi.open.ac.uk/projects/dip/resources/ISWC2006/swese2006Final.pdf

[NIBL 05] "Events and service-oriented architecture: The OASIS Web Services Notification specifications," P. Niblett and S. Graham, IBM Systems Journal, Volume 44, Number 4, 2005, http://www.research.ibm.com/journal/sj/444/niblett.pdf

[SELV 06] "Achieve semantic interoperability in a SOA," M.Selvage, D.Wolfson, B.Zurek and E.Kahan, IBM, June 2006, http://www-128.ibm.com/developerworks/webservices/library/ws-soa-seminterop.html[YAO 00] "Semantically-Based Active Document Collection Templates for Web Information Management Systems", In-Young Ko, Robert Neches and Ke-Thia Yao, In Proceedings of the International Workshop on the Semantic Web, Sep 2000, Lisbon, Portugal. http://www.ics.forth.gr/isl/SemWeb/proceedings/session2-3/html_version/index.html

[SHAR 06] Lightweight Synchronization of Ontologies, Master's Thesis, Arun Sharma, Mat. No. 262119, July, 2006, RWTH Aachen University, ftp://ftp.inrialpes.fr/pub/exmo/reports/diplom-sharma.pdf

[ZHAN 04] "Ontology-Driven Web Services Composition Techniques, Ruoyan Zhang, Master Theses, University of Georgia, 2004, http://webster.cs.uga.edu/~budak/ruoyan_thesis.pdf

[ZHAO 07] "Semantic Matching Across Heterogeneous Data Sources," Huimin Zhao, Communications of the ACM Jan 2007.

[???? 06] "Intelligent Healthcare Monitoring based on a Semantic Interoperability Platform", IST- 27074 SAPHIRE, SAPHIRE – Deliverable 3.1.1 Survey of the State of the Art, Information Society Technologies, Feb 2006, http://www.srdc.metu.edu.tr/webpage/projects/saphire/deliverables/SAPHIRED3.1.1_SOA_0.5.doc