

ENTERPRISE APPLICATIONS OF SEMANTIC WEB: THE SWEET SPOT OF RISK AND COMPLIANCE

Amit Sheth

Semagix, Inc., Athens, GA, USA www.semagix.com

Abstract

Semantic Web is in the transition from vision and research to reality. In this early state, it is important to study the technical capabilities in the context of real-world applications, and how applications built using the Semantic Web technology meet the real market needs. Beyond push from research, it is the market pull and the ability of the technology to meet real business needs that is a key to ultimate success of any technology. In this paper, we discuss the market of Risk and Compliance which presents unique market opportunity combined with challenging technical requirements. We discuss how the Semantic Web technology with an ontology driven approach is especially well suited to support the demanding requirements of the applications in this market. We also discuss the capabilities of a commercial semantic technology that has origins in academic research, as it is utilized in a significant Risk and Compliance application deployed at large financial institutions. Core capabilities of this technology include the ability to develop and maintain focused but large populated ontologies, automatic semantic metadata extraction supported by disambiguation techniques, ability to process heterogeneous information and provide semantic integration combined with link identification and analysis through rule specification and execution, as well as organization and domain specific scoring and ranking. These semantic capabilities are coupled with enterprise software capabilities which are necessary for success of an emerging technology for meeting the needs of demanding enterprise customers.

Keywords: Semantic Web technology, Enterprise Applications, Risk and Compliance, Ontology driven Information Systems, Semantic Metadata, Link Analysis, Rule Processing, Risk Scoring, Customer Identification and Risk Assessment Solution

1. INTRODUCTION

The Semantic Web¹ has arrived. We have early applications that are now functioning and deployed in scientific research as well as industry [Miller2005][Sheth 2005b][Lee 2005]. We also have SW language standards such as RDF and OWL, and we have some stealth applications leading to the pervasive use of enablers such as associating metadata in RDF with digital content over mobile networks/devices and use of metadata in RDF for specifying and validating content license². We also have some early experiences that show where Semantic Web demonstrates clear value and significant differentiation so that we can chart its broader adoption. Two cases stand out in this context: bioinformatics applications in the scientific research arena, and risk and compliance applications in industry. In this paper, we focus on the latter.

Semantics relate to the meaning and use of data. So naturally, characteristics of a domain plays an important role in determining whether a Semantic Web technology is a natural fit for applications and can help address challenges in that domain. Today, ontology is at the heart of any significant Semantic Web technology and solution. Hence a key feature that would make a semantic technology appropriate is the ability to create and manage a large populated ontology for addressing the application requirements. An ontology populated with the domain knowledge provides a critical differentiator for Semantic Web technology in solving problems where other technologies would significantly suffer due to the lack of it. We take the position that semantic technologies that utilize ontology and core technical capabilities such as knowledge representation, entity identification, disambiguation and reasoning that exploit relationships, is of primarily commercial interest for now, whether or not they already use contemporary Semantic Web language standards such as OWL. Albeit the use of standards, especially RDF/RDFS, is highly desirable for interoperability, reuse, commercialization and market adoption reasons³.

While the technical considerations make a technology appropriate to solving a problem, no less important is the non-technical, business issue of market pull or readiness of the businesses to accept new technologies. Unique market circumstances create new opportunities and raise the needs for new applications, which can often break the lethargy or resistance in adopting new technologies and solutions. Again in this case, the risk and compliance market

¹ W3C Semantic Web Activity <http://www.w3.org/2001/sw/>

² Creative Commons License RDF validator: <http://validator.creativecommons.org/>

³ We term the semantic technology that also uses contemporary Semantic Web languages and standards, namely RDF and OWL, as Semantic Web technology. However, for this paper, we will not seek to make significant distinction between the two.

has the external impetus to look for new solutions that traditional technologies do not adequately solve.

This paper deals with the discussions on the needs in the risk and compliance market that uniquely positions the Semantic Web technology as the most appropriate technology, and further gives insights into some of the key technical requirements for which a semantic approach is ideally suited. In brief, this paper seeks to explore or answer the following questions: What are the requirements and characteristics of the risk and compliance market that makes it well suited for Semantic Web technology? What are the technical capabilities of a suitable Semantic Web technology for addressing demanding and unique requirements for applications in this market?

Section 2 characterizes the market in terms of applications. Section 3 focuses on unique requirements for analytics, especially in finding links between heterogeneous data and ontological knowledge. Section 4 discusses key reasons why Semantic Web technology is an excellent fit to address the requirements. Section 5 discusses technical capabilities of a commercial Semantic Web technology. Section 6 briefly describes one application case study.

2. NEW OPPORTUNITIES AND CHALLENGES IN RISK MANAGEMENT AND COMPLIANCE MARKET

There is an unprecedented interest in the risk and compliance applications, especially in financial and government sectors. Two events and circumstances indeed shaped the corresponding market:

(a) September 11, 2001 and ensuing focus on intelligence analysis and fighting terrorism, leading to the USA Patriot Act of 2001.

(b) Corporate scandals and the need for better financial controls and corporate governance resulting from increased regulatory vigor, leading to the *Patriot Act of Finance*, the Sarbanes Oxley Act of 2002.

Correspondingly, many direct and indirect applications have appeared or are being developed. Here are just a few:

Identity and Risk Management: Know Your Customer (KYC) or Customer Identification Program (CIP) applications which the financial organizations are required to perform as part of the Patriot Act section 326 provisions and corresponding European Union regulations

Security Screening: Airport Security Screening or Passenger Threat Assessment applications, to determine if a passenger is directly or indirectly related to any known black listed entities (counties, organizations, people, etc)

and other security and prevention applications needed to support homeland security [Avant et al 2002]

Regulatory Compliance: applications supporting governance and accounting, linking data and processes to comply with the provisions of the Sarbanes Oxley Act [Ruh 2004]

Fraud Prevention: Anti-Money Laundering (AML) application⁴, for example, to help avoid risks associated with doing business with customers (current or potential) who might have links with black listed entities, as required by the Patriot Act as well as European Union's Money Laundering Directive

Financial Crimes Enforcement: such as enforcement of the provisions of Section 314 of the Patriot Act⁵ requiring identification and collection of evidences related to hawala operation involving a sanctioned country, arms trafficking, alien smuggling resulting in fatalities, international criminal network involved in identity theft and wire fraud, and others

Background checks and clearance: for obtaining or renewing security clearances for government jobs, the agencies need to perform substantial background checks on existing and potential employees

Authorized Information Access: for compliances with regulations such as Executive Order on Access to Classified Information⁶) or "need to know" support ensuring that employees access only that information which are necessary to perform their assignments [Aleman et al 2005]

The factors that make the business opportunity for developing risk and compliance applications for financial and government sectors more attractive include the following:

- the institutions are largely unprepared and ill-equipped to deal with the spate of significant new regulations resulting from unexpected circumstances
- the time available to implement a compliance process is in months rather than years, that the risk of non-compliance results in unacceptably high risk (i.e., the solution is *an aspirin, not a vitamin*), and
- the amount of effort involved or time for performing the required compliance activity practically argues for an automated process rather than a manual process.

A risk and compliance process usually span a number of information and knowledge driven activities, including

- identifying reliable information,

⁴ http://www.semagix.com/solutions_circas.html

⁵ FinCEN's 314(a) Fact Sheet, Financial Crimes Enforcement Network, <http://www.fincen.gov>

⁶ Executive Order on Access to Classified Information
<http://www.fas.org/sgp/clinton/co12968.html>

- converting it to a usable form,
- comprehensively analyzing it with respects to mandated and optional objectives,
- identifying relevant actionable information, and
- promptly providing information or action directives to those who need it most, document the results and following up with actions varying from notification, prevention to enforcement.

The problem facing users of risk and threat assessment solutions is that, the information that powers such systems is derived from multiple sources--typically have to be sourced both sourced internally and externally, and is heterogeneous. The challenge becomes how to drive information relevance in much focused domains and then score that information, making it available consistently and in a timely manner.

Information is the key to practically all risk and compliance processes. Following observations outline the complexity of any information processing support for vast majority of applications we outlined above.

- The type of information spans data in its raw form or factual information, as well as domain knowledge and policy descriptions
- Information (data and knowledge) is distributed with the enterprise and information providers, as well as across the open Web. Furthermore, there are different levels of autonomy and control over information sources, varying from internal and proprietary, licensed and subscribed, government and non-government agency supplied as well as open unrestricted information sources.
- Information is heterogeneous in format (unstructured in different file and application specific formats, semi-structured including static and dynamic web pages, and structured including traditional databases)
- Information is often of poor quality and of varying reliability (“Data is difficult to access, and even when it is accessible tends to be dirty or downright inaccurate” [Butler 2005])
- Information is static, time sensitive and dynamic (e.g., news and reports are made available any time), knowledge changes (a new hawala scheme is identified, an organization is added to a black list, policy is updated).

Traditional search techniques do a poor job in supporting risk and compliance applications because of the lack of context, often returning irrelevant or too much information, and without proper ranking or prioritizing. To address this problem space, there is a need to move up the continuum from pure data, to traditional search, to intelligent search utilizing metadata, semantic categorization and finally custom ontologies.

3. BEYOND SEARCH -- TO ANALYTICS VIA INTEGRATION

It is important to note that performing a good search (even when dealing with all varieties of information above) is not sufficient, and that analytical capabilities are critical for these class of applications, without which humans would be inundated with lots of irrelevant information and would not be able to implement policy or regulation uniformly across the organization. Thus the organizations who started with providing their employees the ability to crawl data sources or launch search queries against multiple web sites and data sources have quickly realized that they cannot scale effectively. However, to effectively carry out analysis, we need to integrate heterogeneous multi-source information. In other words, applications encompass search, integration as well as analytics in a highly complex information system. In this context, risk and compliance applications impose much more demanding requirements compared to a vast majority of traditional IT applications that address well defined problems in well controlled environments with limited types of information. Thus, compared to mainstream applications such as inventory management, customer relationship management, order fulfillment and human resource management, risk and compliance applications share more characteristics with the new breed of applications such as business intelligence and knowledge discovery, while not being limited to already integrated (e.g., warehoused) internal and structured data sources.

Analysis of heterogeneous information in these applications involves linking information conveyed by separate independent sources. Furthermore, identifying what is an interesting, important or material link (relationship) is the key. For example, EU Third Money Laundering Directive requires that banks formally introduce a “risk sensitive” approach to customer identification. Also necessary is the ability to focus on critical insight and drill down to arbitrary levels of detail, and translate the insight or discovery into action.

Beyond these unique challenges, these applications do share requirements posed on other enterprise applications, such as ability to do process request in batch mode, scale to millions of documents and gigabytes or terabytes of data, maintain and provide provenance of information, support the workflow that can be adapted to suite organizational structures as well as changing regulatory directives, recording in the process each critical activity for auditing, and so on.

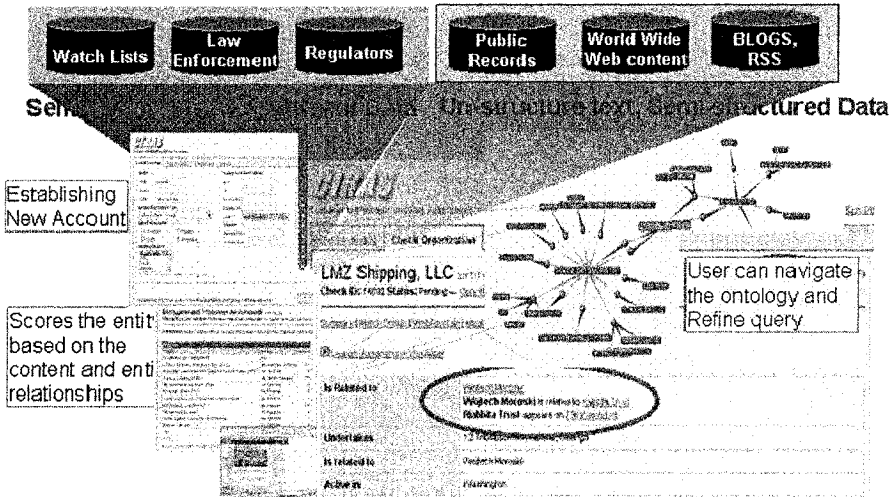


Figure 1: Schematic of a Fraud Prevention application showing heterogeneous information and ontology driven analysis

4. THE APPEAL OF SEMANTICS, ONTOLOGIES AND SEMANTIC WEB TECHNOLOGIES

There are several conceptual and fundamental reasons why semantics, ontologies and the Semantic Web technology are quite possibly the best match for risk and compliance applications.

Relationships are at the heart of semantics. For example, RDF, which is considered as a baseline (representation language) for the Semantic Web, treats relationships as the first class object, which more traditional data representation (e.g., relational model or XML) does not. For a risk and compliance application, linking relevant entities and information is at the heart of required analysis. So the Semantic Web technology is well suited to support this requirement.

It is well known that a syntactic approach grossly fails to make heterogeneous information useful, and that syntactic metadata adds very limited value. A semantic approach is necessary to integrate heterogeneous information. It is very difficult to directly analyzing heterogeneous information, so a more appealing approach is to create semantic metadata which describes information at a more uniform level of abstraction, is domain specific and contextually relevant (as supported by an ontology).

At a more fundamental level, identification or extraction of semantic metadata require two core capabilities: entity recognition/identification (recognizing an object of interest, such as name, organization, event, etc.) and

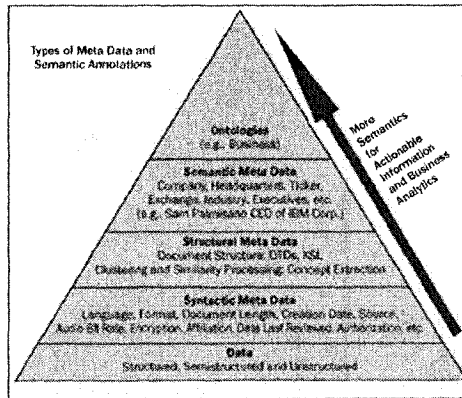


Figure 2: Metadata Semantics (From Syntax to Semantics) [Sheth 2003]

semantic disambiguation (are two objects with the same syntax -- name or description -- also same in the real world or are they different? If the ontology knows of two Bob Smiths, who does the mention of “Bob Smith” in a text refer to? Is Tiger Woods mentioned in the marketing context or the golf context?). Disambiguation is also a critical capability necessary to help build large populated ontologies, as well as deal with dirty data or conflicting information. These capabilities are important building blocks of any Semantic Web technology for enterprises.

Ontology is at the heart of the Semantic Web approach. Ontologies populated with domain knowledge become the key differentiator and enabler for core capabilities that are made possible by what we call explicit semantics (based on formal languages and domain knowledge), compared to implicit semantics (often based on statistical and learning techniques). Ontology and semantic metadata also play a critical role in defining and using context. Context enables scoring and ranking of the most important information and the analysis in help building a 360 degree perspective on an object of interest.

5. TECHNICAL CAPABILITIES OF AN ENTERPRISE SEMANTIC WEB TECHNOLOGY/PLATFORM

We first briefly discuss semantic capabilities, followed by the enterprise software capabilities, both of which are a necessary part of an enterprise grade semantic technology.

5.1 Semantic Capabilities

Earlier we described an excellent match between a semantic approach with the requirements of risk and compliance applications. The corresponding application development lifecycle is depicted in Figure 3. A Semantic Web technology needs to support the following features and capabilities.

Design ontology schema: Ontologies necessary to support most enterprise applications are highly focused. They may be partly based on industry metadata standard but often require customization with respect to coverage and depth. We have not found a practical technology to automatically design such ontologies. So the only practical solution is to use a graphical ontology design tool.

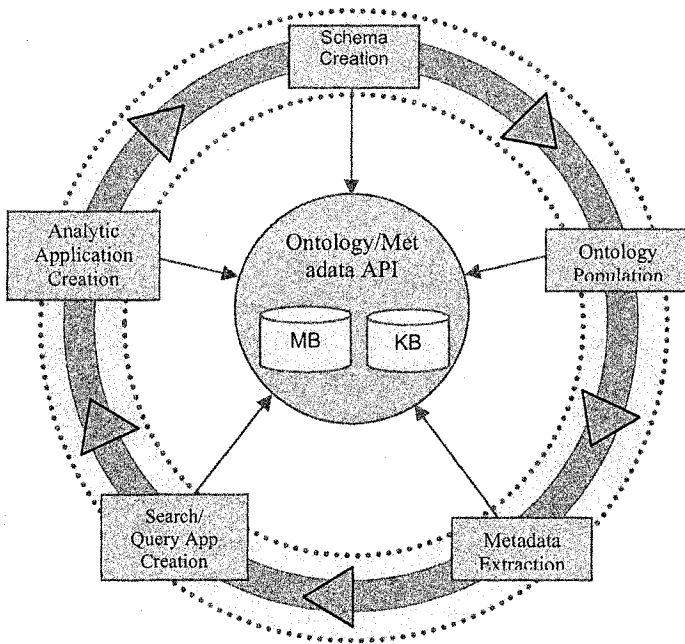


Figure 3: Semantic Application Development Lifecycle

Automatically Populate ontology with domain knowledge: Finding ontology for an enterprise application that is populated with less than a million facts (assertions, entities and relationship instances) is more an exception than a rule. Occasionally, ontology sizes approach 10 million instances. Often data (typically factual information) to populate an ontology is extracted from several trusted knowledge sources (usually data creators/aggregators to provide licensed or subscription based data, such as WorldCheck or Factiva). While knowledge sources provide structured or semi-structured information,

high quality disambiguation techniques including rules that exploit provenance and trustworthiness of data, is critical for the success of automation necessary for such scales. Often it becomes necessary to use specialized disambiguation techniques and tools for matching or comparing names of persons and organization, addresses, and other types of objects. It is interesting to note, as an aside, that this approach to development of ontologies is significantly different than the social and consultative committee oriented process that is used in the development of some of the important biology ontologies and knowledgebase, such as GO and UMLS. The latter takes many years of committee effort and many million of dollars. Most ontologies for supporting industry applications need to be developed in less than three months, and are narrower in scope or coverage (focusing on an application or a class of applications). Human involvement in commercial ontology development is some times indirect - it is in the creation and curation of high quality data provided by knowledge sources, but this cost is shared across many enterprises that license or subscribe such data.

Freshness of ontology

Most customers require ontologies to be updated at the frequency ranging from daily to weekly. Occasionally substantial portions of an ontology may need to be refreshed and repopulated.

Ontology Browsing and Visualization

While software identifies actionable information or provides initial insight, it is often necessary for humans to browse, validate and drill down information. Furthermore it is necessary to be able to easily see original source of information or raw data, as well as traverse related interlinked data and information.

Semantic Metadata Extraction from Heterogeneous Information

A broad variety of heterogeneous information as discussed earlier needs to be processed to extract the semantics with the help of an ontology, resulting in semantic annotation or semantic metadata extraction. Although third party tools are available to deal with proprietary file formats and text conversion, processing unstructured data presents the most challenge. Automatic classifications of unstructured data can improve search, but otherwise have little value in analyzing information. Our experience shows that statistical and learning techniques (including clustering, SVM) are of little value by themselves, and that populated ontologies (i.e., a knowledge based approach) provide the most important basis for entity identification/recognition to extract metadata that is of particular interest to the application. Again disambiguation techniques are also important here. Availability of schema or discernable structure in structured and semi-structured data make is somewhat easier to ingest and process them for metadata extraction.

Semantic Query and Rule Processing

To enable analytical processing, the Semantic Web technology needs to provide comprehensive API for manipulating metadata and ontology, supporting the ability to efficiently process graph oriented information (including graph traversal and path computation). A number of research systems exist for RDF data storage and query processing, which are also likely to be part of future commercial systems (given numerous various of RDF query languages, completion and recommendation of SPARQL by W3C will accelerate commercial support). Support of complex queries involving both metadata (of heterogeneous data) and ontologies—for example, find the stories on the competitors of Intel (where metadata indicates the company that a story is about, and competition relationship is available from the ontology)—is especially important. For performance reasons, Semagix Freedom (a semantic application development platform from Semagix [Sheth et al 2002]) also uses main memory query processing techniques, as traditional database query processing does not given adequate performance.

Reasoning and Analytical Processing

Two types of information processing are possible. If a formal representation such as description logic (e.g., OWL) is used, inferencing is possible. However, in the context of risk and compliance application, the predominant requirement for analytic processing translates to graph oriented or link traversal type of processing. Inferencing based on subsumption does not help. Furthermore, analytical processing can be of the investigative type or the discovery type. Majority of analytical processing today is investigative type, and involves specification of rules identifying links, relationships or patterns of interest and importance. Efficient graph traversal and rules processing is thus an important capability needed for today's advanced risk and compliance applications. Discovery type of processing is an important area of research [Anyanwu 2003] and its support is in its infancy in the current commercial Semantic Web technology.

5.2 Enterprise Software capabilities

Semantic Web technology provides the cutting edge capability needed for risk and compliance applications, and in fact offers critical differentiation. At the same time, it is necessary to support capabilities enterprise users require and demand. Among the capabilities needed include both generic capabilities as well as vertical market specific capabilities. Examples of generic capabilities include:

- flexible, intuitive and highly functional user interface,
- user management (users have different levels of authority, some information is only visible to supervisors and some tasks can only be performed by authorized personnel),

- batch processing (that ability to submit a number of application queries that are then broken up into a series of tasks including semantic query),
- session management (many tasks can be interrupted and the ability to resume at a later stage is important),
- scalability (in many respects, include ability to ingest very large amount of data and large files),
- robustness with round the clock processing support (hence minimal maintenance window, and a need for redundancy),
- system monitoring, reporting, single sign-on and security, and
- use of enterprise class platforms and development methodologies.

Additionally, enterprise software also needs to deal with technology, domain and market specific capabilities. Examples of technology specific capabilities are the support for W3C standards such as RDF and OWL for Semantic Web and WSDL and SOAP for service oriented architecture.

Every risk management project and every enterprise has its own definition of what it perceives as risk. Their perception of risk is best conveyed by means of business rules that can define different scenarios and the corresponding score/action if that scenario is true or false. This calls for a comprehensive risk scoring framework that supports risk specifications which often vary drastically across projects. Also necessary is an ability to support flexible workflows that respect organizational constraints and domain or application specific routing of work, including the ability to deal with escalation of cases and exceptions.

Additional examples of domain and market specific capabilities include name normalization, identity verification, etc. One important capability is that of accessing multiple external systems, often providing the same type of service. For example, ID verification and address verification may be performed by one or more external solution providers. If there are more than one ID Verification Services, the system also needs to perform on-the-fly disambiguation of all the query results.

6. CASE STUDY: CIRAS

Regulations like European Money Laundering Directive and Section 326 of the USA Patriot Act require that financial institutions implement an Anti-Money Laundering (AML) solution. When it comes to money laundering, prevention is definitely better than cure. Detecting it after the event is simply too late, and the consequences can be devastating – both financially and in terms of an enterprise's reputation. While meeting compliance requirements and eliminating money laundering, a comprehensive Know Your Customer

(KYC) process is increasingly valuable [Levy 2004], both in terms of push (as governments introduce increasingly stringent regulations demanding that financial institutions know their customers) and pull (since a richer understanding of an organization’s customers creates enormous business opportunities – in terms of modeling new services to the market at large). The Semagix Customer Identification and Risk Assessment Solution (CIRAS) is an example of a comprehensive semantic technology based solution that enables an organization to quickly and easily identify high risk customers, provides comprehensive analysis tools to perform end-to-end knowledge discovery, vastly reducing the compliance risk to the organization.

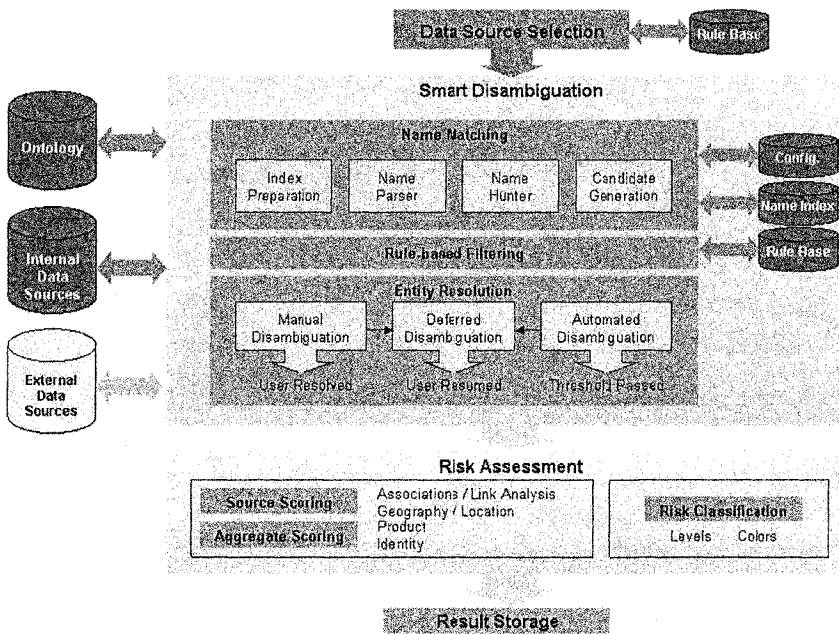


Figure 4: CIRAS KYC process

In order to implement a KYC process successfully, organizations have to bring together vast amounts of very disparate data about their customers. More importantly, though, they need to be able to make sense of all that data and content. Figure 4 shows a schematic of the KYC process engine CIRAS supports.

CIRAS is implemented using the Semagix Freedom semantic application development platform [Sheth et al 2002] with origins in the research at the LSDIS lab of the University of Georgia. Freedom provides the ontology-driven application development platform. Both the semantic and enterprise software capabilities are extensively utilized in realizing the CIRAS KYC

process. The following provide additional context on some of these capabilities:

- Ontology development including knowledge base population and automatic refresh from multiple trusted knowledge sources. This involves use of external sources (as required by the organization) such as WorldCheck, OFAC, and Factiva. These contain names, organizations, aliases, watch-list membership, associations with other individuals, and other information. While ingesting relevant data using extraction agent software to populate the risk ontology, the underlying semantic technology needs to support a comprehensive disambiguation capability, including rule-based techniques. Extraction agents also run periodically as scheduled or on demand to update the ontology based on updates to the knowledge sources.
- Process/analyze wide variety of heterogeneous, multi-source information, including unstructured information (text documents, reports/documents in 150 formats), semi-structured information (Websites, emails), and structured information (databases and XML feeds) for metadata extraction as well as adapters to query data sources on-demand
- Integration with external and third party services such as ID verification (to find if a named entity is that of a recognized real world entity) and custom name matchers using flexible adapters
- Semantic processing capabilities including: entity recognition, entity resolution/disambiguation (covering scenarios such as automated disambiguation (threshold resolved), manual disambiguation (user resolved), deferred disambiguation (user resumed); risk assessment scoring using source scoring (e.g., based on geographical location), aggregate scoring (link analysis and associations), and risk classification using custom rules, provenance, etc.

In summary, unique market conditions, importance of linking and analyzing heterogeneous data, and other advanced technical requirements related to the risk and compliance applications have provided an excellent show case for the emerging Semantic Web technology. Such experiences in building semantic applications using enterprise class software is sure to lead to further successes in many other markets.

Acknowledgements

Special thanks to colleagues at Semagix, specifically Tom Golding for pointers to information on the risk and compliance market, Vernon Lun for

observations on industry and application requirements, and Yash Warke for product information and the figures.

References

[Aleman et al 2005] B. Aleman-Meza, et al, An Ontological Approach to the Document Access Problem of Insider Threat, Proceedings of the IEEE Intl. Conference on Intelligence and Security Informatics (ISI-2005), May 19-20, 2005 (Springer LCNS v3495, Apr 2005, Pages 486 – 491). [OpenURL](#)

[Anyanwu 2003] Kemafor Anyanwu and Amit Sheth, “The ρ Operator: Discovering and Ranking Associations on the Semantic Web,” The Twelfth International World Wide Web Conference, Budapest, Hungary, May 2003, pp. 690-699. <http://lsdis.cs.uga.edu/library/download/AS03-WWW.pdf>

[Avant et al 2002] David Avant, et al., Semantic technology applications for homeland security. In *Proceedings of the Eleventh international Conference on information and Knowledge Management*, McLean, Virginia, USA, November 04 - 09, 2002, ACM Press, pp. 611-613. <http://doi.acm.org/10.1145/584792.584893>

[Butler 2004] Martin Butler, Uncertain Compliance, Information Economics Journal, September 2004, pp. 7-9, http://www.stercomm.co.uk/download/sterling_commerce_IEJ_SEP04.pdf

[Miller 2005] Eric Miller, The Semantic Web is Here, keynote at the Semantic Technology Conference, San Francisco, CA, March 8, 2005 <http://www.w3.org/2005/TALKS/0308-SEMWEB-EM>.

[Lee 2005] Yvonne Lee, Apps Make Semantic Web a Reality, SDTimes, April 1, 2005 <http://www.sdtimes.com/article/story-20050401-05.html>

[Levy 2004] Larry Levy, Semagix AML Solution - Value of a risk approach, Financial Solutions International - Spring 2004. http://www.semagix.com/documents/FIJArticle_FSI9.pdf

[Ruh 2004] William Ruh, The Web of Meaning: The Business Value of the Semantic Web, talk at the W3C 10th Anniversary (W3C10) Symposium, December 1, 2004 <http://www.w3.org/2004/Talks/w3c10-WebOfMeaning/Originals/Ruh.ppt>

[Sheth et al 2002] Amit Sheth, Clemens Bertram, David Avant, Brian Hammond, Krys Kochut, Yashodhan Warke, Managing Semantic Content for the Web, IEEE Internet Computing, July/August 2002, pp. 80-87

[Sheth 2003] Amit Sheth, Semantic Metadata for Next-Gen Enterprise Information Integration, DM Review, July 2003.

[Sheth 2005a] Amit Sheth, From Semantic Search & Integration to Analytics, in Semantic Interoperability and Integration, Dagstuhl Seminar Proceedings 04391, Internationales Begegnungs- und Forschungszentrum (IBFI), Schloss Dagstuhl, Germany, January 1, 2005.

[Sheth 2005b] Amit Sheth, Semantic Technology in the Real World: Experiences and Lessons based on Enterprise Semantic Applications, <http://www.semagix.com/documents/AmitSheth-SemTechConf-Mar082005.pdf>

Brief Bio:

Dr. Amit P. Sheth is a co-founder and CTO of Semagix, Inc., a professor of Computer Science and the director of the Large Scale Distributed Information Systems (LSDIS) lab at the University of Georgia. He is the Editor in Chief of the International Journal on Semantic Web and Information Systems. His research has led to two companies, three major commercial products, two patents numerous commercial applications, and over 200 (co-)authored publications. He has (co-)chaired or organized 20 international conferences and workshops, has served on over 100 program committees and serves on five journal editorial boards. He received BE from BITS, Pilani, India, and MS and PhD from the Ohio State University, USA. <http://lsdis.cs.uga.edu/~amit>.