

Christian De Sainte Marie

Program Manager, IBM

Title: Three good reasons to adopt the W3C RIF

Abstract:

The purpose of this talk is to convince you that W3C new recommendation, the rule interchange format (RIF) is easy to understand and fun to implement, and, last but not least, useful in many ways.

First, I will introduce the purpose and the architecture of RIF, present the current set of specifications and give a short tutorial on RIF syntax and semantics: that is the easy part.

Then, I will discuss some implementation issues, and I hope that you will bring in more: that is the fun part.

Finally, I will discuss use cases and usage issues, such as how to combine RIF with various kinds of data and data models: that is the useful part.

I will conclude with a discussion of future prospects and extensions, and that you should implement, deploy and use W3C RIF.

Dr. Jacob Feldman

Founder & CTO, OpenRules

Title: A Simple Inference Framework for Connecting the Dots

Abstract:

Tom Davenport wrote on January 8, 2010 after the notorious “underwear bomber” attack:

“How easy is it to connect the dots? Granted, there were numerous indications of Abdul Mutallab's evil intent. But it would have been difficult to put them together before the flight. Combining disparate pieces of information about people — whether they are customers or terrorists — is akin to solving a complex jigsaw puzzle.”

In this presentation we will introduce a simple yet practical inference framework for the creation and continuing development of various “connecting the dots” systems. At the heart of the framework is an “always running” inference engine that can accept new facts; propagate them through the existing knowledge base; solicit new facts if necessary; and, finally, reach a conclusion by connecting all the facts together. The framework does not invent a new “magic” technology but rather integrates well-proven techniques and expert knowledge in an ingenious manner. These techniques include: a message broker with a pub/sub mechanism, a time manager, rule engines, finite state machines, multiple event channels and such optional components as a CEP engine and a questionnaire builder. A key feature of this framework is its ability to quickly incorporate new event types and new facts along with the supporting state machines and processing rules into a perpetually running system. More importantly, all these operations could be done in whole or in part by subject matter experts directly in Excel with only a limited need for participation by programmers. We will demonstrate the framework using a complex loan origination scenario where new facts about a loan become known as a loan request moves through the loan approval process. The framework development is still in an experimental phase so the author counts on active audience involvement with constructive criticisms and a lively exchange of ideas.

Alex Guazzelli

VP, Analytics, Zementis

Title: Follow Your Rules, but Listen to Your Data

Abstract:

Business Rules are ubiquitous today. They manage the day to day operations of thousands of companies worldwide. From stocking to maintenance, rules are an integral part of the way we do business in the 21st century. This kind of knowledge, known as Expert Knowledge is forged from years of experience, or what turned out to be the “logical thing to do”. However, along with the information age, more and more data started being gathered all over the world about the processes and services we as a society came to benefit from. In this sea of data, predictive algorithms were designed to extract its hidden patterns, i.e. knowledge that is hidden from the human eye. This is known as Data-Driven Knowledge. In an ideal world, business rules and predictive models live side by side benefiting from each other since both encode complementary types of knowledge.

In this talk, I will be presenting ADAPA, a decision management platform which seamlessly merges the power of rules and predictive analytics to offer advanced decisioning capabilities. ADAPA is centered on open standards like the Predictive Model Markup Language (PMML) and available as a service on the Amazon Elastic Compute Cloud. Scalable, secure and cost-effective, anyone can now access the power of intelligent decisions which seamlessly combine predictions and rules to enhance business decisions.

Rolando Hernandez

Founder, Agility Alliance
Knowledge Engineer, Visible Knowledge
CEO & Enterprise Rules Architect, BIZRULES

Title: Using Enterprise Architecture and Knowledge Engineering to Build Mobil's Lube Knowledge Base and Knowledge Supply Chain

Abstract:

The most valuable business assets are knowledge and wisdom.

Their loss can be devastating: Three years after British Petroleum's top corrosion expert retired, the job was still open because his expertise was hard to replace. That's when pipeline corrosion caused BP's huge oil spill in Alaska.

Businesses compete on knowledge. Corporations often claim that their people are their biggest asset. If that is really true, then the experts with critical knowledge and wisdom must be the most valuable business assets of all.

Mobil Oil took a different approach to managing knowledge and intellectual capital. Mobil developed a global program to manage and automate knowledge using expert rules, business rule engines, and knowledge bases. Mobil created a knowledge supply chain to preserve, share, and automate their most valuable assets: knowledge, wisdom, and intellectual capital.

Every company has a knowledge base and knowledge supply chain. They may exist only in the minds of your top experts. Mobil wanted to transform tacit invisible knowledge into explicit visible knowledge that could be managed and automated.

Mobil developed a process to capture knowledge and model how experts think: how they solve problems, make decisions, give expert advice, offer recommendations, configure products, troubleshoot equipment failures, stop oil spills, fight fires, and so on. EA models were used for knowledge transfer and to preserve knowledge. Knowledge-based systems and rule-based systems were used for knowledge automation. This preserved Mobil's most valuable asset and strengthened the knowledge supply chain.

In this session we will present a case study examining how Enterprise Architecture, Knowledge Engineering and Business Rule Engines helped retain and make available the knowledge of world-class experts at Mobil. We'll show how EA and KE helped create value by bringing together process, technology, and human expertise to acquire, model, and automate critical knowledge about Mobil Oil product recommendation, equipment troubleshooting, and problem diagnosis.

This session will show how Enterprise Architecture and Knowledge Engineering was used to retain, preserve, and share the knowledge of world-class experts at Mobil Oil. Join us to discuss the lessons learned and value created by building the Mobil Lube Knowledge Base and Knowledge Supply Chain. We will show you how you can preserve your most valuable asset. The lessons learned apply to any big company in today's global knowledge-based economy.

Hal Hildebrand

Architect, Oracle Corp

Title: Building a Distributed Rules Engine: an Experience Report

Abstract:

The problems, pitfalls and successes in building a functional distributed rules engine using Jess is described. The system, called C0, is an autonomic system for managing, monitoring and provisioning a large scale distributed system. This talk describes the architecture, the challenges faced in distributed rules processing, as well as the performance of the system under real world load scenarios.

David Holz

Director, Grindwork Corporation

Title: When Truth Goes Bad

Abstract:

Logically confusing cases can crop up in rule usage far easier than introductory materials tend to admit.

What a user expects of rule engine behavior, and what a rule engine expects of its user, differs significantly until the user has run the gauntlet of initiation by debugging various edge cases they didn't even know existed. This routinely happens in cases that the user rightfully expected to be straightforward.

This talk covers some of the messier aspects of dealing with rule-driven truth; how to think about and handle truth duplication, correctly fire side effects, and integrate new facts that might conflict with prior inferences. We take a step back and think about what a logical user's generally simple expectation is in these cases, and what patterns must be used to correctly communicate "gotcha" situations to various types of rule engines.

Stan Belport

Belport Consulting

Title: Using a Declarative Rule Server for Cyber Security Compliance Management Suite

Abstract:

Belport Consulting's Stan Belport talks about developing the CSCMS using declarative rules in place of hand-writing server-side code directly. CSCMS provides nuclear power generating stations with the filing and maintenance needs for cyber security compliance for the Nuclear Regulatory Commission as well as various state and other federal agencies.

Mack Mackenzie

CTO, Starview Technology, Inc.

Title: Emergence of a Unified edBPM Platform:

Abstract:

Discussing the core requirements and semantics of both rule-based event processing systems and workflow business process management systems and show how unifying them can enable powerful next-generation use cases like real-time datamining and emergent pattern detection.

Carole-Ann Matignon

President & CEO, Sparkling Logic, Inc.

Title: <not yet supplied>

Abstract:

Many practitioners fumble when they start new projects. They struggle to find the right methodology, they sometimes apply design patterns that are plain wrong for rules technology. As a result, projects often fail to deliver the expected result may it be \$\$ or in terms of performance or maintainability. Practitioners, on occasion, flee the technology when new projects are decided.

The purpose of this talk is to address the practical challenges of new projects: how to start on the right foot, how to reduce the risk of using a “wrong design”.

Dr. Karen Meyer

Principal Scientist, SRI

Director of the Intelligent Mixed-initiative Planning and Control Technologies (IMPACT) program within the AI Center at SRI International

Title: Learning by Demonstration Technology for the Masses

Abstract:

While the concept of learning by demonstration has been around for many years, recent advances in artificial intelligence have led to a resurgence of work in the area. We describe the development and application of learning by demonstration technology to support user creation of automated rules and procedures for a rich collaborative planning environment that is in widespread use by the U.S. Army. User feedback and evaluation results show that the technology can be used effectively by the target user community and that it can improve the speed and quality of performance for a range of critical tasks.

Hafedh Mili

University of Quebec in Montreal

Title: Towards a Business Rule Capability Maturity Model

Abstract:

The business rules approach provides many advantages over more traditional development methodologies, including:

- 1) forcing business to make the business rules under which the organization (presumably) operates explicit
- 2) providing a common language for business and IT
- 3) providing much enhanced traceability between the business intent, and IT system implementation
- 4) enhancing the agility of the IT infrastructure.

The successful adoption of the business rules approach requires an agile approach to tackle a limited functional area within the context of a pilot or showcase application to be used as learning ground. However, such a SWAT team approach does not lead to a successful and mature implementation. Indeed, it does not scale, and its benefits, even within a single application, do not last. In this talk, we present a methodological roadmap going from the initial opportunistic SWAT team approach to a leveraged and optimized approach, within the broader context of a business rule capability maturity model (BRCMM). In particular, we will relate our BRCMM to SEI's CMMI and to recent work on enterprise architecture.

Chandra B. Mouleeswaran

Consultant, Cupertino, CA

Title: Integrating Classic Rule Systems with Modern Scalable Architectures

Abstract:

This work presents a systematic approach to addressing real time problem-solving by integrating classic rule-system technologies and algorithms with the modern scalable computing architectures. In particular, we are interested in domains which have both symbolic knowledge and vast amounts of numeric data.

Several artificial intelligence (AI) techniques presented in this work have always existed in literature. For some time, AI lost its luster in most commercial engineering applications and appeared only in gaming. As the amount of domain data became explosive, many of the classic AI methods have become very pertinent to modern applications. Rule engines have become popular again, and are used quite widely in several applications.

Two major themes are presented in this work. Firstly, building a rule system needs careful attention to the process of knowledge acquisition, the nature and chaining of rules. Secondly, rule systems will grow and update in two major directions: acquisition of new knowledge and preservation/aging of old knowledge. It is not sufficient to quickly integrate over-the-shelf tools without following proper methodologies. Otherwise, maintainability and robustness of the system cannot be guaranteed.

Using a self-directed investor as a vehicle for understanding, this work presents how each piece of information in the decision-making process can be mapped and a system built using available open source tools. Both scalability and improving performance over time are addressed. An integrated framework will be demonstrated that minimizes the time taken for equity research and handling of vast amount of facts making the entire process more efficient and reliable. The problem-solving methods used in this example easily carry over to other domains such as Sports, Education etc. where rules-of-thumb are known to exist.

James Owen

Consultant, Fort Worth TX

Title: Extending Rulebased Forecasting

Abstract:

Beginning with early agriculture and continuing through the stock market today, forecasting has been considered a “black art” associated with a grizzled old statistician tossing chicken bones on the floor and mumbling something in an arcane language known only to himself and his apprentice. Actually, forecasting, like rulebased systems themselves, is part art but mostly mathematical analysis of historical facts to enable us to predict what might happen in the future. Using a rulebased system is relatively new in that it extracts the logical rules from the control and math routines. This paper is an extension of an earlier work in 1992 by Armstrong, Collopy and Adya. This paper will examine how constantly varying weights on various measurements can lead to better forecasting and how a rulebased makes that task much easier and simpler.

Carlos Serrano-Morales

CTO, Sparkling Logic, Inc.

Title: Extracting Knowledge and Rules From Big Data and Large User Communities

Abstract:

The gradual disappearance of expert systems and knowledge management systems from the public eye, and the significant success enjoyed by big data supported statistical approaches such as those used by Google, have contributed to a current focus on data mining and knowledge extraction through machine learning. At the same time, user feedback is progressively sought out as part of applications, and the collective input thus provided significantly enriching the data produced.

This talk will take a look at some of the approaches used to extract knowledge from both big data and large user communities. It will also touch on how that knowledge can be operationalized through rules engines, and the architectural implications of supporting the corresponding development and management activities.

James Taylor

CEO & Principal Consultant, Decision Management Solutions

Title: Decision Services Need More Than Rules

Abstract:

If you are not using business rules to automate business decisions then you should be. And there are some things you really need to know:

- Decisions matter more than rules do
- Execution is less important than management
- Your data knows where (some of) your rules are
- Data depth can improve its width
- You can't use the same rules on every transaction

This session assumes you know what a rule engine is and how it works. You'll get a quick primer on Decision Management and concrete advice on why decisions matter to rules engines, suggestions on how to use data mining and predictive analytics, and an illustration of why adaptive control matters to anyone doing rules.

Paul Vincent

CTO Business Rules and CEP, TIBCO Software

Title: Practices and Experiences of Rules in CEP Applications

Abstract:

Complex Event Processing has emerged as a recent “technology area” that provides an event-focus for software systems. However CEP is not itself a single technology – it exploits and adapts many existing technologies such as stream processing, model-driven engineering, ... and rules.

As in any other form of IT, “rules” of various forms are regular features in CEP systems. Their roles include direct pattern detection (for example using ECA and Rete-based production rules), as well as representing the “decisions” downstream of the detection of some pertinent business event. The use of any particular rule technology (or not) is of course dependent on the use cases being implemented – for example, stream processing remains a favoured technique in algorithmic trading against market data streams, whereas rule processing is the favourite for more generic business event processing applications in domains such as logistics, transport, telecommunications and marketing.

In this session we look at the

- Typical use cases for CEP, to provide a basis for the motivation for rule usage.
- Types of rules (by role: pattern detection and decision, and by technology: ECA, inference and decision rule execution engines) used in CEP technologies and why, with their associated properties and common alternatives,
- Rule characteristics by type around latency, throughput and reliability, and the various roles they fulfil in a CEP system.
- Some of the extensions made to rule technologies such as production rules to better support event processing needs.
- A review of some customer use cases and design patterns around the use of rules in CEP.

Barbara von Halle

Managing Partner, Knowledge Partners International

Title: The Decision Model for Business Rules and Logic: An Enterprise Analysis Technique

Abstract:

Decades ago, the Relational Model changed the way we perceive and manage data. Now, in 2010, the Decision Model is doing the same for business decisions and rules. Business analysts and business SMEs are using it today to represent business logic and rules as a rigorous, predictable normalized structure. Developers are translating that structure into code in a straightforward manner. Consequently, for the first time, a business's logic is explicit with its own unique look-and-feel. Inconsistencies, gaps, and redundancies become obvious. This shift is having a profound effect on certain industries already and new software is emerging.

The timing is intriguing because Decision Management itself is gaining momentum, according to Harvard Business Review (Tom Davenport, "Make Better Decisions," HBR, Nov 2009).

This presentation is based on the recent book, *The Decision Model: A Framework for Business Logic Linking Business and Technology* (von Halle and Goldberg, 2009, Taylor & Francis LLC, ISBN 1420082817). It also includes lessons learned from *The Decision Model in practice*.

"This book can become one of the classic books of a new era in computing that will have much traction in the next few years"

- Opher Etzion, Ph.D., IBM Master Inventor.

The attendee will:

- Witness the inherent structure, integrity and normalization principles lie hidden in business logic today (as they were hidden in data years ago)
- See how the Decision Model delivers business rules and logic as a tangible and manageable business asset
- Know how Decision Models have instigated creative business thinking and decisions
- Recognize business decisions comprise a new dimension and model
- Follow the step-by-step creation of a Decision Model using an agile iterative approach.
- Envision the integration of the Decision Model, visualization, and data models into requirements practices.

Luke Voss

Founder, Mindviews Labs LLC

Title: User-Mediated Rule Engine Execution

Abstract:

The majority of discussed rule engine case studies involve scenarios where the rule engine is expected to operate without direct user intervention. However, there are scenarios where it is useful to include a user in the rule engine execution loop. One case where user interaction can be useful is in the operation of expert systems that have not yet reached maturity. Such novice systems do not possess the expertise necessary to solve the desired problems independently. Lacking for explicit rules to drive the rule engine execution to completion, human expert users can be included in the process to provide solutions to intermediate steps in the rule execution flow. This paper will discuss the practical approaches to goal formulation, goal seeking, and knowledge base modification to allow user-mediated direction of rule engine execution in the context of a rule-based expert system for spacecraft system design.

Peter Voss

Founder and CEO, Adaptive A.I. Inc. and Smart Action Company, LLC

Title: Artificial General Intelligence Engines, and Modes of Learning

Abstract:

The field of Artificial General Intelligence (AGI) concerns itself with systems that can learn a wide range of novel tasks. In contrast to Narrow AI approaches, AGI engines' core competency is the acquisition of knowledge and skills. Supplemented by built-in capabilities such as feature detection, inference and memory, these systems can potentially bootstrap their overall 'intelligence'. This talk includes a review of several years' research, culminating in the commercial application of this technology. Various rule-based knowledge acquisition and reasoning systems were utilized in this work.

Andrew Waterman

Luis Garcia Barrios

El Colegio de La Frontera Sur, México

Title: Rule Capture and Dynamic Analysis from Divergent Stakeholders by Means of Simulations, Rule Based Gaming and Rule Repositories

Abstract:

One of the greatest problems surrounding social ecological modeling is that of discovering and cataloguing the rules that divergent groups use to manage a natural resource. This problem is similar in nature to that of “knowledge capture” in the realm of expert systems. However, in the case of natural resources, there are often many experts; from different scales and often with divergent opinions. Ostensibly, most of the rules that govern a resource are captured in regulations or laws regulating a specific natural locality. In practice, however, this is not the case, as local users have their own way of exploiting the resource which may or may not conform to these regulations. These rules are often what truly drive the manner which a natural resource, or collection of natural resources (a reserve) is managed in the field.

Companion modeling is a methodology that seeks to allow the true expression of the social and economic rules that govern a natural resource. To this end, role playing games and multi-agent-simulations have been shown to allow greater understanding of the rules which communities use to access (and govern) resources as well as the effects of their exploitative strategies over time. Although experiments built around such work have captured social interactions around specific problems; there is a lack of detail around the social processes that have led to the mutation of proposed or understood rules from one stakeholder group to the next. In our experiments, we have been party to this process of multi-stakeholder mutation as seen during game play and implemented during game development; in particular, during our work with Manantiales.

As this work has occurred through active, *en vivo*, role-play; mutations to rules proposed by ourselves (researchers) or other stakeholders (campesinos, government functionaries, conservationists) have not been captured in a quantitative context. Instead, changes have been incorporated into the game/model as part of the development process and play has continued with these same and different stakeholders; making rule creation and mutation a social amalgam that effects the play and understanding of our agroecological model.

This paper proposes a method for recording such changes by the following means: 1) the development and application of domain specific languages (DSLs) to the description of social/economic rules which govern natural resources and that can be easily understood by all stakeholders; 2) the use of such DSLs for the implementation of social and economic rules that govern Manantiales and other game/simulations that we develop in our work in the field, and 3) the use of source control techniques available in modern rule repositories to compare mutations across time in a quantitative manner. The idea of a “base game” or simulation independent of social rules is also proposed; by which stakeholders could use a sample DSL to create social and economic rules to govern a resource’s exploitation by communities of interest.

We are interested in discussing how such methods can make use of methodologies and processes developed by the business community in terms of editing, maintaining and testing business rules from repositories over time. We shall speak to the possibility of basing our DSLs on RuleML, a rules

based variant of UML, and what possible consequences such an approach might have while working with such divergent populations as those that make up the stakeholders in the Tablón watershed.

George Williamson

Union Pacific Railroad

Title: Rules-Based Applications: Design Decisions

Abstract:

Decisions governing the operation of the Union Pacific Railroad have traditionally been made by human experts following various legal, regulatory, contractual, and policy-based business rules. Automation of these decisions in an IT system naturally lends itself to the utilization of rules-based technologies. This presentation will examine some of the key factors and lessons-learned that have played a major role in the design and development of UP's rules-based applications. Topics include:

- What role rules-based technology plays in the IT toolbox
- How business rules are categorized based on complexity
- How scalability is achieved
- How the system enables business users to maintain the rules without IT involvement
- What role does rules-based technology play in real-time, event-driven processing
- How all these factors combine to determine architectural decisions

High-level use cases of rules-based applications at Union Pacific will also be described to demonstrate how rules-based technology might be applied to solve real-world problems.

We are still awaiting abstracts from:

Karen Mayer

NB: We need titles for a number of presentations - see above.

Abstracts are included as comments in the Title column

There are 17 presentation slots, not including the keynote by Prof. John Laird