

Ontologies as a New Cost Factor in Enterprise Integration

Report Workshop 3/Workgroup 2

H. Ted Goranson¹, (Ed.), Bei-tseng Chu², Michael Gruninger³,
Nenad Ivezic³, Sern Kulvatunyou³, Yannis Labrou⁴, Ryusuke Masuoka⁵,
Yun Peng⁶, Amit Sheth⁷, and David Shorter⁸

¹Old Dominion University, USA; ²University of North Carolina, USA; ³National Institute of Standards and Technology, USA; ⁴FIPA and PowerMarket, USA; ⁵Fujitsu Laboratories, USA; ⁶University of Maryland Baltimore Campus, USA; ⁷University of Georgia, USA; ⁸IT Focus, UK, tedg@sirius-beta.com

Abstract: see Quad Chart on page 2

1 INTRODUCTION

The following Quad-Chart (Table 1) summarizes the work of the group. It identifies the approach taken to address the issues of infrastructures for virtual enterprises exploiting agent technology and proposes future work on ontologies and thereby addresses the issue of model complexity and costs.

2 BACKGROUND

The benefits of modeling are widely recognized. They are significant enough to have driven widespread implementation and the support of substantial theoretic, implementation and user communities. Initially, the benefits of modeling were in identifying and clarifying a process so that it could be formally improved. These models formed the basis of a reusable knowledge store on processes and the science behind process engineering. Since then, two major extensions to this start greatly extended the benefits.

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: [10.1007/978-0-387-35621-1_43](https://doi.org/10.1007/978-0-387-35621-1_43)

Table 1: Working Group Quad-Chart

EI3-IC Workshop 3 <i>Interoperability of business processes and enterprise models</i>	Workgroup 2: <i>Ontologies as a new cost factor in enterprise integration</i>	2002-February-6/8 Gaithersburg MD USA
<p>Abstract: The workgroup focused on key barriers to enterprise modeling for process and system optimization. Overcoming these long-lived barriers requires some new approaches and the workgroup settled on the introduction of ontologies. Several problems and new approaches were explored. Some reasoned speculations resulted, together with proposals for testing their validity. The problems concerned lowering the cost of building and changing models; building and using component libraries; linking to non-process (like data and product) models; and furthering of the agent notions of prior workshops.</p>	<p>Major problems and issues:</p> <ul style="list-style-type: none"> – The cost of enterprise modeling is too high - to model, store, maintain, validate, change, and to inter-operate with other models. – The modeling function has unpredictable costs. – To accomplish the type of integration needed the models are too complex to use. – The latest trend in modeling and interoperability is the need to create ontologies. We need to determine what they are, are not, and how we make them share information. 	
<p>Approach:</p> <ul style="list-style-type: none"> – Study the model and ontology relationship with respect to cost, effectiveness, and ease of use – Follow the promising thread of the previous two workshops regarding knowledge bases and autonomous agents – Consider the difficulty of representing product and process information and moving that data among a product's life-cycle phases. – Use a paradigm in future work that assumes that systems and their models will be able to self-integrate and share information among each other once a context is defined. 	<p>Results</p> <ul style="list-style-type: none"> – Model complexity is due to the semantic content that overloads models. Move semantics into ontology. <p>Future work:</p> <ul style="list-style-type: none"> – Study the role of ontology to reduce the complexity of models – Assess the cost versus benefits as the amount of formal rigor is added to ontologies – Explore the process of separating the semantic aspects of models from the non-semantic aspects – Introduce ontologies into an agent's use of knowledge needed to accomplish its functions. – Explore a security scheme that doesn't prevent sharing information but requires an ontology to assign a context code to each parcel of information to be shared. 	

One extension is the combination of the process and the model into a control model. The model contains the extra functionality to apply its intrinsic algorithmic representation of the process to actually do the work or control the work of the process. The other development is so-called "enterprise integration," where the processes are interlinked in a global framework. Then many of the engineering techniques can be applied at the system level

for optimizations and benefits otherwise impossible. The benefits of process modeling in these two contexts are significant.

This sketch of the history indicates a problem. The early process modeling methodologies were relatively lightweight. Usually they were targeted for use by process owners as a tool to represent their own activities and to engineer them. Over time, as the utility of these models have grown, and the benefits have ballooned, the modeling tools and methodologies have become increasingly complex, formal and costly.

The problem is now not a matter of benefits. These are very large. But so are the costs and risks of modeling. As a result, modeling is not widely employed and enterprise modeling by integration even less so.

- It costs too much to do the modeling in the first place. The tools and methods are so complex and constrained now that two experts must be involved: the process expert and the modeler.
- It costs too much to develop, maintain and practically use libraries or repositories of processes, including components associated with control and best practices.
- It costs too much to validate the model in terms of the interactions among others.
- It costs too much to adapt a model once conditions change. This includes the re-engineering, reintegration and revalidation.
- It costs too much to integrate with other methods. Some methods impinge on the control function, such as the selection of components and their validation in a software engineering sense. Some other methods are “outside” the process scope, such as strategic, financial/legal and marketing concerns.

With each of these costs are attendant risks that increase the worst-case cost. So not only are costs high, but they are unpredictable as well. The group determined that this cost issue is the greatest single barrier to gleaning the benefits of enterprise integration. Additionally, the group concluded that the root of the problem was in “overloading” the modeling function. On the one hand, modeling methods and tools should be intuitive, flexible, easy to use and experiment with. They should be as diverse as the applications, which is to say very diverse and domain-specific. But on the other hand, they should be extremely rigorous and standardized. But accommodating the latter, the former is compromised.

The group proposed taking a look at separating these two needs — allowing the actual modeling tools to be lighter weight, closer to the problem, easier to tinker with, more diverse and specialized. The technique is to propose the use of ontologies for the formal rigor and multiple benefits of standardization. Those benefits are primarily found in ease of integration and enabled component libraries for reuse.

3 ONTOLOGIES

Ontologies are a relatively recent evolution in the modeling and representation family tree. As with many animals in this family, definitions differ. Most definitions share two features: ontologies are formal, explicit conceptualizations and their role is as abstract specification of shared concepts. The latter characteristic places it in the position of empowering the collaboration of models and modeling activities — ontologies are to models in a rough equivalence to what models are to the real world. Except ontologies generally attempt to define a complete domain, and models often focus on elements within a domain.

To support this role, ontologies are generally more abstract, formal and axiomatic than models. Some ontologies are more quick and dirty, and the boundary between ontological engineering and modeling (or knowledge engineering) is fuzzy. But in general, ontologies support the work of representation (in models and the like) and models support work on the real world. As noted, the group explored the possibilities of overcoming the cost barriers of enterprise process models by apt use of ontologies. The general idea is to remove a lot of the restrictions and overhead from the modeling layer into the ontology layer.

Expected benefits will occur on the modeling layer as noted. But greater efficiencies and capabilities are expected at the ontology layer as well. There is universal concurrence in the group that overhead associated with the mechanics of ontologies for the expected needs will be less than similar functions currently supported at the model level. These needs are centered on knowledge capture, but include reuse needs.

4 SUGGESTED ACTIONS AND PROJECTS

The discussion of the workgroup was centered on a number of high value problems. These were considered in turn, and specific actions identified for tests, research or further discussion. Each of these is presented below, in no particular order.

4.1 Project 1: Component Reuse and Tweaking

Problem: This project began with a discussion of the relative advantages of ontologies. The group noted that ontologies are proliferating in apparently much the same way as modeling and representation methods. Some in the larger community question whether anything fundamental is solved by the introduction of ontologies. Perhaps the diversity and need for harmonization

will be just as great but effectively complicated by an additional layer, so the speculation goes.

The group noted that the diversity of ontologies is fundamentally different than that of models and model methods. The latter has two primary dimensions of diversity, a “vertical dimension” of different levels of rigor, and a “horizontal dimension” which reflects the different needs of various domains. By definition, ontologies don’t have this horizontal diversity.

The question then focused on the problem of the proliferation of ontologies in this vertical dimension. At the highest level of formal rigor are approaches like PSL (Process Specification Language); toward the middle of the spectrum is DAML (DARPA Agent Markup Language); toward the bottom are many web-oriented taxonomies. The workgroup believes there is a simple cost-benefit function at work here: the “cheaper” approaches provide accessible benefit. One presumes that the market uses the cheapest approach that solves the problem, and that general perceptions are that the benefits of additional rigor rise less fast than their costs.

The workgroup believes that the cost/benefit curve is reversed, that the benefits of additional rigor grow faster than the additional burden they impose. But this has not been yet demonstrated. A project is proposed that would do just this: audit some benefits of rigor in the context of costs. If the benefits of additional rigor are shown to be much larger than the additional cost, some collapse in the ontological spectrum of diversity are anticipated.

Approach: The project has a narrow scope. An existing set of many integrated models of processes is presumed to exist. The project does not care how they can to be in this stable, functional state. The focus is on two processes, illustrative of many in the enterprise. To be interesting, these are presumed to be in different virtual enterprise components (meaning different companies), they involve some transactions/interactions that involve “soft” elements (trust, uncertainty, tacit knowledge), and the two models were created using different modeling approaches.

The following action will be compared in three modes: without ontological support; with “medium-weight” ontological support; and with PSL ontological support. Suppose that the first model (M1) must be changed in response to a change in the real world; either a change from “below” (within the process) or “above” (from the larger enterprise context). M1 uses some means to peruse a component library supported by an “intermediate” ontology, O0. The mechanism of this perusal is not in the scope of this project, and has been broken out as a separate project below (see “Component Pattern Strategies”).

M1 selects a model component that is a near match for its need. O1 provides a means for mapping that component into M1, identifying the tweak needed, and supporting the tweaking itself. O1 and O2 support a similar

process in M2, which must adapt in response to the changed M1. Then, validation of model interoperability occurs, again supported by the ontology level.

Expected Benefits: The experiment specifically evaluates:

- The cost of making the model changes
- The cost of revalidating the model within itself for consistency
- The cost of revalidating the interoperability of the two models

Specifically excluded from the exercise is validating the correctness of the model(s) in representing the real world. That was judged to be a relatively straightforward process, which will not be affected by the introduction of ontology. Instead, the focus is on the validated interoperability of the models. There is some expectation that some substantial portion of the integration now supported by model frameworks will be better handled by this mechanism.

Second order insights expected from the experiment are:

- Suggestions for a new paradigm for enterprise model integration
- Insights into enterprise component libraries and reuse strategies
- Indications of research needs
- Perspectives on standards tasks
- Links between process model integration and popular parallel trends, many web-based
- Better understanding of dynamic, multidimensional configuration management of large systems
- Possibly guidance on new lightweight, intuitive model freedoms

The experiment is expected to be explored as a “paper” experiment before being fielded on a test bed.

4.2 Project 2: Component Pattern Strategies

Problem: This second project is a companion to the one just described. It focuses on the characterization of the various constructs that are evident in existing and developing enterprise-modeling methods and tools, recognizing that the collection of constructs evident in each method/tool are in some sense coherent. They reflect a particular intellectual approach to the problem of enterprise modeling, and hence the form in which they are characterized and represented to the end-user (the enterprise modeler) reflects this view of how the world is. In this sense the ‘constructs’ represent potential stepping-stones between the putative common ontological underpinning whose cost-benefit will be assessed in Project 1, and the enterprise model to be constructed.

The project focuses on how such stepping stones might be deployed in the (probably Web-enabled) tools that a modeler might use to access compo-

ment libraries, to choose components from different libraries derived from various methods and tools, and to advise on tweaking strategies to support model functionality and interoperability. This approach will complement the lingua franca approach being developed in the UEML project for an Enterprise Model Interchange Format (EMIF).

Approach: Therefore this project focuses on the boundary between M1 and component libraries (or frameworks) through support of the ontologies. The basic idea is that dividing the “old” modeling function into two layers (one formal, one lightweight) is not a trivial job, and some attention must be given to intuitive access of higher level functions. In a functional sense, an intermediate or bridging layer is proposed. The purpose of the layer is to provide a formal metaphor for indexing components and for presenting and managing that index to modelers.

It should also provide placeholders for the propagation of constraints derived from the ontology (or construct to construct mappings), e.g. valid and necessary relationships, restrictions on involved objects and/or their attributes.

This notion grew in importance as the group worked through some scenarios. Given the existing investments in and fragmented usage of existing methods and tools, it is clearly necessary to find ways of increasing tool adoption and encouraging a stronger enterprise modeling community to emerge. This is the issue that a stepping-stone-based component pattern strategy will seek to address. The group also reviewed a collection of de facto and de jure standards that will guide this effort.

Expected Benefits: The specific aims of the project are to:

- Define the functionality of such an interface, including the possible reuse of metaphors from software and system engineering (and similar)
- Explore the use of constraint representation and propagation to support semantic decomposition, semantic aggregation, and more ambitiously semantic mappings
- Explore the direct applicability of existing and emerging standards to support this interface. Certainly to be included will be the standards for model constructs (ENV 12204, under revision), EDOC, BPMI.

A test scenario is likely to result in an additional activity of the project noted above. That scenario may be centred on M1 locating component fragments from two independent (stepping-stone) repositories, combining these in some predictive or computer-assisted way, and then tweaking as already described.

4.3 Project 3: Extending the Previous Workshops' Approach

Problem: Workgroup 1 from the first workshop (in Paris) proposed a new perspective on enterprise integration that focused on integration model frameworks to support knowledge management. In this view, “registering” a model in a context was conflated with “situating” an element of information in the enterprise context as knowledge. Some promising notions were developed.

Workgroup 1 from the second workshop (in Singapore) started from this notion and explored a more rigorous approach involving agents. Key notions from that work involved rationalizing transactions among agents in the virtual enterprise context with process state management and process adaptability. The combined insights of the two workshops showed some promise toward a new approach to agent based virtual enterprise system optimization.

This project extends the work of those two groups by introducing ontologies to the existing notion of actor-model pairs (See “A Merged Future for Knowledge Management and Enterprise Modeling Agents” and “Advanced Virtual Enterprises: Needs and an Approach” in this volume for more detail.) A fine level of granularity is defined by the constraint that an actor can support only one message operation, the messages being based on standard speech act transactions. A set of meta-actors can change processes.

A high value problem concerns the means by which these meta-actors can operate. In the early vision of the Singapore workshop, these processes were allowed to be ad hoc. But in the context of the present workshop the suggestion was to employ a component library as outlined in the two projects above. This project extends the work of those two in the specific case of the agent model.

Approach: What is new in this project is the introduction of global process ontologies for agent evolution. An interesting feature is the discrimination between component libraries for processes that do work and the processes that improve those. The group expects that this distinction will ease the problem of identifying the nodes where standardization can support the entire system. In particular, one set of “pattern” indices might be maintained for process components and another set of indices for:

- Means of ordering searches among components in libraries
- Measuring algorithmic fit between needs and those components
- Drawing from a set of tweaking, assembling and validation processes on those components to meet the requirements
- Perform all of this by autonomous agents working at the bottoms-up level

Expected Benefits: The result of the project is expected to be an agent-based self-evolving index strategy that supports the knowledge management paradigm of the Paris workshop in a virtual enterprise context.

4.4 Project 4: Interfacing with Product Life Cycle Information

Problem: This project arose from interaction between the two workgroups. Workgroup 1 focused on the global problem of integrating the integration frameworks of product and process information through the product and enterprise life cycles. This was seen as a significant, important problem that has been vexing enterprises for decades. Piecemeal solutions appear to be making the problem worse.

The combined workgroups suggested that the new element of ontologies might provide a new tool. Several approaches were suggested. One that seems promising was to consider product data in the context of process data. This is a departure from the current approach, and especially in the case of the U. S. Department of Defense. That default approach posits a substantial distance between the two, in part because the military sees product data as a manufactured item itself that is bought and delivered. Product data carries certain legal assumptions, like ownership and portability that are not as explicit in the case of process information.

Process modeling has other fundamental differences from product data modeling. Modeling of processes is much more difficult, owing to the need for control, and the explicit management of state. Moreover, processes tend to be more self-referential, for instance processes are expected to evolve more opportunistically than products and the means by which a process adapts is a process itself.

In the combined group session, some suggested that the theoretic foundations related to process representation have developed much more rapidly than similar foundations on the data side. This was not a universally held position, but arguments in favor cited the advances in situation logics, the widespread implementations of agents and the new techniques of software engineering that have no counterparts on the data side.

Approach: The approach on the product modeling side has been centered on creating standards that are at the same “semantic level” as the data. The PDES/STEP standard is of this type, but it seems an unwieldy solution that is not widely employed, and does not gracefully extend to process information. The suggested approach adopts two new elements:

- To leverage ontologies where possible — in fact, to move as much of the focus of standardization efforts from the domain-specific to the generic. There is growing appreciation for this approach already.

- To “situate” product data in the context of process data. This is the novel notion of the workshop. Generally, every element of product data, especially data that is maintained in parallel with the product, exists to support a process. But the group offers a more specific notion of bonding, resulting from the overarching notion of enterprise integration.

A primary motivation behind enterprise integration is to provide a philosophy for creating enterprise-wide frameworks. These frameworks tend to use normal transaction boundaries as leverageable division within the framework. The information passed at these transaction boundaries corresponds well to the information, which concerns the product data community. Therefore, one would expect process-oriented integration to show promise for integrating all the information in the enterprise.

This project will:

- Take a fresh look at the problems of life cycle product data integration by defining the actual needs. This would go well beyond the simple call for interoperability and define the particulars of why data needs to interoperate, in what context, and for what ends.
- Examine the possibilities of emerging process integration frameworks to address this requirement. Particular attention will be given to the emerging notion of ontology-supported, scalable component-based frameworks as outlined in the workshop. Some detailed attention will be paid to context-based use strategies, and new notions of integration-supported knowledge management. (The benefits of knowledge management will be added to the benefits of “old-fashioned” integration in the problem statement.)
- Develop an approach of process-context based life cycle integration that is a strong balance of addressing needs and leveraging existing efforts in the standards community
- Run some “paper experiments” to evaluate the relative benefits with a focus on hidden infrastructure costs, lowered costs of adoption and increased benefits (particularly scalability)

Expected Benefits: The project is not expected to result in an action plan for life cycle product and process information integration and knowledge management.

4.5 Project 5: Supporting a New Paradigm for Trusted Systems

Problem: The workgroup recognized that the new approach to enterprise integration would not only address traditional concerns, but also provide new types of benefits not previously recognized. Problems of security came up,

as they always do. The basic problem is that integration is all about appropriate sharing. Current paradigms of security are based on engineered denial of sharing. Usually, the integration framework operates behind a wall of multilevel access gates, which are designed and imposed relatively independently.

The topic arose as a desire to engineer security needs into the integration strategy. In the course of discussing this approach, a radical new approach was incubated. This project was defined to take that relatively speculative notion and to explore it in some detail to discover associated practicalities, boundaries, benefits and constraints.

Approach: The underlying integration paradigm of the workgroup is one based on dynamic, context-sensitive integration. Perhaps this will be agent-supported. Certainly, it will be ontology-driven. In this view, local context and (transaction-oriented) reference frameworks apply to incoming information. It is registered in the local context for local use, and in the system context for global analyses and optimization.

In other words, models are “decoded” for their value in the global context of the enterprise. The advanced case has individual models that may be structured in different ways, so a different model’s perspective very literally needs the global framework to “decode” it. In the normal case, the models are obfuscated by the accidents of domain preference and expediency. But the models could be obfuscated deliberately. The group proposed just this.

This new paradigm is not based on denial. Everything is shareable. Some information is modeled using situation or state-dependent ontologies. “Sense” can only be made of each component by situating it (via ontology-based integration). The net effect is that no piece of information needs to be protected, but no one can use a piece without being in the proper context. That context may include familiar certification mechanisms but can allow other types of possibly more robust protections. The obfuscating methods may include familiar numeric processing, but certainly will support additional or alternative methods.

This project is more speculative than the others noted. The method of the project is to convene a workshop of experts in logic (especially situation theory), non-deterministic abstraction (category theorists), ontologies, process engineering and advanced thinkers in enterprise integration.

Expected Benefits: The goal of the project is to explore the idea, map possibilities and define a research roadmap. A secondary goal is to increase the benefits of system-level, ontology-based integration by adding in support for trusted systems.