Improving Information Quality and Consistency for Modeling and Simulation Activities

Roy Scrudder
Applied Research Laboratories
The University of Texas at Austin
Austin, Texas
scrudder@arlut.utexas.edu

Dr. W. Henson Graves, Tom Tiegen, Chris Johnson
Lockheed Martin Aeronautics Company
Joint Strike Fighter Program
Fort Worth, Texas
henson.graves@lmco.com, tom.s.teigen@lmco.com, john.c.johnson@lmco.com

Steve Hix
Paradigm Technologies, Inc.
Arlington, VA
hix@pti-usa.com

James W. Hollenbach
Simulation Strategies, Inc.
Washington, D.C.
jimh@simstrat.com

ABSTRACT

Modeling and simulation (M&S) has taken its place as a key enabler in all phases of today's military systems development and fielding, from research and development through training and operations planning. The quality and value of M&S activities are dependent not only on the quality of M&S software, but also on the information that drives models and simulations. When M&S is employed in a large-scale enterprise, data dependencies among models and simulations emerge. These range from sharing data among models and simulations to using model and simulation outputs to drive other models and simulations.

In the Joint Strike Fighter (JSF) Program, M&S is a key enabler to all activities, from design through training. Hundreds of models, simulations, and modeling environments (such as computer-aided design systems) enable a web of interrelated activities and require a corresponding information flow. The JSF pioneered an information management approach utilizing metadata to facilitate the understanding and appropriate use of data. Traditional metadata management approaches focus on describing the content and format of individual information resources. The unique aspect of the JSF information management approach is the addition of metadata that captures the lineage of data, allowing information traceability from external sources and through the chain of M&S activities.

This information management approach has been incorporated in the Resource Access System (RAS), which is part of the overall JSF product data management system. RAS is used to locate, access, and register data about the JSF weapons system and other systems with which it interacts. RAS uses metadata for information location and retrieval and for validation of data integrity conditions. We defined an Extensible Markup Language (XML)-based data interchange format that is used to package information resource metadata along with the data content and share these across activities.

ABOUT THE AUTHORS

Roy Scrudder is a program manager at the Applied Research Laboratories, Signal and Information Sciences Laboratory, The University of Texas at Austin. He holds a BS Degree in Mathematics from the University of Tennessee. Mr. Scrudder has over 25 years' experience in systems analysis and design with the last 15 years working in support of various data engineering and data management projects in the defense community. He led the metadata modeling efforts for the JSF program.

W. Henson Graves is a Technical Fellow at Lockheed Martin Aeronautics and is the Architect for Joint Strike Fighter information management. He holds a BS degree in Mathematics from Tulane University in New Orleans and a Ph.D. in Mathematics from McMaster University in Hamilton, Ontario. He held a Post Doctoral Fellowship in Mathematical Physics at Queens University in Kingston, Ontario and was a Research Associate at Stanford University in Palo Alto. He has spent over 10 years in realizing the SBA vision at LM Space Systems and at LM Aero. Dr. Graves was Deputy Program Manager and technical lead at LM Space Systems for the Defense Advanced Research Projects Agency (DARPA) Simulation Based Design Program and was formerly Professor of Computer Science at San Jose State University.

Tom Teigen holds degrees in Information Management Systems from Tarleton State University and Agricultural Engineering from the University of Minnesota. He has spent the last ten years designing and implementing database applications in a number of industries.

Chris Johnson is a Senior Staff Software Engineer at LM Aero. He holds BS and MS degrees in Mechanical Engineering from Virginia Polytechnic Institute and State University (Va. Tech) in Blacksburg, Virginia. He has spent 18 years at LM, first at Missiles & Fire Control in Orlando, Florida, and most recently at LM Aero in Fort Worth, Texas. He has worked on aircraft, rotorcraft, missile, avionics, and DARPA programs performing structural analyses and developing tool and data integration frameworks. He has co-authored papers on multi-disciplinary optimization and served in developer and lead roles for design/analysis tools currently in use on aircraft programs at LM Aero.

Steve Hix is the JSF Authoritative Systems Database project manager with Paradigm Technologies, Inc. He holds a BS in Economics from Rollins College in Winter Park, FL, an MBA from Golden Gate University in San Francisco, CA, and a Certificate in Systems Engineering from Old Dominion University in Norfolk, VA. Mr. Hix spent over 26 years in the aviation field as an air traffic controller in the Air Force and FAA and as a navigator and electronic warfare officer in a variety of Air Force aircraft including 12 combat missions in a B-52 during Desert Storm. The last three years he has been involved in database development in the M&S community.

James W. Hollenbach is an independent consultant to government and industry in the areas of modeling and simulation (M&S), systems engineering and enterprise integration. His company, Simulation Strategies, Inc., currently supports the Joint Strike Fighter Program, the Office of the Secretary of Defense and NASA. He chairs the M&S Committee of the National Defense Industrial Association (NDIA) System Engineering Division. Jim previously served in the U.S. Navy, as a naval aviator, air wing strike leader, EA-6B squadron commanding officer and DoD Acquisition Professional. His last active duty assignment was as Director of the Defense Modeling and Simulation Office from 1994 to 1998, where he led development of the DoD and NATO M&S Master Plans and oversaw development of the High Level Architecture for distributed simulation. He holds a BS in Mechanical Engineering from Rutgers University, an MS in Aeronautical Systems from the University of West Florida and is a graduate of the Canadian Forces Command and Staff College and the U.S. Defense Systems Management College.

Improving Information Quality and Consistency for Modeling and Simulation Activities

Roy Scrudder
Applied Research Laboratories
The University of Texas at Austin
Austin, Texas
scrudder@arlut.utexas.edu

Dr. W. Henson Graves, Tom Tiegen, Chris Johnson
Lockheed Martin Aeronautics Company
Joint Strike Fighter Program
Fort Worth, Texas
henson.graves@lmco.com, tom.s.teigen@lmco.com, john.c.johnson@lmco.com

Steve Hix
Paradigm Technologies, Inc.
Arlington, VA
hix@pti-usa.com

James W. Hollenbach
Simulation Strategies, Inc.
Washington, D.C.
jimh@simstrat.com

INTRODUCTION

We live in an information age, and the international defense community has come to recognize the value of information as a resource. The value of information, like that of any other resource, is contingent upon its availability, usability, and the trust that the user is willing to place in the quality of the information. One of the most fundamental challenges facing any program is how to capture and maintain information during program execution. The volume of data on a large program can be in the terabytes and almost certainly is stored in multiple formats in a distributed collection of repositories.

The US Department of Defense emphasizes this focus on information as a resource through its Net-centric Data Strategy (DoD CIO, 2003b). This strategy provides a vision of data that enables effective decisions and data that is "advertised and available for users and applications when and where they need it." A key element of the Net-centric Data Strategy is a focus on metadata (data describing data). The Net-centric Data Strategy provides a shift away from the traditional metadata focus on data standards, where the goal is for information systems to use the same internal representation of data, to a focus on metadata to facilitate the discovery and assessment of information resources.

In this paper, we will describe an approach to managing information resources using metadata. This approach has been described in a series of Simulation Interoperability Workshop papers (Hollenbach & Harnett, 2000; Hollenbach & Graves, 2002; Graves, Barnhart, & Hollenbach, 2003; Scrudder, et. al., 2003). This approach is consistent with the management of information as described in the JSF Modeling and Simulation Support Plan (JSFPO, 2000). We extend the scope of metadata outlined in the DoD Net-centric Data Strategy. The extensions are necessary when a large number of highly related information resources exist. In such an environment, additional metadata is necessary to characterize the lineage (sometimes referred to as pedigree or provenance) of data. An M&S-rich development environment, such as exists on the JSF program, exemplifies this need. Additional metadata is also necessary to fully understand when a given information resource is appropriate for a given use. For data to be appropriate for use with a given model or simulation, it must be consistent with the underlying assumptions embedded in the model algorithms. This idea is a key concept in validation as described in DoD verification, validation, and accreditation (VV&A) guidance (USD (AT&L), 2001).

The information management and metadata concepts described in this paper were developed through the Joint Strike Fighter (JSF) program, specifically for the management of data associated with the product development process. The JSF Program will produce the affordable, next-generation strike aircraft weapon system for the U.S. Navy, Air Force and Marines along with a number of international partners. Three variants will be built – conventional takeoff and landing, short takeoff and vertical landing, and carrier-capable. The JSF program is a prime example of a simulation based acquisition program (U.S. DoD Acquisition Council, 1998; Frost & Thomen, 1999; Lutz & Keane, 1999; Karangelen, 1998; U.S. DoD Acquisition Council, 2000), where literally hundreds of models, simulations, and modeling environments are in use, all of which rely on and produce information resources.

Figure 1 depicts the objective process for managing information across JSF product development activities, including modeling, simulation, and analysis activities. This figure provides an overview of the range of

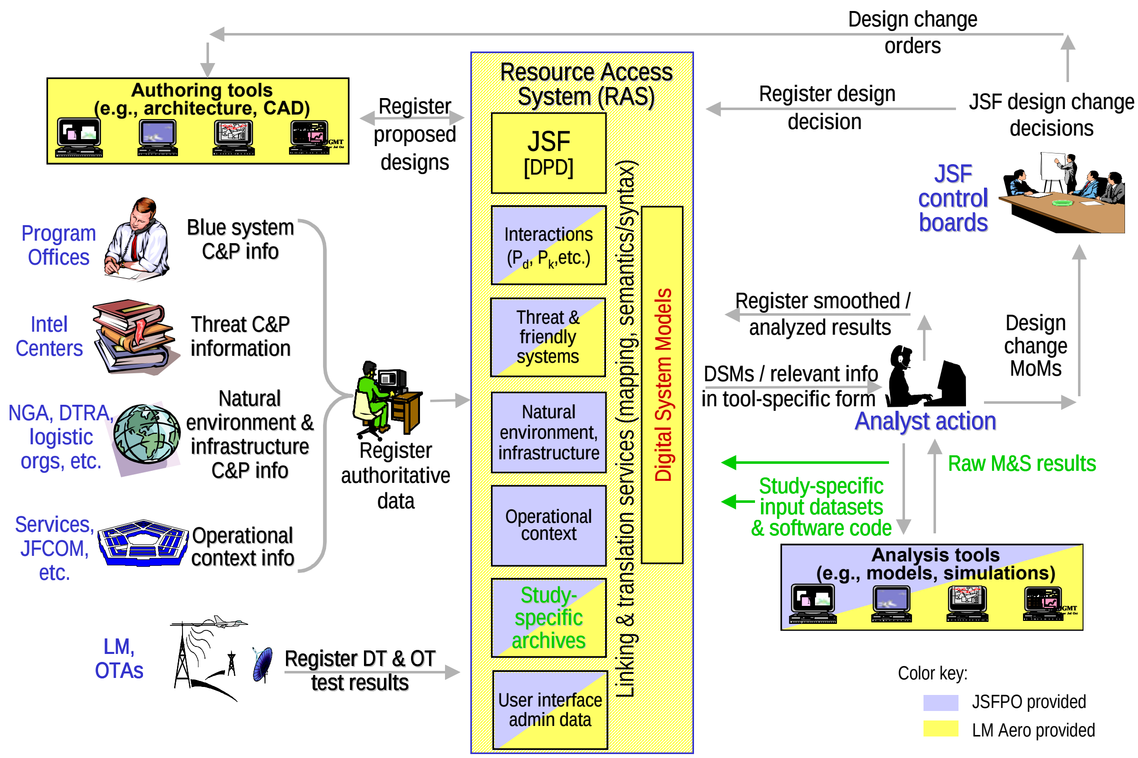

Figure 1. Notional JSF Product Development Information Flow

information necessary to support these activities in a major program such as the JSF. To address the full lifecycle of a system such as JSF, data about much more than just the air system under design and construction must be managed. As can be seen in Figure 1, the scope of data includes data about other systems with which the JSF will operate (both friendly and hostile) and large quantities of data about the environment in which the JSF will operate. This figure also shows the importance of managing the tool-specific datasets used by and produced by modeling, simulation, and analysis activities.

The data needs and outputs of these models, simulations, and tools are tremendous, with significant overlaps and dependencies among the data produced and consumed by the tools. A glimpse into this complexity can be seen in Figure 2, which shows just one example of an analytical data "food chain" across models and simulations used by the JSF program. The types of data and sources of data cover a wide range. Data must be obtained from external authoritative sources (e.g., for threat data and environmental data) as well as produced internally (for JSF product data). The formats of data cover the gamut from structured files of text, to complex binary data, to relational databases.

Although acquisition is the primary focus of the JSF program at this stage, the metadata concepts developed and described here have much broader potential applicability. The same challenges of managing large numbers and a wide variety of information resources pertain to any activity that utilizes M&S, whether it be acquisition, testing, or training. The same metadata and information management concepts that will suffice for a large scale acquisition program such as JSF will prove valuable for managing information resources in smaller scope efforts.

Existing Metadata Standards

Our goal was to use existing metadata standards where they were available and to establish new metadata representations only where no standard existed. We found that standards were plentiful for describing types of information resources, and even for describing individual instances of information resources. The primary information resource metadata standards incorporated into our metadata specification are the Dublin Core Metadata Element Set (NISO, 2001) and the DoD Discovery Metadata Standard (DoD CIO, 2003a), which draws from the Dublin Core and extends it substantially. In the area of VV&A, we made extensive use of the Data Quality Templates from the DoD VV&A Recommended Practices Guide (DMSO, 2001). However, when it came to metadata to describe the lineage of information resources, we

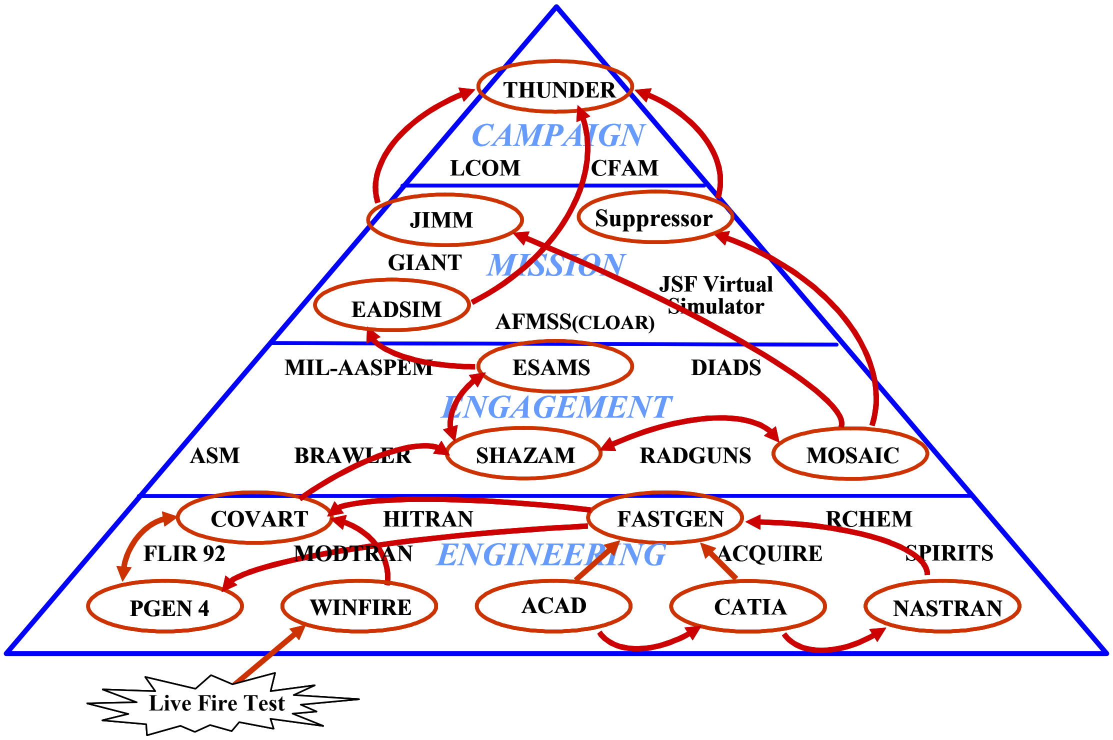

Figure 2. Vulnerability Data Example of Data Dependencies Across Models

found that there was a lack of standards and recommended practices.

Documenting the Metadata

The results of our effort are documented in a formal metamodel specification. A metamodel is a model of metadata classes, attributes, and relationships. The metamodel specification is comprised of the defined metadata terms, a specification of their relationships, and a data interchange format for the transmittal of metadata (optionally along with the content of information resources). In addition to textual definitions of all the metadata terms, this documentation provides a rigorous specification of the metadata structures in the form of an IDEF1X entity relationship model (NIST, 1994). In addition to the IDEF1X model, the metamodel specification provides a tabular listing of the metadata. Although it does not contain all of the relationship information of the IDEF1X diagrams, this tabular representation does allow us to describe additional information including what portions of the metamodel are mandatory and which of the attributes are displayed to the user of a system implementing the metamodel. The data interchange format, known as the Information Set Interchange Format (ISIF), is an Extensible Markup Language (XML) Schema.

METAMODEL ELEMENTS

In the following discussion, we will define each of the major concepts in our metamodel. Each of the concepts will be described, and a high-level IDEF1X data model will depict the relationships between those concepts. Because of space limitations, only the major concepts will be shown (as labeled boxes) in the data model. The full metamodel has a set of descriptors (attributes in IDEF1X terms) that describe each concept. We will discuss some of those in the following paragraphs. The relationships are depicted as lines between the boxes. A dot on the end of a line indicates that there may be many instances of the concepts shown at that end for each instance of the concept shown at the other end. The IDEF1X notation in the diagrams has been extended to show where additional point of contact information (shown as ovals labeled POC) can exist. A point of contact may be either an individual or an organization. Attributes of a point of contact would include a name, phone, address, and email.

Core Elements

The three core elements of the metamodel are information resource, information resource class, and information navigation node. These three elements form the foundation upon which the rest of the metamodel is constructed. Figure 3 shows these three concepts and the relationships among them.

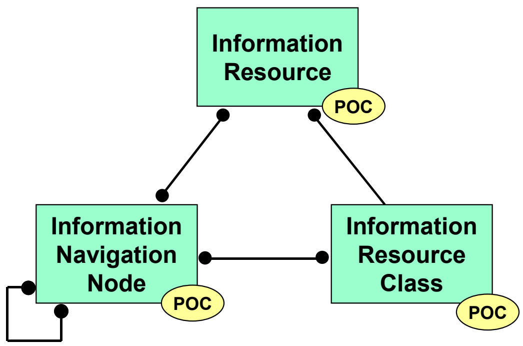

Figure 3. Core Metamodel Elements

Information Resource

An information resource is an individual configuration-managed item composed of data along with the metadata that describes the data. The data content of an information resource may be a set of parametric values (ranging from a single value to an entire dataset of values), an algorithmic representation, an executable model, or a document. While an information resource may be an entire dataset, it may just as easily be used to represent some "slice" (some portion of one or more tables) in a relational database. In this abstract sense, an information resource is declared where a common set of metadata exists. This is a key concept in the metamodel that allows one to apply the metamodel to the management of complete datasets, as well as portions of datasets (down to individual parameter values if appropriate). The individual attributes of an information resource include its name, description, storage location, version identification, purpose, status, and access constraints.

Information Resource Class

An information resource class is a group of information resources that have common properties, relationships, and meaning (semantics). This abstraction allows the user to declare the common characteristics and relationships for entire groups of information resources, such that they do not have to be repeated for each instance of that class (i.e., the individual information resources). The individual attributes of an information resource class include its name and description. Each information resource is a member of a single information resource class. In effect, the information resource class defines a template for a collection of resources. An example of an information resource is type of simulation input file, e.g., a Brawler aircraft file. The information resources associated with this information resource class would be the individual Brawler aircraft files managed by a program.

Information Navigation Node

To support the user navigating a set of information resources, the metamodel contains the concept of an information navigation structure. An information navigation structure is a collection of nodes and the navigation paths between those nodes. As can be seen in Figure 3, each information navigation node may have a set of associated information resources. Likewise, there is a constrained set of information resource classes that are associated with an information navigation node. We refer to the combination of an information navigation structure and an identification of the information resource classes associated with each node in the navigation structure as an information model. While we envision that the primary path for navigation for most information models would be a tree-based hierarchy, the metamodel also allows (by the nature of the many-to-many relationships between information navigation nodes) for additional cross connections between nodes. A common implementation of this pattern is a Web site with a hierarchical set of pages that contain hyperlinks between some subset of the pages. It is important to note that the information model does not dictate where information resources are stored; it provides a way of organizing them for discovery and navigation. The location of an information resource is described as a metadata attribute of the information resource.

Gauging Data Quality and Appropriateness for Use

While the attributes of an information resource may contain a broad statement of the purpose of the resources and the limitations on its use, a more rigorous specification of these concepts is required. This is a major extension beyond what is found in the existing metadata specifications cited earlier. Figure 4 shows how these concepts are represented in the metamodel, and the following paragraphs describe them in more detail.

Data Quality

A quality assessment may be associated with an information resource (or a set of information resources). This quality assessment is the result of VV&A activities and may contain both quantitative

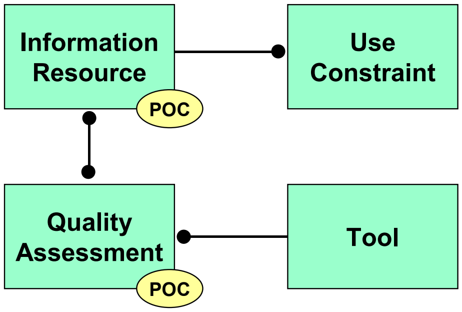

Figure 4. Data Quality and Use Constraints

and qualitative descriptors of data quality. This quality assessment may be in the context of a specific tool. As we will see later, *tool* is an abstraction that includes models and simulations as well as any other tool that may consume or produce data. This association of information resources and models is consistent with the VV&A concept that a model and the data used in a model cannot be judged as valid in isolation. A quality assessment may result in the production of an information resource (e.g., a data quality assessment report that would be managed and described just as any other information resource).

Information Resource Use Constraints

The use constraint portion of the metamodel establishes those conditions under which use of an information resource is appropriate. A use constraint may specify applicability over a ranges of dates, to specific variants of a subject system (e.g., to only the carrier version of the JSF), to a certain operational context, or to a limited range of any specified parameter. An example of this last constraint type would be the case in which a dataset was only applicable for a range of altitudes. An operational context constraint would specify the range of scenario conditions under which an information resource is valid (e.g., only for major regional conflict in a given geographic area). Any information resource may have a combination of use constraints applied.

Lineage of Information Resources

To fully understand the appropriateness of use of an information resource, one must know the sources of information that went into the creation of that information resource. While some existing metadata standards provide the ability to link information resources together, in terms of one resource being the source of another, that information is sometimes lacking. This is almost always the case in the M&S domain. Knowing that a dataset was used to produce other datasets is often not enough; one needs to know the models or simulations that were used to create the dataset. This is potentially the difference between simply obtaining data from an authoritative source and obtaining the *right* data from that source. In addition there are sometimes specifics about the model or simulation run (such as setup parameters or termination conditions) that ultimately determine the appropriateness of the information resource(s) created by the run for other uses.

To meet these needs, the sections of the metamodel shown in Figure 5 were created. The central concept in this portion of the metamodel is a process execution. The information resource source concept in the metamodel is used to capture what process execution(s) resulted in the creation of an information resource. The process execution input portion of the metamodel is used to record what information resources were the inputs to a process execution. For example, if a simulation execution used three simulation input files and produced two output files, the inputs would be recorded as process execution inputs and the two outputs would be recorded as information resource sources.

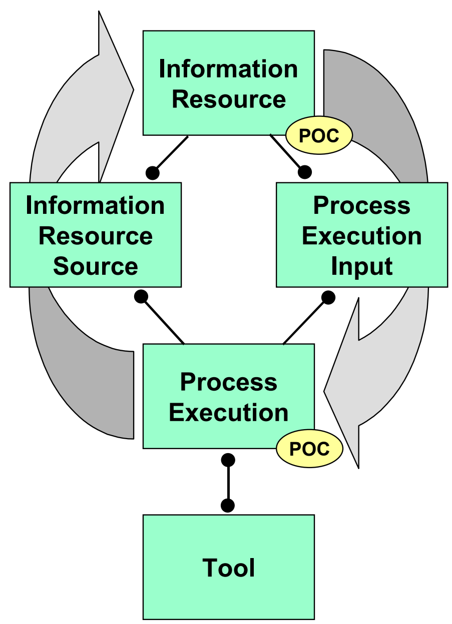

Figure 5. Information Resource Lineage

As mentioned earlier, the metamodel concept of *tool* is an abstraction that includes models and simulations. However, other tools must be considered, such as tools used for data preparation and those used to perform data reduction and analysis. In most cases, the metamodel will be populated with data about process executions that represent model or simulation runs. In this case, the tool associated with the process execution is a specific model or simulation version. However, process execution can also be used to capture information about other processes, such as data preparation from external authoritative sources or data analysis to produce final results.

Assessing Data Consistency

Using the set of source and input relationships defined through the process execution metadata, one can construct a complete model of the lineage of information resources. This is especially useful when the production of information resources from raw data through final analyzed results goes through many steps (process executions), such as was illustrated in Figure 2. This information becomes extremely powerful over time when information resources change. The metadata records can be examined to determine the impact of those changes, identifying where model runs should potentially be re-executed sequentially to produce new end results. Figure 6 shows an example of an information resource lineage tree that could be captured using the metamodel. In this example, we can see that a change in information resource IR_1 has a potential impact on other information resources. In this case, process executions PE_1 and PE_3 potentially need to be rerun to reproduce information resources IR_5 and IR_8.

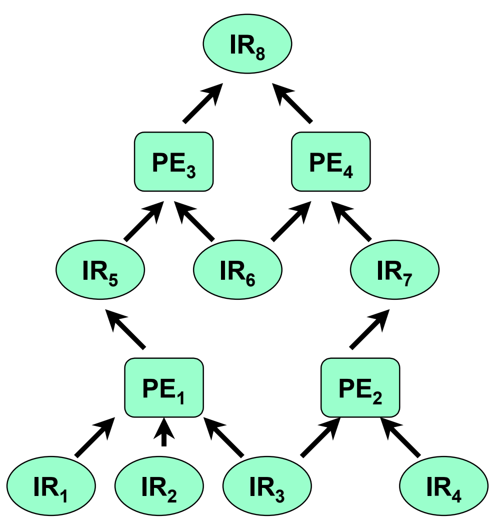

Figure 6. Lineage Example

This example also illustrates how the metamodel is key to addressing data coherency issues. In Figure 6, we see that information resource IR_3 is used in processes PE_1 and PE_2. Thus, we must make sure that a consistent version of the information resource is used in both processes. If inconsistent versions of IR_3 are used for the two process executions, the result is information resources produced higher in the production tree that are not coherent. In this example, the validity of information resource IR_8 is at stake. A concrete example would be to consider where CAD information is used as an input to two analysis chains that eventually produce radar cross section and aerodynamic performance data. The lineage could then merge back together if this radar cross section and aerodynamic performance data were inputs to a higher level engagement model. If different CAD files were used to produce the radar cross section and aerodynamic performance data, then the engagement results would be invalid.

Templates for Lineage

Many modeling, simulation, and analysis processes are not one-time events, but rather recur over time with different sets of data. Training events are a good example of this pattern, where the same set of tools are used to prepare data, conduct the training, and analyze training benefits, but the data used and produced will vary for each training audience.

To support this pattern, the metamodel incorporates the concept of process templates as shown in Figure 7.

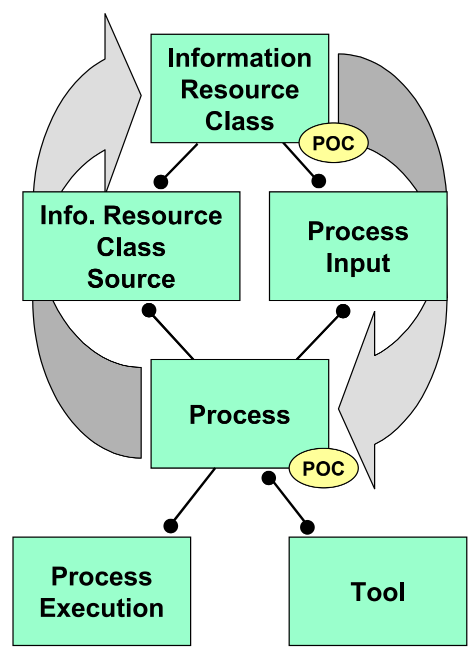

Figure 7. Process Template

The process construct can be viewed as a class, where process executions represent instances of that class. Whereas individual information resources are the inputs and outputs of process executions, information resource classes are the inputs and outputs of processes.

This establishes the "normal" production process for information resources of a given information resource class. Once this metadata is populated, the metamodel user has a powerful quality control capability with which to gauge the completeness of metadata for individual information resources. If one knows that a given process uses one set of information resource classes as inputs and another as outputs, the metadata for individual information resources of those types and individual process executions can be examined and important questions can be addressed:

- If a process execution has been recorded, have all the expected types of inputs and outputs been recorded?

- If an information resource has been recorded has the process execution and all the inputs that were used to produce it been recorded?

This metamodel can also be used to prompt the user for the complete set of metadata when new entries are made. For example, when a process execution is entered, the user can be prompted to enter metadata about the appropriate input and output information resource, based on their process-information resource class relationships.

Access Control Metadata

The final section of the metamodel addresses the permissions for groups of users to modify both metadata about information resources as well as the information resource data content. Figure 8 shows this portion of the metamodel. For all of the concepts shown on the right-hand side of the diagram (information navigation node through process execution), the metamodel contains attributes that capture whether a member of user group has the ability to add, modify, delete, and have visibility of the items of that type. For information resources, the semantics also include whether the user group members have the ability to add, modify, delete, and have visibility of the data content for the information resource. The attributes also establish which user group members have the ability to form associations between the categories of items, e.g., to assign information resources to information navigation nodes.

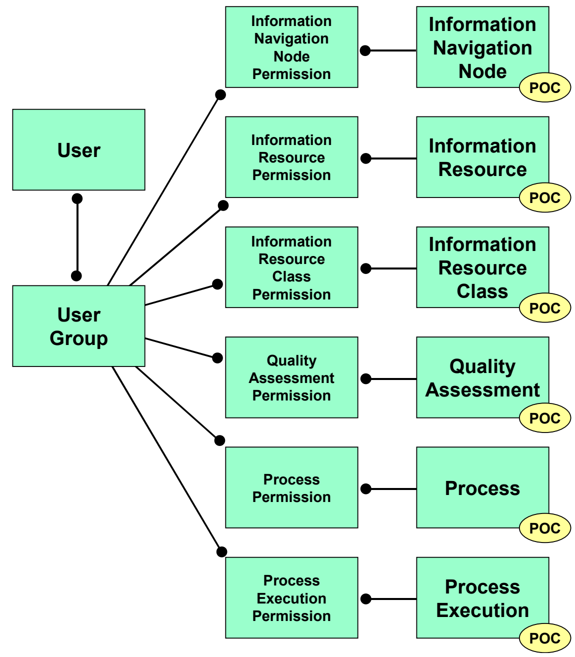

Figure 8. Access Control Metadata

IMPLEMENTING THE METAMODEL

Establishing a metamodel has little value, other than as an academic exercise, until it is actually implemented to manage information. The JSF program, specifically the product developers at Lockheed Martin Aeronautics Company (LM Aero), developed the Resource Access System (RAS) to manage the wide range of product data, and other related data required for analysis, design, production, and testing of the JSF system. RAS is in the early stages of deployment. The RAS uses the metamodel described in this paper as its basis for managing metadata. The multiple repositories that store JSF enterprise information have a wide range of data and metadata storage implementations. These repositories vary greatly in their underlying data models. Figure 9 is a top level depiction of the use of an information model by the RAS system to locate and access data residing in multiple servers on a network. This architecture enables integration of separate repositories and multiple individual information models. As can be seen in Figure 9, the RAS provides a Web-based portal to the metadata and data content managed by a wide range of applications, including the Metaphase product data management (PDM) system, the DOORS requirements database, the PVCS configuration management tools, the JSF Digital Library (JDL), and the threat systems database.

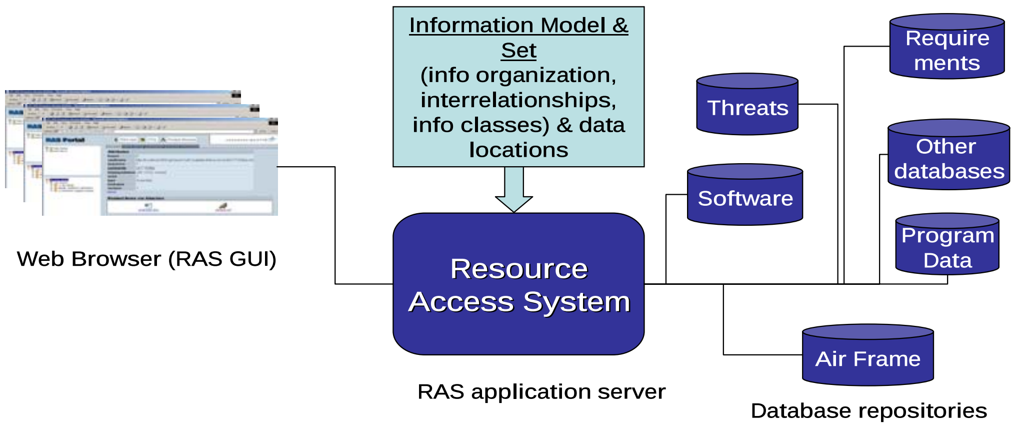

Figure 9. The JSF Resource Access System

The RAS system provides a uniform interface to the configuration-managed data within the JSF program. The interface is consistent, and is based on the metamodel, regardless of the storage location of the data and the system used to natively manage the data. Users are able to access the information resource data with searches based on other metadata attributes such as the information resource class or releasing organization.

The information management reality within most enterprises is that data is distributed and managed in multiple systems. RAS is positioned to connect and bridge individual data servers with three options for maintaining the metadata specified by the metamodel. The metadata can be managed in the individual PDM federates, it can be managed within the RAS application, or the responsibility can be shared among these systems. The obvious advantage to storing and managing all the metadata in the RAS application is that it is a single point of implementation. The alternative of managing all the metadata in the PDM federates requires modifying multiple PDM systems. The current version of the RAS system implements the core of the metadata specification, and alternative architectures are still being evaluated. A necessary (but not sufficient) condition for maintaining metadata consistency across a distributed enterprise is a common metamodel (as we have defined here) and a transmittal format for sharing metadata among information systems.

TRANSMITTING DATA AND METADATA

Since data is typically managed in a distributed manner, a transport mechanism must exist. To elevate this transmittal from the level of data (numbers and values without context) to information, the metadata must accompany the data. In some cases (such as the transmittal of an information model), the transport of only metadata is required. We use the term *information set* to represent this concept of a set of metadata, potentially accompanied by the data content. An information set is a collection that contains a set of related information resources (both data and metadata), information resource metadata, information resource classes, and/or an information model. An information set may take one of three forms:

Content-containing Information Set - an information set that contains both the data component and the metadata component of any information resources it contains. This set may optionally contain an information model that is associated with the information resources.

Content-referencing Information Set - an information set that contains only the metadata component of any information resources it describes. It references the data component of the information resources through the location rather than by inclusion. This set may optionally contain an information model that is associated with the information resources described by metadata.

Information Model Information Set - an information set that contains only an information navigation structure and, optionally, associated information resource classes.

To support the transmittal of all three types of information sets, we defined a single XML-based Information Set Interchange Format (ISIF). XML is a machine-readable mark-up language. Markup tags are used to describe the nature of the data and to control application processing. XML facilitates machine interpretation and the cost-effective development and maintenance of translators. XML documents facilitate separation of transaction protocols, such as storage and retrieval, from data content, and content from format. The ISIF is defined using the XML Schema language. Information sets compliant with ISIF are represented as XML files. XML validation processing can be used to validate that an XML document satisfies the ISIF schema and determine if integrity constraints are met. The RAS makes extensive use of XML processing to validate data integrity conditions. The use of XML for data and metadata representation does not preclude including data in non-XML tool-specific formats as part of data content.

CONCUSIONS

We have defined an approach for managing information resources using metadata. These metadata concepts extend traditional metadata standards by addressing the lineage of data that exists in most enterprises and by providing additional metadata to characterize data quality. The JSF program has successfully demonstrated that these metadata concepts can be implemented to manage information resources of a wide variety of types and resident in a wide variety of repositories. Although the JSF experience is focused on product development, these metadata concepts have much broader potential applicability.

REFERENCES

Defense Modeling and Simulation Office (DMSO) (2001). *VV&A Recommended Practices Guide, Millennium Edition*, http://vva.dmso.mil.

Frost R. & Thomen, D. (1999). "Simulation Based Acquisition, the Road Map," 1999 Spring Simulation Interoperability Workshop.

Graves, W., Barnhart M., & Hollenbach, J. (2003), "JSF Authoritative Modeling Information System (JAMIS) Architecture," 2003 Spring Simulation Interoperability Workshop.

Hollenbach, J., Bishop S., & Graves, H. (2002) "JSF Modeling Information Management", Fall 2002 Simulation Interoperability Workshop.

Hollenbach, J., & Hartnett, R. (2000), "The Joint Strike Fighter (JSF) Distributed Product Description," Fall 2000 Simulation Interoperability Workshop.

Joint Strike Fighter Program Office (JSFPO) (2000), "Joint Strike Fighter (JSF) Modeling & Simulation Support Plan (MSSP), Version 5.2," (not available for public release).

Karangelen, N. (1998). "The Simulation Based Acquisition Vision – A Brief Tutorial," proceedings of NDIA Workshop on Simulation Based Acquisition.

Lutz R. & Keane, J. (1999). "An Architecture for Simulation Based Acquisition," 1999 Spring Simulation Interoperability Workshop.

National Information Standards Organization (NISO) (2001). *The Dublin Core Metadata Element Set*, ANSI/NISO X39.85-2001.

National Institute of Standards and Technology (NIST) (1994). *Integration Definition for Information Modeling (IDEF1X)*, Federal Information Processing Standard Publication 184.

Scrudder, R., Graves, W., Teigen, T., Johnson, C., Hix, S., & Hollenbach, J. (2003). "The Critical Role of Metadata in JSF Development," 2003 Fall Simulation Interoperability Workshop.

Undersecretary of Defense, Acquisition, Technology, and Logistics (USD (AT&L)) (2001) *DoD Modeling and Simulation (M&S) Verification, Validation, and Accreditation (VV&A)*, DoD Instruction 5000.61.

U.S. Department of Defense Acquisition Council (1998). *A Roadmap for Simulation Based Acquisition – Report of the Joint Simulation Based Acquisition Task Force.*

U.S. Department of Defense Chief Information Officer (DoD CIO) (2003a). *Department of Defense Discovery Metadata Standard.*

U.S. DoD Acquisition Council and SBA Industry Steering Group, (2000). *Simulation Based Acquisition (SBA) Definition.*

U.S. Department of Defense Chief Information Officer (DoD CIO) (2003b). *Department of Defense Net-Centric Data Strategy*.