BIOGRAPHICAL INFORMATION

Ken Newberry
Operations Applications Manager
TVA, Electric Systems Operations

Specific Responsibilities

Electric System Operations (ESO) is responsible for safely managing the power grid and continually improving the real-time operation and control of TVA's transmission and generation resources. ESO software applications consist of SCADA/EMS and many business applications to support Asset Scheduling, Operations Management, Transmission Operations and Transmission Reliability. Ken Newberry manages a group of 12 software developers responsible for supporting all non SCADA business applications in ESO.

Past Experience

Ken joined TVA in 1981 working in TVA's Information Services division as a Project Leader before joining Electric Systems Operations in 1995. Ken leads a team which is developing an ESO software development methodology founded upon the Rational Unified Process (RUP) and incorporating data modeling, business vocabulary management using the Common Information Model (CIM), and data movement or transformation using an ETL tool. The final solution will allow ESO to share information among other TVA organizations using an information bus (possibly Web Services) and to be able to easily integrate with other utilities and vendor products via the CIM.

Educational Information

B.S., Computer Science, University of Florida, USA

BIOGRAPHICAL INFORMATION

Greg Robinson
Consultant, Application Integration For Utilities
Xtensible Solutions

Specific Responsibilities

Greg Robinson assists utilities with the planning and implementation of Model-Driven Integration (MDI) infrastructures, helping them to leverage and drive industry standards to their benefit. Systematic integration framework planning is performed to organize the objectives of utilities' disparate organizations into cohesive master plans. Plans and implementations encompass systems performing asset management (AM/FM/GIS), network management (SCADA/EMS/DMS), engineering analysis and design, outage management, work management, maintenance and inspection, metering (interval meters for retail access), regulatory reporting, and customer information. Using various middleware technologies, these utilities are then able to automate the flow of information to support business processes through a common semantic infrastructure. The resulting solutions serve to minimize total cost of ownership while lowering risks associated with implementing and using information technology assets.

Past Experience

After joining Harris Corporation in 1981, Greg led several software development projects for electric utility energy management and distribution automation systems. He joined KEMA Consulting in mid-1994 and provided consulting services that encompassed operational business systems requirements analysis and design, architecture development, business case development, project management, and training. In January 1998, he became a founder of and consultant for Xtensible Solutions, which specialized in utility integration services. In March 2002, KPMG Consulting (now BearingPoint) acquired Xtensible Solutions. With BearingPoint's blessing, Xtensible Solutions started offering integration services again in March 2003.

Educational Information

B.S., Electrical Engineering, Georgia Institute of Technology, USA
M.S., Business Administration, Florida Institute of Technology, USA

Professional Memberships

GITA
Institute of Electrical and Electronics Engineers, Power Engineering Society
IEC TC57 Power System Control and Associated Communications:

- Working Group 14: System Interfaces for Distribution Management, Convenor
- Working Group 13: Energy Management System Control Center API, US Representative.
- Working Group 11: Substation Control and Protection Interfaces, Past US Representative

THE SEMANTIC INTEGRATION FRAMEWORK AT TVA

Greg Robinson
Xtensible Solutions
762 Loggerhead Island Drive
Satellite Beach, FL 32937

Ken Newberry
TVA
1101 Market Street, MR BG
Chattanooga, TN 37402-2801

ABSTRACT

Previous efforts to automate business processes at TVA were being hampered by incongruent data supplied by a myriad of applications and data stores. This paper describes a set of methodology, tools, and standards that TVA has implemented into a semantic-based integration infrastructure. As this approach turns many data sources into one coherent body of information, it is possible to create new application functionality based on consistent data, even though this data may be coming from multiple sources.

INTRODUCTION

A common problem in the industry, TVA (Tennessee Valley Authority) found itself in a situation where efforts to automate and manage business processes were debilitated by incongruent data from applications supporting the planning, constructing, maintaining, and operating of generation and transmission assets. As there are many tools available to bridge the gaps between the disparate technologies, the main show-stopper for large scale integration was that data resides in thousands of incompatible formats and cannot be systematically managed, integrated, or cleansed.

This paper describes an integrated set of methodology, tools, and standards that TVA has implemented to create and maintain a semantic-based integration infrastructure for its Electric System Operations (ESO). As ESO is responsible for the reliability of TVA's network, the Model-Driven Integration (MDI) methodology starts by mapping TVA business processes and roles to the NERC (North American Electric Reliability Council) Functional Model. For information exchange, the utility industry standard Common Information Model (CIM) and associated XML schemas from the IEC (International Electrotechnical Commission) are used as a basis for an inter-application vocabulary. Key for scaling up and maintaining the infrastructure, TVA is using a Business Vocabulary Management (BVM) tool for managing interfaces with applications and databases. A few project artifacts are used to describe the steps TVA follows to automate a business process, with emphasis being placed on how data exchanged among applications and database interfaces is cataloged and mapped to the CIM-based business vocabulary.

DEVELOPING TVA's 'INTEGRATED' INTEGRATION INFRASTRUCTURE

A key goal of TVA's integration infrastructure is to achieve a semantic framework that ensures the same term means the same thing when exchanged among applications, data stores, and portals. The infrastructure should eliminate most of the errors and ambiguities that arise with different interpretations about the intended meaning of information exchanged between applications. For example:

- Confounding conflicts, where information appears to have the same meaning, but does not.
- Scaling conflicts, where different reference systems measure the same value.
- Naming conflicts, where naming schemes differ significantly. See Reference [2].

Furthermore, the infrastructure needs to provide automated support for data management, data integration and data quality that transcends individual applications, data stores, portals and other tools. To be successful, this infrastructure must be manageable by developers with only moderate knowledge of industry standards such as the CIM, XML, and the like. A strategic goal is to turn the many sources of data into a coherent body of information so that new application functionality can be based on consistent data even though this data may be coming from multiple disparate sources.

To achieve these goals, TVA has implemented the Model-Driven Integration (MDI) process depicted in Figure 1.

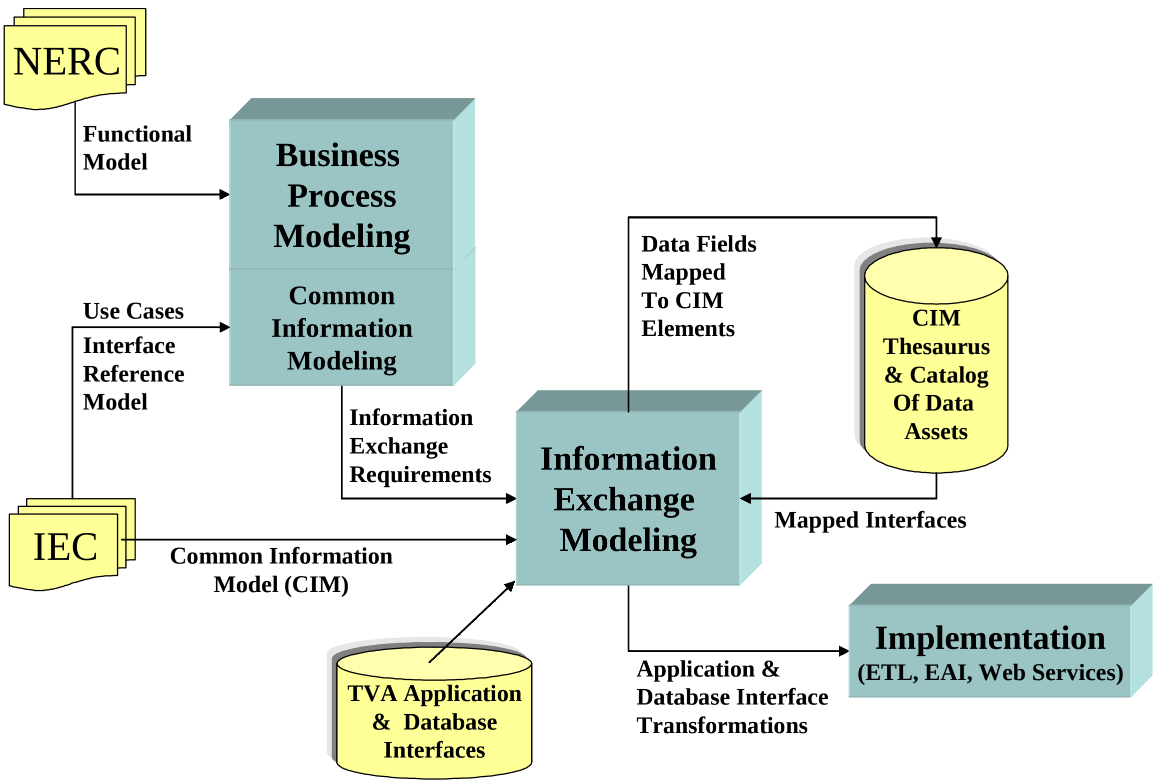

Figure 1: Overview of TVA's Model Driven Integration Process

The first block of the diagram depicts the beginning of the process, "Business Process Modeling & Common Information Modeling", which articulates how business processes are modeled and mapped to industry standard models. As will be described further in the next section, applicable TVA business processes are organized in alignment with the NERC Functional Model. Industry practices, in the form of use cases available from various sources, such as the IEC, are used as appropriate.   Associated activity diagrams identify information flows among organizations and functional systems supporting those organizations. As TVA's systems are continuously being updated and changed, it was decided that the information flows should be defined using a logical design first. At a later stage in the process, the physical systems are mapped to this design. This is accomplished by mapping current and future systems to the functional components of the IEC Interface Reference Model (IRM). The resulting logical design specifies interface requirements that a given application must satisfy when performing the role of a functional component. By using a logical design that is mapped to physical implementations, business process logic is not impacted whenever systems are upgraded or replaced.

The next major portion of the process is described around the second block of the diagram, “Information Exchange Modeling”. Now that information flows – each referred to as a message type - are clearly understood in the standards-based common language, effort at this stage is on making these flows function properly in TVA’s environment. This is accomplished through a business vocabulary management process where interfaces to TVA applications and data stores are mapped to the message types. The output of this process is automatically generated XML stylesheet language transformations (XSLT) and Java transformations, which are used in implementation by adapters for applications and data stores. This feeds into the implementation stage, which has minimal description as it is beyond the scope of this paper.

Arranged in accordance with the major stages shown in Figure 1, the standards and enabling technologies used to establish TVA’s Business Vocabulary are depicted in Figure 2.

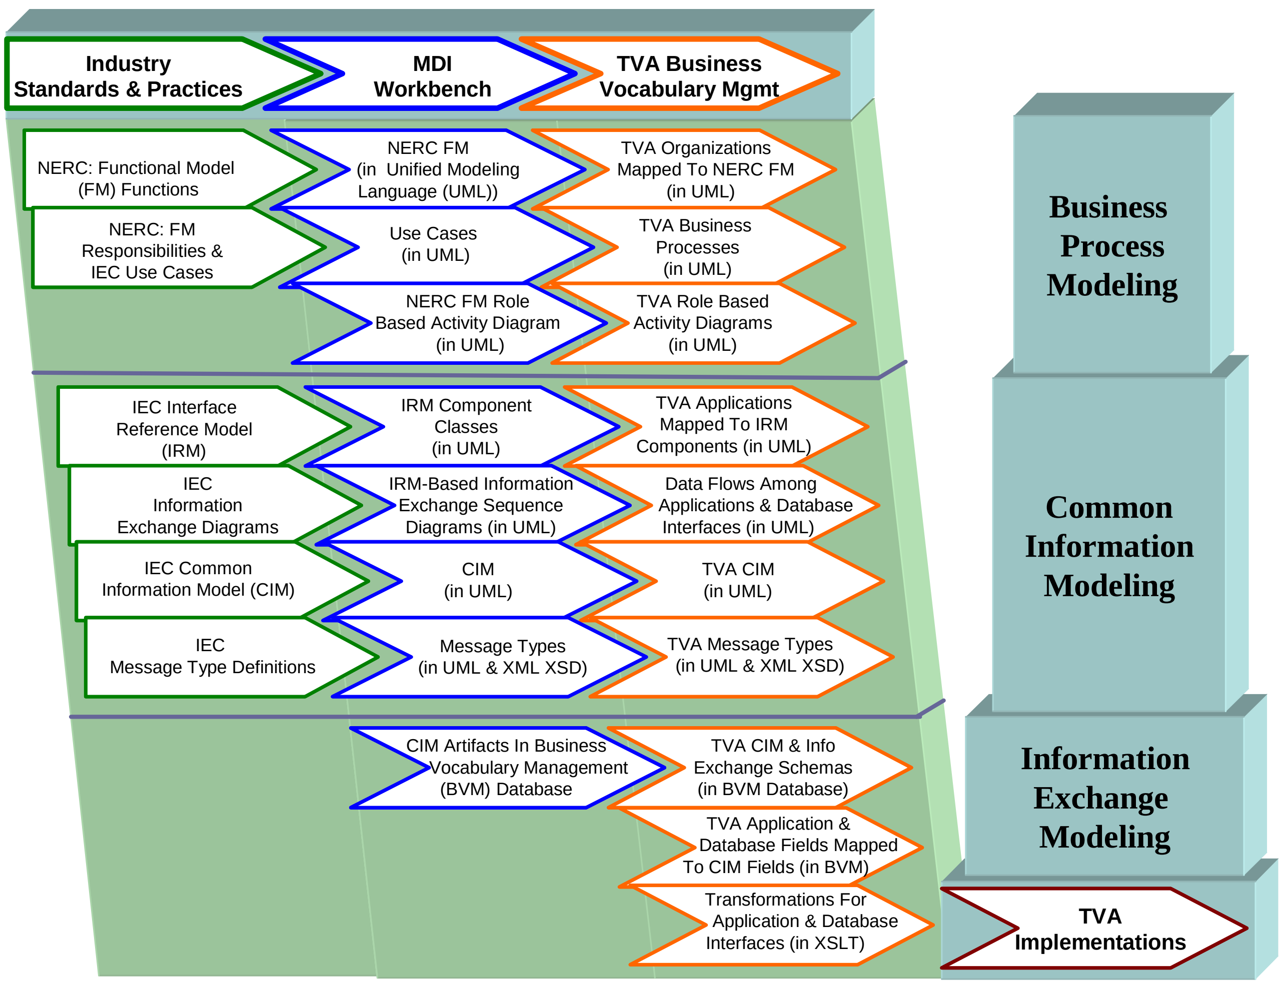

Figure 2: Jump-starting Business Vocabulary Management at TVA

Industry standards and practices are organized for use into a workbench, which is leveraged by TVA to develop and maintain its inter-application business vocabulary. Due to the end-to-end nature of the MDI methodology, no single tool provides all of the needed functionality. Therefore the MDI Workbench is used to integrate tools and pre-defined models for implementing the MDI Methodology. For TVA's integration environment, Rational Rose is used for UML modeling, the Business Vocabulary Management tool is supplied by Contivo, the MDI Workbench is supplied by Xtensible Solutions, and the tool for extraction, transformation, and loading (ETL) is supplied by Data Junction. How these standards, models, and tools are used at TVA to accomplish business vocabulary management is the subject of the remaining sections of this paper.

Business Process Modeling

The initial focus of Business Process Modeling is on how TVA organizations and processes map to the NERC Functional Model (FM). The FM provides a framework on which NERC reliability standards are based. Without specifying particular organizational structures, the FM defines functions involving reliability that are performed by various entities (e.g., Control Areas, Regional Transmission Organizations, Independent System Operators, Generator Owners, Distribution Providers, Purchasing-Selling Entities). Referring to Figure 3, observe that there are seventeen functions in the model. For each of these functions, tasks are defined that must be performed by an entity that claims responsibility for performing that function. An organization that registers with NERC for performing a given function has responsibility for ensuring that all tasks of that function are performed in accordance with associated Reliability Standards.

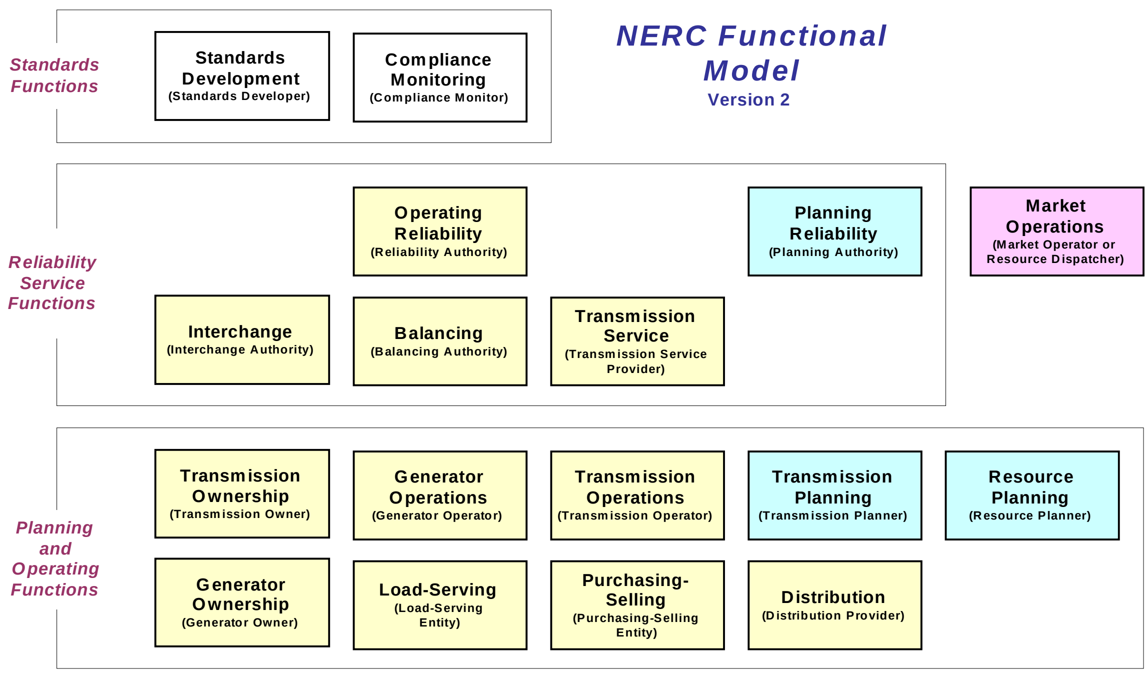

Figure 3: Overview of the NERC Functional Model

An example is shown in Figure 4, where use cases (represented as ovals) relevant to the business actor 'Balancing Authority' (represented as a stick figure) are summarized in a UML diagram. Color coding of use cases indicates the applicable section of NERC requirements for the modeled function: making deals (blue, n/a in Figure 4), ahead of time (green), real time (red), after the hour (rose), or compliance (gray). According to NERC, the Balancing Authority "integrates resource plans ahead of time, and maintains load-interchange-generation balance within a Balancing Authority Area and supports Interconnection frequency in real time." Use cases such as these are used to drive application specific information flow and exchange requirements. First, these NERC standard business process models indicate the industry requirements for information exchange and business interaction. After looking at an example use case, the focus will shift to how TVA realizes NERC requirements in the context of its general business processes.

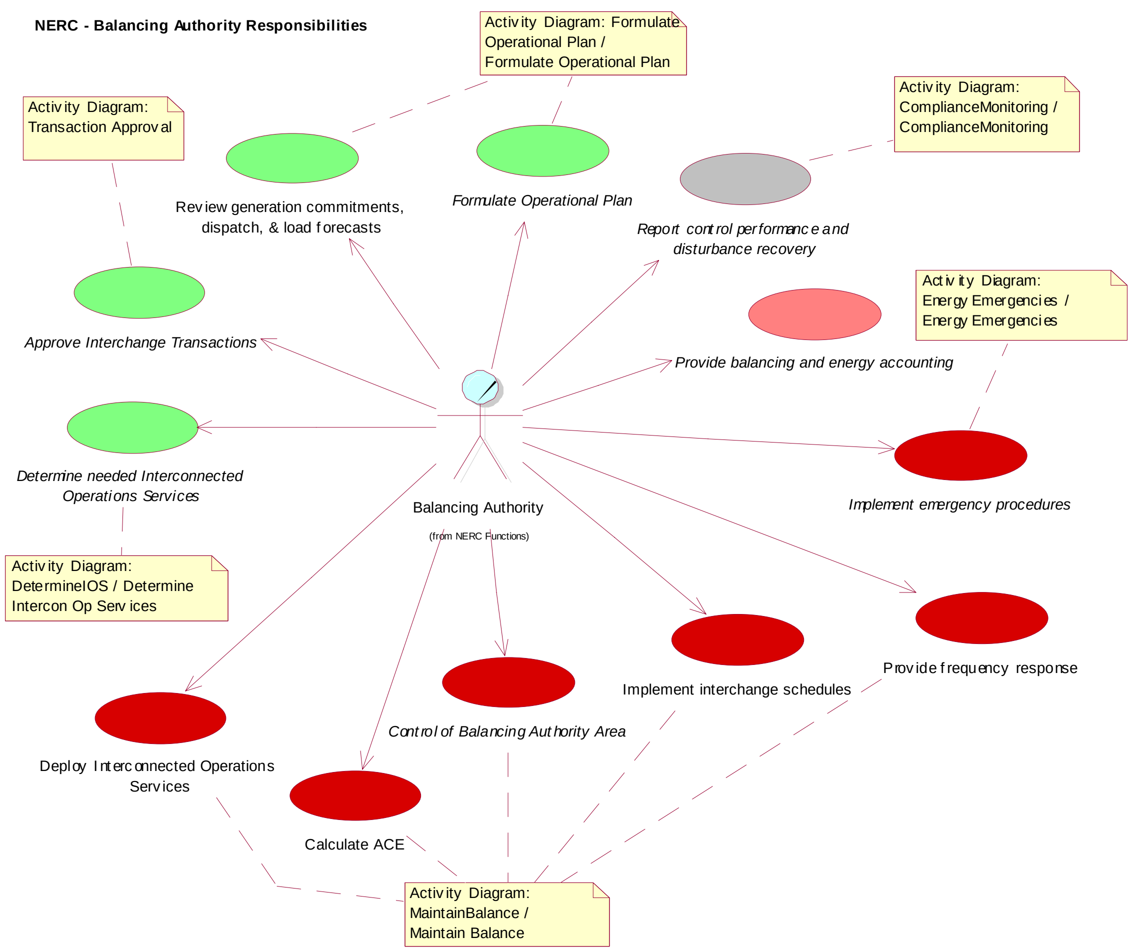

Figure 4: Use Case Diagram for the Balancing Authority

In the UML tool, selection of the use case, "Formulate Operational Plan", provides the description, "1. Compiles load forecasts from Load-Serving Entities. 2. Receives operational plans and commitments from Generator Operators within the Balancing Authority Area." The flow of activities performed in this process is depicted in the associated activity diagram, shown in Figure 5.

In the first swim lane of this activity diagram, designated for Interchange Authority activities, the first activity for the Interchange Authority to perform is to, "Provide Approved/Valid, Balanced Interchange Transactions". As a consequence of this activity, a message type entitled "Interchange Authority: Energy Transaction, IEC verb show", is passed from the Interchange Authority to the Balancing Authority. The contents of generic message types such as this are specified in CIM-based XML schemas (example schemas may be viewed at www.wg14.com).

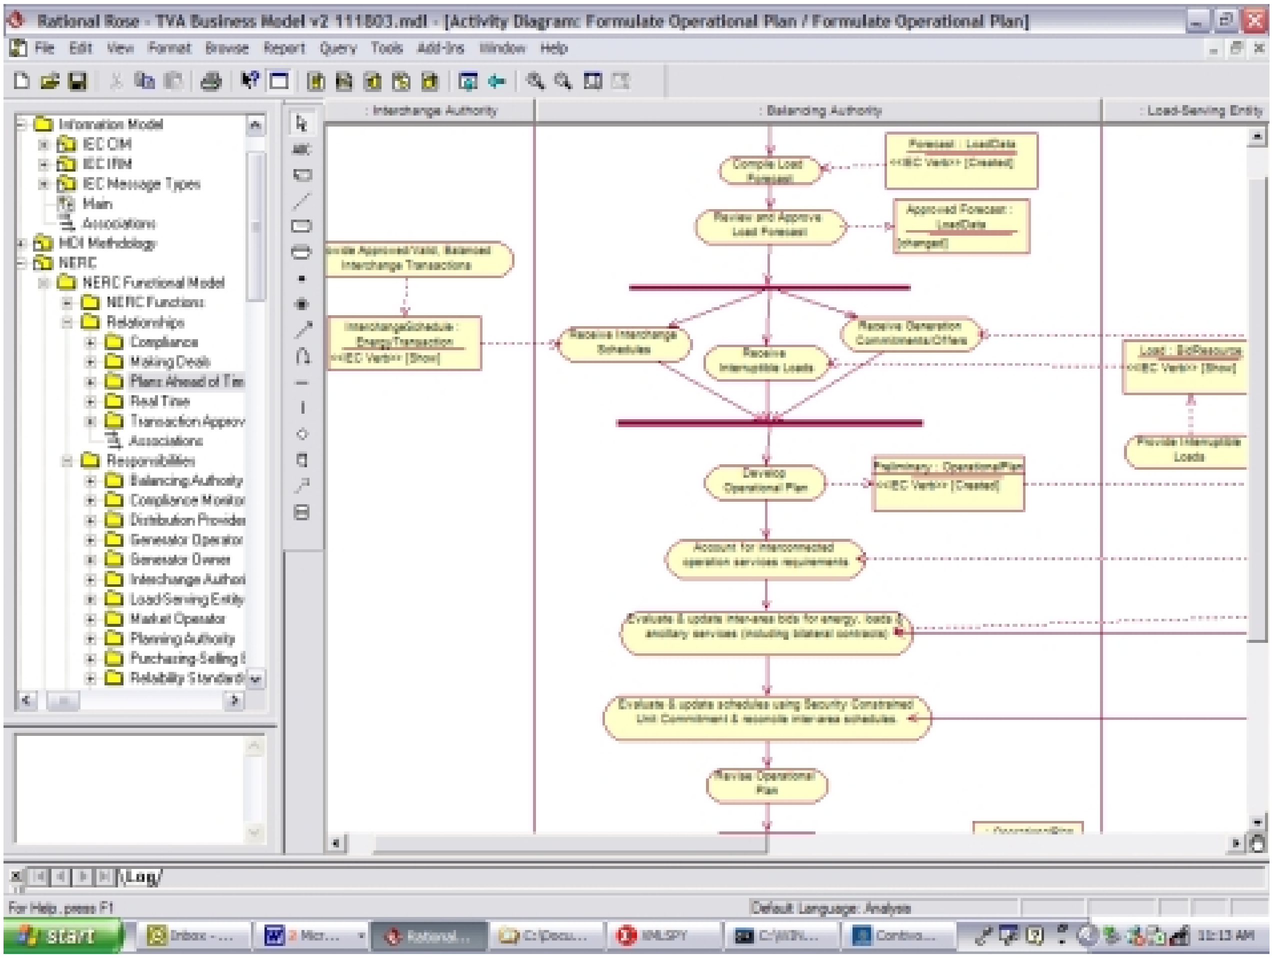

Figure 5: Activity Diagram for the 'Formulate Operational Plan' Use Case of the Balancing Authority.

As depicted in Figure 6 below, TVA models how it realizes NERC requirements in the context of its general business processes. TVA organizations are modeled as business actors and where applicable, mapped to NERC functions. Then TVA business processes are modeled as use cases, and where applicable, mapped to NERC use cases. When the project team discussed processes with each organization, NERC artifacts were used to understand which NERC functions that organization performed and who performed them. This, along with leveraging existing TVA process descriptions and industry use cases, provided a good check list for the project team to ensure that the first draft models articulated all significant business functions. Leveraging an existing framework helped the project team stay on track ("not lose the forest for the trees") while accomplishing this work in a relatively short period of time.

Once the general processes were understood, TVA's use cases were articulated in two types of TVA activity diagrams, 'Role-based' and 'IRM-based.' Both types of activity diagrams model information being exchanged in support of TVA processes, but the swim lanes of the former represent TVA organizations, whereas, the swim lanes of the latter represent abstract functional components of the IEC Interface Reference Model (IRM). The Role-based activity diagram completes the Business Process Modeling stage of MDI.

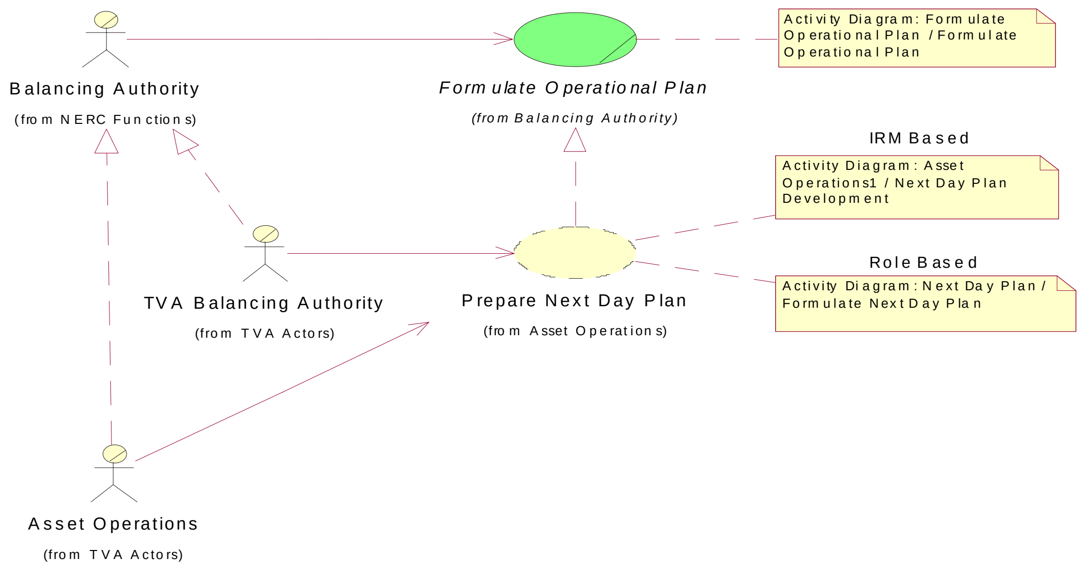

Figure 6: Overview of TVA's Realization of Generic Formulate Operation Plan

Common Information Modeling

For information exchange solutions to transcend product variations and continuous technology changes, stable interface points must be defined for cohesive chunks of application functionality. To that end, the IEC provides an Interface Reference Model (IRM), which is a logical breakdown of utility generation, transmission, and distribution functions as they relate to information technology. The IEC intends for the standard IRM to be extended (custom tailored) to meet an individual utility's needs, such as TVA has done (refer to Figure 7).

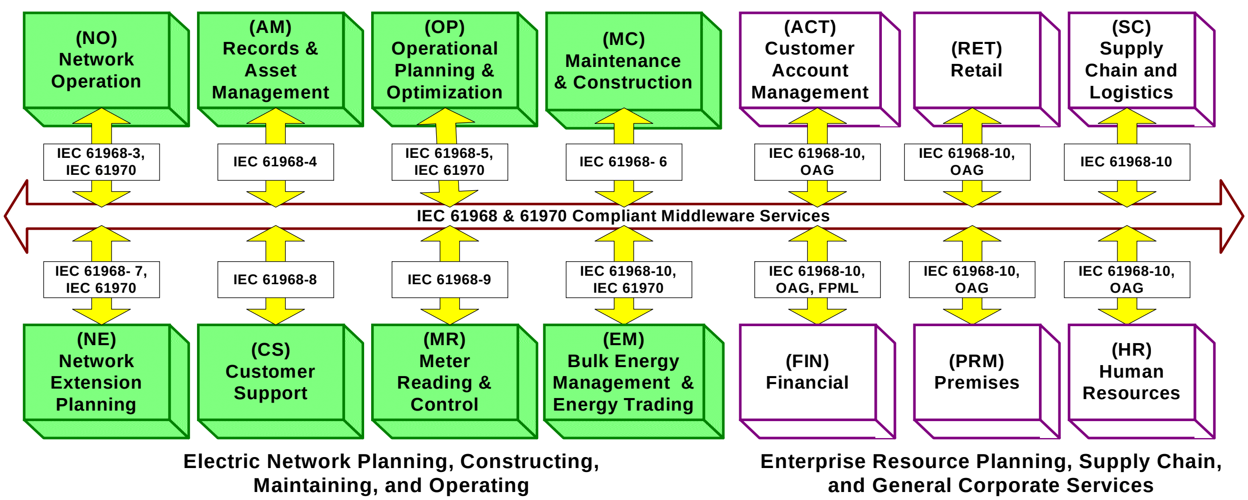

Figure 7: Top Level (Business Function Level) of TVA's IRM

Since the IRM was defined by industry representatives (i.e., utilities, product vendors, integrators, and consultants), mapping the IRM's abstract components to product interfaces is reasonably straightforward.

Understanding when and how these components are intended to be used facilitates the specification of requirements and evaluation of products during the procurement process.

In IRM-based activity diagrams, business processes are modeled based upon IRM abstract components. IRM business events are identified for information exchange purposes and linked to the IEC Message Types (which are developed through the aggregation of CIM classes). The order in which message types are exchanged among components, each component being mapped to the TVA application performing its function, is then shown in Event Sequence Diagrams. Refer to Figure 8 for an example.

One of the main purposes of the Common Information Model (CIM) is to provide a common language for describing exactly what is being exchanged among publishers and subscribers of various categories of information. Real world objects along with their attributes and relationships are defined in an information model using UML - before being used in a message type definition. This ensures that an attribute is used consistently regardless of its use in various types of messages. Once the system boundaries, messages to be exchanged and their contents are defined for a business, then that organization can leverage and extend the standard CIM to define its own information model. This information model then clearly defines the semantics of the messages exchanged among that organization's applications and data stores. For further information on how the CIM is extended for message type definition, refer to Reference [3].

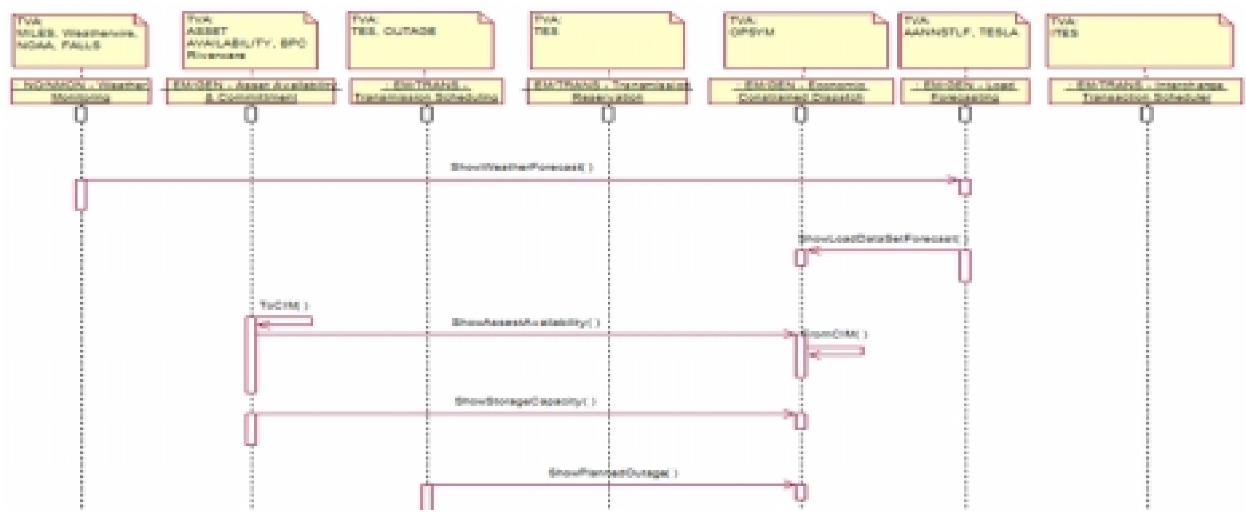

Figure 8: Sequence Diagram for Next Day Plan Development

Information Exchange Modeling

As the next stage begins, TVA has already expressed information exchange requirements in terms of CIM-based message types. The order in which information flows among applications and data stores in business processes is understood. This Information Exchange Modeling stage is to specify how these message types are transformed into the unique interfaces of each application and data store.

TVA's goal for this step is to streamline laborious processes, such as entering the interface for each application and database into cumbersome spreadsheets. This inefficient process is exacerbated since technicians must re-represent each mapping as a custom coded transformation through one of several single-purposed middleware tools. Instead, TVA sought a way whereby analysts could map interfaces and then rely on tools to auto-generate the appropriate transformations, which could then be used across their various middleware choices. This goal is largely achieved through the use of a Business Vocabulary Management (BVM) tool, which enables analysts to map an application's specific interface to meaningful "concepts" rather than low level data fields. Just as a thesaurus is used to find multiple words with similar meaning, the BVM tool maintains associations between CIM concepts and individual elements of the CIM Message Types and the various application and database interfaces. This enables TVA to automate the design of data transformation between applications through model reuse and collaboration. As shown in Figure 9, a graphical interface is used to map fields in CIM-based XML schemas to fields in TVA application interfaces or data stores.

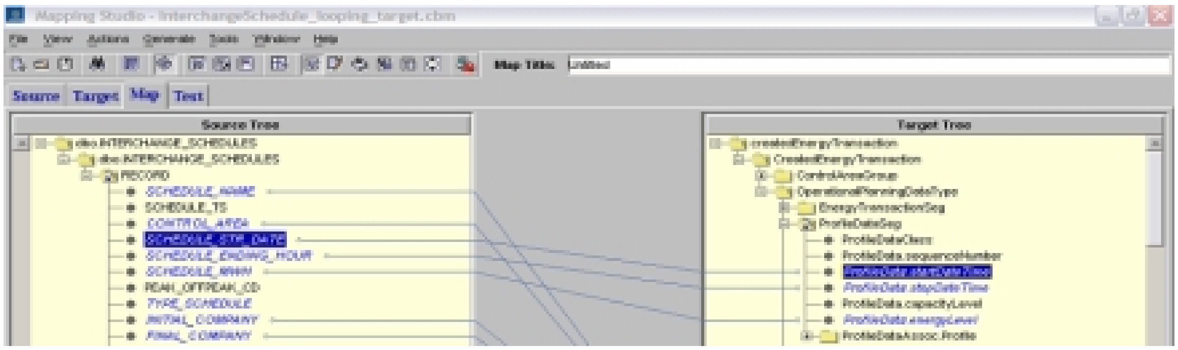

Figure 9: Mapping A Database View to a CIM-based Message Type

Once the map is complete, the BVM tool generates transformation code, which is then used as a roadmap for information flows in a run-time environment. To facilitate this process, the workbench is used to populate a semantic dictionary containing TVA's CIM-based enterprise vocabulary, a Thesaurus with synonyms that match business concepts, and a Rules Dictionary that governs the data transformation. These databases are self-learning and grow with each deployment. MDI uses XSLT for information transformation. XSLT is a language designed to transform one XML document into another, changing both its schema and content in the process.

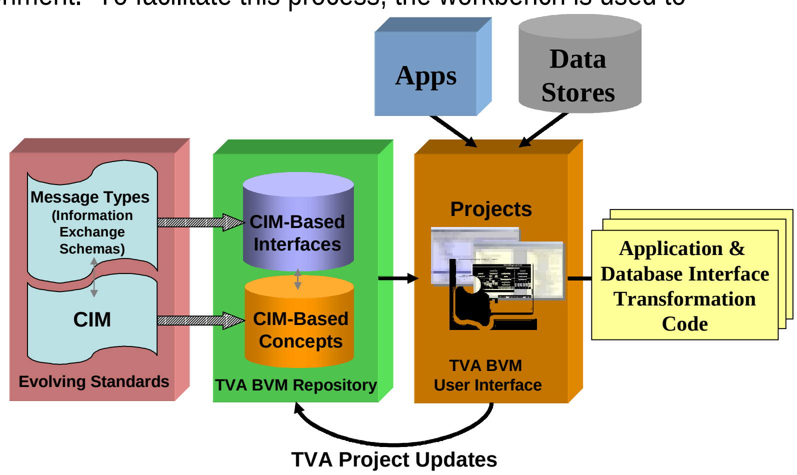

Figure 10: Overview of TVA Business Vocabulary Management

Executing Data Transfer

While the BVM tool generates the XSLT for the data transformation, an XSLT itself will not transfer any data. The final step is to connect the source and sink data sources and execute the XSLT. Rather than doing this manually, TVA uses an ETL (Extraction, Translation, Loading) tool to perform this function.

A Subtle Impact On Application Procurement

Noteworthy is the subtle, but important, impact this integrated infrastructure has had on TVA's application procurement priorities. While an increasing number of application vendors offer CIM-based interfaces in their products, it is recognized that some form of adaptation to TVA's local environment will always be required. ESO also recognizes that forcing alignment of vendor schedules and interfaces to TVA requirements drives up vendor prices and TVA's effort to manage them. All other considerations being equal, TVA will continue to prefer products that support CIM-based interfaces. However, with the level of automation support described in the aforementioned MDI process, it is now more timely and cost effective for TVA (vs. the vendor) to mange interfaces between vendor products and TVA's integration infrastructure. Doing so accelerates delivery of product upgrades, avoids premium charges for custom development during implementation and maintenance, and minimizes risk of broken features.

CONCLUSION

Initial usage of TVA's evolving model-driven integration infrastructure shows much promise. Now that the basic infrastructure has been sorted out, effort has shifted to helping project teams take maximum advantage of it. While this is no simple task, the benefits are well worth it:

- The reusable methodology, business process models, common information models, and information exchange models articulate how business objectives (e.g., NERC reliability requirements) are implemented and provide end-to-end requirements traceability.
- As more projects leverage the infrastructure, more data becomes available as part of one coherent body of information. This makes it possible to create new application functionality based on consistent data even though this data may be coming from multiple disparate sources.
- Faster integration of application systems and information is improving TVA's ability to react quickly to business changes while still providing the right information to people when they need it.
- Improved ability to integrate business processes across commercial-off-the-shelf (COTS) applications is lowering dependence on individual vendors.
- Re-use and adoption by future projects is encouraged by an extensible and adaptable design. There is one repository of all source, target, and CIM-based messages to ensure consistency and leverage, leading to implementation of fewer, simpler, more consistent interfaces.

ACKNOWLEDGEMENTS

The authors thank:

- Mr. Russell Robertson, Senior Manager of Control Systems, Electric Systems Operations, for his vision and sponsorship. Without it, TVA would likely not have performed the work described in this paper.
- Joe Zhou and Dan Martin of Xtensible Solutions, whom have contributed many of the ideas and results discussed in this paper.

REFERENCES

[1] Robinson, Greg, "Key Standards For Utility Enterprise Application Integration (EAI)," GITA 2002, Enterprise Application Integration Seminar.

[2] "NERC Reliability Functional Model, Function Definitions and Responsible Entities," NERC, Version 2 (Draft 12: October 23, 2003).

[3] Dietz, Janet and Greg Robinson, "PacifiCorp's Integration Bus Is Helping To Implement Retail Access," DistribuTECH 2002, Conference Proceedings.

[4] Committee drafts of IEC 61968, IEC TC57 Working Group 14 (refer to http://www.wg14.com), 2003.