The Missing Piece in Achieving Interoperability – a Common Information Model (CIM)-Based Semantic Model

David Becker
Manager, Control Center Technologies
Electric Power Research Institute (EPRI)
3420 Hillview Avenue
Palo Alto, CA 94304-1395
dbecker@epri.com

Terrence L. Saxton
Vice President, Special Projects
Xtensible Solutions
18125 23rd. Avenue North
Plymouth, MN 55447
tsaxton@xtensible.net

Keywords: Common Information Model (CIM), Semantic Model, Model-Driven Integration (MDI), Business Semantics, Integration Framework, Interoperability Testing

Abstract

The IEC 61968/70 Common Information Model (CIM) standards lay the foundation for an enterprise semantic model to achieve interoperability. Key aspects are discussed, including the importance of defining standards boundaries at the right level of abstraction to ensure adoption and continued use in the face of changing information infrastructures and systems, and how unique business contexts based on country and enterprise practices can be incorporated without over-defining the abstract information model standard. The importance of focusing on interfaces for application of semantic model standards and especially for testing for interoperability and compliance is stressed as well as the role of EPRI in extending the CIM into new areas where interoperability is needed and in interoperability and compliance testing to ensure products comply with CIM standards. The key role of profiles and messaging standards to establish interface contracts are explained as well as a related standard, the Generic Interface Definition (GID) for defining interface services.

1. THE NEED FOR SEMANTIC MODELS

The missing piece in most interoperability frameworks is agreement on a semantic model, which is arguably the most strategically important piece of any interoperability solution. This holds true whether one is dealing with system interfaces, field device data reporting, or human interfaces. The need to ensure understanding and avoid confusion in interpreting data while at the same time facilitating the sharing of data among distributed independently-developed applications is common to all enterprises.

1.1. Current Approaches to Achieving Interoperability

Most interoperability frameworks found at utilities today either were built from the ground up as new system interfaces were identified or designed around some type of Enterprise Service Bus (ESB). In either case, the resulting integration framework is defined primarily by the physical connectivity solutions adopted, with information integration typically being handled on a case-by-case basis by the project teams responsible for the particular system interfaces involved. This type of information integration requires unique mappings between every pair of system interfaces, resulting in transformation logic that resides either in a centralized ESB server or at system interfaces.

While Service Oriented Architectures (SOAs) are a step in the right direction by providing a common set of services for information exchange that are independent individual systems involved in the exchanges, they do nothing in and of themselves to address the information integration issues.

1.2. The Role of Semantic Models

At the other end of the spectrum are Model Driven Integration (MDI) frameworks based on a common semantic model that provides the starting point for all information exchanges. That is, any file or message payload defined for the exchange of information between two systems will contain data elements derived directly from a common semantic model, thus ensuring information is integrated regardless of the source of the data. This leads to the adoption of an adapter architecture which provides the transformation logic to map from proprietary data representation to a common model representation in an adapter between each system and the enterprise bus. The big advantage of this approach is that each system has only one mapping (i.e., native to common model), facilitating information sharing, since the source of the data no longer defines the semantics and syntax of the data.

1.3. The Business Case for a Common Semantic Model

While the importance of a common semantic model to system integration cannot be underestimated, the real business value comes from the composite business intelligence and decision support applications that it enables. These applications require data from a variety of sources, but without a common semantic model, they cannot be counted on to deliver on their promise of improved quality of decision making.

1.4. The Need for Enterprise Information Management

With such clear business advantages, it would seem like adopting an interoperability framework based on a common semantic model would be obvious. However, the reality is that it takes advanced planning at the enterprise level to make it a reality. This involves several inter-related efforts:

1. Definition and adoption of an appropriate reference architecture that embraces the notion of a common semantic model.
2. Development of an enterprise semantic model.
3. Establishment of a governance policy for the management and maintenance of this model as well as methodologies to create information exchange models that are based on it.
4. Organization of IT resources to assist individual projects in implementing the policies and procedures necessary to implement system interfaces based on the model. Without strong incentives from the enterprise level, individual project managers will find it difficult to enforce its use due to vendor push back citing increased cost over continued use of proprietary interfaces.

In current industry thinking these are all necessary ingredients of Enterprise Information Management (EIM) plan, which is defined by Gartner as “An organizational commitment to structure, secure and improve the accuracy and integrity of information assets, to solve semantic inconsistencies across all boundaries, and support the technical, operational and business objectives within the organization's enterprise architecture strategy.” The key to successful implementation of EIM is having a plan in place before any of these individual efforts are undertaken. A well thought-out plan will provide clear boundaries between the various roles and responsibilities as well as a methodology for definition of the reference architecture and enterprise semantic model. It should also identify the role of standards in these activities.

The remainder of this paper deals exclusively with points 1 and 2 above. Concepts presented are loosely based on References 1 and 2 with regard to the layered architecture and bridging from UML to OWL to other sources of information, respectively. However, it is of critical importance that all aspects of an EIM strategy be kept in view if the benefits of a common semantic model are to be realized.

1.5. The Role of Standards

Standards can play a vital role in several areas:

1. Definition of a layered reference architecture, clarifying the boundaries between standards in each layer.
2. Provision of a vertical industry information model that can be a key part of an enterprise information model
3. Definition of generic services for information exchange
4. Definition of profiles for the services and semantics for specific information exchanges between business functions

Fortunately for the utility industry, standards addressing these areas have been developed under the initial sponsorship of EPRI. The IEC 61968/70 series of standards define a Common Information Model (CIM), a set of generic services, a set of profiles and message definitions for information exchange. The CIM standards have been developed, managed, and extended by and for utilities, vendors and consultants to ensure completeness and acceptance. While a standard can never address all the information needs of a utility enterprise, it can provide a starting point, and if managed properly, it can be extended via private extensions and later via adoption into the standard.

The layered reference architecture referred to above is the focus of a concentrated effort in IEC Technical Committee 57, Power System Management and Associated Information Exchange, to provide a structure for the deployment of these standards as well as to provide guidelines for the development of standards for the individual layers.

2. STRATEGY FOR BUILDING AN ENTERPRISE SEMANTIC MODEL

One of the key aspects of a successful strategy in building an ESM is to define a reference architecture or framework to show how the various pieces that comprise the ESM all come together to provide model driven integration solutions.

2.1. A Layered Architecture

Figure 1 illustrates a three layered reference architecture that provides clear boundaries between the functions provided in each layer. This reference architecture is useful both for guiding the development of standards for each layer as well as for the development of an ESM within a particular utility. Regarding the standards-related use, this architecture embraces concepts that are currently being adopted into the CIM standards to provide a more stable yet flexible set of standards that can be adapted to a variety of environments.

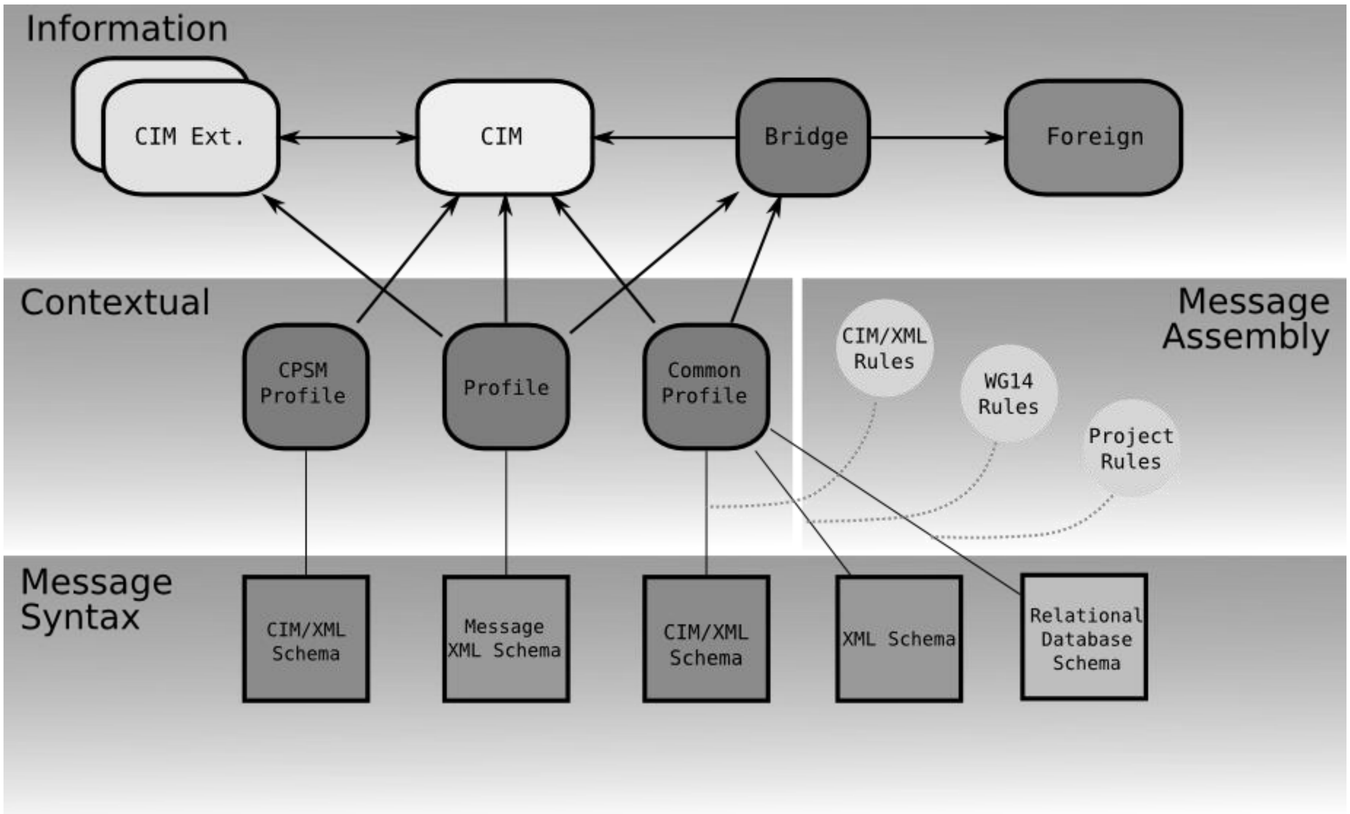

Figure 1, ESM Reference Architecture

The individual layers comprising this reference architecture are:

Information Layer – This layer includes the CIM but provides for the reality that there are other sources of information as well as the CIM that need to be taken into consideration when creating CIM-based ESM. These different models/standards and ways of bridging them together comprise the Information layer.

Contextual Layer – This layer formally recognizes that only a subset of the models in the Information Layer are needed for any particular interface or message definition. The Profile standards defined in this layer:

- Define a subset of the models in the Information layer needed for a particular business purpose as well as constraining those model elements to address specific business needs, and
- Provide a way to incorporate model elements from the different information sources in the Information layer in addition to the CIM.

Message Syntax Layer – This layer provides the rules for implementing the Profiles in the Contextual layer in various technologies.

An important feature of this layered architecture is that there are clear boundaries defined between the information models in the Information Layer and the business context in the Contextual Layer. Without this distinction the current CIM has suffered from an “identity crisis” trying to be an information model that also incorporates business context in a non-uniform way. The tension is created by trying to have the CIM be both general enough to be used in any application while being as specifc and constrained as possible to include descriptions more useful to an application in a specific business context.

It’s not possible to satisfy both objectives in an information model, although attempts have been made unknowingly to do just that. This has resulted in an unnecessary “stirring of the CIM pot”, leading to some changes in the CIM information model that could have been avoided. Separating the information model from the business context permits the CIM to stay more general and stable, while permitting new

Profiles to be defined to apply restrictions needed for a specific business context.

Each layer is described in more detail in the following subsections.

2.2. Information Layer

The important architectural features enabled by the layer are described in the following subsections.

2.2.1. Multiple Sources of Information and Metadata

In the current CIM standards, the CIM in UML is the only recognized source of metadata for defining XML messages or files. Although it is possible to extend the CIM with private extensions, and in fact is expected, the goal has been to eventually incorporate those extensions into a later revision of the CIM UML model if the extensions prove to be generally accepted. In any case, the standard CIM UML model with private extensions is the only recognized source for creating a semantic model as the basis for a model-driven architecture.

The Information Layer in the future reference architecture vision, on the other hand, embraces the notion that there are other sources of metadata that a utility enterprise needs to include in its semantic model without trying to make it a part of the CIM standard. Conceptually, some kind of a Bridge, as shown in Figure 1, is needed to create links to these other metadata, similar to the way associations between classes in UML link different parts of the UML model. Whether or not this Bridge becomes the subject of future standards is unclear.

These other information models denoted as Foreign sources in the diagram could include models from other standards bodies or industry consortiums, such as Geography Markup Language (GML). Other possible sources include other TC57 standards, such as the IEC 61850 Substation Automation standards. In fact, this is a very powerful way of achieving harmonization of the 61968/70 CIM-based standards with the 61850 standards. Rather than trying to change these standards to be the same in the Information Layer where there is overlap, the differences can be resolved in the Contextual Layer by making it possible to include attributes from both sets of standards in a Profile, as elaborated more completely in the Contextual Layer section below.

2.2.2. Abstract General Purpose Information Models

Recognizing the Information Layer as separate and distinct from the Contextual Layer has other benefits as well. The CIM can now be thought of as purely an abstract information model that is general enough to be used in a variety of business contexts. So for example, when defining an attribute describing a generator control mode, the CIM can simply provide a string data type. In the Contextual Layer, the string can be replaced with an enumeration that is appropriate for the country where the CIM is being used. This has the advantage of making the generator control mode in the CIM reusable in many different contexts as well as providing a standard way to constrain the permissable values in a particular business context. This has the benefit of providing for the possibility of validity checking of the instance data to ensure only one of the permitted values is used in an information exchange implementation that includes this attribute.

Another problem this addresses is caused by the use of inheritance in the CIM model. Attributes that are inherited from a parent class have only a general purpose name. In the Contextual Layer the name can be changed to include some reference to the specialized class where it is being used, so that in a particular message payload or file in the Implementation Layer, it will be clear what object the attribute applies to.

2.3. Contextual Layer

The Contextual Layer provides for the definition of Profiles to define a subset of the information models contained in the Information Layer that are needed in a specific business context. Business context or constraints are also applied in this layer. This notion embraces many of the concepts described in the UN/CEFACT Core Components Technical Specification (CCTS) (see Reference 1). Profiles may also incorporate the identification of services to be used for information exchange.

2.3.1. Profile as a Subset of the CIM

The notion of Profiles is not new. For example, the CPSM (Common Power System Model) Profile shown in Figure 1 is currently used to define the subset of classes and attributes that are needed to exchange power system models between RTO/ISOs for maintaining network models of neighboring regions. The CPSM Profile is then used to create the message syntax to be used in actual implementations, in this case to define an RDF/XML schema for the generation of CIM-based XML files or messages. This profile has been standardized as draft IEC 61970-452 and is equivalent to a Platform Independent Model (PIM) as defined in the ONG Model Driven Architecture (MDA). An RDF/XML schema implementation of this profile has also been standardized as IEC 61970-501 and draft 61970-552-4.

2.3.2. Profiles and Multiple Information Sources

In the new vision shown in the diagram, the concept of a Profile has been substantially expanded, so that a Profile can apply a business context to a subset of metadata from multiple information models via the Bridge concept. As shown in Figure 1, the Profile object in the Contextual Layer incorporates metadata from the CIM, private extensions to the CIM, and via the Bridge, other information models as well. The key is to maintain traceability back to the source to facilitate long term management and maintenance of the Profiles as new versions of the information model standards are published.

2.3.3. PIMs and PSMs

Another important concept embodied in the Profiles is the notion that they represent a Platform Independent Model (PIM) of an information exchange or interface, thus creating a clear boundary between the Contextual Layer and the Implementation Layer, where there may be multiple technology implementations of that profile. The standards in the Implementation Layer then are the Platform Specific Models (PSMs). So it can be seen that the future TC57 layered architecture embraces the MDA concepts of PIMs and PSMs.

In Figure 1, the Common Profile object as shown can be implemented in several technologies, each with its own syntax, including RDF/XML schema, XML schema, and a relational database schema. This implies that a Profile must be specified at a high enough level of abstraction to allow it to be implemented in various, different technologies.

2.4. Message Syntax Layer

This layer includes standards for concrete implementations of information exchanges and interfaces to the level of specificity required for achieving interoperability between products/applications/systems from different suppliers. These standards also form the basis for compliance testing to validate system interfaces. As such, they must be technology specific.

Since these PSM standards are based on the PIMs in the Contextual Layer, it is important that they include clear rules for how they are derived from the PIMs. For example, there are several XML Schema structures that can be generated from a single Profile definition – each one correct but different, and not interoperable. So it is important that the PSM standards also include rules for creating the PSM from the PIM. For example, as shown in Figure 1, three different PSMs may be derived from the Common Profile. Each has a different set of rules that must be defined. For generating CIM/XML files based on RDF Schema, rules are defined by WG13 to define the subset of and extensions to the RDF Schema elements as defined by W3C to be used. These are incorporated in a standard so that there is one accepted way of using RDF schema to create the file metadata. Similarly, for the XML Schemas defined by WG14 for message exchange between distribution systems, a set of rules is needed to define how the XML schemas are to be derived from the common profile. As a last example, a project may define a new technology mapping to a relational database with its own set of rules outside the standards arena.

2.5. Concrete Messages and the Three Layer Architecture

Figure 2 illustrates how this three-layered architecture all fits together to define a concrete message for information exchange based on the CIM. Note that this illustration shows only the CIM as a source of the information metadata, but the concepts apply regardless of the information source.

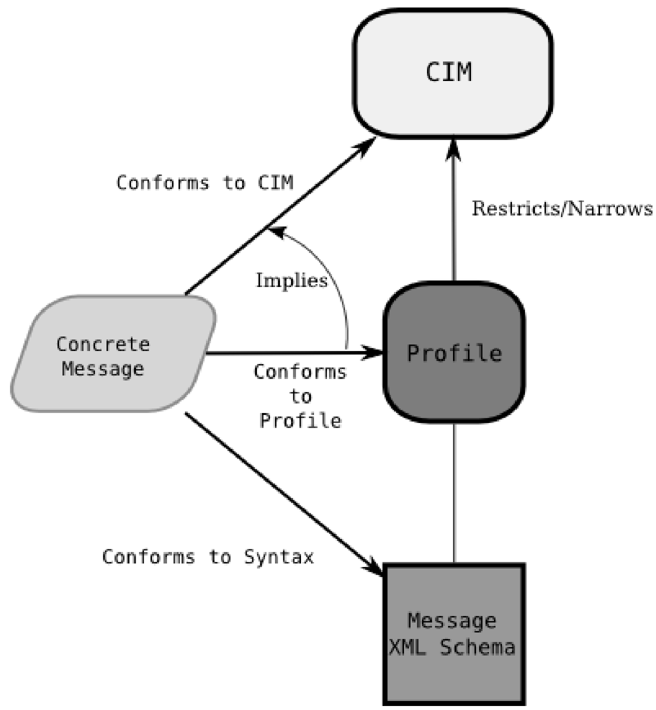

Figure 2, Concrete Message Generation

The CIM is shown as the source of the information metadata used in the message. The Profile defines the subset of the CIM that is to be used in the message, thus restricting the CIM to only those parts needed for the particular business process and information exchange in view. It also adds business context to that subset of the CIM to take the CIM from a general purpose application-independent information model to a semantic model that better represents the specific business context and is thus application-dependent. However, at this point the Profile is still abstract (i.e., technology neutral). The Message XML Schema is then generated from the Profile and CIM following the standard rules for mapping to XML Schema, when the desired concrete message is to be an XML document.

As may be seen in the diagram, the concrete message needs to conform to the CIM standards at three points:

1. The CIM for the information metadata
2. The Profile for the business context restrictions
3. The Message XML Schema for the message syntax

With this view of conformance in view, compliance testing can be better understood. This is an area that is currently not well defined from an architectural perspective, i.e., how to test for compliance with CIM standards, or even more basic, what is the meaning of compliance with CIM standards. Figure 2 illustrates where compliance is necessary to achieve interoperability and claim “compliance with the CIM.”

2.6. Service Model and Interfaces

Interoperability is really about interfaces. An important part of an interface are the services used to exchange information with other systems. SOA and Web services provide a robust services environment but are independent of content. Underlying the reference architecture discussed above are another part of the CIM related standards known as the Generic Interface Definition (GID) services that are part of the IEC 61970 series of standards. The GID includes CIM-aware standards for access to complex data structures, for high speed data exchange, for historical data access, for publishing and subscribing. CIM-aware means data can be browsed and accessed based the CIM representation of the data in view. When combined with specific concrete message payloads based on the CIM, they define an interface that can be tested for interoperability and standards compliance.

3. CONCLUSION

The authors believe that the concepts and supporting standards presented represent the next logical step in the evolution of the CIM standards to help achieve interoperability between the variety of systems used by electric utility transmission and distribution. However, the concepts presented apply equally well to a variety of other domains of application within the scope of the GridWise Architecture framework.

4. REFERENCES AND ACKNOWLEDGEMENTS

Key concepts for this paper were based on working drafts developed in IEC TC57. In particular, the authors want to acknowledge the contributions of Jean-Luc Sanson (EDF), Xiaofeng Wang (Xtensible Solutions), and Arnold deVos (Langdale Consulting). The layered reference architecture described in this paper is also loosely based on concepts described in the following references.

1. *UN/CEFACT ebXML Core Components Technical Specification (CCTS).* From UN/CEFACT, United Nations Centre for Trade Facilitation and Electronic Business. 11-December-2002. Version 1.90.
2. *Ontology Definition Metamodel (ODM)* Sixth Revised Submission to OMG/ RFP ad/2003-03-40, 5 June 2006.

Biography

David Becker is the Manager of Control Center Technologies, Electric Power Research Institute (EPRI) in Palo Alto, California. He is responsible for formulating research strategies for power system control centers. Mr. Becker's current projects include implementation of new technologies incorporating standards, system concepts, and competitive designs to provide solutions for the electric system. Mr. Becker has spent over 15 years guiding the development the CIM, which is now an international standard IEC 61970. Mr. Becker also was sponsor for development of the IEC 60870 Inter-Control Center Communications Protocol (ICCP) and a short- term load forecaster called EPRI ANNSTLF. This forecaster is widely used in the U.S. and abroad. Before joining EPRI in 1993, Mr. Becker worked for Pacific Gas and Electric Company (PG&E) in various operational and managerial roles. He is a Senior Member of IEEE, a USA delegate to IEC TC57 WG13, which is responsible for the CIM standards, and is active in numerous task forces and working groups related to system operations and control centers. Mr. Becker holds a BSEE from Lafayette College in Easton, Pennsylvania, and an MS Engineering Management from the University of Santa Clara, in Santa Clara, California.

Terry Saxton is Vice President and a founder of Xtensible Solutions, a company offering professional consulting services to the international utility industry in the development of Enterprise Information Management (EIM) strategies and frameworks based on the Common Information Model (CIM) and related standards. Mr. Saxton is Convener of IEC TC57 WG13 responsible for the CIM and other international standards for energy management system interfaces. He manages projects for EPRI dealing with the CIM, most recently extending the CIM to support planning applications. Mr. Saxton has many years of experience in the analysis, design, development, and implementation of a wide range of system integration solutions for electric utilities and the US Department of Defense. Prior to starting Xtensible Solutions, Mr. Saxton worked for BearingPoint, KEMA Consulting, Siemens Power Systems Control, Honeywell, Information Exchange Systems, and Bell Telephone Laboratories. He received an MSEE from the Massachusetts Institute of Technology (MIT) in Cambridge, Massachusetts, USA, and BSEE and BS Math with Honors from California State Polytechnic University.