

Model Driven Architecture (MDA)

Document number ormsc/2001-07-01

Architecture Board ORMSC¹

July 9, 2001

1. Edited by Joaquin Miller and Jishnu Mukerji. The following have made significant contribution to this document in course of its development: Carol Burt, Desmond DSouza, Keith Duddy, William El Kaim, William Frank, David Frankel, Sridhar Iyengar, Joaquin Miller, Jeff Mischkin, Jishnu Mukerji, Jon Siegel, Richard Soley, Sandy Tyndal-Biscoe, Axel Uhl, Andrew Watson and Bryan Wood.

1 Why the MDA

1.1 MDA and the Mission of the Object Management Group

The OMG's mission is to help computer users solve integration problems by supplying open, vendor-neutral interoperability specifications. The Model Driven Architecture™ (MDA™) is OMG's next step in solving integration problems.

1.2 The OMG's Specification History

The work of the OMG has been driven since 1990 by a clear architectural statement that has not changed much since it was first designed. The Object Management Architecture (OMA) provided the vision and roadmap for the problem that the OMG has always addressed: the problem of integration. Having created the CORBA (Common Object Request Broker Architecture) interoperability standards, the OMG has in the past used them almost exclusively as the basis for creating standards for use in particular application domains. However, since 1997 the scope of the organization has broadened significantly. In 1997 the OMG issued several important specifications that are not CORBA based, including the Unified Modeling Language™ (UML™) and the Meta Object Facility™ (MOF™), and later XML Metadata interchange (XMI™) and the Common Warehouse Metamodel (CWM™).

Vendors have implemented these specifications widely over the past ten years, most notably OMG's flagship CORBA specification. When every system component supports an OMG-standardized interface, the task of creating a multi-vendor solution (as most are) is greatly eased. Organizations polled in a recent analyst survey confirmed this by ranking CORBA compliance as *the* most important consideration in

choosing middleware for application integration². More recent OMG interoperability specifications like Common Warehouse Metamodel are expected, in the long term, have equal impact.

1.3 The Emergence of New Kinds of Standards

When the OMG was issuing only CORBA-oriented standards, the manner in which its various standards fit together was quite well understood, and clearly mapped by the OMA. The emergence of new kinds of standards that are foundational in nature, and their potential use in defining other standards and fitting them together, necessitates that we expand our vision of the OMG's architecture.

There are limits to the interoperability that can be achieved by creating a single set of standard programming interfaces. Computer systems have lives measured in decades, and not all ancient systems written in obsolete programming languages can be modified to support standards. Furthermore, the increasing need to incorporate Web-based front ends and link to business partners who may be using proprietary interface sets can force integrators back to the low-productivity activities of writing glue code to hold multiple components together. When these systems in their turn need modifying and integrating with next year's hot new technology (and they all will) the result is the kind of maintenance nightmare all computer users fear.

1.4 The Evolution to Model Driven Architecture

This paper is a statement by the OMG Architecture Board (AB) of the expanded vision necessary to support interoperability with specifications that address integration through the entire systems life cycle: from business modeling to system design, to component construction, to assembly, integration, deployment, management, and evolution. This vision is embodied in the OMG's Model Driven Architecture™ (MDA™). It describes how the MDA defines the relationships among OMG standards and how they can be used today in a coordinated fashion, and how the approach helps in the creation, maintenance and evolution of standards. It is important to realize that the MDA is a proposal to *expand* and not replace the OMA, to provide a roadmap and vision that will include and integrate all of the work done to date, and to point the way to future integration standards.

OMG has already specified integration with external specifications (such as XML) and proprietary interface sets (such as Microsoft's DCOM). The MDA approach incorporates this existing work and promises more support for rapidly and effectively creating new specifications that integrate multiple interface standards, from both inside and outside the organization. It is an evolutionary step from how OMG works at the moment, but one that will offer great benefits to those using the OMG interoperability framework. It will also indirectly help all system integrators.

².'A surprising 70 percent of respondents cited CORBA compliance as 'important' or 'very important' to integration, outpacing every other factor in the survey, including core functions such as integration of legacy applications with distributed systems and corporate intranets.'-- Summary of responses from 547 organizations asked to rate middleware selection criteria in the context of application integration in 'Middleware: what end users are buying and why' Gartner Group, February 1999

1.5 How the MDA Integrates Standards

An important aspect of some of the latest OMG standards is that they greatly advance the art, science, and scope, of modeling. The combined power of these model-driven standards forms the basis of a compelling approach to long-lived architectures for distributed, component-based systems. This approach broadens considerably the usability of the standards by:

1. embracing CORBA, J2EE, XML, .NET and other technologies;
2. improving portability of applications by allowing the same model to be realized on multiple platforms through auxiliary mapping standards, or through point mappings to specific platforms
3. improving integration based on models of relationships *across* different domain application, and component interfaces allowing interoperability based on semantically-rich interrelationships.

2 The Model Driven Architecture

2.1 Introduction

The MDA defines an approach to IT system specification that separates the specification of system functionality from the specification of the implementation of that functionality on a specific technology platform. To this end, the MDA defines an architecture for models that provides a set of guidelines for structuring specifications expressed as models.

The MDA approach and the standards that support it allow the same model specifying system functionality to be realized on multiple platforms through auxiliary mapping standards, or through point mappings to specific platforms, and allows different applications to be integrated by explicitly relating their models, enabling integration and interoperability and supporting system evolution as platform technologies come and go.

2.2 Basic concepts

2.2.1 Models

In the MDA, a model is a representation of a part of the function, structure and/or behavior of a system³.

A specification is said to be formal when it is based on a language that has a well-defined form (“syntax”), meaning (“semantics”), and possibly rules of analysis, inference, or proof for its constructs. The syntax may be graphical or textual. The semantics might be defined, more or less formally, in terms of things observed in the world being described (e.g. message sends and replies, object states and state changes, etc.), or by translating higher-level language constructs into other constructs that have a well-defined meaning. The optional rules of inference define what unstated

3. We use *system* here in the system-theoretic sense to include not only software.

properties you can deduce from the explicit statements in the model. In the MDA, a specification that is not *formal* in this sense, is not a model. Thus a diagram with boxes and lines and arrows that does not have behind it a definition of the meaning of a box, and the meaning of a line and of an arrow is not a model—it is just an informal diagram. An MDA model must be paired unambiguously with a definition of the modeling language syntax and semantics, as provide by the MOF.

Note that under this definition, source code is a model that has the salient characteristic that it can be executed by a machine. Similarly, a set of IDL interfaces is a model that can be used with any CORBA implementation and that specifies the signature of operations and attributes of which the interfaces are composed. A UML-based specification is a model whose properties can be expressed graphically via diagrams, or textually via an XML document.

2.2.2 *Abstraction, Refinement and Viewpoint*

The term *abstraction* is used in the MDA in the sense defined in the Reference Model of Open Distributed Processing (RM-ODP)⁴ Part 2: the suppression of irrelevant detail.

It is useful to characterize models in terms of the abstraction criteria that were used to determine what is included in the model. A model that is based on specific abstraction criteria is often referred to as a *model from the viewpoint defined by those criteria*, or in short as a *view* of the system.

Another common use of the word *abstraction* is in the phrase *a model at a higher level of abstraction*, which is used to characterize a model in which more of the details of the system are elided as compared to *a model at a lower level of abstraction*.

Specifically, some pairs of models are in a *refinement* relationship in which one - the *abstraction* - is more abstract than the other - the *realization*. The refinement relation is itself described using a model, defining *abstraction* observations in terms of *realization* observations while maintaining certain guarantees of the *abstraction*. For example, the *realization* observations might be SOAP communications and replies, while the *abstraction* elides differences of SOAP vs. GIOP.

Note that models that simply describe the system from different viewpoints – *of legal drinking age* vs. *resident of Maine* – cannot necessarily be related by refinement.

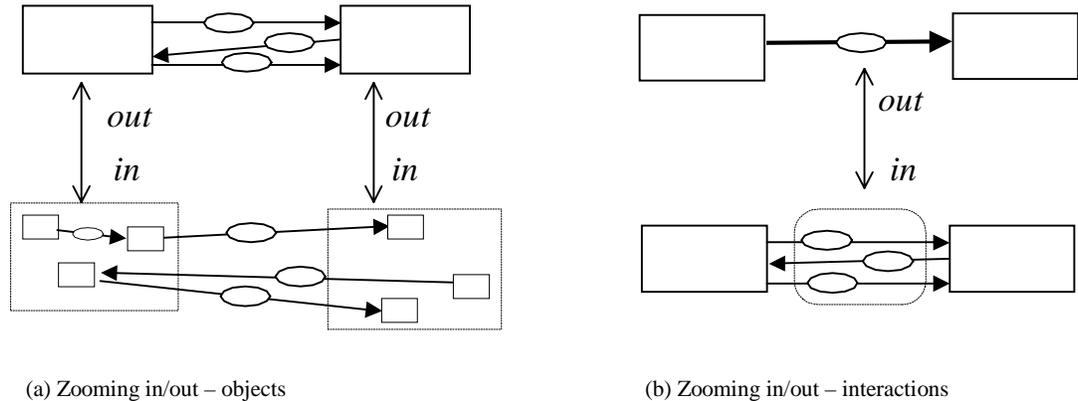
Furthermore in common usage, separate viewpoints sometimes corresponds to different levels of abstraction such that a larger amount of detail from a viewpoint may be elided to obtain a model from that viewpoint, or a view, at a higher level of abstraction, while little or no detail may be elided to obtain a model or view, at a low level of abstraction for those pairs of models. It is essential to note, however, that some models of a system from different viewpoints will not be related by refinement e.g. “of legal drinking age” vs. “resident of Maine” as views of some person.

While the MDA provides guidelines on the architecture of models, the choice of viewpoints to be used in a system specification is a modeling choice. For example, A specific set of five viewpoints is defined in the RM-ODP Part 3.

4. ISO Standard 10746-2

2.2.3 “Zooming” in and out

The MDA leverages UML models as abstractions, different viewpoints, and different levels of abstraction. In particular, as illustrated in Figure 1, the MDA permits:



The symbol  represents an interaction

Figure 1. “Zooming” in and out

- “Zooming” out of a model showing a complex network of objects (e.g. infrastructure objects like stubs, binders, etc.), to get a simplified model with fewer large-grained objects and attributes; and correspondingly zooming in to see those details.
- “Zooming” out of a model of a detailed interaction protocol (e.g. a platform-specific protocol for log-in and authentication) to get a simplified model with a single more abstract interaction with the same overall effect; and correspondingly zoom in to see those detailed interactions.

MDA permits zooming in to result in multiple alternative models (and vice-versa) by separating the refinement model that relates the detailed and simplified views. The MDA needs this flexibility to deal with the modelling of multiple platform specific implementations of the same system functionality; it is more flexible than the 1-1 mappings sometimes assumed by “decomposition” and “composition” models, which are a special case of refinement.

2.2.4 Platform and Implementation Language Environment

In the MDA, the term *platform* is used to refer to technological and engineering details that are irrelevant to the fundamental functionality of a software component. Consider, for example, a formal definition of an operation that transfers funds from a checking account to a savings account. The fundamental functionality of this operation is that a specified amount is subtracted from a designated checking account and added to a designated savings account, with a constraint that the two accounts must belong to the same customer. This functionality remains invariant regardless of whether the operation is performed by a CORBA object, an Enterprise Java Beans, or a SOAP operation⁵.

Thus, a *platform-independent model* is a formal specification of the structure and function of a system that abstracts away technical detail⁶. By this definition a SOAP specification of the funds transfer operation would be *platform-specific*. A specification that depends on interfaces to artifacts of CORBA, like the ORB, Object Services or GIOP/IOP would be an example of a *platform-specific model*.

Note that platforms themselves also have a specification (e.g. of component constructs, such as facet and receptacles, ports, and connectors, and services, such as directory and transactions); and an implementation (the platform component constructs are realized by some refinement e.g. receptacles as some IDL interface pattern, connectors between event sources and sinks as a particular adapter pattern, services implemented in some implementation language). A PSM is expressed in terms of the *specification model* of the target platform. CORBA itself is implemented on an infrastructure, which could properly be referred to as a implementation language platform. However, to avoid confusion, we use the term *implementation language environment* to refer to such infrastructures in the MDA. Thus, analogous to the dichotomy established for platforms, CORBA specifications are *implementation language environment independent*, whereas artifacts like stubs, skeletons and the ORB implemented in a specific language are *implementation language environment specific*.

2.3 Models in the MDA

2.3.1 Architecture for MDA Models

The MDA separates certain key models of a systems, and brings a consistent structure to these models. Figure 2 on page 7 shows that models of different systems are structured explicitly into *Platform Independent Models* (PIMs), and *Platform Specific Models* (PSMs). How the functionality specified in a PIM is realized is specified in a platform-specific way in the PSM, which is derived from the PIM via some transformation.

The PIMs provide formal specifications of the structure and function of the system that abstracts away technical details. A Platform Independent Component View describes computational components and their interactions in a platform-independent manner. These components and interfaces, in turn, are a way of realizing some more abstract information system or application, which itself helps realize a computation-independent Business Model⁷ OMG standards are specified in terms of a PIM and, normally, one or more PSMs, all in UML.

In addition, the MDA defines consistent relationships across these models. As shown in Figure 2, for a given system there are cross-model *refinement* correspondences (the large double-headed arrows in the figures) between business model, platform-

5. In fact, there are many partial specifications of transfer: *s1* = move funds; *s2* = notify IRS if > 100K moved; *s3* = *s1* & *s2*; *s4* = *s3* & failure_spec; *s5* = CORBA version of *s4*. Platform independence separates out *s5* from the rest.

6. Note that we do not say *the* platform-independent model. There could be models of a system from multiple viewpoints, at multiple levels of abstraction, all of which are platform-independent e.g. RM-ODP enterprise, computational, and engineering viewpoints.

independent components, and platform-specific components. Similarly, across two different systems to be integrated (e.g. in EAI), interactions may be specified at the platform-specific, platform-independent, and even business model levels of abstraction.

There could be more and less abstract models within each of the business, PIM, or PSM levels. The PIM (including the Business Model) and the PSM are viewpoint specifications of the system for the OMG standard. They are closely related to specific sets of viewpoints defined by the RM-ODP Part 3.

The computation independent business model is one in which the Computational (c.f. RM-ODP) details are hidden or as yet undetermined. It is sometimes referred to as the business domain model.

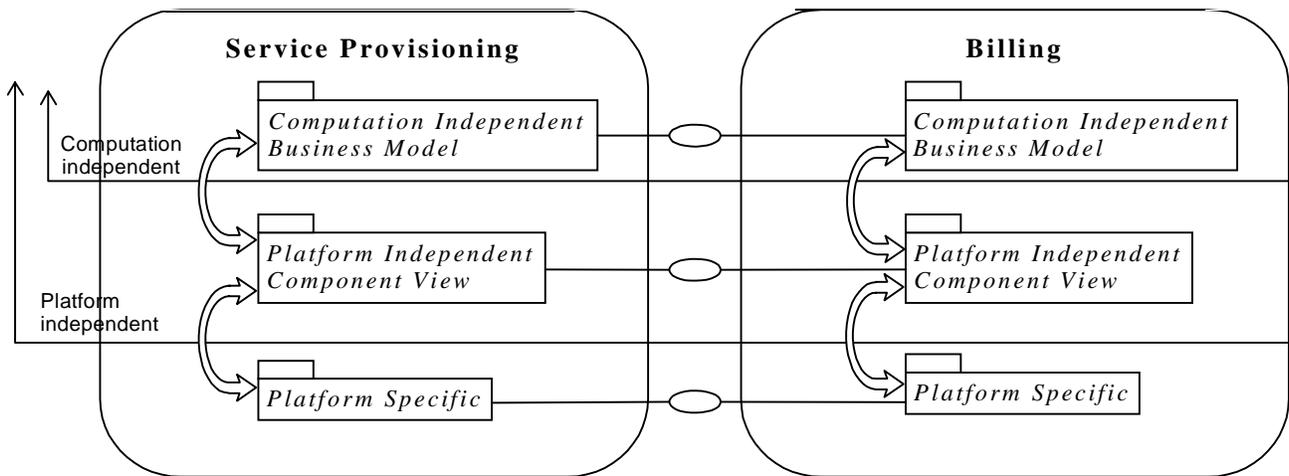


Figure 2. Consistent Model Separations and Relationships in MDA

Abstracting out the fundamental precise structure and behaviour of a system in the PIM from implementation specific concerns in the PSMs has three important benefits:

1. It is easier to validate the correctness of the model uncluttered by platform-specific semantics. For example, PSMs have to use the platform concepts of exception mechanisms, parameter types (including platform-specific rules about objects references, value types, semantics of call by value, etc.), and component model constructs; the PIM does not need these distinctions and can instead use a simpler, more uniform model.

7. This computation-independent description is sometimes referred to as a domain model. While it need not be explicitly present in a particular usage of the MDA scheme, MDA accommodates it consistently in the same overall architecture.

2. It is easier to produce implementations on different platforms while conforming to the same essential and precise structure and behavior of the system. Even further, the business model defines business goals and policies in a computation-independent manner.
3. Integration and interoperability across systems can be defined more clearly in platform-independent terms, then mapped down to platform specific mechanisms.

Where appropriate generic mappings or patterns can be shared across multiple applications, it may be possible to automatically transform a PIM, perhaps after annotating it with some platform information, to different target PSMs either fully or partially. The optimal transformation depends on QoS and other requirements, simplified in Figure 3 to a single shared *PIM to CORBA* generic correspondence mapping. Figure 3 applies regardless of whether the mapping is automated or manual.

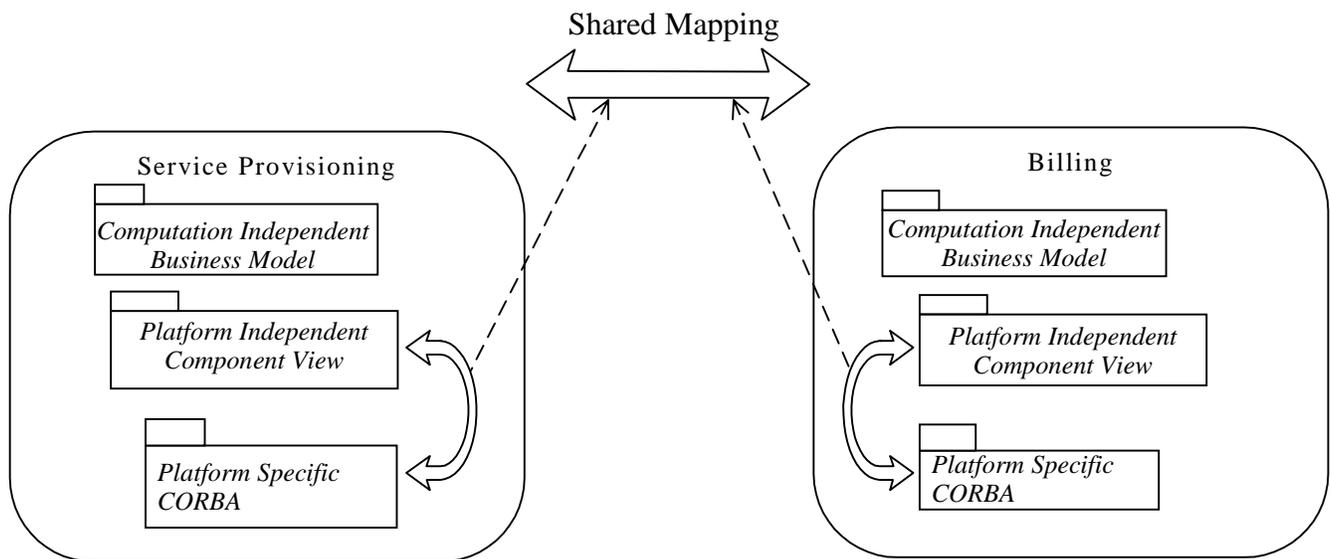


Figure 3. Shared Patterns of PIM <-> PSM Mappings

Figure 4 shows the structure, of specifications, realizations, and refinements from business domain to code, covering any granularity of system, sub-system, etc. A given “platform” is a large (infrastructure) component with a specification, including services available to other assembled components. A given usage of MDA may choose to exploit this to a lesser extent. Recurring patterns occur across all levels, such as between PIM and PSM. These form an important part of the architectural style.

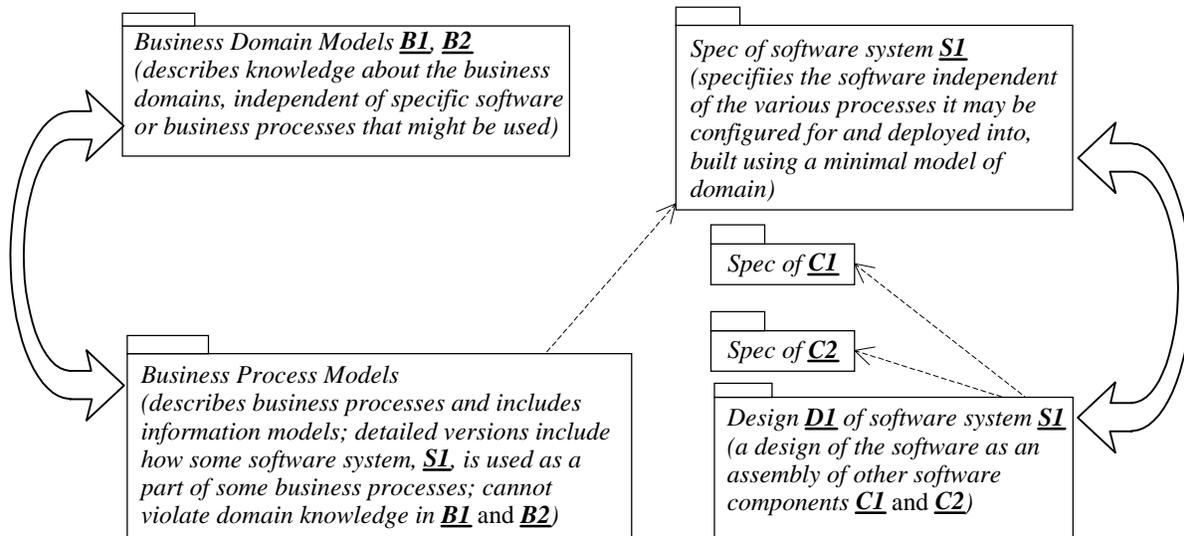


Figure 4. Structure of specification, refinement, and design.

2.3.2 Platform Independent Models in UML

UML models are declarative models, as are IDL-based object models, Java interfaces, and Microsoft IDL interfaces. However, UML models differ from these other kinds of declarative models in some important ways.

First, UML has been defined using core UML modeling concepts and this enhances the power of MDA.

Secondly, UML models can be expressed textually as well as graphically.

Finally, UML models can be semantically much richer than models expressed in the other declarative model languages mentioned above, which can express syntax (i.e. signature) but very little about constraints on usage and behavior such as:

- Static invariants constraints on combinations of attributes.
- Pairs of pre and post-conditions for specifying operations.
- Whether a single-valued parameter is allowed to be null.
- Whether an operation has side effects.
- Whether subtypes of some supertype are disjoint or form a partition.
- Patterns of specifications, designs and refinements.

Static invariants and pre/post conditions are particularly important features of an approach to rigorous software engineering called *contract based design*. UML did not invent the concept of contract based design, but it has very good support for it. UML defines a formal assertion language called *Object Constraint Language* (OCL) that facilitates specification of certain constraints. While contract based design does not

eliminate the need for informal textual explanations of interface usage, it can significantly reduce dependence on them. The UML allows formalization of the vocabulary otherwise left imprecise in interface specifications, as an abstract yet precise model of the state of the object providing that interface and of any parameters exchanged.

Some current OMG specifications including UML, MOF and CWM specifications already use UML and OCL for specifying constraints

Specifying constraints formally rather than in free form text reduces ambiguity in specifications and thus makes life easier for implementers in three important respects:

1. It provides the programmer with more precise instructions, thus lessening the extent to which the programmer has to guess at the designer's intention or track down the designer to find out what to do.
2. It decreases the amount of work required to get different implementations of the same specification working together, or to integrate implementations of two specifications whose models are unambiguously related.
3. The formal specification provides a foundation for defining conformance tests for different implementations.

A model that specifies interfaces to this much precision is profoundly different in character from one that does not. The MDA approach allows OMG specifications to rise to this level of rigor.

2.3.3 *Platform Specific Models in UML*

In section 2.3.1 it is stated that a PSM is expressed in UML, however, since UML is independent of middleware technologies, it is not obvious to the casual observer how to harness this power to express a PSM.

For example, in order to transform a PIM into a CORBA PSM certain decisions need to be made. Are the UML classes represent CORBA interfaces, valuetypes, structs, and unions? If so how does one make it clear that a particular UML class is an interface as opposed to a valuetype? If not, what do the classes represent and how would they be useful?

Such decisions can be defined by a UML profile, which is a set of extensions to UML using the built-in extension facilities of UML, stereotypes and tagged values. Stereotypes label a model element to denote that the element has some particular semantics.

The UML Profile for CORBA, adopted in 2000, specifies how to use UML in a standard way to define CORBA IDL interfaces, structs, unions, etc. For example, it defines stereotypes named *CORBAInterface*, *CORBAValue*, *CORBAStruct*, *CORBAUnion*, etc. that are applied to classes to indicate what the class is supposed to represent. In the graphical UML notation a stereotype is delimited by angle brackets, as illustrated in Figure 5.

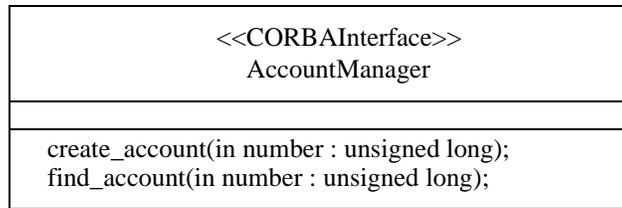


Figure 5. A Stereotype Indicating that the Class Represents a CORBA Interface

Figure 6 is a fragment of the specification of a CORBA interface that uses the semantic power of UML to formalize an invariant rule. The invariant rule cannot be formally specified in IDL, and thus we consider this model to be a semantically enhanced CORBA specification




Figure 6. A Fragment of a Semantically Enhanced CORBA Specification

The model fragment in Figure 6 corresponds to the IDL shown in Figure 7, assuming that UML attributes map directly to exposed attributes in CORBA interfaces.

```

interface Account {
    attribute short number;
    attribute float balance;
};
  
```

Figure 7. IDL--By Nature Semantically Thin

Thus, with the UML Profile for CORBA, CORBA-based specifications can be made much more complete than is possible with IDL only. The normative English text that specifies rules such as the allowable range of the Account number today, will be

replaced by formal invariant rules expressed in terms of the UML Profile for CORBA in the near future. The ORBs of today need only understand the IDL; they do not need to understand the formal specification of behavior and constraints in the more precise specification any more than they need to understand informal specification of behavior and constraints in English text.

Similarly, UML profiles can be defined for other platforms, providing the essential tools for constructing PSMs.

The technology is in place to proceed in this direction. The main barrier is that there is a gap in knowledge of how to use the technology, and there is a lack of universal availability of appropriate tools.

2.3.4 Mappings of Models

Figure 8 below presents a metamodel description of MDA. PIM, PSM and mapping techniques are based on metamodel expressed preferably with OMG core technologies like MOF, CWM or UML.




Figure 8. MDA Metamodel Description

In MDA, one of the key feature of the whole approach is the notion of mapping. A mapping is a set of rules and techniques used to modify one model in order to get another model. Mappings are used for transforming:

-
1. **PIM to PIM.** This transformation is used when models are enhanced, filtered or specialized during the development lifecycle without needing any platform dependent information. One of the most obvious mapping is the analysis to design models transformation. PIM to PIM mappings are generally related to model refinement.
 2. **PIM to PSM.** This transformation is used when the PIM is sufficiently refined to be projected to the execution infrastructure. The projection is based on the platform characteristics. Describing these characteristics should be done using a UML description (and eventually a profile for describing common platform concepts). Going from a logical component model to a commercial existing component model (like EJB for J2EE platform or CCM for CORBA platform) is a kind of PIM to PSM mapping.
 3. **PSM to PSM.** This transformation is needed for component realization and deployment. For example, component packaging is done by selecting services and preparing their configuration. Once packaged, the components delivery could then be done by specifying initialization data, target machines, container generation and configuration, etc. PSM to PSM mapping are generally related to platform dependent model refinement.
 4. **PSM to PIM.** This transformation is required for abstracting models of existing implementations in a particular technology into a platform-independent model. This procedure often resembles a "mining" process that is hard to be fully automated. It may be supported by tools, though. Ideally, the result of this mapping will match the corresponding PIM to PSM mapping.

UML profiles have an important role to play in MDA, since MDA leverage usage of mappings between models. To implement the mapping, one needs to know the metamodels of the input and output models and their mapping rules. Transformation rules execution can be done inside UML tools (for example using scripting) or by external tools (for example, by working on XMI files). We present in Figure 9 on page 14 an example of PIM to PSM mapping using the OMG profile for Enterprise Distributed Object Computing (EDOC). A component architecture is described using the EDOC profile. The EDOC profile specifies also standard mappings between EDOC and standard component models, like EJB. The EDOC to EJB mapping is not a profile *per se*, but could be associated to the profile notion, because you know the input and the output model and metamodel. So specifying standard rules and techniques to go from one model to the other is then possible.

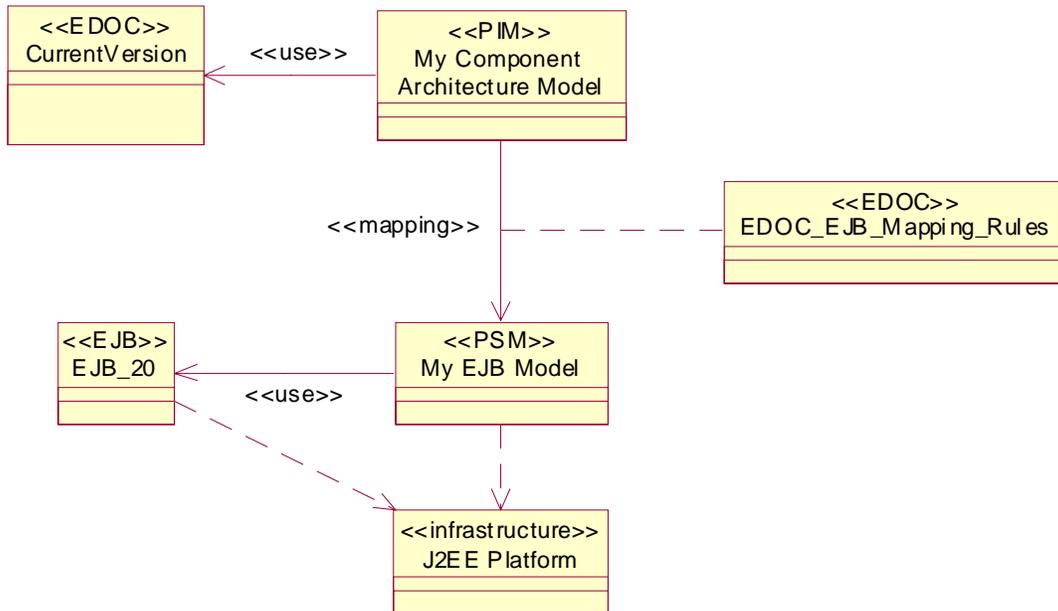


Figure 9. MDA Metamodel Example

2.3.5 Platform Independent and Platform Specific UML Models

There are multiple ways to transform a PIM expressed using UML into a corresponding PSM expressed using UML

1. A human could study the platform-independent model and manually construct a platform-specific model, perhaps manually constructing the one-of refinement mapping between the two.
2. A human could study the platform-independent model and utilize models of known refinement patterns to reduce the burden in constructing the PSM and the refinement relation between the two.
3. An algorithm could be applied to the platform-independent model and create a skeleton of the platform-specific model to be manually enhanced by hand, perhaps using some of the same refinement patterns in 2.
4. An algorithm could create a complete platform-specific model from a complete platform-independent model, explicitly or implicitly recording the refinement relation for use by other automated tools.

Note that the above list does not address the production of executable code from a platform-specific model.

There also are variations in which the platform-specific model is not a semantically rich UML model but, rather, is expressed via a language such as IDL, Java interfaces, XML, etc.

Fully automated transformations are feasible in certain constrained environments. The degree to which transformations can be automated is considerably enhanced when the following conditions are obtained:

- There is no legacy to take into account
- The model that serves as input to the transformation is semantically rich
- The transformation algorithms are of high quality

It is much easier to generate executable code for structural features (attributes, certain associations and similar properties) of a model rather than behavioral features (operations) because the behavior of property getters and setters are quite simple.

Automation of transformations is more tractable when the transformation is parameterized, i.e. a human has a pre-defined set of options to select from, to determine how the transformation is performed. For example, a system that transforms a UML model to an XML DTD could allow some control over how a UML class's attributes are transformed, giving a human a chance to choose to put them in an ATTLIST or to put each attribute in a separate ELEMENT.

Some of these points are illustrated by a CORBA-specific model, corresponding to the fragment of a platform-independent model shown earlier in Figure 10. Note that this model fragment has been annotated with a stereotype to denote that Account is an entity as opposed to a process⁸. The model fragment has also been enhanced to indicate that the Account number constitutes a unique identifier for Account instances. These annotations are platform-independent, but they can directly help select an appropriate PSM pattern

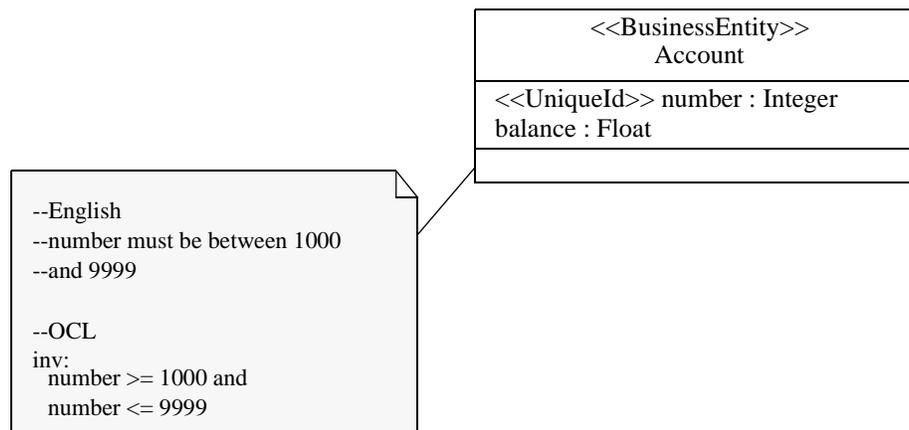


Figure 10. Platform-Independent Model, Annotated with Stereotypes

Figure 11 shows a CORBA-specific UML model constructed from this fragment. We are not taking a stand that this is *the* proper way to construct a CORBA solution from the platform-independent UML fragment. It is simply an example of how one *might* do so either manually or algorithmically. The stereotypes '<<BusinessEntity>>' and '<<UniqueId>>' could themselves indicate the systematic application of this pattern,

⁸In the field of distributed business component-based software it is now widely understood that the entity-process distinction is crucial to building scalable systems.

when interpreted in the context of the CORBA Profile. Note that `Session::BaseBusinessObject` is an element defined in the OMG Task and Session Service. The logic of the construction uses the enhancements to the platform-independent model that indicate that `Account` is a business entity and that the `Account` number constitutes an `Account`'s unique identity. It reflects a commonly used pattern of specifying one interface exposed by entity instances and the other exposed by a manager of the entity instances, including factory finder operations that use the unique identifier as a selector. The entity instance interface has an attribute that provides a reference to the instance manager.

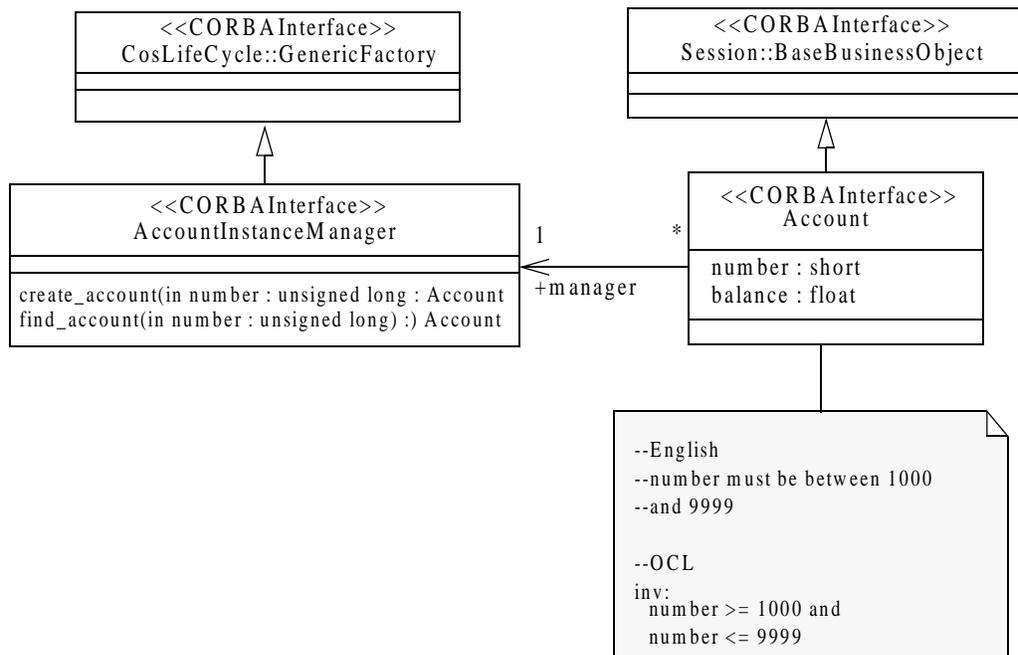


Figure 11. A CORBA-Specific UML Model Derived from the PIM expressed using UML

Figure 12 contains the IDL that expresses the same CORBA-specific solution. Of course the IDL is semantically thin. Its formal constructs do not and cannot express the invariant rule about the account number range. Furthermore, the IDL does not and cannot formally indicate whether a well-formed `Account` instance is allowed to have a null manager reference. On the other hand, the CORBA-specific UML model makes it clear, via the multiplicity of 1 on the *manager* end of the association between `Account` and `AccountInstanceManager`, that a well-formed account must have a non-null manager reference. If the solution designer intended to allow a null manager reference, then the multiplicity would be 0..1 in the CORBA-specific UML model, but the IDL would be the same as in Figure 12 below.

```

interface AccountInstanceManager : CosLifeCycle::GenericFactory {
    Account create_account (in unsigned short number);
    Account find_account (in unsigned short number);
};

interface Account : Session::BaseBusinessObject {
    attribute AccountInstanceManager manager;
    attribute short number;
    attribute float balance;
};

```

Figure 12. IDL Corresponding to the CORBA-Specific UML Model

2.3.6 Packages and Viewpoints

UML provides an important modeling element that is relevant for separating viewpoints, levels of abstraction, and refinements in the context of MDA – the UML *package*, a UML construct for grouping model elements. A package can *import* other packages, making elements from the imported package available for its use. Packages help understand the MDA since the models being interrelated for integration are typically in different packages.

Models of a system from two different viewpoints unrelated by refinement (or models of two distinct interfaces of a component, or models of two different systems to be integrated) would be defined in two separate packages as in Figure 13(a). The interrelationships between the two would be defined in a model correspondence (large arrow), defining the integration of those two viewpoints or systems at the current level of abstraction. Similarly, models of the same system at different levels of abstraction are defined in two packages as in Figure 13(b); the *refinement* model relating them would be another model correspondence. Refinement is a relation between model elements that can be in separate packages, and not between the packages themselves

Naturally, the model correspondences could use recurring patterns from a shared generic model correspondence. Also, since one can “zoom” in and out of any granularity of object or component; the *realization* package P6 assembles some subcomponents by importing their specification packages.

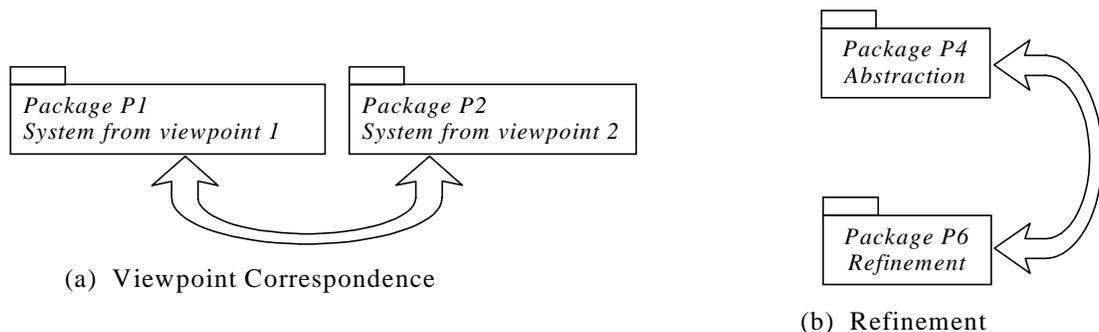


Figure 13. Model Correspondences (a) Viewpoints and (b) Refinement

2.3.7 Traceability

The relationships between elements of the PIM expressed using UML and CORBA-specific UML models can only be partly specified in UML 1.4. Figure 14 illustrates loosely that the AccountInstanceManager CORBA interface is a refinement, at a different level of abstraction, of the number attribute in the PIM that constitutes Account's unique identifier. <<refine>> is a standard UML stereotype of dependency. Note that UML namespaces are used to distinguish between the two Account elements. The platform-independent Account class is contained in a namespace called *PlatformIndependent*. The desired namespace separation can be achieved by putting the two models in separate UML *Packages*, as suggested in Figure 13 on page 17




Figure 14. Tracing Between Elements of the Models

An example of the refinement relationship across levels of abstraction is shown using plain notes to describe the mapping in Figure 15 on page 19⁹.

9. It is generally recognized that UML's facilities for relating models at different levels of abstraction are rudimentary and need expansion. UML's features do not adequately support powerful zoom-in and zoom-out capabilities. The UML 2.0 RFPs call for facilities that address this gap, to enable a more complete definition of the refinement relationship across levels of abstraction, such as the one shown in Figure 15 on page 19.

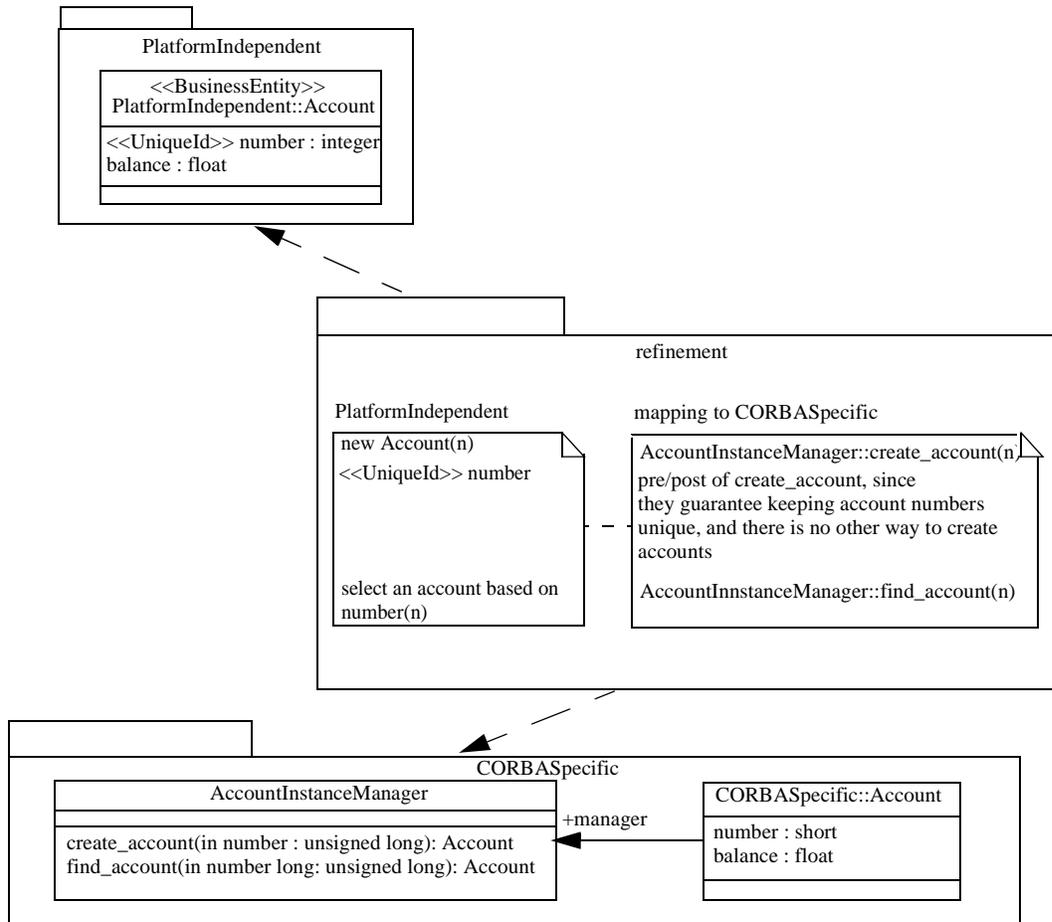


Figure 15. Refinement Relation between PIM and PSM levels of Abstraction

The above set of examples highlights the value of the MDA. The architect/modeler focuses on creating the platform independent architectural and business model of the application. The middleware/e-services designer uses UML profile for CORBA to model the platform specific aspects of the system so that CORBA interface generation can be automated. The programmer then uses the UML model as well as the IDL interfaces to augment whatever generated code exists with additional code to complete the value added business logic for the service. The explicit mapping relations across models are used to enable potential automation of PSM generations as well as for easier integration. All these artifacts are traced and versioned. This allows MDA to be used by systems administrators, architects, designers as well as developers

3 OMG Standards and the MDA

3.1 Overview

Figure 16 illustrates how OMG standards fit together in MDA. It provides an overall framework within which the roles of various OMG and other standards can be uniquely identified. This section looks at the overall picture in Figure 16. The following sections look at some issues related to specific OMG standards.

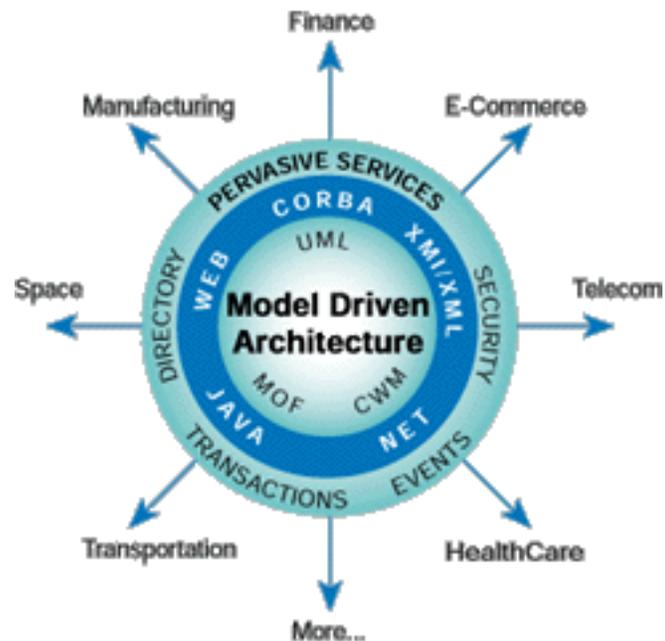


Figure 16. OMG's Model Driven Architecture

3.1.1 The Core

The core of the architecture, at the center of the figure, is based on OMG's modeling standards: UML, the MOF and CWM. The core comprises a number of UML profiles. One will represent Enterprise Computing with its component structure and transactional interaction; another will represent Real-Time computing with its special needs for resource control; more will be added to represent other specialized environments but the total number will be small. Each UML profile represents the common features of all of the middleware platforms appropriate for its category of computing, but will be independent of any specific platform.

Whether your ultimate target is CCM, EJB, MTS, or some other component or transaction-based platform, the first step when constructing an MDA-based application will be to create a PIM of the application expressed in UML using the appropriate UML profile. Platform specialists will transform this application model into one

targeted to a specific platform such as CCM, EJB, or COM+. Standard mappings will allow tools to automate some of the transformation. In Figure 16 on page 20, these target platforms occupy the thin ring surrounding the core.

Among these target platforms, CORBA occupies a special role in OMG's standards activities as the target platform of choice when an implementation language independent target platform with the interoperability implied by the use of IIOP is required. The other target platforms are admissible, but with the understanding that additional work would be involved to obtain the level of out of the box interoperability that one gets when CORBA is used as the target platform.

The PSM faithfully represents both the business and technical run-time semantics of the application. It's still a UML model, but is expressed (because of the transformation step) in a dialect (i.e. a profile) of UML that precisely mirrors technical run-time elements of the target platform. The semantics of the platform-independent original model are carried through into the platform-specific model.

Generally, standards in the Core are developed and managed by the Platform Technical Committee (PTC) at the OMG.

3.1.2 *Pervasive Services*

All applications, independent of their context (e.g. for the enterprise or the Internet, or for embedded computing) rely on some or all of a set of essential services. The list varies somewhat depending on the source but typically includes Directory services, Event handling, Persistence, Transactions, and Security.

When these services are defined and built on a particular platform, they necessarily take on characteristics that restrict them to that platform, or ensure that they work best there. To avoid this, OMG will define such services as *pervasive services* at the PIM level in UML. Only after the features and architecture of a pervasive service are fixed will platform-specific definitions be generated for all of the middleware platforms supported by the MDA.

At the abstraction level of a platform-independent business component model, pervasive services are visible only at a very high level (similar to the view the component developer has in CCM or EJB). When the model is mapped to a particular platform, code will be generated (or dynamically invoked) that makes the calls to the native services of those platforms.

In Figure 16 on page 20, the Pervasive Services are shown as a ring around the outside of the diagram to emphasize that they're available to all applications, in all environments. True integration requires a common model for directory services, events and signals, and security. By extending them to a generalized model, implementable in the different environments and easily integrated, the MDA becomes the basis of our goal of universal integration: the *global information appliance*. In Figure 16, Pervasive services specify services that an infrastructure component offers to other assembled components.

Pervasive Services standardized at the OMG fall in the grey area where some are standardized through the PTC and some through the Domain Technical Committee (DTC).

3.1.3 *OMG Domain Specifications*

A sizeable percentage of OMG activity is focused on standardizing services and facilities in specific vertical markets through Domain Task Forces (DTFs) of the DTC. Initially these specifications consisted of interfaces written in OMG IDL with accompanying semantic description in English text. Standardizing components at a platform level, in terms of standards such as CORBA, is certainly a viable contribution to solving the integration and interoperability problem, but the MDA offers much more.

A well-conceived service or facility is always based on an underlying semantic model that is independent of the target platform. However, the model may not be distilled explicitly, and this is the case with OMG's domain specifications because the model for virtually every one is not expressed separately from its IDL interfaces. Since their models are hidden, these services and facilities have received neither the recognition nor the widespread implementation and use that they deserve outside of the CORBA environment, especially considering the quality of their underlying models. Extending these implied models outside of CORBA just makes sense. The Healthcare Resource Access Decision Facility, already implemented in Java and EJB in addition to CORBA, is an example. There are more.

Thus, in order to maximize the utility and impact of OMG domain facility specifications in the MDA, they will be in the form of normative PIMs expressed using UML, augmented by normative PSMs expressed using UML and interface definitions for at least one target platform. The common basis in the MDA will promote partial generation of implementation code as well, but implementation code of course will not be standardized.

The DTFs of the DTC produce standard frameworks for standard functions in their application space. For example, a Finance DTF standard for an accounts receivable facility might include a PIM expressed using UML, a CORBA-specific UML model and IDL interfaces and a Java-specific UML model and Java interfaces. XML DTDs or schema generated via XMI-based mapping rules could be included as well. All of these artifacts would be normative. Such a standard would have broad impact, in that the platform-independent model would be useful even in middleware environments other than those targeted by the platform-specific parts of the specification. Since accounts receivable is an Enterprise Computing application, the normative, platform-specific artifacts would be derived at least partially via standard mappings of the Enterprise Computing profile to the platforms.

Today OMG has ten DTFs with several more "in the chute." Rather than show them all in a static diagram, only a representative sample is shown in Figure 16 on page 20 where they appear as rays emanating from the center.

3.1.4 *Transparencies and Qualities of Service*

Environments with specific hardware and software attributes – Scalability, Real-Time operation, Fault Tolerance, or Embedded characteristics – may be modeled as well. By defining UML representations for these environments or, OMG will extend the MDA to support and integrate applications with these desirable characteristics.

3.2 **System Lifecycle - MOF, UML, CWM and XMI**

IT systems have historically been developed, managed and integrated using a range of methodologies, tools and middleware and there appears to be no end to this innovation. What we have seen in the last few years, especially as a result of efforts at OMG and W3C is a gradual move to more complete semantic models as well as data representation interchange standards. OMG contributions include CORBA, UML, XMI, MOF and CWM. W3C contributions include XML, XML Schema, and the ongoing work of the XML-P working group. These technologies can be used to integrate more completely the value chain (or life cycle) when it comes to developing and deploying component based applications for various target software architectures.

The life cycle of an application can vary dramatically depending on whether we are building a new application from scratch or just surgically adding a wrapper to an existing application. The cost of enhancement and maintenance of an application as well as the cost of integrating new applications with existing applications far exceeds the cost of initial development. In addition the application life cycle itself can be quite complex, involving several vendors in each of the life cycle phases. Hence the need for information interchange and interoperability between tools and middleware provided by different vendors (a very common situation in enterprises today) is critical.

The MDA supports many of the commonly used steps in model driven component based development and deployment. A key aspect of MDA is that it addresses the complete life cycle covering analysis and design, programming (testing, component build or component assembly) and deployment and management. An example is the way in which UML, XMI, MOF and CWM affect the interchange of information between tools and applications.

The MDA core is based on OMG technologies (MOF, UML, CWM). These technologies are used to describe PIMs. A PIM can be refined n-times until the desired system description level is obtained. Then, the infrastructure is taken into account and the PIM is transformed into a PSM. Then, again, PSMs are refined as many times as needed. Figure 17 on page 24, presents the development lifecycle of the system, emphasizing the models used and the way they are typically refined.

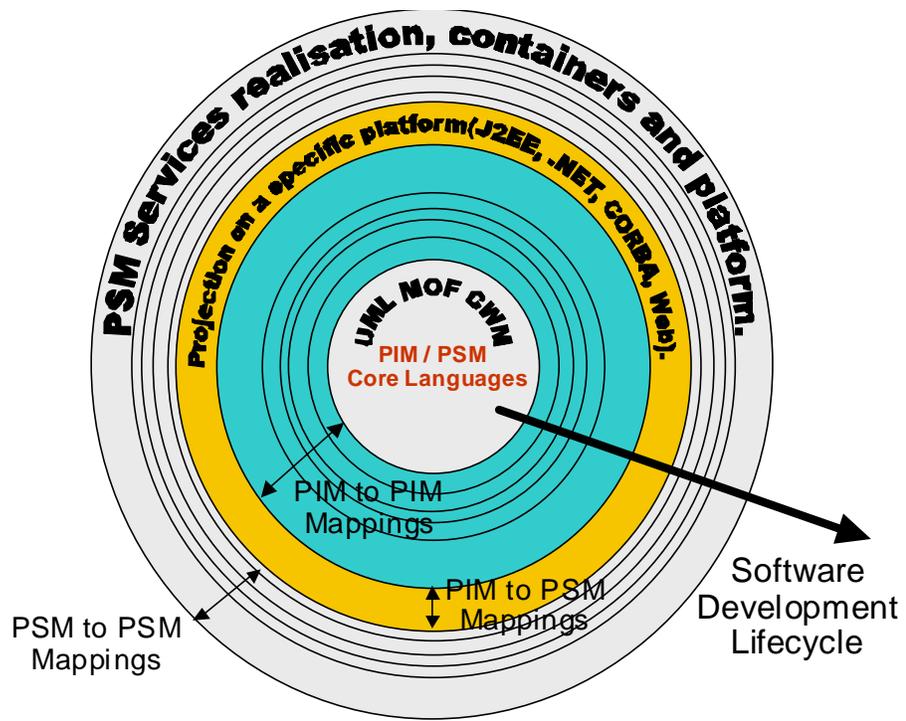


Figure 17. System Lifecycle and the MDA Approach

3.2.1 UML (Unified Modeling Language)

UML addresses the modeling of architecture, objects, interactions between objects, data modeling aspects of the application life cycle, as well as the design aspects of component based development including construction and assembly. Note that UML is powerful enough to be used to represent artifacts of legacy systems. Artifacts captured in UML models (in terms of Classes, Interfaces, UseCases, Activity Graphs etc.) can be easily exported to other tools in the life cycle chain using XMI.

A number of UML profiles (for CORBA, EJB, EDOC etc.) are at various stages of standardization (UML profile for CORBA is adopted). These are critical links that bridge the UML community (model based design and analysis) to the developer community (Java, VB, C++ developers), middleware community (CORBA, EJB, SOAP developers) etc. Additional profiles focused on systems and application management are needed

3.2.2 XMI (XML Metadata Interchange)

XMI is a standard interchange mechanism used between various tools, repositories and middleware. XMI can also be used to automatically produce XML DTDs (and soon XML Schemas) from UML and MOF models, providing an XML serialization mechanism for these artifacts. XMI has been used to render UML artifacts (using the

UML XMI DTD), data warehouse and database artifacts (using the CWM XMI DTD), CORBA interface definitions (using the IDL DTD), and Java interfaces and Classes (using a Java DTD).

XMI, which marries the world of modeling (UML), metadata (MOF and XML) and middleware (UML profiles for Java, EJB, IDL, EDOC etc.) plays a pivotal role in the OMG's use of XML at the core of the MDA. In essence XMI adds *Modeling* and *Architecture* to the world of XML.

3.2.3 *MOF (Meta Object Facility)*

MOF provides the standard modeling and interchange constructs that are used in MDA. Other standard OMG models, including UML and CWM, are defined in terms of MOF constructs. This common foundation provides the basis for model/metadata interchange and interoperability, and is the mechanism through which models are analyzed in XMI. MOF also defines programmatic interfaces for manipulating models and their instances spanning the application lifecycle. These are defined in IDL and are being extended to Java.

3.2.4 *CWM (Common Warehouse Metamodel)*

CWM is the OMG data warehouse standard. It covers the full life cycle of designing, building and managing data warehouse applications and supports management of the life cycle. It is probably the best example to date of applying the MDA paradigm to an application area.

Historically, the integration between the development tools and the deployment into the middleware framework, has been weak. This is now beginning to change by using key elements of the MDA – specific models and XML DTDs that span the life cycle, and profiles that provide mappings between the models used in various life cycle phases.

3.3 Implementation Language Independent Models and IDL

A set of IDL modules containing specifications of IDL interfaces, valuetypes and other datatypes is a declarative syntactic model of a system. Such a model can be used to reason about the validity or lack thereof of relationships among the entities specified using the rules of relationship among IDL declared entities like containment, inheritance etc. An IDL specification is an object model that can be implemented on a CORBA platform that will implicitly verify the syntactic validity of any attempt to use any part of the system.

However, such a specification does not contain much formal information about the meaning of the operations of the interfaces or of the elements of the datatypes declared, nor about the constraints that apply to them. In traditional CORBA specifications such information has been included in a normative but informal description in English.

In this approach, an IDL compiler can be used to statically verify syntactic correctness of the model. An ORB can verify syntactic correctness of attempts to use parts of the system dynamically. However, there is no automatic way of verifying the constraints and functionality that appears in the specifications in informal descriptions

IDL was not designed to express a rich set of relationships among entities. The description of relationships between different parts of a system is also to a large extent informal, and hence prone to multiple interpretations. Traditionally, descriptions of relationships among CORBA Services, and indeed among different artifacts that constitute a CORBA service, appeared in the form of informal text. In more recent specifications (e.g. POA), the use of UML to more completely describe the model has brought additional rigor to the specifications.

3.3.1 Platform and Language Environment Independence of IDL Specified Models

IDL itself is not tied to any specific language environment or platform. This is what made it possible for ISO to adopt IDL as a standard without any specific reference to CORBA. Indeed there are many systems in this world which use IDL to specify the syntactic model of the system but do not use CORBA as the underlying platform. While OMG has not standardized any such usage of IDL with alternative platforms, there are broadly deployed instances in the industry of such use. However, it should be noted that in spite of being platform and language environment independent, IDL specified models are restricted to expressing only the syntax of the interactions, i.e. operation signatures.

OMG has chosen to use IDL together with the CORBA platform (ORB and language mappings) as a reasonable package of facilities to standardize. This facilitates algorithmic construction of skeletons of portable components of the system for a specific language environment, from language independent specifications, using an IDL compiler. The big win from this is *portability of specifications* from one language environment to another, as well as *portability of implementations* among different instances of the same language environment.

Additionally, given specifications of the exact syntax of interaction between objects that constitute the system, it is also possible to automatically generate the syntactic form that is carried on a wire that connects the two communicating objects. OMG has standardized on GIOP/IOP as the standard means of conveying communication between IDL declared objects deployed on a CORBA platform. Again, IDL, and even the CORBA platform, does not preclude use of other means of communication between objects. Indeed, it is quite possible for two CORBA objects to communicate with each other using DCOM or SOAP on the wire. But the adoption of a single means of interoperation ensures *interoperability of implementations*.

3.3.2 Extensions to IDL to Capture Additional Information

Various attempts have been made to extend IDL to capture richer structural and behavioral information and to automatically generate implementation artifacts for a given platform that enforces the constraints as specified in the richer specification. A recent example of this is the Components extension of IDL together with the XML

based deployment descriptors, which facilitates specification of entire systems in terms of its constituent components, their interactions and deployment characteristics. However, it should be noted, that all such extensions so far have been point solutions, without paying much attention to a general model for specifying such extensions.

A model defined in the UML profile for CORBA (see Figure 11 on page 16 and Figure 12 on page 17) provides an alternative representation of an IDL model. The two models are different representations of the same system. In fact, there is precisely one IDL representation that can be derived from a model represented using the UML profile for CORBA. The UML model may, however, provide additional information (such as cardinality) that cannot be represented in an IDL model today. Appropriate extensions to IDL, that allow representation of these additional relevant concepts, would make it possible to map a model expressed in the CORBA profile of UML to an equivalent IDL model in a reversible fashion. That is, one would be able to reconstruct the corresponding UML from the equivalent IDL, without loss of information. The ability to “round-trip” the transformation in this way would allow designers and architects to work in the technology that they are comfortable with (UML or IDL) and algorithmically generate the alternative representation for the specification.

3.4 CORBA, CORBA Services and GIOP

The OMG standard platform consists of the specifications commonly referred to as CORBA, CORBA services, and GIOP/IIOP. A more complete overview of these and the underlying Object Management Architecture can be found in the document [Ref: Discussion of the OMA formal/00-06-04?]

3.4.1 Standard CORBA Platform and Bridging to Other Platforms

The general philosophy behind the CORBA platform standards has been to adopt a single set of standards within a broader framework that allows alternatives if there is such a need. The standard interoperability framework recognizes such possibilities and explicitly defines interoperability domains in which different standards apply, and how bridges can be specified to enable objects in different such domains to communicate with each other, thus making it possible to construct systems that span multiple domains.

This framework has been successfully used to specify bridges between the CORBA platform with GIOP/IIOP based communication and the COM/DCOM platform and communication domain in an existing OMG standard. More recently an inter-domain bridge between the CORBA Component Model and the EJB Component Model has also been adopted as a standard. This shows the tremendous versatility of the CORBA and associated interoperability framework.

3.4.2 Portability and Bridging for Domain Specific Facilities

The problem of bridging from one platform to another becomes considerably simpler if the two platforms in question share a common model at a higher level of abstraction. It is fortunate that most broadly deployed distributed computing environments happen to share such a common model, although never formally expressed as such, thus making construction of bridges among them feasible.

As the basic CORBA platform and associated CORBA services specifications became rich enough to support the building of domain specific facilities the need for expressing the underlying model in a formal way, at an appropriate level of abstraction, has been felt more acutely. This is somewhat analogous to the need that motivated IDL based specifications, but at a higher level of abstraction. A formally specified model is useful because:

- It facilitates creation of compatible platform specific models/specifications corresponding to the same platform-independent model and hence implementations that are easier to bridge together.
- It provides a common reference model and vocabulary with unambiguous meaning thus reducing the chances of miscommunication among system designers and builders.
- It facilitates standardization of more precisely specified designs and patterns, thus allowing for *portability of design*, and makes it easier to support *interoperability among different realizations of the same design on different platforms*.

Thus, given the experience gained working on CORBA systems specifications and of bridges to other similar platforms, it is a natural step for OMG to adopt standardized means of expressing richer formal models at appropriate levels of abstraction, from multiple viewpoints.

3.5 Standards Development

The OMG's UML Profile for EDOC RFP process currently underway will define standard stereotypes for platform-independent models that will provide standard ways to use UML to express the structure, behavior and constraints of components in a platform-independent fashion. It will contain proof of concept mappings to the CORBA Component Model and to EJB. A follow-on RFP is expected that will call for standardization of these mappings¹⁰.

The EDOC profile will not be entirely applicable to all domains. For example, real time applications are likely to require a different profile for platform-independent modeling with their own mappings to platforms.

The Sun Microsystems Java Community Process is currently defining a UML Profile for EJB (JSR #26) to support the declarative specification of EJB-specific solutions. This profile can be considered a peer of the UML Profile for CORBA in that it supports platform-specific modeling. Several members of the OMG Architecture Board are members of the expert group defining the profile as are others active in the OMG.

10. The EDOC profile also addresses business process component modeling and is likely to be integrated sooner or later with another standard in progress, the UML Profile for Event Driven Architectures in EAI. When these standards are finalized we may end up with more than one level of platform-independent model so that a platform-independent business process model could be mapped either to an event-based EAI model or to a more standard component model, where the EAI and component models are still platform-independent.

The MOF-IDL mapping defines an algorithm for transforming any arbitrary MOF-compliant metamodel to a set of IDL interfaces. The generated IDL defines the interfaces for CORBA objects that can represent models in a distributed repository.

The IDL generated from the CWM defines the interfaces for CORBA objects representing specific data models in a repository. The IDL generated from the UML defines the interfaces for CORBA objects representing specific UML models in a repository.

Similarly, the IDL generated from the IR metamodel defines the interfaces for CORBA objects representing specific CORBA object models in a repository.

It is generally agreed that the MOF-IDL mapping is in need of upgrading¹¹. Realistically we will probably have to accept the fact that for the foreseeable future, the automatically generated transformation from PIM to PSM will have to be enhanced by humans. As we gain more experience we will be able to define various standard patterns and allow them to be selected in some way.

There are additional issues regarding evolution of interfaces in a backward compatible interoperable fashion. An interface that is evolved in a UML or MOF model without consideration for backwards compatibility will most likely not result in a newer version of the interface that is backward compatible with the older version when deployed in a given platform. There may need to be additional enhancements to modeling standards that allow specification of platform specific restrictions to the ways in which interfaces can be evolved, so they continue to be usable by older clients.

3.6 Conformance Testing

To support this effort, the OMG must also concentrate extra effort on conformance testing and certification of products (branding). While OMG has been involved in the past with various testing & branding efforts for its standards, the expanded role of the OMG must be built on rock-solid testing, certification and branding. In many cases these efforts will depend on strong relationships with outside organizations with relevant expertise. Focusing on this problem is critical to the success of OMG's expanded role.

4 Conclusion

MDA is OMG's next step in solving integration problems through open, vendor-neutral interoperability specifications. MDA is an evolution of the OMA that addresses integration and interoperability spanning the life cycle of a system from business modeling and design, to component construction, assembly, integration, deployment,

11. The problem is that the generated interfaces are not efficient in distributed systems.

Firstly, the mapping predates CORBA valuetypes and thus does not make use of them. Secondly, a class with N attributes is always mapped to a CORBA interface with N separate getter/setter operations. In a distributed system one would want to group attributes based upon use cases, cache attribute values, or implement other optimizations to reduce the number of distributed calls.

management and evolution. It is built upon the experience gained in creating standards for implementation language independent models in CORBA and the development of the IDL, UML, MOF, CWM and XMI standards.

The MDA defines an architecture for structuring models that effectively separates concerns relevant for integration, interoperability, and portability. It exploits the ability to “zoom” in and out of any model of a system, exposing or eliding details of objects and interactions, with an architecture that uniformly separates specifications from their realizations. It uses these to enable the standardization of platform-independent models which can then be realized in implementation on multiple platforms with traceability between the platform independent models and the platform specific realizations of the models in implemented form.

5 Glossary

This glossary is included for the convenience of the reader. It explains how these terms are used in this document.

Abstraction: An abstraction is a description of something that omits some details that are not relevant to the purpose of the abstraction; the converse of “refinement”. The notion of abstraction is used in MDA in the sense defined in the Reference Model of Open Distributed Processing (RM-ODP) Part 2: the suppression of irrelevant detail.

Execution Environment: An execution environment depends on hardware and software infrastructure and is realized by one or several platforms.

Infrastructure: An infrastructure (also called computing infrastructure) is a set of software pieces or hardware pieces assumed to be already present by some stakeholders when he/she develops a software artifact.

Mapping: Set of rules and techniques used to modify one model in order to get another model. Mapping are used for PIM to PIM, PIM to PSM, PSM to PSM and PSM to PIM transformations.

Model: A model is a representation of part of the function, structure and/or behavior of a system.

Platform Independent Models (PIM): The PIM provide formal specifications of the structure and function of the system that abstracts away technical details.

Platform Specific Models (PSM): A PSM is expressed in terms of the specification model of the target platform. PSM have to use the platform concepts of exception mechanisms, parameter types (including platform-specific rules about objects references, value types, semantics of call by value, etc.), and component model.

Platform: A platform is a software infrastructure implemented with a specific technology (Unix platform, CORBA platform, Windows platform) on specified hardware technology.

Refinement: A refinement is a more detailed description that conforms to another (its abstraction). Everything said about the abstraction still holds, perhaps in a somewhat different form, in the refinement. In MDA, we would like to permit that more detailed descriptions are built in a systematic way from abstract ones.

Software Infrastructure: A software infrastructure is a computing infrastructure comprising software artifact with tools and guidelines to use them.

View: A representation of a whole system from the perspective of a related set of concerns [IEEE STD 1471-2000]. Views are not necessarily orthogonal, but each view generally contains specific information. In MDA, a view is a collection of models that represent one aspect of an entire system. A view applies to only one system, not to generalizations across many systems. The IEEE concept of "view" is equivalent to the RM-ODP concept of "viewpoint specification", which is the specification of a system from a given viewpoint.

Viewpoint: A specification of the conventions for constructing and using a view. A pattern or template from which to develop individual views by establishing the purposes and audience for a view and the techniques for its creation and analysis [IEEE STD 1471-2000]. The IEEE concept of "viewpoint" is equivalent to the RM-ODP concept, which is "a form of abstraction achieved using a selected set of architectural concepts and structuring rules, in order to focus on particular concerns within a system".