

Distributed Management Task Force, Inc.



**COMMON INFORMATION MODEL (CIM)
INFRASTRUCTURE SPECIFICATION**

**DSP0004
Version 2.3 Final
October 4, 2005**

DSP0004 Status: Final Standard

Copyright © 1997-2005 Distributed Management Task Force, Inc. (DMTF). All rights reserved.

DMTF is a not-for-profit association of industry members dedicated to promoting enterprise and systems management and interoperability. Members and non-members may reproduce DMTF specifications and documents for uses consistent with this purpose, provided that correct attribution is given. As DMTF specifications may be revised from time to time, the particular version and release date should always be noted.

Implementation of certain elements of this standard or proposed standard may be subject to third party patent rights, including provisional patent rights (herein "patent rights"). DMTF makes no representations to users of the standard as to the existence of such rights, and is not responsible to recognize, disclose, or identify any or all such third party patent right, owners or claimants, nor for any incomplete or inaccurate identification or disclosure of such rights, owners or claimants. DMTF shall have no liability to any party, in any manner or circumstance, under any legal theory whatsoever, for failure to recognize, disclose, or identify any such third party patent rights, or for such party's reliance on the standard or incorporation thereof in its product, protocols or testing procedures. DMTF shall have no liability to any party implementing such standard, whether such implementation is foreseeable or not, nor to any patent owner or claimant, and shall have no liability or responsibility for costs or losses incurred if a standard is withdrawn or modified after publication, and shall be indemnified and held harmless by any party implementing the standard from any and all claims of infringement by a patent owner for such implementations.

For information about patents held by third-parties which have notified the DMTF that, in their opinion, such patent may relate to or impact implementations of DMTF standards, visit <http://www.dmtf.org/about/policies/disclosures.php>.

=====

Terminology

The key phrases and words MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY and OPTIONAL in this document are to be interpreted as described in the IETF's RFC 2119. The complete reference for this document is: "Key words for Use in RFCs to Indicate Requirement Levels", IETF RFC 2119, March 1997 (<http://www.ietf.org/rfc/rfc2119.txt>).

=====

Abstract

The DMTF Common Information Model (CIM) Infrastructure is an approach to the management of systems and networks that applies the basic structuring and conceptualization techniques of the object-oriented paradigm. The approach uses a uniform modeling formalism that together with the basic repertoire of object-oriented constructs supports the cooperative development of an object-oriented schema across multiple organizations.

A management schema is provided to establish a common conceptual framework at the level of a fundamental typology both with respect to classification and association, and with respect to a basic set of classes intended to establish a common framework for a description of the managed environment. The management schema is divided into these conceptual layers:

- Core model—an information model that captures notions that are applicable to all areas of management.
- Common model—an information model that captures notions that are common to particular management areas, but independent of a particular technology or implementation. The common areas are systems, applications, databases, networks and devices. The information model is specific enough to provide a basis for the development of management applications. This model provides a set of base classes for extension into the area of technology-specific schemas. The Core and Common models together are expressed as the CIM schema.
- Extension schemas—represent technology-specific extensions of the Common model. These schemas are specific to environments, such as operating systems (for example, UNIX[†] or Microsoft Windows[†]).

[†] Other product and corporate names may be trademarks of other companies and are used only for explanation and to the owners' benefit, without intent to infringe.

Contents

1	Introduction and Overview	1
1.1	CIM Management Schema	1
1.1.1	Core Model	1
1.1.2	Common Model	2
1.1.3	Extension Schema	2
1.2	CIM Implementations	2
	Figure 1-1 Four Ways to Use CIM	3
1.2.1	CIM Implementation Conformance	4
2	Meta Schema	5
2.1	Definition of the Meta Schema	5
	Figure 2-1 Meta Schema Structure	6
	Figure 2-2 Reference Naming	8
	Figure 2-3 References, Ranges, and Domains	9
	Figure 2-4 References, Ranges, Domains and Inheritance	9
2.2	Property Data Types	10
2.2.1	Datetime Type	10
2.2.2	Indicating Additional Type Semantics with Qualifiers	11
2.3	Supported Schema Modifications	11
2.3.1	Schema Versions	12
2.4	Class Names	13
2.5	Qualifiers	14
2.5.1	Meta Qualifiers	14
2.5.2	Standard Qualifiers	14
2.5.3	Optional Qualifiers	23
2.5.4	User-defined Qualifiers	26
2.5.5	Mapping MIF Attributes	26
2.5.6	Mapping Generic Data to CIM Properties	27
3	Managed Object Format	29
3.1	MOF usage	29
3.2	Class Declarations	29
3.3	Instance Declarations	29
4	MOF Components	30
4.1	Keywords	30
4.2	Comments	30
4.3	Validation Context	30
4.4	Naming of Schema Elements	30
4.5	Class Declarations	30
4.5.1	Declaring a Class	31
4.5.2	Subclasses	31
4.5.3	Default Property Values	31
4.5.4	Class and Property Qualifiers	31
4.5.5	Key Properties	34
4.6	Association Declarations	35
4.6.1	Declaring an Association	35
4.6.2	Subassociations	36
4.6.3	Key References and Properties	36
4.6.4	Object References	36
4.7	Qualifier Declarations	37
4.8	Instance Declarations	37
4.8.1	Instance Aliasing	38
4.8.2	Arrays	38
4.9	Method Declarations	39
4.10	Compiler Directives	40

4.11 Value Constants.....	41
4.11.1 String Constants.....	41
4.11.2 Character Constants.....	41
4.11.3 Integral Constants.....	41
4.11.4 Floating-Point Constants.....	42
4.11.5 Object Ref Constants.....	42
4.11.6 NULL.....	42
4.12 Initializers.....	42
4.12.1 Initializing Arrays.....	42
4.12.2 Initializing References Using Aliases.....	43
5 Naming.....	44
5.1 Background.....	44
Figure 5-1 Definitions of instances and classes.....	45
Figure 5-2 Exporting to MOF.....	47
Figure 5-3 Information Exchange.....	47
5.1.1 Management Tool Responsibility for an Export Operation.....	47
5.1.2 Management Tool Responsibility for an Import Operation.....	47
5.2 Weak Associations: Supporting Key Propagation.....	48
Figure 5-4 Example of Weak Association.....	48
5.2.1 Referencing Weak Objects.....	49
5.3 Naming CIM Objects.....	49
Figure 5-5 Object Naming.....	50
5.3.1 Namespace Path.....	50
Figure 5-6 Namespaces.....	51
5.3.2 Model Path.....	51
5.3.3 Specifying the Object Name.....	51
6 Mapping Existing Models Into CIM.....	53
6.1 Technique Mapping.....	53
Figure 6-1 Technique Mapping Example.....	53
Figure 6-2 MIF Technique Mapping Example.....	54
6.2 Recast Mapping.....	54
Figure 6-3 Recast mapping.....	54
6.3 Domain Mapping.....	56
6.4 Mapping Scratch Pads.....	56
7 Repository Perspective.....	57
Figure 7-1 Repository Partitions.....	57
7.1 DMTF MIF Mapping Strategies.....	58
7.2 Recording Mapping Decisions.....	58
Figure 7-2 Homogeneous and Heterogeneous Export.....	59
Figure 7-3 Scratch Pads and Mapping.....	59
Appendix A MOF Syntax Grammar Description.....	61
Appendix B CIM META SCHEMA.....	67
Appendix C Values for UNITS Qualifier.....	74
Appendix D UML Notation.....	76
Appendix E Glossary.....	79
Appendix F Unicode Usage.....	82
F.1 MOF Text.....	82
F.2 Quoted Strings.....	82
Appendix G Guidelines.....	83
G.1 Mapping of Octet Strings.....	83
G.2 SQL Reserved Words.....	83
Appendix H Embedded Object Qualifier.....	86
H.1 Encoding for MOF.....	86
H.2 Encoding for CIM-XML.....	87
Appendix I Schema Errata.....	88
Appendix J References.....	90

Appendix K Change History.....91
Appendix L Ambiguous Property and Method Names93

Table of Figures

Figure 1-1 Four Ways to Use CIM	3
Figure 2-1 Meta Schema Structure	6
Figure 2-2 Reference Naming	8
Figure 2-3 References, Ranges, and Domains	9
Figure 2-4 References, Ranges, Domains and Inheritance	9
Figure 5-1 Definitions of instances and classes	45
Figure 5-2 Exporting to MOF.....	47
Figure 5-3 Information Exchange.....	47
Figure 5-4 Example of Weak Association.....	48
Figure 5-5 Object Naming	50
Figure 5-6 Namespaces	51
Figure 5-7 Technique Mapping Example	53
Figure 5-8 MIF Technique Mapping Example	54
Figure 5-9 Recast mapping.....	54
Figure 6-1 Repository Partitions.....	57
Figure 6-2 Homogeneous and Heterogeneous Export	59
Figure 6-3 Scratch Pads and Mapping.....	59

1 Introduction and Overview

This section describes the many ways in which the Common Information Model (CIM) can be used. It provides a context in which the details described in the later chapters can be understood.

Ideally, information used to perform tasks is organized or structured to allow disparate groups of people to use it. This can be accomplished by developing a model or representation of the details required by people working within a particular domain. Such an approach can be referred to as an information model. An information model requires a set of legal statement types or syntax to capture the representation, and a collection of actual expressions necessary to manage common aspects of the domain (in this case, complex computer systems). Because of the focus on common aspects, the DMTF refers to this information model as CIM, the Common Information Model.

This document describes an object-oriented meta model based on the Unified Modeling Language (UML). This model includes expressions for common elements that must be clearly presented to management applications (for example, object classes, properties, methods and associations). This document does not describe specific CIM implementations, APIs, or communication protocols – those topics will be addressed in a future version of the specification. For information on the current core and common schemas developed using this meta model, contact the Distributed Management Task Force.

Throughout this document, elements of formal syntax are described in the notation defined in [7], with these deviations:

Each token may be separated by an arbitrary number of white space characters, except where stated otherwise (at least one tab, carriage return, line feed, form feed or space).

The vertical bar (“|”) character is used to express alternation, rather than the virgule (“/”) specified in [7].

1.1 CIM Management Schema

Management schemas are the building blocks for management platforms and management applications, such as device configuration, performance management, and change management. CIM is structured in such a way that the managed environment can be seen as a collection of interrelated systems, each of which is composed of a number of discrete elements.

CIM supplies a set of classes with properties and associations that provide a well-understood conceptual framework within which it is possible to organize the available information about the managed environment. It is assumed that CIM will be clearly understood by any programmer required to write code that will operate against the object schema, or by any schema designer intending to make new information available within the managed environment.

CIM itself is structured into these distinct layers:

- Core model—an information model that captures notions that are applicable to all areas of management.
- Common model—an information model that captures notions that are common to particular management areas, but independent of a particular technology or implementation. The common areas are systems, applications, networks and devices. The information model is specific enough to provide a basis for the development of management applications. This schema provides a set of base classes for extension into the area of technology-specific schemas. The Core and Common models together are referred to in this document as the CIM schema.
- Extension schemas—represent technology-specific extensions of the Common model. These schemas are specific to environments, such as operating systems (for example, UNIX or Microsoft Windows).

1.1.1 Core Model

The Core model is a small set of classes, associations and properties that provide a basic vocabulary for analyzing and describing managed systems. The Core model represents a starting point for the analyst in determining how to extend the common schema. While it is possible that additional classes will be added to the Core model over time, major re-interpretations of the Core model classes are not anticipated.

47 **1.1.2 Common Model**

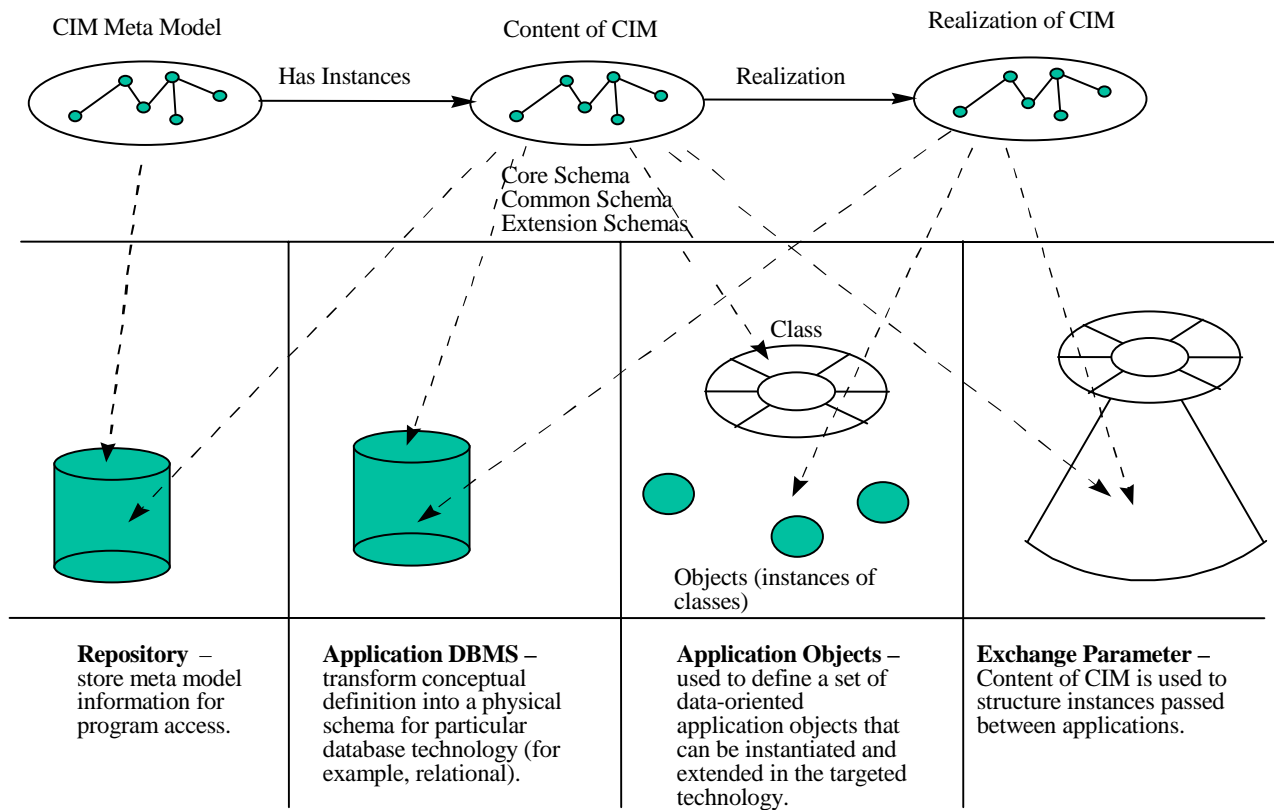
48 The Common model is a basic set of classes that define various technology-independent areas. The areas are
49 systems, applications, networks and devices. The classes, properties, associations and methods in the Common
50 model are intended to provide a view of the area that is detailed enough to use as a basis for program design and, in
51 some cases, implementation. Extensions are added below the Common model in platform-specific additions that
52 supply concrete classes and implementations of the Common model classes. As the Common model is extended, it
53 will offer a broader range of information.

54 **1.1.3 Extension Schema**

55 The Extension schemas are technology-specific extensions to the Common model. It is expected that the Common
56 model will evolve as a result of the promotion of objects and properties defined in the Extension schemas.

57 **1.2 CIM Implementations**

58 CIM is a conceptual model that is not bound to a particular implementation. This allows it to be used to exchange
59 management information in a variety of ways; four of these ways are illustrated in Figure 1-1. It is possible to use
60 these ways in combination within a management application.



61
62

63

Figure 1-1 Four Ways to Use CIM

64 As a repository (see the Repository Perspective section for more detail), the constructs defined in the model are
 65 stored in a database. These constructs are not instances of the object, relationship, and so on; but rather are
 66 definitions for someone to use in establishing objects and relationships. The meta model used by CIM is stored in a
 67 repository that becomes a representation of the meta model. This is accomplished by mapping the meta-model
 68 constructs into the physical schema of the targeted repository, then populating the repository with the classes and
 69 properties expressed in the Core model, Common model and Extension schemas.

70 For an application DBMS, the CIM is mapped into the physical schema of a targeted DBMS (for example,
 71 relational). The information stored in the database consists of actual instances of the constructs. Applications can
 72 exchange information when they have access to a common DBMS and the mapping occurs in a predictable way.

73 For application objects, the CIM is used to create a set of application objects in a particular language. Applications
 74 can exchange information when they can bind to the application objects.

75 For exchange parameters, the CIM—expressed in some agreed-to syntax—is a neutral form used to exchange
 76 management information by way of a standard set of object APIs. The exchange can be accomplished via a direct set
 77 of API calls, or it can be accomplished by exchange-oriented APIs which can create the appropriate object in the
 78 local implementation technology.

79 **1.2.1 CIM Implementation Conformance**

80 The ability to exchange information between management applications is fundamental to CIM. The current
81 mechanism for exchanging management information is the Management Object Format (MOF). At the present
82 time,¹ no programming interfaces or protocols are defined by (and hence cannot be considered as) an exchange
83 mechanism. Therefore, a CIM-capable system must be able to import and export properly formed MOF constructs.
84 How the import and export operations are performed is an implementation detail for the CIM-capable system.

85 Objects instantiated in the MOF must, at a minimum, include all key properties and all properties marked as
86 required. Required properties have the REQUIRED qualifier present and set to TRUE.

¹ The standard CIM application programming interface and/or communication protocol will be defined in a future version of the CIM Infrastructure specification.

87 2 Meta Schema

88 The Meta Schema is a formal definition of the model. It defines the terms used to express the model and its usage
89 and semantics (see Appendix B).

90 The Unified Modeling Language (UML) is used to define the structure of the meta schema. In the discussion that
91 follows, italicized words refer to objects in the figure. The reader is expected to be familiar with UML notation (see
92 <http://www.rational.com/uml>) and with basic object-oriented concepts in the form of classes, properties, methods,
93 operations, inheritance, associations, objects, cardinality and polymorphism.

94 2.1 Definition of the Meta Schema

95 The elements of the model are Schemas, Classes, Properties and Methods. The model also supports Indications and
96 Associations as types of Classes and References as types of Properties.

97 A *Schema* is a group of classes with a single owner. Schemas are used for administration and class naming. Class
98 names must be unique within their owning schemas.

99 A *Class* is a collection of instances that support the same type: that is, the same properties and methods.

100 Classes can be arranged in a generalization hierarchy that represents subtype relationships between Classes. The
101 generalization hierarchy is a rooted, directed graph and does not support multiple inheritance.

102 Classes can have Methods, which represent the behavior relevant for that Class. A Class may participate in
103 Associations by being the target of one of the References owned by the Association. Classes also have instances (not
104 represented in Figure 2-1).

105 Each instance provides values for the Properties associated with the instance's defining Class. An Instance does not
106 carry values for any other Properties or Methods that aren't defined in (or inherited by) its defining Class. An
107 Instance may not redefine the Properties or Methods defined in (or inherited by) its defining Class.

108 Instances are not Named Elements, and cannot have Qualifiers associated with them. (Qualifiers MAY be
109 associated with the Instance's Class, however, as well as the Properties and Methods defined in, or inherited by, that
110 Class.) Instances are also not permitted to attach new Qualifiers to Properties, Methods, or Parameters, as the
111 association between Qualifier and Named Element is not restricted to the context of a particular Instance.

112 A *Property* assigns values used to characterize instances of a Class. A Property can be thought of as a pair of Get
113 and Set functions that, when applied to an object,² return state and set state, respectively.

114 A *Method* is a declaration of a signature (that is, the method name, return type and parameters), and, in the case of a
115 concrete Class, may imply an implementation.

116 A *Trigger* is recognition of a state change (such as create, delete, update, or access) of a Class instance, and update
117 or access of a Property.

118 An *Indication* is an object created as a result of a Trigger. Because Indications are subtypes of Class, they can have
119 Properties and Methods, and be arranged in a type hierarchy.

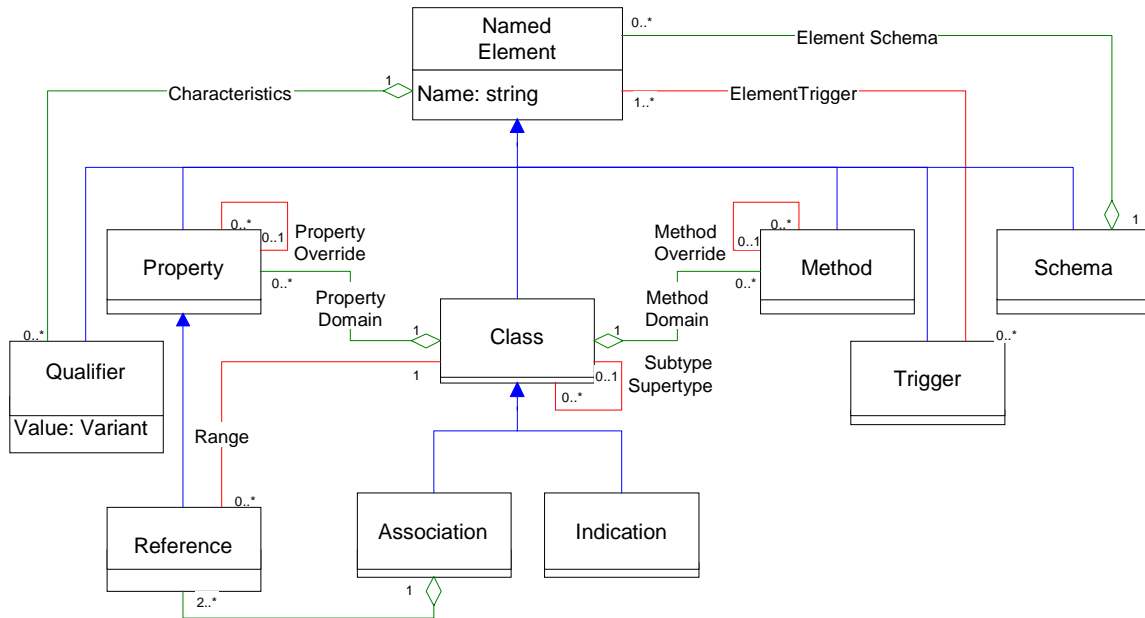
120 An *Association* is a class that contains two or more References. It represents a relationship between two or more
121 objects. Because of the way Associations are defined, it is possible to establish a relationship between Classes
122 without affecting any of the related Classes. That is, addition of an Association does not affect the interface of the
123 related Classes. Associations have no other significance. Only Associations can have References. An Association
124 cannot be a subclass of a non-association Class. Any subclass of an Association is an Association.

125 *References* define the role each object plays in an Association. The Reference represents the role name of a Class in
126 the context of an Association. Associations support the provision of multiple relationship instances for a given
127 object. For example, a system can be related to many system components.

² Note the equivocation between "object" as instance and "object" as class. This is common usage in object-oriented literature and reflects the fact that in many cases, operations and concepts may apply to or involve both classes and instances.

128 Properties and Methods have reflexive associations that represent Property and Method overriding. A Method can
 129 override an inherited Method, which implies that any access to the inherited Method will result in the invocation of
 130 the implementation of the overriding Method. A similar interpretation implies the overriding of Properties.

131 *Qualifiers* are used to characterize Named Elements (for example, there are Qualifiers that define the characteristics
 132 of a Property or the key of a Class). Qualifiers provide a mechanism that makes the meta schema extensible in a
 133 limited and controlled fashion. It is possible to add new types of Qualifiers by the introduction of a new Qualifier
 134 name, thereby providing new types of meta data to processes that manage and manipulate classes, properties and
 135 other elements of the meta schema. See below for details on the qualifiers provided.



136

137

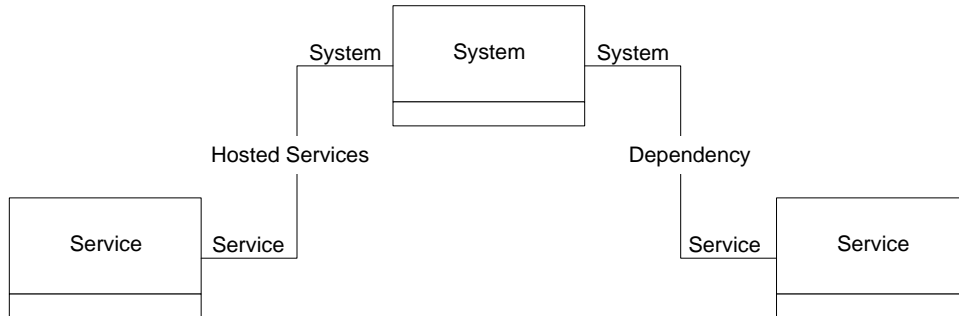
Figure 2-1 Meta Schema Structure

138 Figure 2-1 provides an overview of the structure of the meta schema. The complete meta schema is defined by the
 139 MOF found in Appendix B. The rules defining the meta schema are:

- 140 1. Every meta construct is expressed as a descendent of a Named Element.
- 141 2. A Named Element has zero or more Characteristics. A Characteristic is a Qualifier that
 142 characterizes a Named Element.
- 143 3. A Named Element can trigger zero or more Indications.
- 144 4. A Schema is a Named Element and can contain zero or more classes. A Class must belong to only
 145 one schema.
- 146 5. A Qualifier Type (not shown in Figure 2-1) is a Named Element and must be used to supply a type
 147 for a Qualifier (that is, a Qualifier must have a Qualifier Type). A Qualifier Type can be used to
 148 type zero or more Qualifiers.
- 149 6. A Qualifier is a Named Element and has a Name, a Type (intrinsic data type), a Value of this type,
 150 a Scope, a Flavor and a default Value. The type of the Qualifier Value must agree with the type of
 151 the Qualifier Type.
- 152 7. A Property is a Named Element and has only exactly one Domain: the Class that owns the
 153 Property. The Property is applicable to Instances of the Domain (including Instances of subclasses
 154 of the Domain), and not to any other Instances.

- 155 8. A Property can have an Override relationship with another Property from a different class. The
 156 Domain of the overridden Property must be a supertype of the Domain of the overriding Property.
 157 For non-Reference Properties, the type associated with the overriding Property MUST agree with
 158 (be the same as) the type of the overridden Property.
- 159 9. The Class referenced by the Range association (Figure 2-4) of an overriding Reference must be
 160 the same as, or a subtype of, the Class referenced by the Range associations of the Reference
 161 being overridden.
- 162 10. The Domain of a Reference must be an Association.
- 163 11. A Class is a type of Named Element. A Class can have instances (not shown on the diagram) and
 164 is the Domain for zero or more Properties. A Class is the Domain for zero or more Methods.
- 165 12. A Class can have zero or one supertype, and zero or more subtypes.
- 166 13. An Association is a type of Class. Associations are classes with an Association qualifier.
- 167 14. An Association must have two or more References.
- 168 15. An Association cannot inherit from a non-association Class.
- 169 16. Any subclass of an Association is an association.
- 170 17. A Method is a Named Element and has exactly one Domain: the Class that owns the Method. The
 171 Method is applicable to Instances of the Domain (including Instances of subclasses of the
 172 Domain), and not to any other Instances.
- 173 18. A Method can have an Override relationship with another Method from a different Class. The
 174 Domain of the overridden Method must be a superclass of the Domain of the overriding Method.
- 175 19. A Trigger is an operation that is invoked on any state change, such as object creation, deletion,
 176 modification or access, or on property modification or access. Qualifiers, Qualifier Types and
 177 Schemas may not have triggers. The changes that invoke a trigger are specified as a Qualifier.
- 178 20. An Indication is a type of Class and has an association with zero or more Named Triggers that can
 179 create instances of the Indication.
- 180 21. Every meta-schema object is a descendent of a Named Element and, as such, has a Name. All
 181 names are case-insensitive. The rules applicable to Name vary, depending on the creation type of
 182 the object.
- 183 A Fully-qualified Class Names (that is, the Class name prefixed by the schema name) are
 184 unique within the schema. (See the discussion of schemas later in this section).
- 185 B Fully-qualified Association and Indication Names are unique within the schema (implied
 186 by the fact that Associations and Indications are subtypes of Class).
- 187 C Implicitly-defined Qualifier Names are unique within the scope of the characterized
 188 object (that is, a Named Element may not have two Characteristics with the same Name).
 189 Explicitly defined Qualifier Names are unique within the defining Namespace. An
 190 implicitly-defined Qualifier must agree in type, scope and flavor with any explicitly-
 191 defined Qualifier of the same name.
- 192 D Trigger names must be unique within the Property, Class or Method to which the Trigger
 193 applies.
- 194 E Method and Property names must be unique within the Domain Class. A Class can inherit
 195 more than one Property or Method with the same name. Property and Method names can
 196 be qualified using the name of the declaring Class.

197 F Reference Names must be unique within the scope of their defining Association.
 198 Reference Names obey the same rules as Property Names. Note that Reference names are
 199 not required to be unique within the scope of the related Class. In such a scope, the
 200 Reference provides the name of the Class within the context defined by the Association.



201

202


Figure 2-2 Reference Naming

203 It is legal for the class System to be related to Service by two independent Associations (*Dependency* and *Hosted*
 204 *Services*, each with roles *System* and *Service*). It would not be legal for *Hosted Services* to define another Reference
 205 *Service* to the Service class, since a single association would then contain two references called *Service*.

206 22. Qualifiers are Characteristics of Named Elements. A Qualifier has a Name (inherited from Named
 207 Element) and a Value. The Value is used to define the characteristics of the Named Element. For
 208 example, a Class might have a Qualifier with the Name “Description,” the Value of which is the
 209 description for the Class. A Property might have a Qualifier with the Name “Units,” which has
 210 Values such as “Bytes” or “KiloBytes.” The Value can be thought of as a variant (that is, a value
 211 plus a type).

212 23. Association and Indication are types of Class; as such, they can be the Domain for Methods,
 213 Properties and References (that is, Associations and Indications can have Properties and Methods
 214 in the same way as a Class does). Associations and Indications can have instances. The instance
 215 of an Association has a set of references that relate one or more objects. An instance of an
 216 Indication represents the occurrence of an event, and is created because of that occurrence—
 217 usually a Trigger. Indications are not required to have keys. Typically, Indications are very short-
 218 lived objects used to communicate information to an event consumer.

219 24. A Reference has a range that represents the type of the Reference. For example, in the model of
 220 PhysicalElements and PhysicalPackages, there are two References: ContainedElement, which has
 221 PhysicalElement as its range and Container as its domain, and ContainingElement, which has
 222 PhysicalPackage as its range and Container as its domain.



223

224

Figure 2-3 References, Ranges, and Domains

225

25. A Class has a subtype-supertype association that represents substitutability relationships between the Named Elements involved in the relationship. The association implies that any instance of a subtype can be substituted for any instance of the supertype in an expression, without invalidating the expression.

226

227


228

229

Revisiting the Container example: Card is a Subtype of PhysicalPackage. Therefore, Card can be used as a value for the Reference ContainingElement (that is, an instance of Card can be used as a substitute for an instance of PhysicalPackage).

230

231



232

233

Figure 2-4 References, Ranges, Domains and Inheritance

234

A similar relationship can exist between Properties. For example, given that PhysicalPackage has a Name property (which is a simple alphanumeric string), Card Overrides Name to a name of alpha-only characters.

235

236

The same idea applies to Methods. A Method that overrides another Method must support the same signature as the original Method and, most importantly, must be substitutable for the original Method in all cases.

237

238

26. The Override relationship is used to indicate the substitution relationship between a property or method of a subclass and the overridden property or method inherited from the superclass. This is the opposite of the C++ convention in which the superclass property or method is specified as virtual, with overriding occurring thereafter as a side effect of declaring a feature with the same signature as the inherited virtual feature.

239

240

241

242

27. The number of references in an Association class defines the arity of the Association. An Association containing two references is a binary Association. An Association containing three references is a ternary Association. Unary Associations (Associations containing one reference) are not meaningful. Arrays of references are not allowed. When an association is sub-classed, its arity cannot change.

243

244

245

246

28. Schemas provide a mechanism that allows ownership of portions of the overall model by individuals and organizations who are responsible for managing the evolution of the schema. In any given installation, all classes are mutually visible, regardless of schema ownership. Schemas have a universally unique name. The schema name is considered part of the class name. The full class name (that is, class name plus owning schema name) is unique within the namespace and is referred to as the fully-qualified name (see Section 2.4).

247

248

249

250

2.2 Property Data Types

Property data types are limited to the intrinsic data types, or arrays of such. Structured types are constructed by designing new classes. If the Property is an array property, the corresponding variant type is simply the array equivalent (fixed or variable length) of the variant for the underlying intrinsic type. There are no subtype relationships among the intrinsic data types uint8, sint8, uint16, sint16, uint32, sint32, uint64, sint64, string, boolean, real32, real64, datetime, and char16.

This table contains the intrinsic data types and their interpretation:

INTRINSIC DATA TYPE	INTERPRETATION
uint8	Unsigned 8-bit integer
sint8	Signed 8-bit integer
uint16	Unsigned 16-bit integer
sint16	Signed 16-bit integer
uint32	Unsigned 32-bit integer
sint32	Signed 32-bit integer
uint64	Unsigned 64-bit integer
sint64	Signed 64-bit integer
string	UCS-2 string
boolean	Boolean
real32	IEEE 4-byte floating-point
real64	IEEE 8-byte floating-point
datetime	A string containing a date-time
<classname> ref	Strongly typed reference
char16	16-bit UCS-2 character

258

2.2.1 Datetime Type

The datetime type is used to specify a timestamp (point in time) or an interval. If it specifies a timestamp, it allows to preserve the timezone offset. In both cases, datetime allows to specify varying precision of the date and time information.

Datetime uses a fixed string-based format. The format for timestamps is: `yyymmddhhmmss.mmmmmmsutc`

where the meaning of each field is:

- `yyyy` is a 4 digit year
- `mm` is the month within the year (starting with 01)
- `dd` is the day within the month (starting with 01)
- `hh` is the hour within the day (24-hour clock, starting with 00)
- `mm` is the minute within the hour (starting with 00)
- `ss` is the second within the minute (starting with 00)
- `mmmmmm` is the microsecond within the second (starting with 000000)
- `s` is a "+" or "-", indicating that the value is a timestamp, and indicating the sign of the UTC (Universal Coordinated Time; for all intents and purposes the same as Greenwich Mean Time) correction field. A "+" is used for time zones east of Greenwich, and a "-" is used for time zones west of Greenwich.
- `utc` is the offset from UTC in minutes (using the sign indicated by `s`).

Since datetime contains the time zone information, it is possible to reconstruct the original time zone from the value. However, this also makes it possible to specify the same point in time using different UTC offsets by adjusting the hour and possibly the minutes fields accordingly.

279 For example, Monday, May 25, 1998, at 1:30:15 PM EST would be represented as: 19980525133015.0000000-300.

280 An alternative representation of the same point in time would be: 19980525183015.0000000+000.

281 The format for intervals is: ddddddhhmmss.mmmmm:000

282 where the meaning of each field is:

- 283 • dddddd is the number of days
- 284 • hh is the remaining number of hours
- 285 • mm is the remaining number of minutes
- 286 • ss is the remaining number of seconds
- 287 • mmmmm is the remaining number of microseconds
- 288 • ":" is indicating that the value is an interval
- 289 • 000 (the UTC offset field) is always zero for interval properties

290 For example, an interval of 1 day, 13 hours, 23 minutes, 12 seconds and 0 microseconds would be:
291 0000001132312.000000:000.

292 For both timestamps and intervals, the field values MUST be zero-padded so that the entire string is always the same
293 25-character length.

294 For both timestamps and intervals, fields that are not significant MUST be replaced with asterisk ("*") characters.
295 Not significant fields are those that are beyond the resolution of the data source. This is used to indicate the
296 precision of the value and can only be done for an adjacent set of fields, starting with the least significant field
297 (mmmmm), and continuing to more significant fields. The granularity of using asterisks is always the entire field,
298 except for the mmmmm field where the granularity is single digits. The UTC offset field MUST NOT contain
299 asterisks.

300 For example, if an interval of 1 day, 13 hours, 23 minutes, 12 seconds, and 125 milliseconds was measured with a
301 precision of 1 millisecond, the format would be: 0000001132312.125***:000.

302 2.2.2 Indicating Additional Type Semantics with Qualifiers

303 Since "counter" and "gauge" types (as well as many others) are actually simple integers with specific semantics, they
304 are not treated as separate intrinsic types. Instead, qualifiers must be used to indicate such semantics when properties
305 are being declared (note the example below merely suggests how this may be done; the qualifier names chosen are
306 not part of this standard):

```
307     class Acme_Example
308     {
309         [counter]
310         uint32 NumberOfCycles;
311         [gauge]
312         uint32 MaxTemperature;
313         [octetstring, ArrayType("Indexed")]
314         uint8 IPAddress[10];
315     };
```

316 Implementers are permitted, for documentation purposes, to introduce arbitrary qualifiers in this manner. The
317 semantics are not enforced.

318 2.3 Supported Schema Modifications

319 This is a list of supported schema modifications, some of which, when used, will result in changes in application
320 behavior. Changes are all subject to security restrictions; in particular, only the owner of the schema, or someone
321 authorized by the owner, can make modifications to the schema.

322 A class can be added to or deleted from a schema.

323 A property can be added to or deleted from a class.

324 A class can be added as a subtype or supertype of an existing class.

325 A class can become an association as a result of the addition of an Association qualifier, plus two or more
326 references.

- 327 A qualifier can be added to or deleted from any Named Element.
- 328 The Override qualifier can be added to or removed from a property or reference.
- 329 A method can be added to a class.
- 330 A method can override an inherited method.
- 331 Methods can be deleted, and the signature of a method can be changed.
- 332 A trigger may be added to or deleted from a class.
- 333 In defining an extension to a schema, the schema designer is expected to operate within the constraints of the classes
- 334 defined in the Core model. With respect to classification, it is recommended that any added component of a system
- 335 be defined as a subclass of an appropriate Core model class. It is expected that the schema designer will address the
- 336 following question to each of the Core model classes: “Is the class being added a subtype of this class?” Having
- 337 identified the Core model class to be extended, the same question should be addressed with respect to each of the
- 338 subclasses of the identified class. This process, which defines the superclasses of the class to be defined, should be
- 339 continued until the most detailed class is identified. The Core model is not a part of the meta schema, but is an
- 340 important device for introducing uniformity across schemas intended to represent aspects of the managed
- 341 environment.

342 **2.3.1 Schema Versions**

343 Schema Versioning is described in Section 4 of DSP0129 (the “DMTF Release Process”) and in Section 2.5.2 where

344 the Version qualifier is discussed. Versioning takes the form m.n.u, where:

- 345 m = major version identifier in numeric form
- 346 n = minor version identifier in numeric form
- 347 u = update (errata or coordination changes) in numeric form

348 The Usage Rules for the Version qualifier (Section 2.5.2) provide additional information.

349 Classes are versioned in the CIM Schemas. A class’ Version qualifier indicates the schema release when the last

350 change to the class occurred. Class Versions in turn dictate the “Schema Version”. A major version change for a

351 class would require that the major version number of the Schema Release be incremented. All class Versions must

352 be at the same level, or a previous level, than the Schema Release. This is because classes and models which differ

353 in minor version numbers MUST be backwards compatible. In other words, valid instances MUST continue to be

354 valid if the minor version number is incremented. Classes and models which differ in major version numbers will

355 not be backwards compatible. This then requires that the major version number of the Schema Release be

356 incremented.

357 The following table describes modifications to the CIM Schemas in FINAL status that would cause a major version

358 number change. (Preliminary Models are allowed to evolve based on implementation experience.) These

359 modifications result in changes in application behavior and/or customer code. Therefore, they force a major version

360 update and are discouraged. (Note that the table is “exhaustive” based on current CIM experience and knowledge.

361 Items may be added as new issues are raised and in response to the evolution of the CIM Standards.)

362 Other alterations (beyond those listed in the table) are considered to be interface-preserving and require that the

363 minor version number be incremented. Updates/errata are NOT classified as having “major” or “minor” impact –

364 but are required to correct errors, or to coordinate across standards bodies.

365 **Changes that Increment the CIM Schema Major Version Number**

366 **(This table may be updated in the future, as experience and extensions require.)**

367

<i>Description</i>	<i>Explanation or exceptions</i>
Class deletion	
Property deletion or data type change	
Method deletion or signature change	

Reorganization of values in an enumeration	The semantics and mappings of an enumeration cannot change, but values can be added in unused ranges as a minor change or update.
Movement of a class “upwards” in the inheritance hierarchy – I.E., the removal of superclasses from the inheritance hierarchy	The removal of superclasses results in the deletion of properties or methods. New classes CAN be inserted as superclasses as a minor change or update. Inserted classes MUST NOT change keys or add “Required” properties.
Addition of Abstract, Indication or Association qualifiers to an existing class	
Change of an association reference “downward” in the object hierarchy (i.e., to a subclass) or to a different part of the hierarchy	The change of an association reference to a subclass could invalidate existing instances.
Addition or removal of a Key or Weak qualifier	
Addition of a Required qualifier	
Decrease in MaxLen, decrease in MaxValue, increase in MinLen or increase in MinValue	Decreasing a maximum, or increasing a minimum invalidates current data. The opposite change (increasing a maximum) results in truncated data, where necessary.
Decrease in Max or increase in Min cardinalities	
Addition or removal of Override qualifier	There is one exception to this – An Override qualifier can be added if a property is promoted to a superclass, and it is necessary to maintain the specific qualifiers and descriptions in the original subclass. In this case, there is no change to existing instances.
Change in the following qualifiers: In/Out, Units	

368 **2.4 Class Names**

369 Fully-qualified class names are in the form <schema name>_<class name>. An underscore is used as a delimiter
 370 between the <schema name> and the <class name>. The delimiter is not allowed to appear in the <schema name>
 371 although it is permitted in the <class name>.

372 The format of the fully-qualified name is intended to allow the scope of class names to be limited to a schema: that
 373 is, the schema name is assumed to be unique, and the class name is only required to be unique within the schema.
 374 The isolation of the schema name using the underscore character allows user interfaces to conveniently strip off the
 375 schema when the schema is implied by the context.

376 Examples of fully-qualified class names:

- 377 • CIM_ManagedSystemElement: the root of the CIM managed system element hierarchy.
- 378 • CIM_ComputerSystem: the object representing computer systems in the CIM schema.
- 379 • CIM_SystemComponent: the association relating systems to their components.
- 380 • Win32_ComputerSystem: the object representing computer systems in the Win32 schema.

381 **2.5 Qualifiers**

382 Qualifiers are values that provide additional information about classes, associations, indications, methods, method
 383 parameters, triggers, properties or references. All qualifiers have a name, type, value, scope, flavor and default
 384 value. Qualifiers cannot be duplicated; there cannot be more than one qualifier of the same name for any given class
 385 or property.

386 The following sections describe meta, standard, optional and user-defined qualifiers. When any of these qualifiers
 387 are used in a model, they must be declared in the MOF file before being used. These declarations must abide by the
 388 details (name, applied to, type) specified in the tables below. It is not valid to change any of this information for the
 389 meta, standard and optional qualifiers. It is possible to change the default values. A default value is the assumed
 390 value for a qualifier when it is not explicitly specified for particular model elements.

391 **2.5.1 Meta Qualifiers**

392 This table lists the qualifiers that are used to refine the definition of the meta constructs in the model. These
 393 qualifiers are used to refine the actual usage of an object class or property declaration within the MOF syntax. These
 394 qualifiers are all mutually exclusive.

QUALIFIER	DEFAULT	TYPE	MEANING
ASSOCIATION	FALSE	BOOLEAN	The object class is defining an association.
INDICATION	FALSE	BOOLEAN	The object class is defining an indication.

395 **2.5.2 Standard Qualifiers**

396 This table is a list of standard qualifiers that all CIM-compliant implementations are required to handle. Any given
 397 object will not have all of the qualifiers listed. It is expected that additional qualifiers will be supplied by extension
 398 classes to facilitate the provision of instances of the class and other operations on the class.

399 It is also important to recognize that not all of these qualifiers can be used together. First, as indicated in the table,
 400 not all qualifiers can be applied to all meta-model constructs. These limitations are identified in the “Applies To”
 401 column. Second, for a particular meta-model construct like associations, the use of the legal qualifiers may be
 402 further constrained because some qualifiers are mutually exclusive or the use of one qualifier implies some
 403 restrictions on the value of another qualifier, and so on. These usage rules are documented in the “Meaning” column
 404 of the table. Third, legal qualifiers are not inherited by meta-model constructs. For example, the MAXLEN qualifier
 405 that applies to properties is not inherited by references.

406 The “Applies To” column in the table identifies the meta-model construct(s) that can use a particular qualifier. For
 407 qualifiers like ASSOCIATION (discussed in the previous section), there is an implied usage rule that the meta
 408 qualifier must also be present. For example, the implicit usage rule for the AGGREGATION qualifiers is that the
 409 ASSOCIATION qualifier must also be present.

410

QUALIFIER	DEFAULT	APPLIES TO	TYPE
ABSTRACT	FALSE	Class, Association, Indication	BOOLEAN
MEANING: Indicates that the class is abstract and serves only as a base for new classes. It is not possible to create instances of such classes.			
AGGREGATE	FALSE	Reference	BOOLEAN
MEANING: Defines the "parent" component of an Aggregation association. Usage Rule: The Aggregation and Aggregate qualifiers are used together – Aggregation qualifying the association, and Aggregate specifying the "parent" reference			
AGGREGATION	FALSE	Association	BOOLEAN
MEANING: Indicates that the association is an aggregation.			

QUALIFIER	DEFAULT	APPLIES TO	TYPE
ARRAYTYPE	"Bag"	Property, Parameter	STRING
<p>MEANING: Indicates the type of the qualified array. Valid values are "Bag", "Indexed" and "Ordered".</p> <p>For a "Bag" array type, no significance is defined for the array index other than as a convenience for accessing the elements of the array. For example, there can be no assumption that the same index returns the same value for every access to the array.</p> <p>For an "Ordered" array type, the array index is significant as long as no array elements are added, deleted, or changed. If this is true, the same index returns the same value for every access to the array.</p> <p>For an "Indexed" array type, the array maintains the correspondence between element position and value.</p> <p>Usage rule: The ArrayType qualifier MUST only be applied to properties and method parameters that are arrays (defined using the square bracket syntax specified in Appendix A).</p> <p>Usage Note: For any of the ArrayTypes, elements in the array may be duplicated.</p>			
BITMAP	NULL	Property, Method, Parameter	STRING ARRAY
<p>MEANING: Indicates bit positions that are significant in a bit map. The bit map is evaluated from the right, starting with the least significant value. This value is referenced as "0". For example, using a uint8 data type, the bits take the form, Mxxx xxxL, where M and L designate the most and least significant bits, respectively. The least significant bits referenced as "0", and the most significant bit is "7". The position of a specific value in the BitMap array defines an index that is used in selecting a string literal from the BitValues array.</p> <p>Usage Rule: The number of entries in the BitValues and BitMap arrays MUST match.</p>			
BITVALUES	NULL	Property, Method, Parameter	STRING ARRAY
<p>MEANING: Provides translation between a bit position value and an associated string. See the description for the BitMap qualifier.</p> <p>Usage Rule: The number of entries in the BitValues and BitMap arrays MUST match.</p>			
COUNTER	FALSE	Property, Method, Parameter	BOOLEAN
<p>MEANING: Applicable only to unsigned integer types. Represents a non-negative integer that monotonically increases until it reaches a maximum value of $2^n - 1$, when it wraps around and starts increasing again from zero. N can be 8, 16, 32 or 64 depending on the datatype of the object that the qualifier is applied to. Counters have no defined "initial" value, and thus, a single value of a Counter has (in general) no information content</p>			
COMPOSITION	FALSE	Association	BOOLEAN
<p>MEANING: Refines the definition of an Aggregation association, adding the semantics of a whole-part/compositional relationship, to distinguish it from a collection or basic aggregation. This is necessary to better map CIM associations into UML where whole-part relationships are considered compositions. The semantics conveyed by Composition align with that of the UML Specification from the OMG. Quoting from Section 3.48 of the V4 UML Specification (September 2001):</p> <p>Composite aggregation is a strong form of aggregation, that requires that a part instance be included in AT MOST ONE composite at a time and that the composite object has sole responsibility for the disposition of its parts.</p> <p>Use of this qualifier imposes restrictions on the membership of the "collecting" object (the whole). Care should be taken when entities are added to the aggregation – since they MUST be "parts" of the whole. Also, if the collecting entity (the whole) is deleted, it is the responsibility of the implementation to dispose of the parts. The behavior may vary with the type of collecting entity whether the parts are also deleted. This is very different from that of a collection, since a collection may be removed without deleting the entities that are collected.</p> <p>Usage Rule: The Aggregation and Composition qualifiers are used together. Aggregation indicating the general nature of the association and Composition indicating more specific semantics of whole-part relationships. This duplication of information is necessary since Composition is a more recent addition to the list of qualifiers – and applications may be built only</p>			

QUALIFIER	DEFAULT	APPLIES TO	TYPE
understanding the earlier Aggregation qualifier.			
DEPRECATED	NULL	Any	STRING ARRAY
<p>MEANING: Indicates that the feature the qualifier is applied to (e.g., class, property, etc.) is tolerated but not recommended, and may be superseded by another Method. Existing instrumentation should continue to support the old feature, so as not to break current applications, and are encouraged to support the new approach as well. Replacement or substitution information is conveyed in the contents of the qualifier's string value. The qualifier serves only as an indication that the Deprecated item is not recommended for use in new development efforts. Existing and new implementations MAY still use this feature, but SHOULD move to the replacement or substitution feature as soon as possible. The Deprecated feature MAY be removed in a future major version release of the CIM Schema (ex, CIM 2.x to CIM 3.0). The qualifier acts inclusively in that if a class is deprecated, all of the properties, references and methods in that class are also deprecated. However, no subclasses or associations that reference that class are deprecated, unless explicitly declared as such. For clarity, all deprecated properties and associations SHOULD be specifically labeled as DEPRECATED.</p> <p>The syntax for the Deprecated qualifier's string value should indicate replacement values (if any exist) as follows: <className> [*(" " <propertyName>) ["." <methodName> ["(" <parameterName> ")"]]]</p> <p>The property is defined as a string ARRAY to allow a single feature to be replaced by multiple other constructs.</p> <p>Usage Rule: The string value for the qualifier indicates what the replacement or substitution feature is for the deprecated construct. If there is no replacement feature, then the string value MUST be set to "No value". When a feature is Deprecated, its description MUST indicate why it was deprecated and how the replacement features (if any are defined) are used. For example, "The property X is deprecated in lieu of the method Y defined in this class. The reason for this change is that the property is actually causing a change of state and requires an input parameter."</p> <p>Note 1: Publishing a feature by deprecating and replacing it results in duplicate representations of the feature. This is of particular concern when deprecated classes are replaced by new classes, and instances may be duplicated. To allow a management application to detect such duplication, implementations SHOULD document in a README, in MOF or by other documentation, how duplicate instances are detected.</p> <p>Note 2: Properties with the KEY qualifier CAN be deprecated but MUST NOT be left blank or NULL, MUST keep the KEY qualifier, and MUST NOT be replaced with a new KEY property. If any of the latter 3 changes occur, this breaks existing instrumentation and applications and hence would not be "tolerated" changes (a requirement for using the DEPRECATED qualifier, as documented in the first sentence of this description). In scenarios where a KEY property should be deprecated AND replaced, only one option is available- it is necessary to deprecate the entire class and therefore its properties, methods, references, etc.</p>			
DESCRIPTION	NULL	Any	STRING
MEANING: Provides a description of a Named Element.			
DISPLAYNAME	NULL	Any	STRING
MEANING: Defines a name that is displayed on UI instead of the actual name of the element.			
DN	FALSE	Property, Parameter, Method	BOOLEAN
MEANING: When applied to a string element the DN qualifier specifies that the string MUST be a distinguished name as defined in Section 9 of X.501 "Information Technology – Open Systems Interconnection – The Directory: Models" and the string representation defined in RFC2253 "Lightweight Directory Access Protocol (v3): UTF-8 String Representation Of Distinguished Names". This qualifier MUST NOT be applied to qualifiers that are not of intrinsic data type string.			
EMBEDDEDOBJECT	FALSE	Property, Parameter, Method	BOOLEAN
MEANING: This qualifier can only be used in conjunction with string-valued features. It indicates that the qualified entity contains an encoding of an instance's data. Please see Appendix H for examples and important details.			

QUALIFIER	DEFAULT	APPLIES TO	TYPE
EXCEPTION	FALSE	Class	BOOLEAN
<p>MEANING: Indicates that the class and all subclasses of this class describe transient Exception information. The definition of this qualifier is identical to the Abstract Qualifier except that this qualifier cannot be overridden. It is not possible to create instances of Exception classes.</p> <p>Usage Rule: The Exception Qualifier is used to denote a class hierarchy that defines transient (i.e., very short-lived) Exception objects. Instances of Exception classes are used to communicate exception information between CIMEntities. The Exception Qualifier cannot be used with the Abstract Qualifier. The subclass of an Exception class, if it exists, MUST be an Exception class.</p>			
EMBEDDEDINSTANCE	NULL	Property, Parameter, Method	STRING
<p>MEANING: The EMBEDDEDINSTANCE qualifier modifies a string property, method parameter or return value to indicate that it contains an embedded instance.</p> <p>The definition of this qualifier is identical to the EMBEDDEDOBJECT Qualifier except that it restricts the set of permissible values for the qualified property, parameter, or method return to instances of the class, or to one of its subclasses, that is designated by the string value.</p> <p>Usage Rule: When the EMBEDDEDINSTANCE qualifier is used, its value MUST be specified and not NULL.</p> <p>The encoding of the instance contained in the string qualified by EMBEDDEDINSTANCE follows the same rules as for EMBEDDEDOBJECT. This allows the same parsing engine that decodes the instance/indication envelope to decode the EMBEDDEDOBJECT string. (As with EMBEDDEDOBJECT, encodings are defined by an implementation and are not the subject of this document.) Please see Appendix H for examples and further explanation.</p>			
EXPERIMENTAL	FALSE	Any	BOOLEAN
<p>MEANING: Represents that the specified element is proposed for inclusion in a future release of the CIM Schemas, but is not currently a part of the standard Schema. The specified element is in an interim state - available for experimentation and implementation experience, but not a part of the standard. For example, if an implementation uses schema from a DMTF Preliminary Standard, the features from the Preliminary Standard should be labeled with the EXPERIMENTAL qualifier.</p> <p>Based on implementation experience, changes may occur to this element in future releases, it may be standardized "as is", or it may be removed.</p> <p>An implementation does not have to support an EXPERIMENTAL feature.</p>			
GAUGE	FALSE	Property, Method, Parameter	BOOLEAN
<p>MEANING: Applicable only to unsigned integer types.</p> <p>Represents an integer that may increase or decrease in any order of magnitude.</p> <p>The value of a Gauge is capped at the implied limits of the property's datatype. If the information being modeled exceeds an implied limit, the value represented is that limit. Values do not wrap. For unsigned integers, the limits are zero (0) to 2^n-1, inclusive. For signed integers, the limits are $-(2^{n-1})$ to $2^{n-1}-1$, inclusive. N can be 8, 16, 32, or 64 depending on the datatype of the property that the qualifier is applied to.</p>			
IN	TRUE	Parameter	BOOLEAN
<p>MEANING: Indicates that the associated parameter is used to pass values to a method.</p>			
KEY	FALSE	Property, Reference	BOOLEAN
<p>MEANING: Indicates that the property is part of the namespace handle (see Section 5.3.1.2 for information about namespace handles). If more than one property has the KEY qualifier, then all such properties collectively form the key (a compound</p>			

QUALIFIER	DEFAULT	APPLIES TO	TYPE
key).			
Usage Rule: Keys are written once at object instantiation and MUST NOT be modified thereafter. Default values are not applied to a KEY-qualified properties. Properties of an array type MUST NOT be qualified with KEY.			
MAPPINGSTRINGS	NULL	Any	STRING ARRAY
MEANING: Mapping strings for one or more management data providers or agents. See Section 2.5.5 and 2.5.6 for more details.			
MAX	NULL	Reference	uint32
MEANING: Indicates the maximum cardinality of the reference (i.e., the maximum number of values a given reference may have for each set of other reference values in the association). For example, if an association relates A instances to B instances, and there MUST be at most one A instance for each B instance, then the reference to A should have a Max(1) qualifier.			
MAXLEN	NULL	Property, Method, Parameter	uint32
MEANING: Indicates the maximum length, in characters, of a string data item. MAXLEN may also be applied to elements with a string array value, in which case it is interpreted as applying to each string in the string array. When overriding the default value, any unsigned integer value (uint32) can be specified. A value of NULL implies unlimited length.			
MAXVALUE	NULL	Property, Method, Parameter	sint64
MEANING: Maximum value of this element.			
MIN	0	Reference	uint32
MEANING: Indicates the minimum cardinality of the reference (i.e., the minimum number of values a given reference may have for each set of other reference values in the association). For example, if an association relates A instances to B instances, and there MUST be at least one A instance for each B instance, then the reference to A should have a Min(1) qualifier.			
MINLEN	0	Property, Method, Parameter	uint32
MEANING: Indicates the minimum length, in characters, of a string data item. When overriding the default value, any unsigned integer value (uint32) can be specified. The default is zero length for the minimum length..			
MINVALUE	NULL	Property, Method, Parameter	sint64
MEANING: Minimum value of this element.			
MODELRESPONDENCE	NULL	Any	STRING ARRAY
MEANING: Indicates a correspondence between two elements in the CIM Schema. The referenced elements MUST be defined in a standard or extension MOF file, such that the correspondence can be examined. If possible, forward referencing of elements SHOULD be avoided.			
Object elements are identified using the following syntax:			
<code><className> [*("(<propertyName> <referenceName>)) ["(" <methodName> ["(" <parameterName> "]"]]]]</code>			
Note that the basic relationship between the referenced elements is a "loose" correspondence – simply indicating that the elements are coupled. And, this coupling MAY be unidirectional. Additional qualifiers MAY be used to describe a tighter			

QUALIFIER	DEFAULT	APPLIES TO	TYPE
<p>coupling.</p> <p>The following list provides examples of several correspondences found in CIM and vendor schemas:</p> <ul style="list-style-type: none"> · A vendor defines an Indication class "corresponding" to a particular CIM property or method – i.e., Indications are generated based on the values or operation of the property or method (in this case, the ModelCorrespondence MAY only be on the vendor's Indication class which is an extension to CIM) · A property provides more information for another - for example, an enumeration has an allowed value of "Other", and another property further clarifies the intended meaning of "Other"; alternately, a property specifies status and another property provides human-readable strings (using an array construct) expanding on this status (in these cases, ModelCorrespondence is found on both properties, each referencing the other – note also that referenced array properties MAY NOT be ordered but carry the default ArrayType qualifier definition of "Bag") · A property is defined in a subclass to supplement the meaning of an inherited property (in this case, the ModelCorrespondence is only found on the construct in the subclass) · Multiple properties taken together are needed for complete semantics – for example, one property might define units, another a multiplier and a third property define a specific value (in this case, ModelCorrespondence is found on all related properties, each referencing all the others) · Multi-dimensional arrays are desired – for example, one array may define names while another the name formats (in this case, the arrays would each be defined with the ModelCorrespondence qualifier, referencing the other array properties or parameters, and would also be indexed - i.e., also carry the ArrayType qualifier with the value "Indexed") <p>The specific semantics of the correspondence are based on the elements themselves. ModelCorrespondence is only a hint or indicator of a relationship between the elements.</p>			
NONLOCAL			
MEANING: This instance-level qualifier, and the corresponding pragma, were removed as an erratum by CR1461.			
NONLOCALTYPE			
MEANING This instance-level qualifier, and the corresponding pragma, were removed as an erratum by CR1461.			
NULLVALUE	NULL	Property	STRING
<p>MEANING: Defines a value that indicates the associated property is NULL (i.e., the property is considered to have a valid or meaningful value.</p> <p>The NullValue qualifier may only be used with string and integer valued properties. When used with an integer type, the qualifier value is a MOF integerValue. The syntax for representing an integer value is:</p> <p>["+" / "-"] 1*<decimalDigit></p> <p>The content, maximum number of digits and represented value are constrained by the datatype of the qualified property.</p> <p>Note this qualifier cannot be overridden as it seems unreasonable to permit a subclass to return a different null value to that of the superclass.</p>			
OCL	NULL	Class, Association, Indication, Method	STRING ARRAY
<p>MEANING: Indicates that the qualified element specifies one or more constraints. These constraints are defined using the Object Constraint Language syntax, as defined by the Open Management Group (OMG), in the UML v2 OCL specification (http://www.omg.org/cgi-bin/apps/doc?ptc/03-10-14.pdf).</p> <p>USAGE: The OCL array contains string values using the syntax defined by UML's OCL 'inv:', 'pre:' and 'post:' clauses (invariants, and method pre and post clauses). The context of the constraint is defined by the qualified element (i.e., the context is the class or method that carries the qualifier).</p>			

QUALIFIER	DEFAULT	APPLIES TO	TYPE
<p>Note:</p> <ul style="list-style-type: none"> 'self' refers to the context 'result' refers to the result of a method A property in a class is defined using the syntax: ("self" <class name>)." <property name> Following an association to a specific role is defined using the syntax: ("self" <class name>)." <association class name> "[" <role name> "]" Various set operations can be performed as a result of following an association: sum, size, isEmpty, notEmpty, includes, excludes, count(<specific object>) These are invoked by specifying "->" and the operation name, after the association and role name are selected as above. A few basic tests exist: oclIsTypeOf(<class name>) for specific class checks oclIsKindOf(<class name>) for class or superclass checking oclIsNew() for operation post condition checking that a new instance is created (i.e., did not previously exist) oclIsUndefined() checking for null values <p>For example, checking that both property x and y cannot be null is specified using the following syntax, defined on a class: OCL {"inv: not(self.x.oclIsUndefined() AND self.y.oclIsUndefined())"}</p>			
OCTETSTRING	FALSE	Property, Parameter, Method	BOOLEAN
<p>MEANING: This qualifier is used to identify the qualified property or parameter as an octet string.</p> <p>When used in conjunction with an unsigned 8-bit integer (uint8) array, the OCTETSTRING qualifier indicates that the unsigned 8-bit integer array represents a single octet string.</p> <p>When used in conjunction with arrays of strings, the OCTETSTRING qualifier indicates that the qualified character strings are encoded textual conventions representing octet strings. The text encoding of these binary values conforms to the following grammar: "0x" 4*(<hexDigit> <hexDigit>). In both cases, the first 4 octets of the octet string (8 hexadecimal digits in the text encoding) are the number of octets in the represented octet string with the length portion included in the octet count (e.g., "0x00000004" is the encoding of a 0 length octet string).</p>			
OUT	FALSE	Parameter	BOOLEAN
<p>MEANING: Indicates that the associated parameter is used to return values from a method.</p>			
OVERRIDE	NULL	Property, Method, Reference	STRING
<p>MEANING: Indicates that the element in the derived class overrides another element (of the same name) defined in the ancestry of the class. The overriding element effectively takes the place of the element being overridden, for instances of the</p>			

QUALIFIER	DEFAULT	APPLIES TO	TYPE
<p>class on which the overriding element appears, and any of its subclasses (unless, of course, the overriding element is itself overridden in one of the subclasses). The flavor of the qualifier is defined as 'Restricted', so that the Override qualifier is NOT repeated in (i.e., inherited by) each subclass. The effect of the override is inherited, but not the identification of the Override qualifier itself. This enables new Overrides in subclasses to be easily located and applied. The value of this qualifier MAY identify the ancestral class whose subordinate element (property or method) is overridden. The format of the string to accomplish this is: [<className> "."] <IDENTIFIER>. If the class name is omitted, the element being overridden is found by searching the ancestry of the class until a definition of an appropriately-named subordinate element (of the same meta-schema class) is found. Usage Rule: The Override qualifier can only refer to elements of the same meta-schema class, e.g., properties can only override properties, etc. Also, it is not allowed to change an element's name or signature when overriding.</p>			
PROPAGATED	NULL	Property	STRING
<p>MEANING: The propagated qualifier is a string-valued qualifier that contains the name of the key that is being propagated. Its use assumes the existence of only one weak qualifier on a reference that has the containing class as its target. The associated property MUST have the same value as the property named by the qualifier in the class on the other side of the weak association. The format of the string to accomplish this is: [<className> "."] <IDENTIFIER></p> <p>Usage Rule: When the PROPAGATED qualifier is used, the KEY qualifier MUST be specified with a value of TRUE.</p>			
READ	TRUE	Property	BOOLEAN
<p>MEANING: Indicates that the property is readable.</p>			
REQUIRED	FALSE	Property, Reference	BOOLEAN
<p>MEANING: Indicates that a non-NULL value is required for the property or reference. Properties of a class that are inherent characteristics of a class and identifying in nature (e.g., domain name, file name, burned-in device identifier, IP address, etc.) that are likely to be useful for applications as query entry points and that are not KEY properties should be REQUIRED properties</p>			
REVISION (DEPRECATED)	NULL	Class, Association, Indication,	STRING
<p>MEANING: DEPRECATED - See VERSION Qualifier Provides the minor revision number of the schema object. Usage Rule: The VERSION qualifier MUST be present to supply the major version number when the REVISION qualifier is used.</p>			
SCHEMA	NULL	Property, Method	STRING
<p>MEANING: The name of the schema that contains the feature.</p>			
SOURCE			
<p>MEANING: This instance-level qualifier, and the corresponding pragma, were removed as an erratum by CR1461.</p>			
SOURCETYPE			
<p>MEANING: This instance-level qualifier, and the corresponding pragma, were removed as an erratum by CR1461.</p>			
STATIC	FALSE	Property, Method	BOOLEAN

QUALIFIER	DEFAULT	APPLIES TO	TYPE
<p>MEANING: For methods indicates that the method is a class method that does not depend on any per-instance data. For properties, indicates that the property is a class variable rather than an instance variable.</p>			
TERMINAL	FALSE	Class, Association, Indication	BOOLEAN
<p>MEANING: Indicate that the class can have no subclasses. If such a subclass is declared the compiler will generate an error. Note this qualifier cannot coexist with the Abstract qualifier. If both are specified the compiler generates an error.</p>			
UNITS	NULL	Property, Method, Parameter	STRING
<p>MEANING: Indicates the units of the associated data item (e.g., a Size data item might have Units of "bytes"). The complete set of standard units is defined in Appendix C.</p>			
VALUEMAP	NULL	Property, Method, Parameter	STRING ARRAY
<p>Defines the set of permissible values for the qualified property, method return or method parameter.</p> <p>The ValueMap qualifier can be used alone, or in combination with the Values qualifier. When used with the Values qualifier, the location of the value in the ValueMap array determines the location of the corresponding entry in the Values array.</p> <p>Where: ValueMap may only be used with string or integer types. When used with a string type, a ValueMap entry is a MOF stringvalue. When used with an integer type, a ValueMap entry is a MOF integervalue or an integervalue range as defined here.</p> <p style="padding-left: 20px;">integervalue range: [integervalue] ".." [integervalue]</p> <p>A ValueMap entry of : "x" claims the value x, "..x" claims all values less than and including x, "x.." claims all values greater than and including x, and, ".." claims all values not otherwise claimed.</p> <p>The values claimed are constrained by the type of the associated property.</p> <p>ValueMap = ("..") is not permitted.</p> <p>If used with a Value array, then all values claimed by a particular ValueMap entry applies to the corresponding Value entry.</p> <p>Example: [Values {"zero&one", "2to40", "fifty", "the unclaimed", "128-255"}, ValueMap {"..1","2..40" "50", "..", "x80.." }]</p> <p>uint8 example: In this example, where the type is uint8, "..1" and "zero&one" map to "0" and "1" "2..40" and "2to40" map to "2" thru "40" ".." and "the unclaimed" map to "41" thru "49" and to "51" thru "127" "0x80.." and "128-255" map to "128" thru "255"</p>			
VALUES	NULL	Property, Method, Parameter	STRING ARRAY
<p>MEANING: Provides translation between an integer values and strings (such as abbreviations or English terms) in the ValueMap array, and an associated string at the same index in the Values array. If a ValueMap qualifier is not present, the Values array is indexed (zero relative) using the value in the associated property, method return type or method parameter. If a ValueMap qualifier is present, the Values index is defined by the location of the property value in the ValueMap.</p>			

QUALIFIER	DEFAULT	APPLIES TO	TYPE
Usage Rule: The number of entries in the Values and ValueMap arrays MUST match.			
VERSION	NULL	Class, Association, Indication	STRING
<p>MEANING: Provides the version information of the object. This is incremented when changes are made to the object.</p> <p>Usage Rule: Starting with CIM Schema 2.7 (including extension schema), the VERSION qualifier MUST be present on each class to indicate the version when the class was last updated.</p> <p>The string representing the version comprises three decimal integers separated by periods, i.e., M.N.U, or, more formally, 1*<decimalDigit> "." 1*<decimalDigit> "." 1*<decimalDigit>, where</p> <p>M - The major version in numeric form that changed the class</p> <p>N - The minor version in numeric form that changed the class</p> <p>U - The update (e.g. errata, patch, ...) in numeric form that changed the class.</p> <p>NOTES:</p> <p>The addition/removal of the <i>Experimental</i> qualifier does not require that the version information be updated.</p> <p>The version change is applicable to only those elements that are local to the class. In other words, the version change of a superclass does not require the version in the subclass to be updated.</p> <p>Examples:</p> <p>Version("2.7.0")</p> <p>Version("1.0.0")</p>			
WEAK	FALSE	Reference	BOOLEAN
<p>MEANING: Indicates that the keys of the referenced class include the keys of the other participants in the association. This qualifier is used when the identity of the referenced class depends on the identity of the other participants in the association. No more than one reference to any given class can be weak. The other classes in the association MUST define a key. The keys of the other classes in the association are repeated in the referenced class and tagged with a propagated qualifier.</p>			
WRITE	FALSE	Property	BOOLEAN
<p>MEANING: Indicates that the modeling semantics of a property support modification of that property by "consumers". This qualifier is only intended to capture modeling semantics and specifically is not intended to address more dynamic characteristics such as provider capability or authorization rights.</p>			

411

412 **2.5.3 Optional Qualifiers**

413 The optional qualifiers listed in this table address situations that are not common to all CIM-compliant
 414 implementations. Thus, CIM-compliant implementations can ignore optional qualifiers since they are not required
 415 to interpret or understand these qualifiers. These are provided in the specification to avoid random user-defined
 416 qualifiers for these recurring situations.

Qualifier	Default	Applies To	Type
ALIAS	NULL	Property, Reference, Method	STRING
MEANING: Establishes an alternate name for a property or method in the schema.			

Qualifier	Default	Applies To	Type
DELETE	FALSE	Association, Reference	BOOLEAN
<p>MEANING: For associations: Indicates that the qualified association MUST be deleted if any of the objects referenced in the association are deleted, AND the respective object referenced in the association is qualified with IFDELETED.</p> <p>For references: Indicates that the referenced object MUST be deleted if the association containing the reference is deleted, AND qualified with IFDELETED, or if any of the objects referenced in the association are deleted AND the respective object referenced in the association is qualified with IFDELETED.</p> <p>Usage Rule: Applications MUST chase associations according to the modeled semantic and delete objects appropriately. <i>Note: This usage rule must be verified when the CIM security model is defined.</i></p>			
EXPENSIVE	FALSE	Any	BOOLEAN
<p>MEANING: Indicates the element is expensive to manipulate and/or compute.</p>			
IFDELETED	FALSE	Association, Reference	BOOLEAN
<p>MEANING: Indicates that all objects qualified by DELETE within the association MUST be deleted if the referenced object or the association, respectively, is deleted.</p>			
INVISIBLE	FALSE	Association, Property, Method, Reference, Class	BOOLEAN
<p>MEANING: Indicates that the element is defined only for internal purposes (for example, as an intermediate value in a calculation or to facilitate association semantics) and should not be displayed or otherwise relied upon.</p>			
LARGE	FALSE	Property, Class	BOOLEAN
<p>MEANING: Indicates the property or class requires a large amount of storage space.</p>			
PROVIDER	NULL	Any	STRING
<p>MEANING: An implementation specific handle to the instrumentation that populates those elements in the schemas that refer to dynamic data.</p>			
PROPERTYUSAGE	“CURRENTCONTEXT”	Property	STRING

Qualifier	Default	Applies To	Type
<p>MEANING: This qualifier allows classification of properties according to how they are intended to be used from the managed elements point of view. They are intended to allow the managed element to convey intent for property usage for the managed element. It is not intended to convey what access CIM has to the properties, i.e.. not all properties classified as configuration are necessarily writeable. Some configuration properties may be maintained by the provider or resource that the managed element represents, and not by CIM. The PropertyUsage qualifier is meant to allow distinguishing between properties that represent attributes of a managed resource vs. capabilities of a managed resource vs. configuration data for a managed resource vs. metrics about or from a managed resource vs. state information for a managed resource. If the qualifiers value is set to CurrentContext (i.e., the default value) then the PropertyUsage qualifiers actual value should be determined by looking at what class the property is placed in. The rules for which classes/subclasses result in which default PropertyUsage values are defined below:</p> <p>Class>CurrentContext PropertyUsage Value Setting > Configuration Configuration > Configuration Statistic > Metric ManagedSystemElement > State Product > Descriptive FRU > Descriptive SupportAccess > Descriptive Collection > Descriptive</p> <p>Usage Rules: The valid values for this qualifier are UNKNOWN, OTHER, CURRENTCONTEXT, DESCRIPTIVE, CAPABILITY, CONFIGURATION, STATE and METRIC.</p> <p>UNKNOWN indicates that the property's usage qualifier has not been determined and set.</p> <p>OTHER indicates that the property's usage is not Descriptive, Capabilities, Configuration, Metric, or State.</p> <p>CURRENTCONTEXT indicates that the PropertyUsage qualifiers value shall be inferred based on the class placement of the property according to the following rules:</p> <ul style="list-style-type: none"> *If the property is in a subclass of Setting or Configuration then the PropertyUsage value of CURRENTCONTEXT should be treated as CONFIGURATION. *If the property is in a subclass of Statistics then the PropertyUsage value of CURRENTCONTEXT should be treated as METRIC. *If the property is in a subclass of ManagedSystemElement then the PropertyUsage value of CURRENTCONTEXT should be treated as STATE. *If the property is in a subclass of Product, FRU, SupportAccess or Collection then the PropertyUsage value of CURRENTCONTEXT should be treated as DESCRIPTIVE. <p>DESCRIPTIVE indicates that the property contains information that describes the managed element. This can be properties like vendor, description, caption, etc. These properties are generally not good candidates for representation in Settings subclasses.</p> <p>CAPABILITY indicates that property contains information that reflects the inherent capabilities of the managed element regardless of its configuration. These are usually specifications of a product. For example, VideoController.MaxMemorySupported=128 would be a capability.</p> <p>CONFIGURATION indicates that the property contains information that influences or reflects the configurational state of the managed element. These properties are candidates for representation in Settings subclasses. VideoController.CurrentRefreshRate would be a configuration value.</p> <p>STATE indicates that the property contains information that reflects or can be used to derive the current status of the managed element.</p> <p>METRIC indicates that the property contains a numerical value representing a statistic or metric reporting performance oriented and/or accounting oriented information for the managed element. This would be appropriate for properties containing counters like 'BytesProcessed'.</p>			
SYNTAX	NULL	Property, Reference, Method, Parameter	STRING

Qualifier	Default	Applies To	Type
MEANING: Specific type assigned to a data item. Usage Rule: Must be used with the SyntaxType qualifier.			
SYNTAXTYPE	NULL	Property, Reference, Method, Parameter	STRING
MEANING: Defines the format of the SYNTAX qualifier. Usage Rule: Must be used with the SYNTAX qualifier			
TRIGGERTYPE	NULL	Class, Property, Method, Association, Indication, Reference	STRING
MEANING: Indicates the circumstances that cause a trigger to be fired. Usage Rule: The trigger types vary by meta-model construct. For classes and associations, the legal values are CREATE, DELETE, UPDATE and ACCESS. For properties and references, the legal values are: UPDATE and ACCESS. For methods, the legal values are BEFORE and AFTER. For indications, the legal values are THROWN.			
UNKNOWN VALUES	NULL	Property	STRING ARRAY
MEANING: Defines a set of values whose presence indicates that the value of the associated property is unknown – that is that the property cannot be considered as having a valid or meaningful value. The conventions and restrictions used for defining unknown values are the same as those applicable to the ValueMap qualifier. Note this qualifier cannot be overridden as it seems unreasonable to permit a subclass to treat as a known value a value that is treated as unknown by some superclass.			
UNSUPPORTED VALUES	NULL	Property	STRING ARRAY
MEANING: Defines a set of values whose presence indicates that the value of the associated property is unsupported – that is that the property cannot be considered as having a valid or meaningful value. The conventions and restrictions used for defining unsupported values are the same as those applicable to the ValueMap qualifier. Note this qualifier cannot be overridden as it seems unreasonable to permit a subclass to treat as a supported value a value that is treated as unknown by some superclass.			

417

418 2.5.4 User-defined Qualifiers

419 The user can define any additional arbitrary named qualifiers. However, it is recommended that only defined
420 qualifiers be used, and that the list of qualifiers be extended only if there is no other way to accomplish a particular
421 objective.

422 2.5.5 Mapping MIF Attributes

423 Mapping Management Information Format (MIF) attributes to CIM Properties can be accomplished using the
424 MAPPINGSTRINGS qualifier. This qualifier provides a mechanism to specify the mapping from DMTF and
425 vendor-defined MIF groups to specific properties. This allows for mapping using either Domain or Recast Mapping.

426 Every MIF group contains a unique identification that is defined using the class string, which is defined as follows:

427 defining body|specific name|version

428 where defining body is the creator and owner of the group, specific name is the unique name of the group and
429 version is a three-digit number that identifies the version of the group definition. In addition, each attribute has a
430 unique numeric identifier, starting with the number one.

431 Therefore, the mapping qualifier can be represented as a string that is formatted as follows:

432 MIF.defining body|specific name|version.attributeid

433 where MIF is a constant defining this as a MIF mapping followed by a dot. This is then followed by the class string
434 for the group this defines, and optionally followed by a dot and the identifier of a unique attribute.

435 In the case of a Domain Mapping, all of the above information is required, and provides a way to map an individual
436 MIF attribute to a particular CIM Property. In the case of the recast mapping, a CIM class can be recast from a MIF
437 group and only the MIF constant, followed by the dot separator followed by the class string, is required.

438 For example, a Domain Mapping of a DMTF MIF attribute to a CIM property would be as follows:

```
439           [MAPPINGSTRINGS { "MIF.DMTF | ComponentID | 001.4" }, READ]
440           SerialNumber = "";
```

441 The above declaration defines a mapping to the SerialNumber property from the DMTF Standard Component ID
442 group's serial number attribute. Because the qualifiers of CIM are a superset of those found in MIF syntax, any
443 qualifier may be overridden in the CIM definition.

444 To recast an entire MIF group into a CIM Object, the mapping string can be used to define an entire Class. For
445 example:

```
446           [MAPPINGSTRINGS { "MIF.DMTF | Software Signature | 002" }]
447           class MicroSoftWord : SoftwareSignature
448           {
449                 ...
450           }
```

451 2.5.6 Mapping Generic Data to CIM Properties

452 In addition to mapping MIF attributes, the MAPPINGSTRINGS qualifier can be used to map SNMP variables to
453 CIM properties. Every standard SNMP variable has associated with it a variable name and a unique object identifier
454 (OID) that is defined by a unique naming authority. This naming authority is a string. This string can either be a
455 name

456 standards body (e.g., "IETF"), a company name (e.g., "Acme") for defining the mappings to a company's private
457 MIB, and/or an appropriate management protocol (e.g., "SNMP"). For the IETF case, the ASN.1 module name, not
458 the RFC number, should be used as the MIB name (e.g., instead of saying RFC1493, the string "BRIDGE-MIB"
459 should be used). This is also true for the case of a company name being used as the naming authority. For the case of
460 using a management protocol like SNMP, the SNMP OID can be used to identify the appropriate SNMP variable.
461 This latter is especially important for mapping variables in private MIBs.

462 It should be noted that the concept of a naming authority for mapping data other than SNMP data into CIM
463 properties could be derived from this requirement. As an example, this can be used to map attributes of other data
464 stores (e.g., directories) using an application-specific protocol (e.g., LDAP).

465 The syntax for mapping MIF attributes as defined in Section 2.5.5 is as follows:

```
466           " MIF.<defining_body | specific_name | version>.attributeid"
```

467
468 The above MIF format can be reconciled with the more general syntax needed to map generic data to CIM
469 properties by realizing that both forms can be represented as follows:

```
470           " <Format>.<Scoping_Name>.<Content> "
```

471
472
473 where:

474 "Format" defines the format of the entry. It has the following values:

476 "MIF" means that the rest of the string is interpreted as MIF data
477 "MIB" means that the rest of the string is interpreted as a variable name of a MIB
478 "OID" means that the rest of the string is interpreted as an OID that is defined by a particular protocol to represent a
479 variable name
480
481 "Scoping_Name" defines the format used to uniquely identify the entry. It has the following values:
482 "defining_body | specific_name | version" is used for MIF mappings
483 "Naming_Authority | MIB_Name" is used for MIB mappings
484 "Naming_Authority | Protocol_Name" is used for protocol mappings that use OIDs to represent a variable name
485
486 "Content" defines the value of the entry. It has the following values:
487 "attributeid" is used for MIF mappings
488 "Variable_Name" is used for MIB mappings
489 "OID" is used for protocol mappings

490
491 Here are two examples of the syntax. The first uses the MIB format and looks as follows:

```
492         [Description(  
493           "OperatingSystem's notion of the local date and time of day"),  
494           MappingStrings {"MIB.IETF | HOST-RESOURCES-MIB.hrSystemDate"}]  
495         datetime LocalDateTime;
```

496
497 The second example uses the OID format and looks as follows:

```
498         [Description(  
499           "OperatingSystem's notion of the local date and time of day"),  
500           MappingStrings {"OID.IETF | SNMP.1.3.6.1.2.1.25.1.2"}]  
501         datetime LocalDateTime;
```

503 **3 Managed Object Format**

504 The management information is described in a language based on Interface Definition Language (IDL) [3] called the
505 Managed Object Format (MOF). This document uses the term MOF specification to refer to a collection of
506 management information described in a manner conformant to the MOF syntax.

507 Elements of MOF syntax are introduced on a case-by-case basis with examples. In addition, a complete description
508 of the MOF syntax is provided in Appendix A.

509 Note: All grammars defined in this specification use the notation defined in [7]; any exceptions are stated with the
510 grammar.

511 The MOF syntax is a way to describe object definitions in textual form. It establishes the syntax for writing
512 definitions. The main components of a MOF specification are textual descriptions of classes, associations,
513 properties, references, methods and instance declarations and their associated qualifiers. Comments are permitted.

514 In addition to serving the need for specifying the managed objects, a MOF specification can be processed using a
515 compiler. To assist the process of compilation, a MOF specification consists of a series of compiler directives.

516 A MOF file can be encoded in either Unicode or UTF-8.

517 **3.1 MOF usage**

518 The managed object descriptions in a MOF specification can be validated against an active namespace (See Section
519 5). Such validation is typically implemented in an entity acting in the role of a Server. This section describes the
520 behavior of an implementation when introducing a MOF specification into a namespace. Typically, such a process
521 validates both the syntactic correctness of a MOF specification, as well as the semantic correctness of such a
522 specification against a particular Implementation. In particular, MOF declarations must be ordered correctly with
523 respect to the target implementation state. For example, if the specification references a class without defining it
524 first, the reference is valid only if the server already has a definition of that class. A MOF specification can be
525 validated for the syntactic correctness alone, in a component such as a MOF compiler

526 **3.2 Class Declarations**

527 A class declaration is treated as an instruction to create a new class. It is a local matter as to whether the process of
528 introducing a MOF specification into a namespace is allowed to add classes or modify classes.

529 If the specification references a class without defining it first, the server must reject it as invalid if it does not already
530 have a definition of that class.

531 **3.3 Instance Declarations**

532 If the specification references a class without defining it first, the server must reject it as invalid if it does not already
533 have a definition of that class.

534 Any instance declaration is treated as an instruction to create a new instance where the object's key values do not
535 already exist, or an instruction to modify an existing instance where an object with identical key values already
536 exists.

537 4 MOF Components

538 4.1 Keywords

539 All keywords in the MOF syntax are case-insensitive.

540 4.2 Comments

541 Comments can appear anywhere in MOF syntax and are indicated by either a leading double slash "//", or a pair of
542 matching "/*" and "*/" sequences.

543 A "/" comment is terminated by carriage return, line feed or by the end of the MOF specification (whichever comes
544 first).

545 For example:

```
546 // This is a comment
```

547 A "/*" comment is terminated by the next "*/" sequence or by the end of the MOF specification (whichever comes
548 first). Comments are not recognized by the meta model and as such, will not be preserved across compilations. In
549 other words, the output of a MOF compilation is not required to include any comments.

550 4.3 Validation Context

551 Semantic validation of a MOF specification involves an explicit or implied namespace context. This is defined as
552 the namespace against which the objects in the MOF specification are validated and the namespace in which they
553 are created. Multiple namespaces typically indicate the presence of multiple management spaces or multiple devices.

554 4.4 Naming of Schema Elements

555 This section describes the rules for naming of schema elements; this applies to classes, properties, qualifiers,
556 methods and namespaces.

557 CIM is a conceptual model that is not bound to a particular implementation. This allows it to be used to exchange
558 management information in a variety of ways, examples of which are described in Section 1. Some implementations
559 may use case-sensitive technologies, while others may use case-insensitive technologies. The naming rules defined
560 in this section are chosen to allow efficient implementation in either environment, and to enable the effective
561 exchange of management information between all compliant implementations.

562 All names are case-insensitive, in that two schema item names are identical if they differ only in case. This is
563 mandated so that scripting technologies that are case-insensitive can leverage CIM technology. (Note, however, that
564 string values assigned to properties and qualifiers are not covered by this rule, and must be treated in a case-sensitive
565 manner).

566 The case of a name is set by its defining occurrence and must be preserved by all implementations. This is mandated
567 so that implementations can be built using case-sensitive technologies such as Java and object databases. (This also
568 allows names to be consistently displayed using the same user-friendly mixed-case format).

569 For example, an implementation, if asked to create class 'Disk', must reject the request if there is already a class
570 'DISK' in the current schema. Otherwise, when returning the name of the class 'Disk', it must return the name in
571 mixed case as it was originally specified.

572 CIM does not currently require support for any particular query language. It is assumed that implementations will
573 specify which query languages are supported by the implementation and will adhere to the case conventions that
574 prevail in the specified language. That is, if the query language is case-insensitive, statements in the language will
575 behave in a case-insensitive manner.

576 For the full rules for schema names see Appendix F, Unicode Usage.

577 4.5 Class Declarations

578 A class is an object describing a grouping of data items that are conceptually related and thought of as modeling an
579 object. Class definitions provide a type system for instance construction.

580 4.5.1 Declaring a Class

581 A class is declared by specifying these components:

- 582 1. The qualifiers of the class. This may be empty, or a list of qualifier name/value bindings separated
583 by commas "," and enclosed with square brackets "[" and "]".
- 584 2. The class name.
- 585 3. The name of the class from which this class is derived (if any).
- 586 4. The class properties, which define the data members of the class. A property may also have an
587 optional qualifier list, expressed in the same way as the class qualifier list. In addition, a property
588 has a data type, and (optionally) a default (initializer) value.
- 589 5. The methods supported by the class. A method may have an optional qualifier list. A method has a
590 signature consisting of its return type, plus its parameters and their type and usage.

591 This sample shows how to declare a class:

```
592         [abstract]
593 class Win32_LogicalDisk
594 {
595     [read]
596     string DriveLetter;
597     [read, Units("KiloBytes")]
598     sint32 RawCapacity = 0;
599     [write]
600     string VolumeLabel;
601     [Dangerous]
602     boolean Format([in] boolean FastFormat);
603 };
```

604 4.5.2 Subclasses

605 To indicate that a class is a subclass of another class, the derived class is declared by using a colon followed by the
606 superclass name.

607 For example, if the class Acme_Disk_v1 is derived from the class CIM_Media:

```
608     class Acme_Disk_v1 : CIM_Media
609     {
610         // Body of class definition here ...
611     };
```

612 The terms Base class, superclass and supertype are used interchangeably, as are Derived class, subclass and subtype.

613 The superclass declaration **must** appear at a prior point in the MOF specification or already be a registered class
614 definition in the namespace in which the derived class is defined.

615 4.5.3 Default Property Values

616 Any properties in a class definition can have default initializers. For example:

```
617     class Acme_Disk_v1 : CIM_Media
618     {
619         string Manufacturer = "Acme";
620         string ModelNumber = "123-AAL";
621     };
```

622 When new instances of the class are declared, then any such property is automatically assigned its default value
623 unless the instance declaration explicitly assigns a value to the property.

624 4.5.4 Class and Property Qualifiers

625 Qualifiers are meta data about a property, method, method parameter, or class, and are not part of the definition
626 itself. For example, a qualifier is used to indicate whether a property value is modifiable (using the WRITE
627 qualifier). Qualifiers always precede the declaration to which they apply.

628 Certain qualifiers are well known and cannot be redefined (see the description of the meta schema). Apart from
629 these, arbitrary qualifiers may be used.

630 Qualifier declarations include an explicit type indicator, which must be one of the intrinsic types. A qualifier with an
631 array-based parameter is assumed to have a type, which is a variable-length homogeneous array of one of the
632 intrinsic types. Note that in the case of boolean arrays, each element in the array is either TRUE or FALSE.

633 Examples:

```
634         Write(true)                               // boolean
635         profile { true, false, true }             // boolean []
636         description("A string")                   // string
637         info { "this", "a", "bag", "is" }         // string []
638         id(12)                                     // uint32
639         idlist { 21, 22, 40, 43 }                 // uint32 []
640         apple(3.14)                               // real32
641         oranges { -1.23E+02, 2.1 }               // real32 []
```

642 Qualifiers are applied to a class by preceding the class declaration with a qualifier list, comma-separated, and
643 enclosed within square brackets. Qualifiers are applied to a property or method in a similar fashion.

644 For example:

```
645         class CIM_Process:CIM_LogicalElement
646         {
647             uint32 Priority;
648             [Write(true)]
649             string Handle;
650         };
```

651 When specifying a boolean qualifier in a class or property declaration, the name of the qualifier can be used without
652 also specifying a value. From the previous example:

```
653         class CIM_Process:CIM_LogicalElement
654         {
655             uint32 Priority;
656             [Write] // Equivalent declaration to Write (True)
657             string Handle;
658         };
```

659 If only the qualifier name is listed for a boolean qualifier, it is implicitly set to TRUE.

660 In contrast, when a qualifier is not specified at all for a class or property, the default value for the qualifier is
661 assumed. Using another example:

```

662         [Association,
663         Aggregation] // Specifies the Aggregation qualifier to be True
664     class CIM_SystemDevice: CIM_SystemComponent
665     {
666         [Override ("GroupComponent"),
667         Aggregate] // Specifies the Aggregate qualifier to be True
668         CIM_ComputerSystem Ref GroupComponent;
669         [Override ("PartComponent"),
670         Weak] // Defines the Weak qualifier to be True
671         CIM_LogicalDevice Ref PartComponent;
672     };
673
674     [Association] // Since the Aggregation qualifier is not specified,
675                 // its default value, False, is set
676     class Acme_Dependency: CIM_Dependency
677     {
678         [Override ("Antecedent")] // Since the Aggregate and Weak
679                                 // qualifiers are not used, their
680                                 // default values, False, are assumed
681         Acme_SpecialSoftware Ref Antecedent;
682         [Override ("Dependent")]
683         Acme_Device Ref Dependent;
684     };

```

685 Qualifiers can be transmitted automatically from classes to derived classes, or from classes to instances, subject to
686 certain rules. The rules behind how the transmission occurs are attached to each qualifier and encapsulated in the
687 concept of the qualifier flavor. For example, a qualifier may be designated in the base class as automatically
688 transmitted to all of its derived classes, or it may be designated as belonging specifically to that class and not
689 transmittable.

690 The former is achieved by using the ToSubclass flavor, and the latter by using the Restricted flavor. These two
691 flavors MUST NOT be used at the same time. In addition, if a qualifier gets transmitted to its derived classes, the
692 qualifier flavor can be used to control whether or not derived classes can override the qualifier value, or whether it
693 must be fixed for an entire class hierarchy. This aspect of qualifier flavor is referred to as override permissions.

694 This is done by using the EnableOverride or DisableOverride flavors, which MUST NOT be used at the same time.
695 If a qualifier does not get transmitted to its derived classes, these two flavors are meaningless and MUST be ignored.

696 Qualifier flavors are indicated by an optional clause after the qualifier and preceded by a colon. They consist of
697 some combination of the key words EnableOverride, DisableOverride, ToSubclass and Restricted, indicating the
698 applicable propagation and override rules. For example:

```

699         class CIM_Process: CIM_LogicalElement
700         {
701             uint32 Priority;
702             [Write(true):DisableOverride ToSubclass]
703             string Handle;
704         };

```

705 In this example, Handle is designated as writable for the Process class and for every subclass of this class.

706 The recognized flavor types are:

PARAMETER	Interpretation	Default
ToSubclass	The qualifier is inherited by any subclass.	ToSubclass
Restricted	The qualifier applies only to the class in which it is declared.	ToSubclass
EnableOverride	If ToSubclass is in effect: The qualifier is overridable.	EnableOverride
DisableOverride	If ToSubclass is in effect: The qualifier cannot be overridden.	EnableOverride
Translatable	Indicates the value of the qualifier can be specified in multiple locales (language and country combination). When Translatable(yes) is specified for a qualifier, it is legal to create implicit qualifiers of the form :	no

PARAMETER	Interpretation	Default
	<p>label_ll_cc</p> <p>where "label" is the name of the qualifier with Translatable(yes), and ll and cc are the language code and country code designation, respectively, for the translated string. In other words, a label_ll_cc qualifier is a clone, or derivative, of the "label" qualifier with a postfix to capture the translated value's locale. The locale of the original value (that is, the value specified using the qualifier with a name of "label") is determined by the locale pragma.</p> <p>When a label_ll_cc qualifier is implicitly defined, the values for the other flavor parameters are assumed to be the same as for the "label" qualifier. When a label_ll_cc qualifier is defined explicitly, the values for the other flavor parameters must also be the same. A "yes" for this parameter is only valid for string-type qualifiers.</p> <p>Example: if an English description is translated into Mexican Spanish the actual name of the qualifier is: DESCRIPTION_es_MX.</p>	

707 **4.5.5 Key Properties**

708 Instances of a class require some mechanism through which the instances can be distinguished within a single
 709 namespace. Designating one or more properties with the reserved qualifier "key" provides instance identification.

710 For example, this class has one property (Volume) which serves as its' key:

```

711     class Acme_Drive
712     {
713         [key]
714         string Volume;
715         string FileSystem;
716         sint32 Capacity;
717     };
    
```

718 In this example, instances of Drive are distinguished using the Volume property, which acts as the key for the class.

719 Compound keys are supported and are designated by marking each of the required properties with the key qualifier.

720 If a new subclass is defined from a superclass, and the superclass has key properties (including those inherited from
 721 other classes), the new subclass **cannot** define any additional key properties. New key properties in the subclass can
 722 be introduced only if all classes in the inheritance chain of the new subclass are keyless.

723 If any reference to the class has the Weak qualifier, the properties that are qualified as Key in the other classes in the
 724 association are propagated to the referenced class. The key properties are duplicated in the referenced class using the
 725 name of the property, prefixed by the name of the original declaring class. For example:

```
726     class CIM_System: CIM_LogicalElement
727     {
728         [Key]
729         string Name;
730     };
731
732     class CIM_LogicalDevice: CIM_LogicalElement
733     {
734         [Key]
735         string DeviceID;
736         [Key, Propagated("CIM_System.Name")]
737         string SystemName;
738     };
739
740     [Association]
741     class CIM_SystemDevice: CIM_SystemComponent
742     {
743         [Override ("GroupComponent"), Aggregate, Min(1), Max(1)]
744         CIM_System Ref GroupComponent;
745         [Override ("PartComponent"), Weak]
746         CIM_LogicalDevice Ref PartComponent;
747     };
```

748 4.6 Association Declarations

749 An association is a special kind of a class describing a link between other classes. As such, they also provide a type
750 system for instance constructions. Associations are just like other classes with a few additional semantics explained
751 below.
752

753 4.6.1 Declaring an Association

754 An association is declared by specifying these components:

755 The qualifiers of the association (at least the ASSOCIATION qualifier, if it doesn't have a supertype). Further
756 qualifiers may be specified as a list of qualifier/name bindings separated by commas ",". The entire qualifier list is
757 enclosed in square brackets "[" and "]").

758 . The association name.

759 The name of the association from which this association is derived (if any).

760 The association references which define pointers to other objects linked by this association. References may also
761 have qualifier lists, expressed in the same way as the association qualifier list. Especially the qualifiers to specify
762 cardinalities of references are important to be mentioned (see 2.5.2. "Standard Qualifiers"). In addition, a reference
763 has a data type, and (optionally) a default (initializer) value.

764 Additional association properties which define further data members of this association. They are defined in the
765 same way as for ordinary classes.

766 The methods supported by the association. They are defined in the same way as for ordinary classes.

767 The following example shows how to declare an association (assuming given classes CIM_A and CIM_B):

```
768     [Association]
769     class CIM_LinkBetweenAandB : CIM_Dependency
770     {
771         [Override ("Antecedent")]
772         CIM_A Ref Antecedent;
773         [Override ("Dependent")]
774         CIM_B Ref Dependent;
775     };
```

776 4.6.2 Subassociations

777 To indicate that an association is a subassociation of another association, the same notation as for ordinary classes is
778 used, i.e. the derived association is declared by using a colon followed by the superassociation name. (An example is
779 provided above.)
780

781 4.6.3 Key References and Properties

782 Instances of an association also require some mechanism through which the instances can be distinguished, implied
783 by the fact that they are just a special kind of a class. Designating one or more references/properties with the
784 reserved KEY qualifier provides instance identification.

785 A reference/property of an association is (part of) the association key if the KEY qualifier is applied.

```
786     [Association, Aggregation]
787     class CIM_Component
788     {
789         [Aggregate, Key]
790         CIM_ManagedSystemElement Ref GroupComponent;
791         [Key]
792         CIM_ManagedSystemElement Ref PartComponent;
793     };
```

794 In principle, the key definition of association follows the same rules as for ordinary classes. Compound keys are
795 supported in the same way. Also a new subassociation **cannot** define any additional key properties/references.

796 If any reference to a class has the WEAK qualifier, the KEY-qualified properties of the other class, whose reference
797 is not WEAK-qualified are propagated to the class. (see subchapter 4.5.5 "Key Properties").

798 4.6.4 Object References

799 Object references are properties whose values are links or pointers to other objects (classes or instances). The value
800 of an object reference is expressed as a string, which represents a path to another object. The path includes:

801 The namespace in which the object resides.

802 The class name of the object.

803 If the object represents an instance, the values of all key properties for that instance.

804 Object reference properties are declared by "XXX ref", indicating a strongly typed reference to objects of the class
805 with name "XXX" (or a derived class thereof). For example:

```
806     [Association]
807     class Acme_ExampleAssoc
808     {
809         Acme_AnotherClass ref Inst1;
810         Acme_Aclass      ref Inst2;
811     };
```

812 In the above declaration, Inst1 can only be set to point to objects of type Acme_AnotherClass.

813 Also see Section 4.12.2 on Initializing References Using Aliases.

814 In associations, object references have cardinalities - denoted using Min and Max qualifiers. Here are examples of
815 UML cardinality notations and their respective combinations of Min and Max values:

UML	MIN	MAX	Required MOF Text*	Description
*	0	NULL		Many
1..*	1	NULL	Min(1)	At least one
1	1	1	Min(1), Max(1)	One
0,1 (or 0..1)	0	1	Max(1)	At most one

816

817 4.7 Qualifier Declarations

818 Qualifiers may be declared using the keyword “qualifier”. The declaration of a qualifier allows the definition of
819 types, default values, propagation rules (also known as Flavors), and restrictions on use.

820 The default value for a declared qualifier is used when the qualifier is not explicitly specified for a given schema
821 element (explicit specification includes when the qualifier specification is inherited).

822 The MOF syntax allows specifying a qualifier without an explicit value. In this case, the assumed value depends on
823 the qualifier type: booleans are true, numeric types are null, strings are null and arrays are empty.

824 For example, the alias qualifier is declared as follows:

```
825     qualifier alias :string = null, scope (property, reference, method);
```

826 This declaration establishes a qualifier called alias. The type of the qualifier is string. It has a default value of null
827 and may only be used with properties, references and methods.

828 The meta qualifiers are declared as:

```
829     Qualifier Association : boolean = false,  
830     Scope(class, association), Flavor(DisableOverride);
```

```
831     Qualifier Indication : boolean = false,  
832     Scope( class, indication), Flavor(DisableOverride);
```

834

835 See Appendix B for the complete list of standard qualifiers.

836 4.8 Instance Declarations

837 This section is specific to instance declarations, Method declarations are covered in Section 4.9. Instances are
838 declared using the keyword sequence "instance of" and the class name. The property values of the instance may be
839 initialized within an initialization block. Any Qualifiers specified for the Instance MUST already be present in the
840 defining Class, and MUST have the same value and flavor(s).

841 Property initialization consists of an optional list of preceding qualifiers, the name of the property and an optional
842 value. Any Qualifiers specified for the Property MUST already be present in the Property definition from the
843 defining Class, and MUST have the same value and flavor(s). Any property values not initialized have default
844 values as specified in the class definition, or (if no default value has been specified) the special value NULL to
845 indicate "absence of value". For example, given the class definition:

```
846     class Acme_LogicalDisk: CIM_Partition  
847     {  
848         [key]  
849         string DriveLetter;  
850         [Units("kilo bytes")]  
851         sint32 RawCapacity = 128000;  
852         [write]  
853         string VolumeLabel;  
854         [Units("kilo bytes")]  
855         sint32 FreeSpace;  
856     };
```

857 an instance of the above class might be declared as:

```
858     instance of Acme_LogicalDisk  
859     {  
860         DriveLetter = "C";  
861         VolumeLabel = "myvol";  
862     };
```

863 The resulting instance would take these property values:

- 864 1. DriveLetter would be assigned the value "C".
- 865 2. RawCapacity would be assigned the default value 128000.
- 866 3. VolumeLabel would be assigned the value "myvol".

867 4. FreeSpace would be assigned the value NULL.

868 For subclasses, all of the properties in the superclass must have their values initialized along with the properties in
869 the subclass. Any property values not specifically assigned in the instance block will have either the default value
870 for the property (if there is one), or else the value NULL (if there is not one).

871 The values of all key properties must be specified in order for an instance to be identified and created. There is no
872 requirement to explicitly initialize other property values. See Section 4.11.6 on behavior when there is no property
873 value initialization.

874 As described in item 21-E of Section 2.1, a class may have, by inheritance, more than one property with a particular
875 name. If a property initialization has a property name which applies to more than one property in the class, the
876 initialization applies to the property defined "closest" to the class of the instance. That is, the property can be
877 located by starting at the class of the instance; if the class defines a property with the name from the initialization,
878 then that property is initialized, otherwise, the search is repeated from the class' direct superclass. See also
879 Appendix L for further discussion of the name conflict issue.

880 Instances of Associations may also be defined. For example:

```
881     instance of CIM_ServiceSAPDependency
882     {
883         Dependent = "CIM_Service.Name = \"mail\"";
884         Antecedent = "CIM_ServiceAccessPoint.Name = \"PostOffice\"";
885     };
886
```

887 4.8.1 Instance Aliasing

888 An alias can be assigned to an instance using this syntax:

```
889     instance of Acme_LogicalDisk as $Disk
890     {
891         // Body of instance definition here ...
892     };

```

893 Such an alias can later be used within the same MOF specification as a value for an object reference property. For
894 more information, see Section 4.12.2 Initializing References using Aliases.

895 4.8.2 Arrays

896 Arrays of any of the basic data types can be declared in the MOF specification by using square brackets after the
897 property identifier. Fixed-length arrays indicate their length as an unsigned integer constant within the square
898 brackets; otherwise, the array is assumed to be variable length. Arrays can be bags, ordered lists or indexed arrays.
899 An array's type is defined by the ARRAYTYPE qualifier, whose values are "Bag", "Ordered" or "Indexed". The
900 default array type is "Bag". Regarding each of the array types:

901 An array of type "Bag" is unordered and multi-valued, allowing duplicate entries.

902 An ordered list ("Ordered") is a special case of a bag, which is multi-valued and allows duplicate entries. It returns
903 the property values in an implementation dependent, but fixed order.

904 An indexed array ("Indexed") maintains the order of the elements, and could be implemented based on an integer
905 index for each of the array values.

906 Note that for the "Bag" array type, no significance is defined for the array index other than a convenience for
907 accessing the elements of the array. For example, there can be no assumption that the same index will return the
908 same value for every access to the array. The only assumption is that a complete enumeration of the indices will
909 return a complete set of values.

910 For the "Ordered" array type, the array index is significant as long as no array elements are added, deleted or
911 changed. In this case the same index will return the same value for every access to the array. If an element is added,
912 deleted or changed, the index of the elements might change according to the implementation-specific ordering
913 algorithm.

914 The "Indexed" array maintains the correspondence between element position and value. Array elements can be
915 overwritten, but not deleted. Indexes start at 0 and have no gaps.

916 The current release of CIM does not support n-dimensional arrays.

917 Arrays of any basic data type are legal for properties. Arrays of references are not legal for properties. Arrays must
918 be homogeneous. Arrays of mixed types are not supported. In MOF, the data type of an array precedes the array
919 name. Array size, if fixed length, is declared within square brackets, following the array name. If a variable length
920 array is to be defined, empty square brackets follow the array name.

921 Arrays are declared using this MOF syntax:

```
922     class A
923     {
924         [Description("An indexed array of variable length"), ArrayType("Indexed")]
925         uint8 MyIndexedArray[];
926
927         [Description("A bag array of fixed length")]
928         uint8 MyBagArray[17];
929     };
```

930 If default values are to be provided for the array elements, this syntax is used:

```
931     class A
932     {
933         [Description("A bag array property of fixed length")]
934         uint8 MyBagArray[17] = {1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17};
935     };
```

936 This MOF presents further examples of "Bag", "Ordered" and "Indexed" array declarations:

```
937     class Acme_Example
938     {
939         char16 Prop1[];           // Bag (default) array of chars, Variable length
940
941         [ArrayType ("Ordered")] // Ordered array of double-precision reals,
942         real64 Prop2[];         // Variable length
943
944         [ArrayType ("Bag")]     // Bag array containing 4 32-bit signed integers
945         sint32 Prop3[4];
946
947         [ArrayType ("Ordered")] // Ordered array of strings, Variable length
948         string Prop4[] = {"an", "ordered", "list"};
949
950         // Prop4 is variable length with default values defined at the
951         // first three positions in the array
952
953         [ArrayType ("Indexed")] // Indexed array of 64-bit unsigned integers
954         uint64 Prop5[];
955     };
```

956 4.9 Method Declarations

957 A method is defined as an operation together with its signature. The signature consists of a possibly empty list of
958 parameters and a return type. There are no restrictions on the type of parameters other than they **must** be one of the
959 data types described in Section 2.2, a fixed or variable length array of one of those types. Method return types
960 defined in MOF must be one of the data types described in Section 2.2. Return types cannot be arrays, but otherwise
961 are unrestricted.

962 Note: Methods are expected, but not required, to return a status value indicating the result of execution of the
963 method. Methods may use their parameters to pass arrays.

964 Syntactically, the only thing that distinguishes a method from a property is the parameter list. The fact that methods
965 are expected to have side-effects is outside the scope of this specification.

966 In this example, Start and Stop methods are defined on the Service class. Each method returns an integer value:

```

967     class CIM_Service: CIM_LogicalElement
968     {
969         [Key]
970         string Name;
971         string StartMode;
972         boolean Started;
973         uint32 StartService();
974         uint32 StopService();
975     };

```

976 In this example, a Configure method is defined on the Physical DiskDrive class. It takes a
977 DiskPartitionConfiguration object reference as a parameter, and returns a boolean value.

```

978     class Acme_DiskDrive: CIM_Media
979     {
980         sint32 BytesPerSector;
981         sint32 Partitions;
982         sint32 TracksPerCylinder;
983         sint32 SectorsPerTrack;
984         string TotalCylinders;
985         string TotalTracks;
986         string TotalSectors;
987         string InterfaceType;
988         boolean Configure([IN] DiskPartitionConfiguration REF config);
989     };

```

990 4.10 Compiler Directives

991 Compiler directives are provided as the keyword "pragma", preceded by a hash (#) character, and followed by a
992 string parameter.

993 The current standard compiler directives are:

compiler Directive	Interpretation
#pragma include()	Has a file name as a parameter. The file is assumed to be a MOF file. The pragma has the effect of textually inserting the contents of the include file at the point where the include pragma is encountered.
#pragma instancelocale()	Declares the locale used for instances described in a MOF file. This pragma specifies the locale when "INSTANCE OF" MOF statements include string or char16 properties, and the locale is not the same as the locale specified by a #pragma locale() statement. The locale is specified as a parameter of the form ll_cc where ll is the language code based on ISO/IEC 639, and cc is the country code based on ISO/IEC 3166.
#pragma locale()	Declares the locale used for a particular MOF file. The locale is specified as a parameter of the form ll_cc, where ll is the language code based on ISO/IEC 639, and cc is the country code based on ISO/IEC 3166. When the pragma is not specified, the assumed locale is "en_US". It is important to note that this pragma does not apply to the syntax structures of MOF. Keywords, such as "class" and "instance", are always in en_US.
#pragma namespace()	This pragma is used to specify a Namespace path.
#pragma nonlocal()	These compiler directives, and the corresponding instance-level qualifiers, were removed as errata by CR1461.
#pragma nonlocaltype()	
#pragma source()	
#pragma sourcetype()	

994 Additional pragma directives may be added as a MOF extension mechanism. Unless standardized in a future CIM
995 Infrastructure specification, such new pragma definitions must be considered vendor-specific. Use of non-standard
996 pragma will affect interoperability of MOF import and export functions.

997 4.11 Value Constants

998 The constant types supported in the MOF syntax are described in the subsections that follow. These are used in
999 initializers for classes and instances, and in the parameters to named qualifiers.

1000 A formal specification of the representation is found in Appendix A, MOF Syntax Grammar Description.

1001 4.11.1 String Constants

1002 A string constant is a sequence of zero or more UCS-2 characters enclosed in double-quotes ("). A double-quote is
1003 allowed within the value, as long as it is preceded immediately by a backslash (\).

1004 For example:

```
1005     "This is a string"
```

1006 Successive quoted strings are concatenated, as long as only white space or a comment intervenes:

```
1007     "This" " " becomes a long string"
1008     "This" /* comment */ " becomes a long string"
```

1009 The escape sequences such as `\n`, `\t` and `\r` are recognized as legal characters within a string. The complete set is:

```
1010     \b      // \x0008: backspace BS
1011     \t      // \x0009: horizontal tab HT
1012     \n      // \x000A: linefeed LF
1013     \f      // \x000C: form feed FF
1014     \r      // \x000D: carriage return CR
1015     \"      // \x0022: double quote "
1016     \'      // \x0027: single quote '
1017     \\      // \x005C: backslash \
1018     \x<hex> // where <hex> is one to four hex digits
1019     \X<hex> // where <hex> is one to four hex digits
```

1020 The character set of the string depends on the character set supported by the local installation. While the MOF
1021 specification may be submitted in UCS-2 form [10], the local implementation may only support ANSI and vice-
1022 versa. Therefore, the string type is unspecified and dependent on the character set of the MOF specification itself. If
1023 a MOF specification is submitted using UCS-2 characters outside of the normal ASCII range, then the
1024 implementation may have to convert these characters to the locally-equivalent character set.

1025 4.11.2 Character Constants

1026 Character and wide-character constants are specified as.

```
1027     'a'
1028     '\n'
1029     '1'
1030     '\x32'
```

1031 Note: Forms such as octal escape sequences (e.g. `'\020'`) are not supported.

1032 Integer values can also be used as character constants, as long as they are within the numeric range of the character
1033 type. For example, wide-character constants must fall within the range 0 to 0xFFFF.

1034 4.11.3 Integral Constants

1035 Integer constants may be decimal, binary, octal or hexadecimal.

1036 For example, these are all legal:

```
1037     1000
1038     -12310
1039     0x100
1040     01236
1041     100101B
```

1042 Note that binary constants have a series of 1 and 0 digits, with a "b" or "B" suffix to indicate that the value is binary.

1043 The number of digits permitted depends on the current type of the expression. For example, it is not legal to assign
1044 the constant 0xFFFF to a property of type uint8.

1045 4.11.4 Floating-Point Constants

1046 Floating point constants are declared as specified by IEEE in Ref. [6].

1047 For example, these are legal:

```
1048         3.14
1049        -3.14
1050       -1.2778E+02
```

1051 The range for floating point constants depends on whether float or double properties are used and must fit within the
1052 range specified for IEEE 4-byte and 8-byte floating point values, respectively.

1053 4.11.5 Object Ref Constants

1054 Object references are simple URL-style links to other objects (which may be classes or instances). They take the
1055 form of a quoted string containing an object path. The object path is a combination of a namespace path and the
1056 model path.

1057 For example:

```
1058         "///./root/default:LogicalDisk.SystemName=\"acme\",LogicalDisk.Drive=\"C\" "  
1059         "///./root/default:NetworkCard=2"
```

1060 An object reference can also be an alias. See Section 4.12.2 for more details.

1061 4.11.6 NULL

1062 All types can be initialized to the predefined constant NULL, which indicates no value has been provided. The
1063 details of the internal implementation of the NULL value are not mandated by this document.

1064 4.12 Initializers

1065 Initializers are used both in class declarations for default values and instance declarations to initialize a property to a
1066 value. The format of initializer values is specified in Section 2 and its subsections.

1067 The initializer value **must** match the property data type. The only exceptions are the NULL value, which may be
1068 used for any data type, and integral values, used for characters.

1069 4.12.1 Initializing Arrays

1070 Arrays can be defined to be of type, "Bag", "Ordered" or "Indexed", and can be initialized by specifying their values
1071 in a comma-separated list (as in the C programming language). The list of array elements is delimited with curly
1072 brackets.

1073 For example, given this class definition:

```
1074         class Acme_ExampleClass
1075         {
1076             [ArrayType ("Indexed")]
1077             string ip_addresses []; // Indexed array of variable length
1078             sint32 sint32_values [10]; // Bag array of fixed length = 10
1079         };
```

1080 this is a valid instance declaration:

```
1081         instance of Acme_ExampleClass
1082         {
1083             ip_addresses = { "1.2.3.4", "1.2.3.5", "1.2.3.7" };
1084
1085             // ip_address is an indexed array of at least 3 elements, where
1086             // values have been assigned to the first three elements of the
1087             // array
1088
1089             sint32_values = { 1, 2, 3, 5, 6 };
1090         };
```

1091 Refer to Section 4.8.2 for additional information on declaring arrays, and the distinctions between bags, ordered
1092 arrays and indexed arrays.

1093 **4.12.2 Initializing References Using Aliases**

1094 Aliases are symbolic references to an object located elsewhere in the MOF specification. They only have
1095 significance within the MOF specification in which they are defined, and are only used at compile time to facilitate
1096 establishment of references. They are not available outside of the MOF specification.

1097 Instances may be assigned an alias as described in Section 4.8.1. Aliases are identifiers which begin with the \$
1098 symbol. When a subsequent reference to that instance is required for an object reference property, the identifier is
1099 used in place of an explicit initializer.

1100 Assuming that \$Alias1 and \$Alias2 have been declared as aliases for instances, and the obref1 and obref2 properties
1101 are object references, this example shows how the object references could be assigned to point to the aliased
1102 instances:

```
1103         instance of Acme_AnAssociation  
1104         {  
1105             strVal = "ABC";  
1106             obref1 = $Alias1;  
1107             obref2 = $Alias2;  
1108         };
```

1109 Forward-referencing and circular aliases are permitted.

1110 **5 Naming**

1111 Because CIM is not bound to a particular technology or implementation, it promises to facilitate sharing
 1112 management information between a variety of management platforms. The CIM Naming mechanism was defined to
 1113 address enterprise-wide identification of objects, as well as the sharing of management information.

1114 1. CIM Naming addresses these requirements:

1115 Ability to locate and uniquely identify any object in an enterprise

1116 Unambiguous enumeration of all objects

1117 Ability to determine when two object names reference the same entity

- 1118 • Location transparency (no need to understand which management platforms proxy other platforms’
 1119 instrumentation)

1120 2. Allow sharing of objects and instance data among management platforms

- 1121 • Allow creation of different scoping hierarchies which vary by “time” (for example, a “current” vs.
 1122 “proposed” scoping hierarchy)

1123 3. Facilitate move operations between object trees (including within a single management platform)

- 1124 • Hide underlying management technology/provide technology transparency for the domain-mapping
 1125 environment

1126 Object name identifiable regardless of instrumentation technology


1127 Allowing different names for DMI vs. SNMP objects requires the management platform to understand how the
 1128 underlying objects are implemented

1129 The KEY qualifier is the CIM Meta-Model mechanism used to identify the properties that uniquely identify an
 1130 instance of a class (and indirectly an instance of an association). CIM Naming enhances this base capability by
 1131 introducing the WEAK and PROPOGATED qualifiers to express situations in which the keys of one object are to be
 1132 propagated to another object.

1133 **5.1 Background**

1134 CIM MOF files can contain definitions of instances, classes or both, as illustrated in this diagram:

1135



1136

1137

Figure 5-1 Definitions of instances and classes

1138 MOF files can be used to populate a technology that understands the semantics and structure of CIM. When a MOF
1139 file is consumed by a particular implementation, there are two operations that are actually being performed,
1140 depending on the file's content. First, a compile or definition operation is performed to establish the structure of the
1141 model. Second, an import operation is performed to insert instances into the platform or tool.

1142 Once the compile and import are completed, the actual instances are manipulated using the native capabilities of the
1143 platform or tool. In other words, in order to manipulate an object (for example, change the value of a property), one
1144 must know the type of platform the information was imported into, the APIs or operations used to access the
1145 imported information, and the name of the platform instance that was actually imported. For example, the semantics
1146 become:

1147 Set the Version property of the Logical Element object with Name="Cool" in the relational database named
1148 LastWeeksData to "1.4.0".

1149 The contents of a MOF file are loaded into a namespace that provides a domain (in other words, a container), in
1150 which the instances of the classes are guaranteed to be unique per the KEY qualifier definitions. The term
1151 namespace is used to refer to an implementation that provides such a domain.

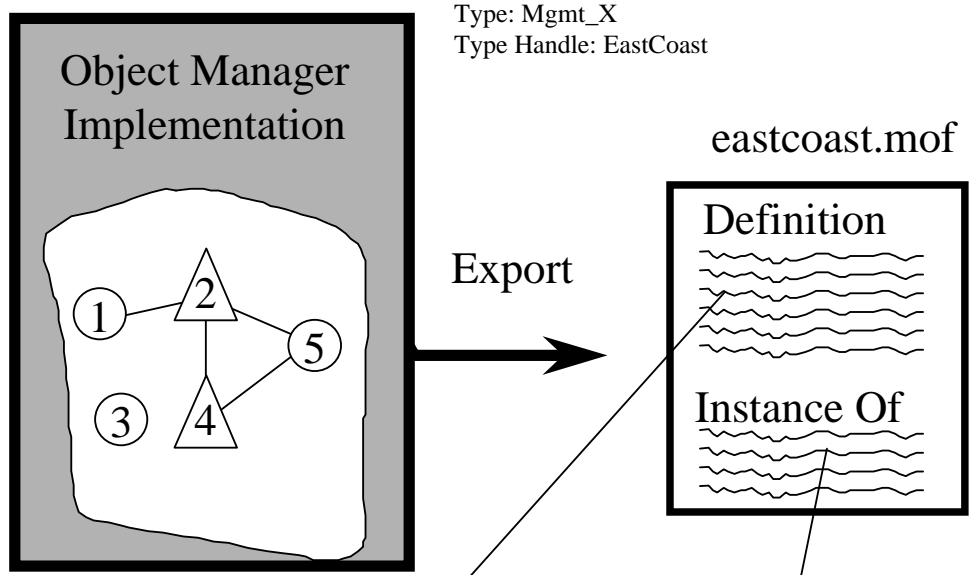
1152 Namespaces can be used to:

1153 Define chunks of management information (objects and associations) to limit implementation resource requirements,
1154 such as database size.

1155 Define views on the model for applications managing only specific objects, such as hubs.

1156 Pre-structure groups of objects for optimized query speed.

1157 Another viable operation is exporting from a particular management platform. Essentially, this operation creates a
1158 MOF file for all or some portion of the information content of a platform.



```
[
class Figs_Circle
{
  [ key ] uint32 Name;
  string Color; };

class Figs_Triangle
{
  [ key ] uint32 Label;
  string Color;
  uint32 Area;
};

[Association] class Figs_CircleToTriangle
{
  Figs_Circle REF ACircle;
  Figs_Triangle REF ATriangle;
};

[Association] class Figs_Covers
{
  Figs_Triangle REF Over;
  Figs_Triangle REF Under;
};
```

```
instance of Figs_Triangle {Label=2 ; Color="Blue";Area=12 };
instance of Figs_Triangle {Label=4 ; Color="Blue";Area=12 };
instance of Figs_Circle { Name=1 ; Color="Blue" };
instance of Figs_Circle { Name=3 ; Color="Blue" };
instance of Figs_Circle { Name=5 ; Color="Blue" };

instance of Figs_CircleToTriangle
{ ACircle = "Circle.Name=1";
  ATriangle = "Triangle.Label=2"; };

instance of Figs_CircleToTriangle
{ ACircle = "Circle.Name=5";
  ATriangle = "Triangle.Label=2"; };

instance of Figs_CircleToTriangle
{ ACircle = "Circle.Name=5";
  ATriangle = "Triangle.Label=4"; };

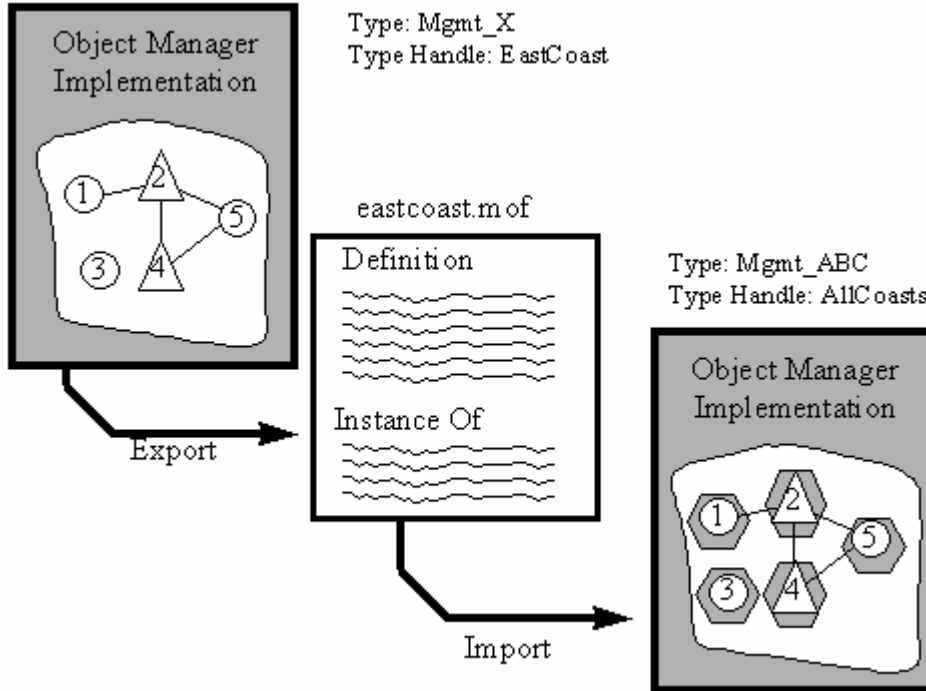
instance of Figs_Covers
{ Over = "Triangle.Label=2";
  Under = "Triangle.Label=4"; };
```

1159

1160

Figure 5-2 Exporting to MOF

1161 For example, information is exchanged when the source system is of type Mgmt_X and its name is EastCoast. The
 1162 export produces a MOF file with the circle and triangle definitions and instances 1, 3, 5 of the circle class and
 1163 instances 2, 4 of the triangle class. This MOF file is then compiled and imported into the management platform of
 1164 type Mgmt_ABC with the name AllCoasts.



1165

1166

Figure 5-3 Information Exchange

1167 The import operation involves storing the information in a local or native format of Mgmt_ABC so its native
 1168 operations can be used to manipulate the instances. The transformation to a native format is shown in the figure by
 1169 wrapping the five instances in hexagons. The transformation process must maintain the original keys.

1170 5.1.1 Management Tool Responsibility for an Export Operation

1171 The management tool must be able to create unique key values for each distinct object it places in the MOF file.

1172 For each instance placed in the MOF file, the management tool must maintain a mapping from the MOF file keys to
 1173 the native key mechanism.

1174 5.1.2 Management Tool Responsibility for an Import Operation

1175 The management tool must be able to map the unique keys found in the MOF file to a set of locally-understood
 1176 keys.

1177 5.2 Weak Associations: Supporting Key Propagation

1178 CIM provides a mechanism to name instances within the context of other object instances. For example, if a
 1179 management tool is handling a local system, then it can refer to the C drive or the D drive. However, if a
 1180 management tool is handling multiple machines, it must refer to the C drive on machine X and the C drive on
 1181 machine Y. In other words, the name of the drive must include the name of the hosting machine. CIM supports the
 1182 notion of weak associations to specify this type of key propagation.

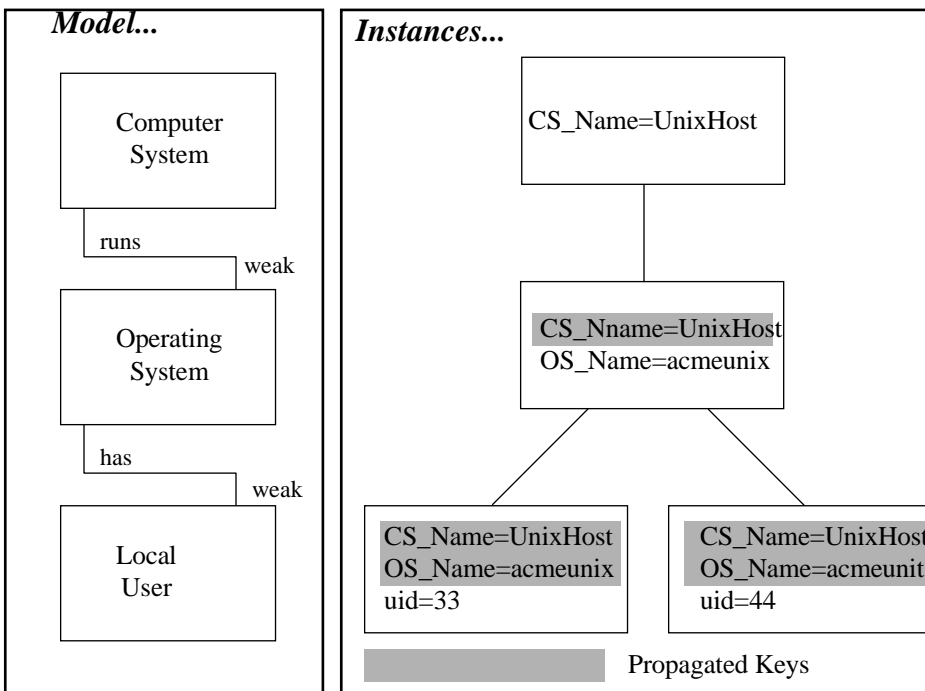
1183 A weak association is defined using a qualifier. For example:

```
1184     Qualifier Weak: boolean = false, Scope(reference), Flavor(DisableOverride);
```

1185 The key(s) of the referenced class includes the key(s) of the other participants in the WEAK association. This
 1186 situation occurs when the referenced class identity depends on the identity of other participants in the association.

1187 Usage Rule: This qualifier can only be specified on one of the references defined for an association. The Weak
 1188 referenced object is the one that depends on the other object for identity.

1189 This figure shows an example. There are three classes: ComputerSystem, OperatingSystem and Local User. The
 1190 Operating System class is weak with respect to the Computer System class, since the runs association is marked
 1191 weak. Similarly, the Local User class is weak with respect to the Operating System class, since the association is
 1192 marked weak.



1193

1194

Figure 5-4 Example of Weak Association

1195 In the context of a weak association definition, the Computer System class is a scoping class for the Operating
 1196 System class, since its keys are propagated to the Operating System class. The Computer System and the Operating
 1197 System classes are both scoping classes for the Local User class, since the Local User class gets keys from both.
 1198 Finally, the Computer System is referred to as a Top Level Object (TLO) because it is not weak with respect to any
 1199 other class. The fact that a particular class is a top-level object is inferred because no references to that class are
 1200 marked with the WEAK qualifier. In addition, Top Level Objects must have the possibility of an enterprise-wide,
 1201 unique key. An example may be a computer's IP address in a company's enterprise-wide IP network. The goal of
 1202 the TLO concept is to achieve uniqueness of keys in the model path portion of the object name. In order to come as
 1203 close as possible to this goal, TLO must have relevance in an enterprise context.

1204 Objects in the scope of another object can in turn be a scope for other objects; hence, all model object instances are
1205 arranged in directed graphs with the Top Level Object's (TLO's) as peer roots. The structure of this graph – in other
1206 words, which classes are in the scope of another given class – is defined as part of CIM by means of associations
1207 qualified with the WEAK qualifier.

1208 **5.2.1 Referencing Weak Objects**

1209 A reference to an instance of an association includes the propagated keys. The properties must have the propagated
1210 qualifier that identifies which class the property originates in and what the name of the property is in that class – for
1211 example

```
1212  
1213     instance of Acme_has  
1214     {  
1215         anOS = "Acme_OS.Name=\"acmeunit\",SystemName=\"UnixHost\"";  
1216         aUser = "Acme_User.uid=33,OSName=\"acmeunit\",SystemName=\"UnixHost\"";  
1217     };  
1218
```

1219 The operating system being weak to system would be declared as:

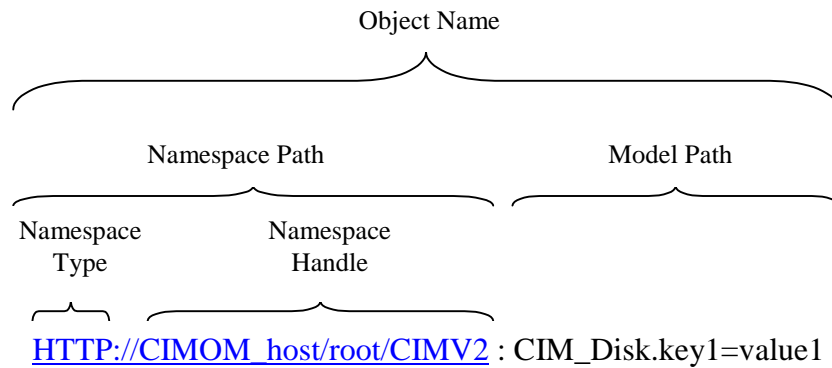
```
1220     Class Acme_OS  
1221     {  
1222         [key]  
1223         String Name;  
1224         [key, Propagated("CIM_System.Name")]  
1225         String SystemName;  
1226     };
```

1227 The user class being weak to operating system would be declared as:

```
1228     Class Acme_User  
1229     {  
1230         [key]  
1231         String uid;  
1232         [key, Propagated("Acme_OS.Name")]  
1233         String OSName;  
1234         [key, Propagated("Acme_OS.SystemName")]  
1235         String SystemName;  
1236     };
```

1237 **5.3 Naming CIM Objects**

1238 Since CIM allows for multiple implementations, it is not sufficient to think of the name of an object as just the
1239 combination of properties that have the KEY qualifier. The name must also identify the implementation that actually
1240 hosts the objects. The object name consists of the Namespace Path, which provides access to a CIM implementation,
1241 plus the Model Path, which provides full navigation within the CIM schema. The namespace path is used to locate a
1242 particular name space. The details of the namespace path are dependent on a particular implementation. The model
1243 path is the concatenation of the class name and the properties of the class that are qualified with the KEY qualifier.
1244 When the class is weak with respect to another class, the model path includes all key properties from the scoping
1245 objects. The following figure shows the various components of an object name. These are described in more details
1246 in the following sections. See the objectName non-terminal in Appendix A for the formal description of object
1247 name syntax.



1248

1249

Figure 5-5 Object Naming**1250 5.3.1 Namespace Path**

1251 A Namespace path references a namespace within an implementation that is capable of hosting CIM objects.

1252 A Namespace path resolves to a namespace hosted by a CIM-Capable implementation (in other words, a CIM
 1253 Object Manager). Unlike the Model Path, the details of the Namespace path are implementation-specific. Therefore,
 1254 the Namespace path provides two pieces of information: it identifies the type of implementation or namespace type,
 1255 and it provides a handle that references a particular implementation or namespace handle.

1256 5.3.1.1 Namespace Type

1257 The namespace type classifies or identifies the type of implementation. The provider of such an implementation is
 1258 responsible for describing the access protocol for that implementation. This is analogous to specifying http or ftp in
 1259 a browser.

1260 Fundamentally, a namespace type implies an access protocol or API set that can be used to manipulate objects.
 1261 These APIs would typically support: (1) generating a MOF file for a particular scope of classes and associations,
 1262 (2) importing a MOF file and (3) manipulating instances. A particular management platform may have a variety of
 1263 ways to access management information. Each of these ways must have a namespace type definition. Given this
 1264 type, there would be an assumed set of mechanisms for exporting, importing and updating instances.

1265 5.3.1.2 Namespace Handle

1266 The Namespace handle identifies a particular instance of the type of implementation. This handle must resolve to a
 1267 namespace within an implementation.

1268 The details of the handle are implementation-specific. It might be a simple string for an implementation that
 1269 supports one namespace, or it might be a hierarchical structure if an implementation supports multiple namespaces.
 1270 Either way, it resolves to a namespace.

1271 It is important to note that some implementations can support multiple namespaces. In this case, the implementation-
 1272 specific reference must resolve to a particular namespace within that implementation.

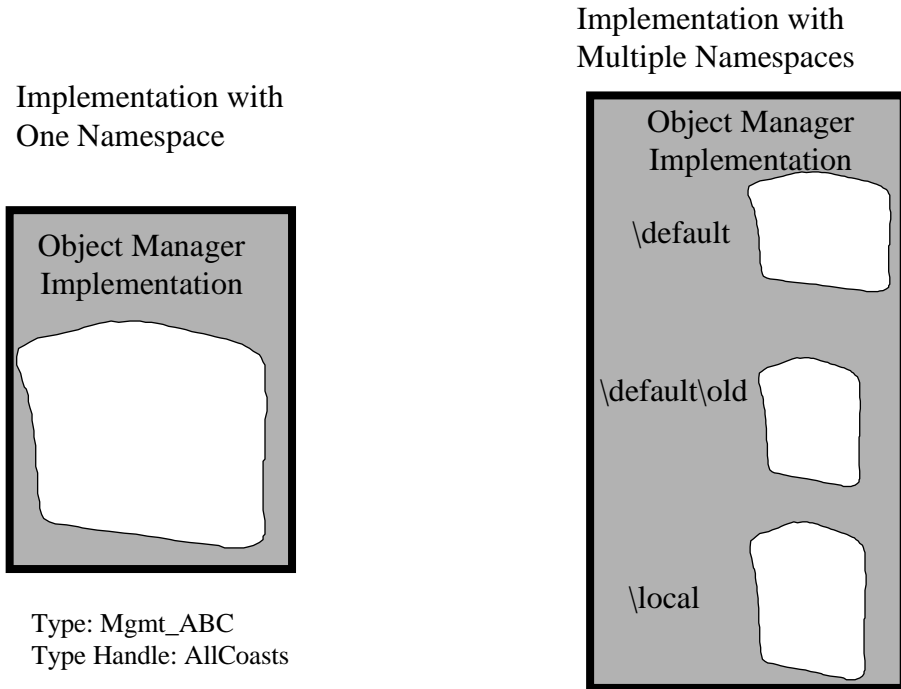


Figure 5-6 Namespaces

1273

1274

1275 There are two important observations to make:

1276 Namespaces can overlap with respect to their contents.

1277 An object in one name space, which has the same model path as an object in another name, space does not guarantee
1278 that the objects are representing the same reality.

1279 5.3.2 Model Path

1280 The object name constructed as a scoping path through the CIM schema is referred to as a Model Path. In the case
1281 of a model path for an instance, the model path is a combination of the key properties names and values qualified by
1282 the class name. It is solely described by CIM elements and is absolutely implementation-independent. It is used to
1283 describe the path to a particular object or to identify a particular object within a namespace. The name of any
1284 instance is a concatenation of named key property values, including all key values of its scoping objects.. When the
1285 class is weak with respect to another class, the model path includes all key properties from the scoping objects.

1286 The formal syntax of Model Path is provided in Appendix A.

1287 The syntax of Model Path is:

1288 <className>.<key1>=<value1>[,<keyx>=<valuex>]*

1289 5.3.3 Specifying the Object Name

1290 There are various mechanisms for specifying the object name details for any class instance or any association
1291 reference in a MOF file.

1292 The model path is specified for object and association differently. For objects (instances of classes), the model path
1293 is the combination of property value pairs that are marked with the KEY qualifier. So the model path for the
1294 following is: "ex_sampleClass.label1=9921,label2=8821". Since the order of the key properties is not significant, the
1295 model path could also be: "ex_sampleClass.label2=8821,label1=9921".

```

1296      Class ex_sampleClass
1297      {
1298          [key]
1299          uint32 label1;
1300          [key]
1301          string label2;
1302          uint32 size;
1303          uint32 weight;
1304      };
1305
1306      instance of ex_sampleClass
1307      {
1308          label1 = 9921;
1309          label2 = "SampleLabel";
1310          size = 80;
1311          weight = 45
1312      };
1313
1314      instance of ex_sampleClass
1315      {
1316          label1 = 0121;
1317          label2 = "Component";
1318          size = 80;
1319          weight = 45
1320      };

```

1321 For associations, a model path is used to specify the value of a reference in an INSTANCE OF statement for an
1322 association. In the following composedof-association example, the model path
1323 "ex_sampleClass.label1=9921,label2=8821" is used to reference an instance of the ex_sampleClass that is playing
1324 the role of a composer.

```

1325      [Association ]
1326      Class ex_composedof
1327      {
1328          [key] composer REF ex_sampleClass;
1329          [key] component REF ex_sampleClass;
1330      };
1331
1332      instance of ex_composedof
1333      {
1334          composer = "ex_sampleClass.label1=9921,label2=\"SampleLabel\"";
1335          component = "ex_sampleClass.label1=0121,label2=\"Component\"";
1336      }
1337

```

1338 An object path for the ex_composedof instance would be (note the handling of double quote characters):
1339 ex_composedof.composer="ex_sampleClass.label1=9921,label2=\"SampleLabel\"",
1340 component="ex_sampleClass.label1=0121,label2=\"Component\""

1341

1342 Even in the unusual case of a reference to an association, the object name is formed the same way:

```

1343      [Association ]
1344      Class ex_moreComposed
1345      {
1346          composedof REF ex_composedof;
1347          . . .
1348      };
1349
1350      instance of ex_moreComposed
1351      {
1352          composedof =
1353          "ex_composedof.composer=\"ex_sampleClass.label1=9921,label2=\\\\"SampleLabel
1354          \\\\"\",component=\"ex_sampleClass.label1=0121,label2=\\\\"Component\\\\"\"";
1355          . . .
1356      };

```

1357 The object name can be used as the value for object references and for object queries.

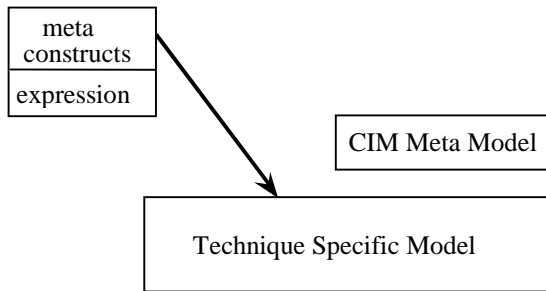
1358 **6 Mapping Existing Models Into CIM**

1359 Existing models have their own meta model and model. There are three types of mapping that can occur between
 1360 meta schemas: technique, recast and domain. Each of these mappings can be applied when converting from MIF
 1361 syntax to MOF syntax.

1362 **6.1 Technique Mapping**

1363 A technique mapping provides a mapping that uses the CIM meta-model constructs to describe the source modeling
 1364 technique’s meta constructs (for example, MIF, GDMO and SMI). Essentially, the CIM meta model is a meta meta-
 1365 model for the source technique.

1366

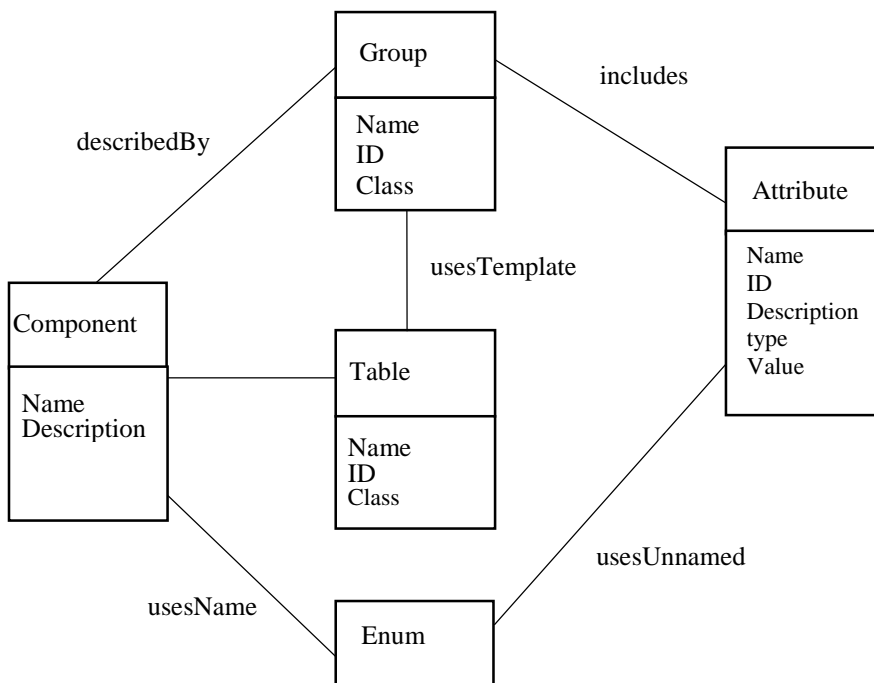


1367

1368 **Figure 6-1 Technique Mapping Example**

1369 The DMTF uses the management information format (MIF) as the meta model to describe distributed management
 1370 information in a common way. Therefore, it is meaningful to describe a technique mapping in which the CIM meta
 1371 model is used to describe the MIF syntax.

1372 The mapping presented here takes the important types that can appear in a MIF file and then creates classes for
 1373 them. Thus, component, group, attribute, table and enum are expressed in the CIM meta model as classes. In
 1374 addition, associations are defined to document how these are combined. Figure 6-2 illustrates the results:



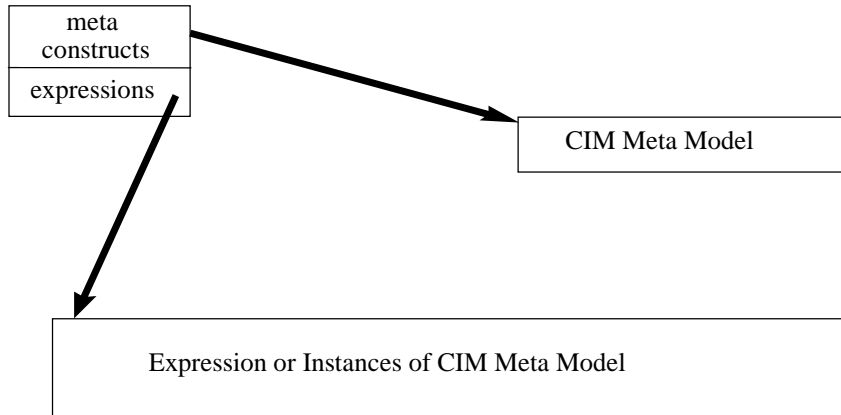
1375

1376

Figure 6-2 MIF Technique Mapping Example**6.2 Recast Mapping**

1377

1378 A recast mapping provides a mapping of the sources' meta constructs into the targeted meta constructs, so that a
 1379 model expressed in the source can be translated into the target. The major design work is to develop a mapping
 1380 between the sources' meta model and the CIM meta model. Once this is done, the source expressions are recast.



1381

1382

Figure 6-3 Recast mapping

1383 This is an example of a recast mapping for MIF, assuming:

```

1384     DMI attributes -> CIM properties
1385     DMI key attributes -> CIM key properties
1386     DMI groups -> CIM classes
1387     DMI components -> CIM classes
  
```

1388 The standard DMI ComponentID group might be recast into a corresponding CIM class:

```

1389     Start Group
1390     Name = "ComponentID"
1391     Class = "DMTF|ComponentID|001"
1392     ID = 1
1393     Description = "This group defines the attributes common to all "
1394                 "components. This group is required."
1395     Start Attribute
1396         Name = "Manufacturer"
1397         ID = 1
1398         Description = "Manufacturer of this system."
1399         Access = Read-Only
1400         Storage = Common
1401         Type = DisplayString(64)
1402         Value = ""
1403     End Attribute
1404     Start Attribute
1405         Name = "Product"
1406         ID = 2
1407         Description = "Product name for this system."
1408         Access = Read-Only
1409         Storage = Common
1410         Type = DisplayString(64)
1411         Value = ""
1412     End Attribute
1413     Start Attribute
1414         Name = "Version"
1415         ID = 3
1416         Description = "Version number of this system."
1417         Access = Read-Only
  
```

```
1418         Storage = Specific
1419         Type = DisplayString(64)
1420         Value = ""
1421     End Attribute
1422     Start Attribute
1423         Name = "Serial Number"
1424         ID = 4
1425         Description = "Serial number for this system."
1426         Access = Read-Only
1427         Storage = Specific
1428         Type = DisplayString(64)
1429         Value = ""
1430     End Attribute
1431     Start Attribute
1432         Name = "Installation"
1433         ID = 5
1434         Description = "Component installation time and date."
1435         Access = Read-Only
1436         Storage = Specific
1437         Type = Date
1438         Value = ""
1439     End Attribute
1440     Start Attribute
1441         Name = "Verify"
1442         ID = 6
1443         Description = "A code that provides a level of verification that the "
1444                     "component is still installed and working."
1445         Access = Read-Only
1446         Storage = Common
1447         Type = Start ENUM
1448             0 = "An error occurred; check status code."
1449             1 = "This component does not exist."
1450             2 = "Verification is not supported."
1451             3 = "Reserved."
1452             4 = "This component exists, but the functionality is untested."
1453             5 = "This component exists, but the functionality is unknown."
1454             6 = "This component exists, and is not functioning correctly."
1455             7 = "This component exists, and is functioning correctly."
1456         End ENUM
1457         Value = 1
1458     End Attribute
1459 End Group
```

1460 A corresponding CIM class might be the following. Note that properties in the example include an ID qualifier to
1461 represent the corresponding DMI attribute's ID. Here, a user-defined qualifier may be necessary.

```

1462     [Name ("ComponentID"), ID (1), Description (
1463         "This group defines the attributes common to all components. "
1464         "This group is required.")]
1465
1466     class DMTF|ComponentID|001 {
1467         [ID (1), Description ("Manufacturer of this system."), maxlen (64)]
1468         string Manufacturer;
1469         [ID (2), Description ("Product name for this system."), maxlen (64)]
1470         string Product;
1471         [ID (3), Description ("Version number of this system."), maxlen (64)]
1472         string Version;
1473         [ID (4), Description ("Serial number for this system."), maxlen (64)]
1474         string Serial_Number;
1475         [ID (5), Description("Component installation time and date.")]
1476         datetime Installation;
1477         [ID (6), Description("A code that provides a level of verification "
1478             "that the component is still installed and working."),
1479             Value (1)]
1480         string Verify;
1481     };

```

1482 6.3 Domain Mapping

1483 A domain mapping takes a source expressed in a particular technique and maps its content into either the core or
 1484 common models, or extension sub-schemas of the CIM. This mapping does not rely heavily on a meta-to-meta
 1485 mapping; it is primarily a content-to-content mapping. In one case, the mapping is actually a re-expression of
 1486 content in a more common way using a more expressive technique.

1487 This is an example of how CIM properties can be supplied by DMI, using information from the DMI disks group
 1488 ("DMTF|Disks|002"). For a hypothetical CIM disk class, the CIM properties are expressed as:

CIM "Disk" property	Can be sourced from DMI group/attribute
StorageType	"MIF.DMTF Disks 002.1"
StorageInterface	"MIF.DMTF Disks 002.3"
RemovableDrive	"MIF.DMTF Disks 002.6"
RemovableMedia	"MIF.DMTF Disks 002.7"
DiskSize	"MIF.DMTF Disks 002.16"

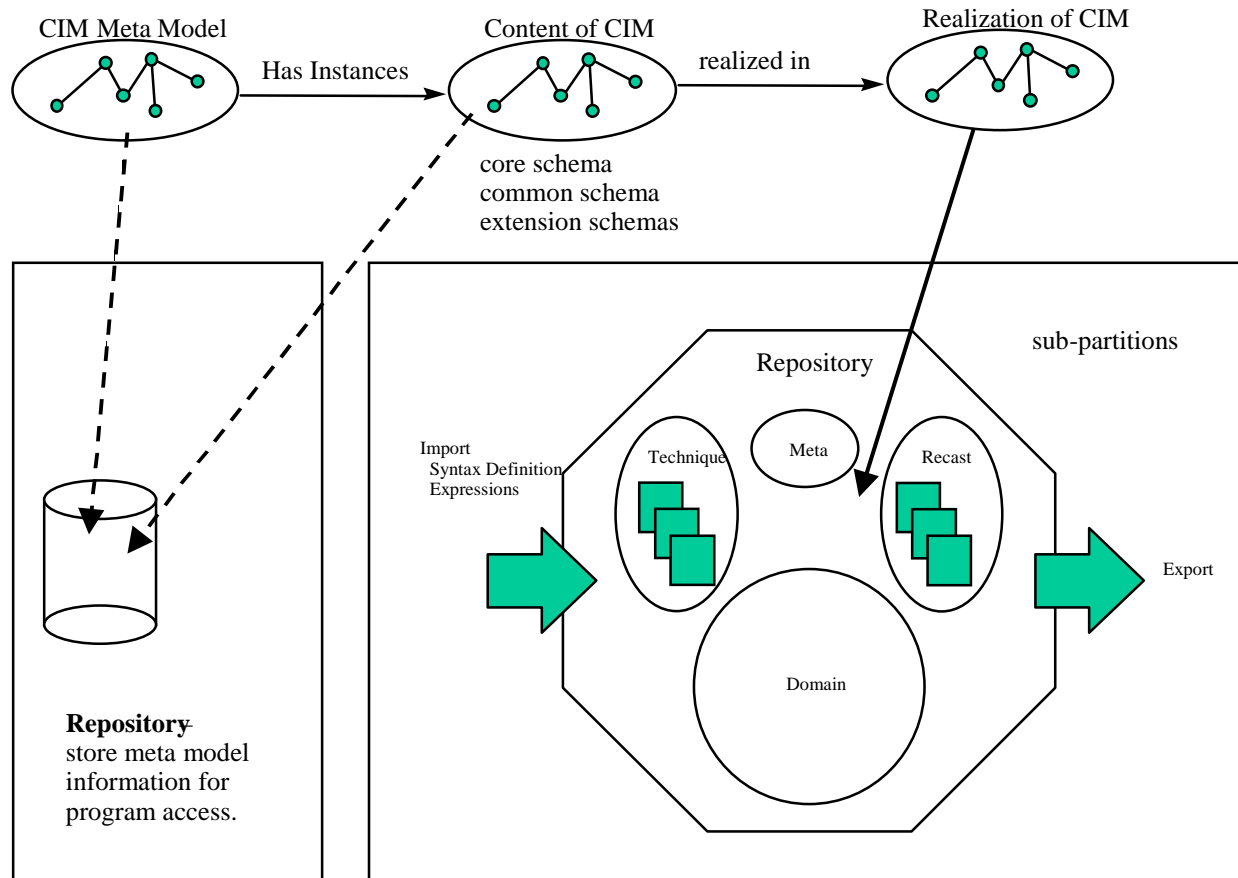
1490 6.4 Mapping Scratch Pads

1491 In general, when the content of models are mapped between different meta schemas, information gets lost or is
 1492 missing. To fill this gap, "scratch pads" are expressed in the CIM meta model using qualifiers, which are actually
 1493 extensions to the meta model (for example, see section 2.5.5 Mapping MIF Attributes and section 2.5.6 Mapping
 1494 Generic Data to CIM Properties). These scratch pads are critical to the exchange of core, common and extension
 1495 model content with the various technologies used to build management applications.

1496 7 Repository Perspective

1497 This section provides a basic description of a repository and a complete picture of the potential exploitation of it. A
 1498 repository stores definitional and/or structural information, and includes the capability to extract the definitions in a
 1499 form that is useful to application developers. Some repositories allow the definitions to be imported into and
 1500 exported from the repository in multiple forms. The notions of importing and exporting definition can be refined so
 1501 that they distinguish between three types of mappings.

1502 Using the mapping definitions in Section 6, the repository can be organized into the four partitions (meta, technique,
 1503 recast and domain).



1504
 1505

1506

Figure 7-1 Repository Partitions

1507 The repository partitions have the following characteristics:

- 1508 • Each partition is discrete. The meta partition refers to the definitions of the CIM meta model. The
 1509 technique partition refers to definitions that are loaded using technique mappings. The recast
 1510 partition refers to definitions that are loaded using recast mappings. The domain partition refers to
 1511 the definitions that are associated with the core and common models, and Extension schemas.
- 1512 • The technique and recast partitions can be organized into multiple sub-partitions in order to
 1513 capture each source uniquely. For example, there would be a technique sub-partition for each
 1514 unique meta language encountered (that is, one for MIF, GDMO, SMI, and so on). In the re-cast
 1515 partition, there would be a sub-partition for each meta language.

- 1516 • The act of importing the content of an existing source can result in entries in the recast or domain
1517 partition.

1518 **7.1 DMTF MIF Mapping Strategies**

1519 Assume the meta-model definition and the baseline for the CIM schema are complete. The next step is to map
1520 another source of management information (such as standard groups) into the repository. The primary objective is
1521 to do the work required to import one or more of the standard group(s).

1522 The possible import scenarios for a DMTF standard group are:

1523 *To Technique Partition:* Create a technique mapping for the MIF syntax. This mapping would be the same for all
1524 standard groups and would only need to be updated if the MIF syntax changed.

1525 *To Recast Partition:* Create a recast mapping from a particular standard group into a sub-partition of the recast
1526 partition. This mapping would allow the entire contents of the selected group to be loaded into a sub-partition of the
1527 recast partition. The same algorithm can be used to map additional standard groups into that same sub-partition.

1528 *To Domain Partition:* Create a Domain Mapping for the content of a particular standard group that overlaps with the
1529 content of the CIM schema.

1530 *To Domain Partition:* Create a Domain Mapping for the content of a particular standard group that does not overlap
1531 with CIM's schema into an extension sub-schema.

1532 *To Domain Partition:* Propose extensions to the content of the CIM schema and then perform Steps 3 and/or 4.

1533 Any combination of these five scenarios can be initiated by a team that is responsible for mapping an existing source
1534 into the CIM repository. There are many other details that must be addressed as the content of any of the sources
1535 changes and/or when the core or common model changes.

1536 Assuming numerous existing sources have been imported using all the import scenarios, now look at the export side.
1537 Ignoring the technique partition, the possible scenarios are:

1538 *From Recast Partition:* Create a recast mapping for a sub-partition in the recast partition to a standard group (that is,
1539 inverse of import 2). The desired method would be to use the recast mapping to translate a standard group into a
1540 GDMO definition.

1541 *From Recast Partition:* Create a Domain Mapping for one of the recast sub-partitions to a known management
1542 model that was not the original source for the content that overlaps.


1543 *From Domain Partition:* Create a recast mapping for the complete content of the CIM to a selected technique (for
1544 MIF, this results in a non-standard group).

1545 *From Domain Partition:* Create a Domain Mapping for the content of the CIM schema that overlaps with the
1546 content of an existing management model

1547 *From Domain Partition:* Create a Domain Mapping for the entire content of the CIM schema to an existing
1548 management model with the necessary extensions.

1549 **7.2 Recording Mapping Decisions**

1550 In order to understand the role of the scratch pad (see Section 6.4) in the repository, it is necessary to look at the
1551 import and export scenarios for the different partitions in the repository (technique, recast and application). These
1552 mappings can be organized into two categories: homogeneous and heterogeneous. The homogeneous category
1553 includes the mapping where the imported syntax and expressions are the same as the exported (for example,
1554 software MIF in and software MIF out). The heterogeneous category addresses the mappings where the imported
1555 syntax and expressions are different from the exported (for example, MIF in and GDMO out). For the homogenous
1556 category, the information can be recorded by creating qualifiers during an import operation so the content can be
1557 exported properly. For the heterogeneous category, the qualifiers must be added after the content is loaded into a
1558 partition of the repository. Figure 7-2, shows the X schema imported into the Y schema, and then being
1559 homogeneously exported into X or heterogeneously exported into Z. Each of the export arrows works with a
1560 different scratch pad.




1561

1562

Figure 7-2 Homogeneous and Heterogeneous Export

1563 The definition of the heterogeneous category is actually based on knowing how a schema was loaded into the
 1564 repository. A more general way of looking at this is to think of the export process using one of multiple scratch pads.
 1565 One of the scratch pads was created when the schema was loaded, and the others were added to handle mappings to
 1566 schema techniques other than the import source (Figure 7-3).



1567

1568

Figure 7-3 Scratch Pads and Mapping

1569 Figure 7-3 shows how the scratch pads of qualifiers are used without factoring in the unique aspects of each of the
 1570 partitions (technique, recast, applications) within the CIM repository. The next step is to put this discussion in the
 1571 context of these partitions.

1572 For the technique partition, there is no need for a scratch pad since the CIM meta model is used to describe the
1573 constructs used in the source meta schema. Therefore, by definition, there is one homogeneous mapping for each
1574 meta schema covered by the technique partition. These mappings create CIM objects for the syntactical constructs of
1575 the schema and create associations for the ways they can be combined (for example, MIF groups include attributes).

1576 For the recast partition, there are multiple scratch pads for each of the sub-partitions, since one is required for each
1577 export target and there can be multiple mapping algorithms for each target. The latter occurs because part of creating
1578 a recast mapping involves mapping the constructs of the source into CIM meta-model constructs. Therefore, for the
1579 MIF syntax, a mapping must be created for component, group, attribute, and so on, into appropriate CIM meta-
1580 model constructs like object, association, property, and so on. These mappings can be arbitrary. As a specific
1581 example, one of the decisions that must be made is whether a group or a component maps into an object. It would be
1582 possible to have two different recast mapping algorithms, one that mapped groups into objects with qualifiers that
1583 preserved the component, and one that mapped components into objects with qualifiers that preserved the group
1584 name for the properties. Therefore, the scratch pads in the recast partition are organized by target technique and
1585 employed algorithm.

1586 For the domain partitions, there are two types of mappings. The first is similar to the recast partition in that some
1587 portion of the domain partition is mapped into the syntax of another meta schema. These mappings can use the same
1588 qualifier scratch pads and associated algorithms that are developed for the recast partition. The second type of
1589 mapping facilitates documenting the content overlap between the domain partition and some other model (for
1590 example, software groups). These mappings cannot be determined in a generic way at import time; therefore, it is
1591 best to consider them in the context of exporting. The mapping uses filters to determine the overlaps and then
1592 performs the necessary conversions. The filtering can be done using qualifiers that indicate a particular set of
1593 domain partition constructs map into some combination of constructs in the target/source model. The conversions
1594 would be documented in the repository using a complex set of qualifiers that capture how to write or insert the
1595 overlapped content into the target model. The mapping qualifiers for the domain partition would be organized like
1596 the recasting partition for the syntax conversions, and there would be scratch pads for each of the models for
1597 documenting overlapping content.

1598 In summary, pick the partition, develop a mapping, and identify the qualifiers necessary to capture potentially lost
1599 information when developing mapping details for a particular source. On the export side, the mapping algorithm
1600 checks to see if the content to be exported includes the necessary qualifiers for the logic to work.

1601 **Appendix A MOF Syntax Grammar Description**

1602 This section contains the grammar for MOF syntax. While the grammar presented here is convenient for describing
1603 the MOF syntax clearly, it should be noted that the same MOF language can also be described by a different, LL(1)-
1604 parseable, grammar. This has been done to allow low-footprint implementations of MOF compilers.

1605 In addition, note these points:

- 1606 1. An empty property list is equivalent to "*".
- 1607 2. All keywords are case-insensitive.
- 1608 3. The IDENTIFIER type is used for names of classes, properties, qualifiers, methods and
1609 namespaces; the rules governing the naming of classes and properties are to be found in section 1
1610 of Appendix F.
- 1611 4. A string Value may contain quote (") characters, provided that each is immediately preceded by a
1612 backslash (\) character.
- 1613 5. In the current release, the MOF BNF does not support initializing an array value to empty (i.e., an
1614 array with no elements). In the 3.0 version of this specification, it is the intention of the DMTF to
1615 extend the MOF BNF to support this functionality as follows:

1616
1617 arrayInitialize = "{ [arrayElementList] }

1618
1619 arrayElementList = constantValue *("," constantValue)

1620
1621 In order to ensure interoperability with V2.x implementations, the DMTF recommends that, where
1622 possible, the value of NULL rather than empty (i.e., "{}") be used to represent the most common
1623 use cases. However, where this may result in confusion or other issues, implementations MAY
1624 use the above v3.0 syntax to initialize an empty array.

1625

```

mofSpecification      = *mofProduction

mofProduction        = compilerDirective |
                      classDeclaration  |
                      assocDeclaration  |
                      indicDeclaration  |
                      qualifierDeclaration |
                      instanceDeclaration

compilerDirective     = PRAGMA pragmaName "(" pragmaParameter ")"

pragmaName            = IDENTIFIER

pragmaParameter       = stringValue

classDeclaration      = [ qualifierList ]
                      CLASS className [ superClass ]
                      "{" *classFeature "}" ";"

assocDeclaration      = "[" ASSOCIATION *( "," qualifier ) "]"
                      CLASS className [ superClass ]
                      "{" *associationFeature "}" ";"

                      // Context:
                      // The remaining qualifier list must not include
                      // theASSOCIATION qualifier again. If the
                      // association has no super association, then at
                      // least two references must be specified! The
                      // ASSOCIATION qualifier may be omitted in
                      // sub-associations.

indicDeclaration      = "[" INDICATION *( "," qualifier ) "]"
                      CLASS className [ superClass ]
                      "{" *classFeature "}" ";"

className             = schemaName "_" IDENTIFIER // NO whitespace !

                      // Context:
                      // Schema name must not include "_" !

alias                 = AS aliasIdentifier

aliasIdentifier        = "$" IDENTIFIER // NO whitespace !

superClass            = ":" className

classFeature           = propertyDeclaration | methodDeclaration

associationFeature     = classFeature | referenceDeclaration

qualifierList          = "[" qualifier *( "," qualifier ) "]"

qualifier              = qualifierName [ qualifierParameter ] [ ":" 1*flavor ]

qualifierParameter     = "(" constantValue ")" | arrayInitializer

flavor                = ENABLEOVERRIDE | DISABLEOVERRIDE | RESTRICTED |
                      TOSUBCLASS | TRANSLATABLE

propertyDeclaration    = [ qualifierList ] dataType propertyName
                      [ array ] [ defaultValue ] ";"

referenceDeclaration    = [ qualifierList ] objectRef referenceName

```

```

        [ defaultValue ] ";"

methodDeclaration    = [ qualifierList ] dataType methodName
                      "(" [ parameterList ] ")" ";"

propertyName        = IDENTIFIER

referenceName        = IDENTIFIER

methodName           = IDENTIFIER

dataType             = DT_UINT8 | DT_SINT8 | DT_UINT16 | DT_SINT16 |
                      DT_UINT32 | DT_SINT32 | DT_UINT64 | DT_SINT64 |
                      DT_REAL32 | DT_REAL64 | DT_CHAR16 |
                      DT_STR | DT_BOOL | DT_DATETIME

objectRef            = className REF

parameterList        = parameter *( "," parameter )

parameter            = [ qualifierList ] (dataType|objectRef) parameterName
                      [ array ]

parameterName        = IDENTIFIER

array                = "[" [positiveDecimalValue] "]"

positiveDecimalValue = positiveDecimalDigit *decimalDigit

defaultValue         = "=" initializer

initializer          = ConstantValue | arrayInitializer | referenceInitializer

arrayInitializer     = "{" constantValue*( "," constantValue)"}"

constantValue        = integerValue | realValue | charValue | stringValue |
                      booleanValue | nullValue

integerValue         = binaryValue | octalValue | decimalValue | hexValue

referenceInitializer = objectHandle | aliasIdentifier

objectHandle         = stringValue
                      // the(unescape)contents of which must form an
                      // objectName; see examples

objectName           = [ namespacePath ":" ] modelPath

namespacePath        = [ namespaceType "://" ] namespaceHandle

namespaceType        = One or more UCS-2 characters NOT including the sequence
                      "://"

namespaceHandle      = One or more UCS-2 character, possibly including ":"
                      // Note that modelPath may also contain ":" characters
                      // within quotes; some care is required to parse
                      // objectNames.

modelPath            = className "." keyValuePairList
                      // Note: className alone represents a path to a class,
                      // rather than an instance

keyValuePairList     = keyValuePair *( "," keyValuePair )

```

```

keyValuePair          = ( propertyName "=" constantValue ) | ( referenceName "="
                        objectHandle )

qualifierDeclaration  = QUALIFIER qualifierName qualifierType scope
                        [ defaultFlavor ] ";"

qualifierName         = IDENTIFIER

qualifierType         = ":" dataType [ array ] [ defaultValue ]

scope                 = "," SCOPE "(" metaElement *( "," metaElement ) ")"

metaElement           = CLASS | ASSOCIATION | INDICATION | QUALIFIER
                        PROPERTY | REFERENCE | METHOD | PARAMETER | ANY

defaultFlavor         = "," FLAVOR "(" flavor *( "," flavor ) ")"

instanceDeclaration   = [ qualifierList ] INSTANCE OF className [ alias ]
                        "{" 1*valueInitializer "}" ";"

valueInitializer      = [ qualifierList ]
                        ( propertyName | referenceName ) "=" initializer ";"
    
```

1626 These productions do not allow whitespace between the terms:

```

schemaName      = IDENTIFIER
                  // Context:
                  // Schema name must not include "_" !
fileName        = stringValue
binaryValue     = [ "+" | "-" ] 1*binaryDigit ( "b" | "B" )
binaryDigit     = "0" | "1"
octalValue      = [ "+" | "-" ] "0" 1*octalDigit
octalDigit      = "0" | "1" | "2" | "3" | "4" | "5" | "6" | "7"
decimalValue    = [ "+" | "-" ] ( positiveDecimalDigit *decimalDigit | "0" )
decimalDigit    = "0" | positiveDecimalDigit
positiveDecimalDigit = "1" | "2" | "3" | "4" | "5" | "6" | "7" | "8" | "9"
hexValue        = [ "+" | "-" ] ( "0x" | "0X" ) 1*hexDigit
hexDigit        = decimalDigit | "a" | "A" | "b" | "B" | "c" | "C" |
                  "d" | "D" | "e" | "E" | "f" | "F"
realValue       = [ "+" | "-" ] *decimalDigit "." 1*decimalDigit
                  [ ( "e" | "E" ) [ "+" | "-" ] 1*decimalDigit ]
charValue       = // any single-quoted Unicode-character, except
                  // single quotes
stringValue     = 1*( "" *stringChar "" )
stringChar      = "\" "" | // encoding for double-quote
                  "\" \" | // encoding for backslash
                  any UCS-2 character but "" or "\"
booleanValue    = TRUE | FALSE
nullValue       = NULL

```

1627 The remaining productions are case-insensitive keywords:

ANY	=	"any"
AS	=	"as"
ASSOCIATION	=	"association"
CLASS	=	"class"
DISABLEOVERRIDE	=	"disableOverride"
DT_BOOL	=	"boolean"
DT_CHAR16	=	"char16"
DT_DATETIME	=	"datetime"
DT_REAL32	=	"real32"
DT_REAL64	=	"real64"
DT_SINT16	=	"sint16"
DT_SINT32	=	"sint32"
DT_SINT64	=	"sint64"
DT_SINT8	=	"sint8"
DT_STR	=	"string"
DT_UINT16	=	"uint16"
DT_UINT32	=	"uint32"
DT_UINT64	=	"uint64"
DT_UINT8	=	"uint8"
ENABLEOVERRIDE	=	"enableoverride"
FALSE	=	"false"
FLAVOR	=	"flavor"
INDICATION	=	"indication"
INSTANCE	=	"instance"
METHOD	=	"method"
NULL	=	"null"
OF	=	"of"
PARAMETER	=	"parameter"
PRAGMA	=	"#pragma"
PROPERTY	=	"property"
QUALIFIER	=	"qualifier"
REF	=	"ref"
REFERENCE	=	"reference"
RESTRICTED	=	"restricted"
SCHEMA	=	"schema"
SCOPE	=	"scope"
TOSUBCLASS	=	"tosubclass"
TRANSLATABLE	=	"translatable"
TRUE	=	"true"

1628 **Appendix B CIM META SCHEMA**

```

1629 // =====
1630 //   NamedElement
1631 // =====
1632     [Version("2.3.0"), Description(
1633         "The Meta_NamedElement class represents the root class for the "
1634         "Metaschema. It has one property: Name, which is inherited by all the "
1635         "non-association classes in the Metaschema. Every metaconstruct is "
1636         "expressed as a descendent of the class Meta_Named Element.") ]
1637 class Meta_NamedElement
1638 {
1639     [Description (
1640         "The Name property indicates the name of the current Metaschema element. "
1641         "The following rules apply to the Name property, depending on the "
1642         "creation type of the object:<UL><LI>Fully-qualified class names, such "
1643         "as those prefixed by the schema name, are unique within the schema."
1644         "<LI>Fully-qualified association and indication names are unique within "
1645         "the schema (implied by the fact that association and indication classes "
1646         "are subtypes of Meta_Class).<LI>Implicitly-defined qualifier names are "
1647         "unique within the scope of the characterized object; that is, a named "
1648         "element may not have two characteristics with the same name."
1649         "<LI>Explicitly-defined qualifier names are unique within the defining "
1650         "schema. An implicitly-defined qualifier must agree in type, scope and "
1651         "flavor with any explicitly-defined qualifier of the same name."
1652         "<LI>Trigger names must be unique within the property, class or method "
1653         "to which the trigger applies.<LI>Method and property names must be "
1654         "unique within the domain class. A class can inherit more than one "
1655         "property or method with the same name. Property and method names can be "
1656         "qualified using the name of the declaring class.<LI>Reference names "
1657         "must be unique within the scope of their defining association class. "
1658         "Reference names obey the same rules as property names.</UL><B>Note:</B> "
1659         "Reference names are not required to be unique within the scope of the "
1660         "related class. Within such a scope, the reference provides the name of "
1661         "the class within the context defined by the association.") ]
1662     string Name;
1663 };
1664
1665 // =====
1666 //   QualifierFlavor
1667 // =====
1668     [Version("2.3.0"), Description (
1669         "The Meta_QualifierFlavor class encapsulates extra semantics attached "
1670         "to a qualifier such as the rules for transmission from superClass "
1671         "to subClass and whether or not the qualifier value may be translated "
1672         "into other languages") ]
1673 class Meta_QualifierFlavor:Meta_NamedElement
1674 {
1675 };
1676
1677 // =====
1678 //   Schema
1679 // =====
1680     [Version("2.3.0"), Description (
1681         "The Meta_Schema class represents a group of classes with a single owner."
1682         " Schemas are used for administration and class naming. Class names must "
1683         "be unique within their owning schemas.") ]
1684 class Meta_Schema:Meta_NamedElement
1685 {
1686 };
1687

```

```
1688 // =====
1689 //   Trigger
1690 // =====
1691     [Version("2.3.0"), Description (
1692       "A Trigger is a recognition of a state change (such as create, delete, "
1693       "update, or access) of a Class instance, and update or access of a "
1694       "Property.") ]
1695 class Meta_Trigger:Meta_NamedElement
1696 {
1697 };
1698
1699 // =====
1700 //   Qualifier
1701 // =====
1702     [Version("2.3.0"), Description (
1703       "The Meta_Qualifier class represents characteristics of named elements. "
1704       "For example, there are qualifiers that define the characteristics of a "
1705       "property or the key of a class. Qualifiers provide a mechanism that "
1706       "makes the Metaschema extensible in a limited and controlled fashion."
1707       "<P>It is possible to add new types of qualifiers by the introduction of "
1708       "a new qualifier name, thereby providing new types of metadata to "
1709       "processes that manage and manipulate classes, properties, and other "
1710       "elements of the Metaschema.") ]
1711 class Meta_Qualifier:Meta_NamedElement
1712 {
1713     [Description ("The Value property indicates the value of the qualifier.")]
1714     string Value;
1715 };
1716
1717 // =====
1718 //   Method
1719 // =====
1720     [Version( "2" ), Revision( "2" ), Description (
1721       "The Meta_Method class represents a declaration of a signature; that is, "
1722       "the method name, return type and parameters, and (in the case of a "
1723       "concrete class) may imply an implementation.") ]
1724 class Meta_Method:Meta_NamedElement
1725 {
1726 };
1727
1728 // =====
1729 //   Property
1730 // =====
1731     [Version( "2" ), Revision( "2" ), Description (
1732       "The Meta_Property class represents a value used to characterize "
1733       "instances of a class. A property can be thought of as a pair of Get and "
1734       "Set functions that, when applied to an object, return state and set "
1735       "state, respectively.") ]
1736 class Meta_Property:Meta_NamedElement
1737 {
1738 };
1739
```

```
1740 // =====
1741 // Reference
1742 // =====
1743 [Version( "2" ), Revision( "2" ), Description (
1744     "The Meta_Reference class represents (and defines) the role each object "
1745     "plays in an association. The reference represents the role name of a "
1746     "class in the context of an association, which supports the provision of "
1747     "multiple relationship instances for a given object. For example, a "
1748     "system can be related to many system components.") ]
1749 class Meta_Reference:Meta_Property
1750 {
1751 };
1752
1753 // =====
1754 // Class
1755 // =====
1756 [Version( "2" ), Revision( "2" ), Description (
1757     "The Meta_Class class is a collection of instances that support the same "
1758     "type; that is, the same properties and methods. Classes can be arranged "
1759     "in a generalization hierarchy that represents subtype relationships "
1760     "between classes.<P>The generalization hierarchy is a rooted, directed "
1761     "graph and does not support multiple inheritance. Classes can have "
1762     "methods, which represent the behavior relevant for that class. A Class "
1763     "may participate in associations by being the target of one of the "
1764     "references owned by the association.") ]
1765 class Meta_Class:Meta_NamedElement
1766 {
1767 };
1768
1769 // =====
1770 // Indication
1771 // =====
1772 [Version( "2" ), Revision( "2" ), Description (
1773     "The Meta_Indication class represents an object created as a result of a "
1774     "trigger. Because Indications are subtypes of Meta_Class, they can have "
1775     "properties and methods, and be arranged in a type hierarchy. ") ]
1776 class Meta_Indication:Meta_Class
1777 {
1778 };
1779
1780 // =====
1781 // Association
1782 // =====
1783 [Version( "2" ), Revision( "2" ), Description (
1784     "The Meta_Association class represents a class that contains two or more "
1785     "references and represents a relationship between two or more objects. "
1786     "Because of how associations are defined, it is possible to establish a "
1787     "relationship between classes without affecting any of the related "
1788     "classes.<P>For example, the addition of an association does not affect "
1789     "the interface of the related classes; associations have no other "
1790     "significance. Only associations can have references. Associations can "
1791     "be a subclass of a non-association class . Any subclass of "
1792     "Meta_Association is an association.") ]
1793 class Meta_Association:Meta_Class
1794 {
1795 };
1796
```

```
1797 // =====
1798 //   Characteristics
1799 // =====
1800     [Association, Version( "2" ), Revision( "2" ), Aggregation, Description (
1801       "The Meta_Characteristics class relates a Meta_NamedElement to a "
1802       "qualifier that characterizes the named element. Meta_NamedElement may "
1803       "have zero or more characteristics." ) ]
1804 class Meta_Characteristics
1805 {
1806     [Description (
1807       "The Characteristic reference represents the qualifier that "
1808       "characterizes the named element." ) ]
1809     Meta_Qualifier REF Characteristic;
1810     [Aggregate, Description (
1811       "The Characterized reference represents the named element that is being "
1812       "characterized." ) ]
1813     Meta_NamedElement REF Characterized;
1814 };
1815
1816 // =====
1817 //   PropertyDomain
1818 // =====
1819     [Association, Version( "2" ), Revision( "2" ), Aggregation, Description (
1820       "The Meta_PropertyDomain class represents an association between a class "
1821       "and a property.<P>A property has only one domain: the class that owns "
1822       "the property. A property can have an override relationship with another "
1823       "property from a different class. The domain of the overridden property "
1824       "must be a supertype of the domain of the overriding property. The "
1825       "domain of a reference must be an association." ) ]
1826 class Meta_PropertyDomain
1827 {
1828     [Description (
1829       "The Property reference represents the property that is owned by the "
1830       "class referenced by Domain." ) ]
1831     Meta_Property REF Property;
1832     [Aggregate, Description (
1833       "The Domain reference represents the class that owns the property "
1834       "referenced by Property." ) ]
1835     Meta_Class REF Domain;
1836 };
1837
1838 // =====
1839 //   MethodDomain
1840 // =====
1841     [Association, Version( "2" ), Revision( "2" ), Aggregation, Description (
1842       "The Meta_MethodDomain class represents an association between a class "
1843       "and a method.<P>A method has only one domain: the class that owns the "
1844       "method, which can have an override relationship with another method "
1845       "from a different class. The domain of the overridden method must be a "
1846       "supertype of the domain of the overriding method. The signature of the "
1847       "method (that is, the name, parameters and return type) must be "
1848       "identical." ) ]
1849 class Meta_MethodDomain
1850 {
1851     [Description (
1852       "The Method reference represents the method that is owned by the class "
1853       "referenced by Domain." ) ]
1854     Meta_Method REF Method;
1855     [Aggregate, Description (
1856       "The Domain reference represents the class that owns the method "
1857       "referenced by Method." ) ]
1858     Meta_Class REF Domain;
1859 };
```

```
1860
1861 // =====
1862 //     ReferenceRange
1863 // =====
1864     [Association, Version( "2" ), Revision( "2" ), Description (
1865         "The Meta_ReferenceRange class defines the type of the reference." ) ]
1866 class Meta_ReferenceRange
1867 {
1868     [Description (
1869         "The Reference reference represents the reference whose type is defined "
1870         "by Range." ) ]
1871     Meta_Reference REF Reference;
1872     [Description (
1873         "The Range reference represents the class that defines the type of "
1874         "reference." ) ]
1875     Meta_Class REF Range;
1876 };
1877
1878 // =====
1879 //     QualifiersFlavor
1880 // =====
1881     [Association, Version( "2" ), Revision( "2" ), Aggregation, Description (
1882         "The Meta_QualifiersFlavor class represents an association between a "
1883         "flavor and a qualifier." ) ]
1884 class Meta_QualifiersFlavor
1885 {
1886     [Description (
1887         "The Flavor reference represents the qualifier flavor to "
1888         "be applied to Qualifier." ) ]
1889     Meta_QualifierFlavor REF Flavor;
1890     [Aggregate, Description (
1891         "The Qualifier reference represents the qualifier to which "
1892         "Flavor applies." ) ]
1893     Meta_Qualifier REF Qualifier;
1894 };
1895
1896 // =====
1897 //     SubtypeSupertype
1898 // =====
1899     [Association, Version( "2" ), Revision( "2" ), Description (
1900         "The Meta_SubtypeSupertype class represents subtype/supertype "
1901         "relationships between classes arranged in a generalization hierarchy. "
1902         "This generalization hierarchy is a rooted, directed graph and does not "
1903         "support multiple inheritance." ) ]
1904 class Meta_SubtypeSupertype
1905 {
1906     [Description (
1907         "The SuperClass reference represents the class that is hierarchically "
1908         "immediately above the class referenced by SubClass." ) ]
1909     Meta_Class REF SuperClass;
1910     [Description (
1911         "The SubClass reference represents the class that is the immediate "
1912         "descendent of the class referenced by SuperClass." ) ]
1913     Meta_Class REF SubClass;
1914 };
1915
```

```
1916 // =====
1917 //   PropertyOverride
1918 // =====
1919     [Association, Version( "2" ), Revision( "2" ), Description (
1920       "The Meta_PropertyOverride class represents an association between two "
1921       "properties where one overrides the other.<P>Properties have reflexive "
1922       "associations that represent property overriding. A property can "
1923       "override an inherited property, which implies that any access to the "
1924       "inherited property will result in the invocation of the implementation "
1925       "of the overriding property. A Property can have an override "
1926       "relationship with another property from a different class.<P>The domain "
1927       "of the overridden property must be a supertype of the domain of the "
1928       "overriding property. The class referenced by the Meta_ReferenceRange "
1929       "association of an overriding reference must be the same as, or a "
1930       "subtype of, the class referenced by the Meta_ReferenceRange "
1931       "associations of the reference being overridden.") ]
1932 class Meta_PropertyOverride
1933 {
1934     [Description (
1935       "The OverridingProperty reference represents the property that overrides "
1936       "the property referenced by OverriddenProperty.") ]
1937     Meta_Property REF OverridingProperty;
1938     [Description (
1939       "The OverriddenProperty reference represents the property that is "
1940       "overridden by the property reference by OverridingProperty.") ]
1941     Meta_Property REF OverriddenProperty;
1942 };
1943
1944 // =====
1945 //   MethodOverride
1946 // =====
1947     [Association, Version( "2" ), Revision( "2" ), Description (
1948       "The Meta_MethodOverride class represents an association between two "
1949       "methods, where one overrides the other. Methods have reflexive "
1950       "associations that represent method overriding. A method can override an "
1951       "inherited method, which implies that any access to the inherited method "
1952       "will result in the invocation of the implementation of the overriding "
1953       "method.") ]
1954 class Meta_MethodOverride
1955 {
1956     [Description (
1957       "The OverridingMethod reference represents the method that overrides the "
1958       "method referenced by OverriddenMethod.") ]
1959     Meta_Method REF OverridingMethod;
1960     [Description (
1961       "The OverriddenMethod reference represents the method that is overridden "
1962       "by the method reference by OverridingMethod.") ]
1963     Meta_Method REF OverriddenMethod;
1964 };
1965
```

```
1966 // =====
1967 //   ElementSchema
1968 // =====
1969     [Association, Version( "2" ), Revision( "2" ), Aggregation, Description (
1970       "The Meta_ElementSchema class represents the elements (typically classes "
1971       "and qualifiers) that make up a schema." ) ]
1972 class Meta_ElementSchema
1973 {
1974     [Description (
1975       "The Element reference represents the named element that belongs to the "
1976       "schema referenced by Schema." ) ]
1977     Meta_NamedElement REF Element;
1978     [Aggregate, Description (
1979       "The Schema reference represents the schema to which the named element "
1980       "referenced by Element belongs." ) ]
1981     Meta_Schema REF Schema;
1982 };
```


1983 Appendix C Values for UNITS Qualifier

1984	The UNITS qualifier specifies the unit of measure in which the associated property is expressed. For example, a
1985	Size property might have Units ("bytes"). Currently recognized values are:
1986	
1987	Bits, KiloBits, MegaBits, GigaBits
1988	< Bits, KiloBits, MegaBits, GigaBits> per Second
1989	Bytes, KiloBytes, MegaBytes, GigaBytes, Words, DoubleWords, QuadWords
1990	Degrees C, Tenths of Degrees C, Hundredths of Degrees C, Degrees F, Tenths of Degrees F, Hundredths of Degrees
1991	F, Degrees K, Tenths of Degrees K, Hundredths of Degrees K, Color Temperature
1992	Volts, MilliVolts, Tenths of MilliVolts, Amps, MilliAmps, Tenths of MilliAmps, Watts, MilliWattHours
1993	Joules, Coulombs, Newtons
1994	Lumen, Lux, Candelas
1995	Pounds, Pounds per Square Inch
1996	Cycles, Revolutions, Revolutions per Minute, Revolutions per Second
1997	Minutes, Seconds, Tenths of Seconds, Hundredths of Seconds, MicroSeconds, MilliSeconds, NanoSeconds
1998	Hours, Days, Weeks
1999	Hertz, MegaHertz
2000	Pixels, Pixels per Inch
2001	Counts per Inch
2002	Percent, Tenths of Percent, Hundredths of Percent
2003	Meters, Centimeters, Millimeters, Cubic Meters, Cubic Centimeters, Cubic Millimeters
2004	Inches, Feet, Cubic Inches, Cubic Feet Ounces, Liters, Fluid Ounces
2005	Radians, Steradians, Degrees
2006	Gravities, Pounds, Foot-Pounds
2007	Gauss, Gilberts, Henrys, MilliHenrys, Farads, MilliFarads, MicroFarads, PicoFarads
2008	Ohms, Siemens
2009	Moles, Becquerels, Parts per Million
2010	Decibels, Tenths of Decibels
2011	Grays, Sieverts
2012	MilliWatts
2013	DBm
2014	<Bytes, KiloBytes, MegaBytes, GigaBytes> per Second
2015	BTU per Hour
2016	PCI clock cycles
2017	<Numeric value> <Minutes, Seconds, Tenths of Seconds, Hundreths of Seconds, MicroSeconds, MilliSeconds,
2018	Nanoseconds>

- 2019 Us³
- 2020 Amps at <Numeric Value> Volts
- 2021 Clock Ticks
- 2022 Packets

³ Standard Rack Measurement equal to 1.75 inches.

2023 **Appendix D UML Notation**

2024 The CIM meta-schema notation is based directly on the notation used in Unified Modeling Language (UML). There
 2025 are distinct symbols for all of the major constructs in the schema, with the exception of qualifiers (as opposed to
 2026 properties, which are directly represented in the diagrams).

2027 In UML, a class is represented by a rectangle. The class name either stands alone in the rectangle or is in the
 2028 uppermost segment of the rectangle. If present, the segment below the segment containing the name contains the
 2029 properties of the class. If present, a third region indicates the presence of methods.

2030 A line decorated with a triangle indicates an inheritance relationship; the lower rectangle represents a subtype of the
 2031 upper rectangle. The triangle points to the superclass.

2032 Other solid lines represent relationships. The cardinality of the references on either side of the relationship is
 2033 indicated by a decoration on either end. The following character combinations are commonly used:

2034 “1” indicates a single-valued, required reference

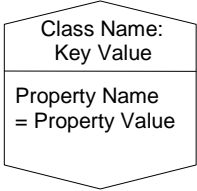
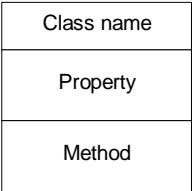
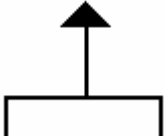
2035 “0..1” indicates an optional single-valued reference

2036 “*” indicates an optional many-valued reference (as does “0..*”)

2037 “1..*” indicates a required many-valued reference

2038 A line connected to a rectangle by a dotted line represents a subclass relationship between two associations.

2039 The diagramming notation and its interpretation are summarized in this table:

Meta element	Interpretation	Diagramming Notation
Object		
Primitive type	Text to the right of the colon in the center portion of the class icon	
Class		
Subclass		

Meta element	Interpretation	Diagramming Notation
Association	1:1 1:Many 1:zero or 1 Aggregation	
Association with properties	link-class with the link-class having the same name as the association, and using normal conventions for representing properties and methods.	
Association with subclass	A dashed line running from the sub association to the super class.	
Property	Middle section of the class icon is a list of the properties of the class.	
Reference	One end of the association line labeled with the name of the reference.	
Method	Lower section of the class icon is a list of the methods of the class.	

Meta element	Interpretation	Diagramming Notation
Overriding	No direct equivalent. Note: Use of the same name does not imply overriding.	
Indication	Message trace diagram in which vertical bars represent objects and horizontal lines represent messages.	
Trigger	State transition diagrams.	
Qualifier	No direct equivalent.	

2040

Appendix E Glossary

Aggregation	A strong form of an <i>association</i> . For example, the containment relationship between a system and the components that make up the system can be called an <i>aggregation</i> . An <i>aggregation</i> is expressed as a <i>Qualifier</i> on the <i>association</i> class. <i>Aggregation</i> often implies, but does not require, that the aggregated <i>objects</i> have mutual dependencies.
Association	A <i>class</i> that expresses the relationship between two other <i>classes</i> . The relationship is established by the presence of two or more <i>references</i> in the <i>association class</i> pointing to the related <i>classes</i> .
Cardinality	A relationship between two classes that allows more than one <i>object</i> to be related to a single <i>object</i> . For example, Microsoft Office* is made up of the software elements Word, Excel, Access and PowerPoint.
CIM	Common Information Model is the schema of the overall managed environment. It is divided into a <i>Core model</i> , <i>Common model</i> and <i>extended schemas</i> .
CIM Schema	The schema representing the <i>Core</i> and <i>Common models</i> . Versions of this schema will be released by the DMTF over time as the schema evolves.
Class	A collection of instances, all of which support a common type; that is, a set of <i>properties</i> and <i>methods</i> . The common <i>properties</i> and <i>methods</i> are defined as <i>features</i> of the <i>class</i> . For example, the <i>class</i> called Modem represents all the modems present in a system.
Common model	A collection of <i>models</i> specific to a particular area, derived from the <i>Core model</i> . Included are the system <i>model</i> , the application <i>model</i> , the network <i>model</i> and the device <i>model</i> .
Core model	A subset of <i>CIM</i> , not specific to any platform. The <i>Core model</i> is set of <i>classes</i> and <i>associations</i> that establish a conceptual framework for the <i>schema</i> of the rest of the managed environment. Systems, applications, networks and related information are modeled as extensions to the <i>Core model</i> .
Domain	A virtual room for object names that establishes the range in which the names of objects are unique.
Explicit Qualifier	A <i>qualifier</i> defined separately from the definition of a <i>class</i> , <i>property</i> or other schema element (see <i>implicit qualifier</i>). <i>Explicit qualifier</i> names must be unique across the entire <i>schema</i> . <i>Implicit qualifier</i> names must be unique within the defining schema element; that is, a given schema element may not have two <i>qualifiers</i> with the same name.
Extended schema	A platform specific <i>schema</i> derived from the Common model. An example is the Win32 <i>schema</i> .
Feature	A <i>property</i> or <i>method</i> belonging to a <i>class</i> .
Flavor	Part of a <i>qualifier</i> specification indicating overriding and <i>inheritance</i> rules. For example, the <i>qualifier</i> KEY has Flavor(DisableOverride ToSubclass), meaning that every subclass must inherit it and cannot override it.
Implicit Qualifier	A <i>qualifier</i> defined as a part of the definition of a <i>class</i> , <i>property</i> or other schema element (see <i>explicit qualifier</i>).
Indication	A type of <i>class</i> usually created as a result of the occurrence of a <i>trigger</i> .
Inheritance	A relationship between two <i>classes</i> in which all the members of the <i>subclass</i> are required to be members of the <i>superclass</i> . Any member of the <i>subclass</i> must also support any <i>method</i> or <i>property</i> supported by the <i>superclass</i> . For example, Modem is a <i>subclass</i> of Device.
Instance	A unit of data. An <i>instance</i> is a set of <i>property</i> values that can be uniquely identified by a <i>key</i> .
Key	One or more qualified class properties that can be used to construct a name. One or more qualified object properties which uniquely identify instances of this object

	in a namespace.
Managed Object	The actual item in the system environment that is accessed by the <i>provider</i> . For example, a Network Interface Card.
Meta model	A set of <i>classes</i> , <i>associations</i> and <i>properties</i> that expresses the types of things that can be defined in a <i>Schema</i> . For example, the <i>meta model</i> includes a <i>class</i> called <i>property</i> which defines the <i>properties</i> known to the system, a <i>class</i> called <i>method</i> which defines the <i>methods</i> known to the system, and a <i>class</i> called <i>class</i> which defines the <i>classes</i> known to the system.
Meta schema	The schema of the meta model.
Method	A declaration of a signature; that is, the method name, return type and parameters, and, in the case of a concrete class, may imply an implementation.
Model	A set of <i>classes</i> , <i>properties</i> and <i>associations</i> that allows the expression of information about a specific domain. For example, a Network may consist of Network Devices and Logical Networks. The Network Devices may have attachment <i>associations</i> to each other, and may have member <i>associations</i> to Logical Networks.
Model Path	A reference to an object within a namespace.
Namespace	An <i>object</i> that defines a scope within which object keys must be unique.
Namespath Path	A reference to a namespace within an implementation that is capable of hosting CIM objects.
Name	Combination of a Namespace path and a Model path that identifies a unique object.
Trigger	The occurrence of some action such as the creation, modification or deletion of an <i>object</i> , access to an <i>object</i> , or modification or access to a <i>property</i> . <i>Triggers</i> may also be fired as a result of the passage of a specified period of time. A <i>trigger</i> typically results in an <i>Indication</i> .
Polymorphism	A <i>subclass</i> may redefine the implementation of a <i>method</i> or <i>property</i> inherited from its <i>superclass</i> . The <i>property</i> or <i>method</i> is thereby redefined, even if the <i>superclass</i> is used to access the object. For example, Device may define availability as a string, and may return the values "powersave", "on" or "off." The Modem <i>subclass</i> of Device may redefine (override) availability by returning "on," "off," but not "powersave". If all Devices are enumerated, any Device that happens to be a modem will not return the value "powersave" for the availability <i>property</i> .
Property	A value used to characterize an instance of a <i>class</i> . For example, a Device may have a <i>property</i> called status.
Provider	An executable that can return or set information about a given <i>managed object</i> .
Qualifier	A value used to characterize a <i>method</i> , <i>property</i> , or <i>class</i> in the <i>meta schema</i> . For example, if a property has the qualifier KEY with the value TRUE, the property is a key for the class.
Reference	Special <i>property types</i> that are references or "pointers" to other instances.
Schema	A management schema is provided to establish a common conceptual framework at the level of a fundamental topology-both with respect to classification and association, and with respect to a basic set of classes intended to establish a common framework for a description of the managed environment. A <i>Schema</i> is a namespace and unit of ownership for a set of classes. <i>Schemas</i> may come in forms such as a text file, information in a repository, or diagrams in a CASE tool.
Scope	Part of a <i>Qualifier</i> specification indicating with which meta constructs the <i>Qualifier</i> can be used. For example, the <i>Qualifier</i> ABSTRACT has Scope(Class Association Indication), meaning that it can only be used with <i>Classes</i> , <i>Associations</i> and <i>Indications</i> .
Scoping Object	Objects which represent a real-world managed element, which in turn propagate keys to other objects.

Signature	The return type and parameters supported by a <i>method</i> .
Subclass	See Inheritance.
Superclass	See Inheritance.
Top Level Object	A class or object that has no scoping object.

2042 **Appendix F Unicode Usage**

2043 All punctuation symbols associated with object path or MOF Syntax occur within the Basic Latin range U+0000 to
2044 U+007F. These include normal punctuators, such as slashes, colons, commas, and so on. No important syntactic
2045 punctuation character occurs outside of this range.

2046 All characters above U+007F are treated as parts of names, even though there are several reserved characters such as
2047 U+2028 and U+2029 which are logically whitespace.

2048 Therefore, all namespace, class and property names are identifiers composed as follows:

2049 Initial identifier characters must be in set S1, where $S1 = \{U+005F, U+0041...U+005A, U+0061...U+007A,$
2050 $U+0080...U+FFEF\}$ [This is alphabetic, plus underscore]

2051 All following characters must be in set S2 where $S2 = S1 \text{ union } \{U+0030...U+0039\}$ [This is alphabetic, underscore,
2052 plus Arabic numerals 0 through 9.]

2053 Note that the Unicode specials range (U+FFF0...U+FFFF) are not legal for identifiers. While the above sub-range of
2054 U+0080...U+FFEF includes many diacritical characters which would not be useful in an identifier, as well as the
2055 Unicode reserved sub-range which has not been allocated, it seems advisable for simplicity of parsers to simply treat
2056 this entire sub-range as 'legal' for identifiers.

2057 Refer to RFC2279, published by the Internet Engineering Task Force (IETF), as an example of a Universal
2058 Transformation Format that has specific characteristics for dealing with multi-octet characters on an application-
2059 specific basis.

2060 **F.1 MOF Text**

2061 MOF files using Unicode must contain a signature as the first two bytes of the text file, either U+FFFE or U+FEFF,
2062 depending on the byte ordering of the text file (as suggested in Section 2.4 of the Unicode specification).

2063 U+FFFE is little endian.

2064 All MOF keywords and punctuation symbols are as described in the MOF Syntax document and are not locale-
2065 specific. They are composed of characters falling in the range U+0000...U+007F, regardless of the locale of origin
2066 for the MOF or its identifiers.

2067 **F.2 Quoted Strings**

2068 In all cases where non-identifier string values are required, delimiters must surround them.

2069 The supported delimiter for strings is U+0027. Once a quoted string is started using the delimiter, the same
2070 delimiter, U+0027, is used to terminate it.

2071 In addition, the digraph U+005C ("\" followed by U+0027 "" constitutes an embedded quotation mark, not a
2072 termination of the quoted string.

2073 The characters permitted within the quotation mark delimiters just described may fall within the range U+0001
2074 through U+FFEF.

2075 **Appendix G Guidelines**

2076 Method descriptions are recommended and must, at a minimum, indicate that method's side-effects (pre- and post-
2077 conditions).

2078 Associations must not be declared as subtypes of classes that are not associations.

2079 Leading underscores in identifiers are to be discouraged and not be used at all in the standard schemas.

2080 As a general rule, it is recommended that class names not be reused as part of property or method names. Property
2081 and method names are already unique within their defining class.

2082 To enable information sharing between different CIM implementations, the MAXLEN qualifier should be used to
2083 specify the maximum length of string properties. This qualifier must **always** be present for string properties used as
2084 keys.

2085 A class that has no ABSTRACT qualifier must define, or inherit, key properties.

2086 **G.1 Mapping of Octet Strings**

2087 Most management models, including SNMP and DMI, support octet strings as data types. The octet string data type
2088 represents arbitrary numeric or textual data. This data is stored as an indexed byte array of unlimited, but fixed size.
2089 Typically, the first N bytes indicate the actual string length. Since some environments only reserve the first byte,
2090 they do not support octet strings larger than 255 bytes.

2091 In the current release, CIM does not support octet strings as a separate data type. To map a single octet string (i.e.,
2092 octets of binary data), it is recommended that the equivalent CIM property be defined as an array of unsigned 8-bit
2093 integers (uint8). The first four bytes of the array contain the length of the octet data: byte 0 is the most significant
2094 byte of the length; byte 3 is the least significant byte of the length. The octet data starts at byte 4. The
2095 OCTETSTRING qualifier may be used to indicate that the uint8 array conforms to this encoding.

2096 In the case where an array of octet strings must be mapped, since arrays of uint8 arrays are not supported, a textual
2097 convention encoding the binary information as hexadecimal digit characters (i.e., 0x<<0-9,A-F><0-9,A-F>>*) is
2098 used for each of the octet strings in the array. The number of octets in the octet string is encoded in the first 8
2099 hexadecimal digits of the string with the most significant digits in the left- most characters of the string. The length
2100 count octets are included in the length count (i.e., "0x00000004" is the encoding of a 0- length octet string.

2101 The OCTETSTRING qualifier is used to qualify the string array.

2102

2103 Example usages of the OCTETSTRING qualifier on a property follows:

2104

2105 [Description ("An octet string"), Octetstring]

2106 uint8 Foo[];

2107 [Description ("An array of octet strings"), Octetstring]

2108 String Bar[];

2109 **G.2 SQL Reserved Words**

2110 It is recommended that SQL reserved words be avoided in the selection of class and property names. This
2111 particularly applies to property names, since class names are prefixed by the schema name, making a clash with a
2112 reserved word unlikely. The current set of SQL reserved words are:

2113 From sql1992.txt:

AFTER	ALIAS	ASYNC	BEFORE
BOOLEAN	BREADTH	COMPLETION	CALL
CYCLE	DATA	DEPTH	DICTIONARY
EACH	ELSEIF	EQUALS	GENERAL
IF	IGNORE	LEAVE	LESS

LIMIT	LOOP	MODIFY	NEW
NONE	OBJECT	OFF	OID
OLD	OPERATION	OPERATORS	OTHERS
PARAMETERS	PENDANT	PREORDER	PRIVATE
PROTECTED	RECURSIVE	REF	REFERENCING
REPLACE	RESIGNAL	RETURN	RETURNS
ROLE	ROUTINE	ROW	SAVEPOINT
SEARCH	SENSITIVE	SEQUENCE	SIGNAL
SIMILAR	SQL EXCEPTION	SQLWARNING	STRUCTURE
TEST	THERE	TRIGGER	TYPE
UNDER	VARIABLE	VIRTUAL	VISIBLE
WAIT	WHILE	WITHOUT	

2114 From sql1992.txt (Annex E):

ABSOLUTE	ACTION	ADD	ALLOCATE
ALTER	ARE	ASSERTION	AT
BETWEEN	BIT	BIT_LENGTH	BOTH
CASCADE	CASCADED	CASE	CAST
CATALOG	CHAR_LENGTH	CHARACTER_LENGTH	COALESCE
COLLATE	COLLATION	COLUMN	CONNECT
CONNECTION	CONSTRAINT	CONSTRAINTS	CONVERT
CORRESPONDING	CROSS	CURRENT_DATE	CURRENT_TIME
CURRENT_TIMESTAMP	CURRENT_USER	DATE	DAY
DEALLOCATE	DEFERRABLE	DEFERRED	DESCRIBE
DESCRIPTOR	DIAGNOSTICS	DISCONNECT	DOMAIN
DROP	ELSE	END-EXEC	EXCEPT
EXCEPTION	EXECUTE	EXTERNAL	EXTRACT
FALSE	FIRST	FULL	GET
GLOBAL	HOURLY	IDENTITY	IMMEDIATE
INITIALLY	INNER	INPUT	INSENSITIVE
INTERSECT	INTERVAL	ISOLATION	JOIN
LAST	LEADING	LEFT	LEVEL
LOCAL	LOWER	MATCH	MINUTE
MONTH	NAMES	NATIONAL	NATURAL
NCHAR	NEXT	NO	NULLIF
OCTET_LENGTH	ONLY	OUTER	OUTPUT
OVERLAPS	PAD	PARTIAL	POSITION
PREPARE	PRESERVE	PRIOR	READ
RELATIVE	RESTRICT	REVOKE	RIGHT
ROWS	SCROLL	SECOND	SESSION
SESSION_USER	SIZE	SPACE	SQLSTATE
SUBSTRING	SYSTEM_USER	TEMPORARY	THEN
TIME	TIMESTAMP	TIMEZONE_HOUR	TIMEZONE_MINUTE
TRAILING	TRANSACTION	TRANSLATE	TRANSLATION

TRIM	TRUE	UNKNOWN	UPPER
USAGE	USING	VALUE	VARCHAR
VARYING	WHEN	WRITE	YEAR
ZONE			

2115 From sql3part2.txt (Annex E):

ACTION	ACTOR	AFTER	ALIAS
ASYNC	ATTRIBUTES	BEFORE	BOOLEAN
BREADTH	COMPLETION	CURRENT_PATH	CYCLE
DATA	DEPTH	DESTROY	DICTIONARY
EACH	ELEMENT	ELSEIF	EQUALS
FACTOR	GENERAL	HOLD	IGNORE
INSTEAD	LESS	LIMIT	LIST
MODIFY	NEW	NEW_TABLE	NO
NONE	OFF	OID	OLD
OLD_TABLE	OPERATION	OPERATOR	OPERATORS
PARAMETERS	PATH	PENDANT	POSTFIX
PREFIX	PREORDER	PRIVATE	PROTECTED
RECURSIVE	REFERENCING	REPLACE	ROLE
ROUTINE	ROW	SAVEPOINT	SEARCH
SENSITIVE	SEQUENCE	SESSION	SIMILAR
SPACE	SQLEXCEPTION	SQLWARNING	START
STATE	STRUCTURE	SYMBOL	TERM
TEST	THERE	TRIGGER	TYPE
UNDER	VARIABLE	VIRTUAL	VISIBLE
WAIT	WITHOUT		

2116 sql3part4.txt (ANNEX E):

CALL	DO	ELSEIF	EXCEPTION
IF	LEAVE	LOOP	OTHERS
RESIGNAL	RETURN	RETURNS	SIGNAL
TUPLE	WHILE		

2117

2118 **Appendix H Embedded Object Qualifier**

2119 The EmbeddedObject qualifier is used motivated by the need to include a specific instance's data in an indication
2120 (event notification) or to capture the contents of an instance at a point in time (for example, to include the
2121 CIM_DiagnosticSetting properties that dictated a particular CIM_DiagnosticResult, in the Result object).

2122 To address this need, it is expected that the next major version of the CIM Specification will include a separate data
2123 type for representing instances (or snapshots of instances) directly. Until that time, the EmbeddedObject qualifier
2124 can be used to achieve an approximately equivalent effect by permitting a CIM object manager (or other entity) to
2125 simulate embedded instances by encoding them as strings when they are presented externally. Clients which aren't
2126 concerned with handling embedded objects may treat properties with this qualifier as they would any other string-
2127 valued property, while clients that wish to realize the capability of embedded objects can extract the embedded
2128 object information by decoding the presented string value.

2129 In order to reduce the parsing burden on embedded object-aware consumers of CIM data, the encoding by which the
2130 embedded object is represented in the string value depends on the protocol or representation being used for
2131 transmission of the containing instance. This makes the string value appear to vary depending on the circumstances
2132 in which it is observed; this is an acknowledged weakness of using a qualifier instead of a new data type.

2133 The DMTF defines only two representations by which CIM data are made available from a CIM object manager,
2134 although it allows for other representations and protocols. The two defined representations are MOF and CIM-
2135 XML; accordingly, the DMTF defines EmbeddedObject string formats only for these two representations. When
2136 other protocols or representations are used, they MUST include particulars on the encoding to be used for
2137 EmbeddedObject-qualified string values if that protocol or representation will be used to communicate with
2138 embedded object-aware consumers of CIM data.

2139 **H.1 Encoding for MOF**

2140 When CIM data including an EmbeddedObject-qualified string value are rendered in MOF by an embedded object-
2141 aware entity (e.g., CIM object manager or client), the embedded object MUST be encoded into string form
2142 following the MOF syntax for the instanceDeclaration nonterminal in the case of embedded instances, or for the
2143 classDeclaration, assocDeclaration, or indicDeclaration nonterminals, as appropriate, in the case of embedded
2144 classes (see Appendix A).

2145 Examples:

```
2146 Instance of CIM_InstCreation {  
2147     EventTime = "20000208165854.457000-360";  
2148     SourceInstance =  
2149         "Instance of CIM_FAN {"  
2150             "DeviceID = \"Fan 1\";"  
2151             "Status = \"Degraded\";"  
2152         }";  
2153 };  
2154
```

```
2155 Instance of CIM_ClassCreation {  
2156     EventTime = "20031120165854.457000-360";  
2157     ClassDefinition =  
2158         "class CIM_Fan : CIM_CoolingDevice {"  
2159             " boolean VariableSpeed;"  
2160             " [Units (\"Revolutions per Minute\")] "  
2161             " uint64 DesiredSpeed;"  
2162         }";  
2163 };
```

2164 **H.2 Encoding for CIM-XML**

2165 When CIM data including an EmbeddedObject-qualified string value are rendered in CIM-XML by an embedded
2166 object-aware entity, the embedded object **MUST** be encoded into string form as either an INSTANCE element (in
2167 the case of instances) or a CLASS element (in the case of classes), as defined in the DMTF CIM-XML specification
2168 (Representation of CIM in XML, DSP0201).

2169 **Appendix I Schema Errata**

2170 It is expected that the CIM Schema (based on the constructs and thinking laid out in this Specification) will evolve.
2171 Evolution occurs for three reasons:
2172

2173
2174 To add new classes, associations, qualifiers, properties and/or methods

2175 To correct errors in the "Final Release" versions of the Schema

2176 To deprecate and update the model (labeling classes, associations, qualifiers, etc. as "not recommended for future
2177 development" and replacing them with "new" constructs)

2178

2179 #1 is easily accommodated and addressed in Section 2.3 of this Specification.

2180 #3 is addressed by the Deprecated qualifier explained in Section 2.5 of this Specification.

2181 #2 is an Errata to the CIM Schemas, after their "Final Release." Some examples of errata

2182 are:

2183 Incorrectly or incompletely defined keys (an array defined as a key property, or incompletely specified propagated
2184 keys)

2185 Invalid subclassing (subclassing an optional association from a weak relationship i.e., a mandatory association,
2186 subclassing a nonassociation class from an association, or subclassing an association but having different reference
2187 names resulting in 3 or more references on an association)

2188 Class references reversed as defined by an association's roles (Antecedent/Dependent references reversed)

2189 Use of SQL reserved words as property names

2190 Violation of semantics (Missing Min(1) on a Weak relationship, contradicting that a Weak relationship is
2191 mandatory)

2192

2193 Errata are a serious matter since the Schema should be correct, but existing implementations must be taken into
2194 account. Therefore, the DMTF will implement the following process (in addition to the normal release process) with
2195 respect to any Schema errata:

2196

2197 (a) As soon as an error is found, it should be raised to the Technical Committee (technical@dmf.org) for review.
2198 Suggestions for correcting the error should also be described, if possible.

2199

2200 (b) The findings of the Technical Committee will be documented in an email to the submitter, within 21 days. These
2201 findings will describe the Committee's decision regarding whether the submission is a valid Errata, the reasoning
2202 behind the decision, the recommended strategy to correct the error, and whether backward compatibility is possible.

2203

2204 (c) If the error is valid, an email will be sent simultaneously (with the reply to the submitter) to all DMTF members
2205 (members@dmf.org). The email will highlight the error, the findings of the Technical Committee and the strategy
2206 to correct the error. In addition, the Committee will indicate what versions of the Schema will be affected i.e., only
2207 the latest, or all schemas after a specific version.

2208

2209 (d) All members are invited to respond to the Technical Committee -- within 30 days regarding impacts to their
2210 implementations due to the "correction strategy." Impacts should be explained as thoroughly as possible, as well as
2211 alternate strategies to correct the error.

2212

2213 (e) If one or more members are affected, then the Technical Committee will evaluate all proposed, alternate
2214 correction strategies. It will choose one of three options:

2215

2216 To stay with the correction strategy proposed in (b)

2217 To move to one of the proposed alternate strategies

2218 To define a new correction strategy based on the evaluation of member impacts.

2219

2220 (f) If an "alternate" strategy is proposed in (e), the Technical Committee may decide to reenter the Errata process,
2221 resuming with Item (c) -- an email to all DMTF members regarding the alternate correction strategy. However, if the
2222 Technical Committee believes that further comment will not raise any new issues, then the outcome of Item (e) will

2223 be declared "final."

2224

2225 (g) If a "final" strategy is decided, this strategy will be implemented via a Change Request to the affected
2226 Schema(s). The Change Request will be written and issued by the Technical Committee. Affected Models and MOF
2227 will be updated, and their introductory comment section flagged to indicate that a correction for an ERRATA was
2228 applied.

2229

Appendix J References

- 2231 [1] Grady Booch and James Rumbaugh, *Unified Method for Object-Oriented Development Document Set*,
2232 Rational Software Corporation, 1996, <http://www.rational.com/uml>
- 2233 [2] HyperMedia Management Protocol, Protocol Encoding, draft-hmmp-encoding-03.txt, February, 1997
- 2234 [3] *Interface Definition Language*, DCE/RPC, The Open Group.
- 2235 [4] Georges Gardarin and Patrick Valduriez, *Relational Databases and Knowledge Bases*, Addison-
2236 Wesley, 1989
- 2237 [5] Coplein, James O., Schmidt, Douglas C (eds). *Pattern Languages of Program Design*, Addison-
2238 Wesley, Reading Mass., 1995
- 2239 [6] IEEE Standard for Binary Floating-Point Arithmetic, *ANSI/IEEE Standard 754-1985*, Institute of
2240 Electrical and Electronics Engineers, August 1985.
- 2241 [7] Augmented BNF for Syntax Specifications: ABNF, RFC 2234, Nov 1997
- 2242 [8] G. Weinberger, *General Systems Theory*
- 2243 [9] *The Unicode Standard*, Version 2.0, by The Unicode Consortium, Addison-Wesley, 1996.
- 2244 [10] Universal Multiple-Octet Coded Character Set, ISO/IEC 10646
- 2245 [11] UCS Transformation Format 8 (UTF-8), ISO/IEC 10646-1:1993 Amendment 2 (1996)
- 2246 [12] Code for the Representation of Names of Languages, ISO/IEC 639:1988 (E/F)
- 2247 [13] Code for the Representation of Names of Territory, ISO/IEC 3166:1988 (E/F)

Appendix K Change History

Version 1	April 09, 1997	First Public Release
Version 1.1	October 23, 1997	Output after Working Groups input
Version 1.2a	November 03, 1997	Naming
Version 1.2b	November 17, 1997	Remove reference qualifier
Version 2.0a	December 11, 1997	Apply pending changes and new metaclass
Version 2.0d	December 11, 1997	Output of 12/9/1997 TDC, Dallas
Version 2.0f	February 16, 1998	Output of 2/3/1998 TDC, Austin
Version 2.0g	February 26, 1998	Apply approved change requests and final edits submitted through 2/26/1998.
Version 2.2	June 14, 1999	Incorporate Errata and approved change requests through 1999-06-08
Version 2.2.1	June 07, 2003	Incorporate Addenda 1
Version 2.2.2	June 07, 2003	Incorporate Addenda 2
Version 2.2.1000	June 7, 2003	00638-Replace Section 2.3.1 of the CIM Spec 00664-Rules for Versioning CIM Classes 00665-Remove Participants Section 00685-Add Units in the CIM Spec App C 00707-New "Composition" Qualifier 00713-Corrections to Qualifiers 00714-Clarification of the BitMap Qualifier 00715-Clarification of the Deprecated qualifier 00727-Required Qualifier for Identification 00729-ALIAS qualifier from Standard to Optional 00731-New Units value - Packets 00762-Correction to Qualifiers 00813-Clarify Semantics of the Write Qualifier 00814-Add syntax to MOF BNF to allow the specification of an empty array 00817-Errata Class Alias Support 00881-Clarify Model Path 00910-New Exception Qualifier 01062-Updates to Experimental and Version 01069-Clarification of ranges in ValueMaps
Version 2.3 Draft	November 05, 2003 Lawrence J. Lamers	00716-New MinLen qualifier
Version 2.3 Draft	November 12, 2003	Updated Appendix K - fixed Table of Contents formatting.
Version 2.3 Draft	January 8, 2004	ARCH00001 - Add RFC2119 terminology to MOF files ARCH00002 CR1133 - Overriding (non-reference) property type ARCH00004 CR1169 - Instances don't have qualifiers ARCH00005 CR1148 - Cleanups for Standard Qualifiers ARCH00006 CR1149 - Datetime cleanup ARCH00008 CR1158 - Parameters can be arrays of references ARCH00009 CR1153 - Formal syntax reference ARCH00010 CR1155 - Instances don't define properties or methods ARCH00011 CR1156 - MODELRESPONSE syntax ARCH00012 CR1158 - NULLVALUE qualifier syntax and wording ARCH00013 CR1159 - MAXVALUE qualifier wording

		ARCH00016 CR1172 - ModelPath syntax
Version 2.3 Draft	February 10, 2004	ARCH00015 CR1167 -Correct DisableOverride flavor on OVERRIDE qualifier ARCH00017 CR1168 - Clarify definitions of Qualifier flavors ARCH00021 CR1237 - Identifying properties
Version 2.3 Draft	July 6, 2004	ARCH00019 CR1174.003 EMBEDDEDOBJECT qualifier.
Version 2.3 Draft	July 7, 2004	ARCH00020.004 CR1461.000 Deprecate use of SOURCE, SOURCETYPE, NONLOCAL, NONLOCALTYPE qualifiers and pragmas. ARCH00022.001 CR1278.000 Model Correspondence Scope ARCH00027.000 CR1289.001 Change name of DSP0004 to CIM Infrastructure Specification ARCH00029.004 CR1390.001 EMBEDDED INSTANCE qualifier ARCH00031.002 INT type; MINVALUE and MAXVALUE usage ARCH00032.001 Usage Rules for DN string arrays
Version 2.3 Draft	October 4, 2005	ARCHCR0048 Remove XML of embedded object example ARCHCR0051 Starting numbers for datetime fields ARCHCR0052 OCL qualifier ARCHCR0059 Remove mention of "instance" with qualifiers ARCHCR0060 Remove inappropriate mention of "alias" ARCHCR0061 Remove incorrect #pragma namespace syntax rule ARCHCR0062 Remove moot qualifier/pragma conflict rule ARCHCR0063 Explain struck-through text for compiler directives ARCHCR0064 Remove "Synchronizing Namespaces" section ARCHCR0065 Remove qualifier defs and Revision qual from Appendix B

2249

2250 **Appendix L Ambiguous Property and Method Names**

2251 Item 21-E of Section 2.1 explicitly allows a subclass to define a property which may have the same name as a
2252 property defined by a superclass, and for that new property NOT to override the superclass' property. (The subclass
2253 may override the superclass' property by attaching an OVERRIDE qualifier; this situation is well-behaved, and not
2254 part of the problem under discussion.)

2255 Similarly, a subclass may define a method with the same name as a method defined by a superclass, without
2256 overriding the superclass' method. This appendix will refer only to properties in what follows, but it's to be
2257 understood that the issues regarding methods are essentially the same, and for any statement about properties, a
2258 similar statement about methods can be inferred.

2259 This same-name capability is provided to allow the superclass to be enhanced or extended in a minor schema change
2260 by one group (the DMTF, in particular), without coordination with, or even knowledge of, the development of the
2261 subclass in another schema by another group. That is, a subclass that has been defined using one version of the
2262 superclass should not become invalid if a subsequent version of the superclass introduces a new property whose
2263 name is the same as a property defined on the subclass. Any other use of the same-name capability is strongly
2264 discouraged, and additional constraints on allowable cases may well be added in future versions of CIM.

2265 It's natural for CIM applications to be written under the assumption that property names alone suffice to uniquely
2266 identify properties. However, applications written this way are at risk of failure if they refer to properties from a
2267 subclass whose superclass has been modified to include a new property, and the new property has the same name as
2268 a previously-existing property defined by the subclass. As an example, suppose we have:

```
2269 [abstract]
2270 class CIM_Superclass
2271 {
2272 };
2273
2274 class VENDOR_Subclass
2275 {
2276     string    Foo;
2277 };
2278
```

2279 Suppose further that there's just one instance of VENDOR_Subclass; a call to
2280 enumerateInstances("VENDOR_Subclass") might produce the following XML result from the CIMOM (assuming it
2281 didn't bother to ask for CLASSORIGIN information):

```
2282 <INSTANCE CLASSNAME="VENDOR_Subclass">
2283     <PROPERTY NAME="Foo" TYPE="string">
2284         <VALUE>Hello, my name is Foo</VALUE>
2285     </PROPERTY>
2286 </INSTANCE>
2287
2288
```

2289 If the definition of CIM_Superclass changes to:

```
2290 [abstract]
2291 class CIM_Superclass
2292 {
2293     string foo = "You lose!";
2294 };
2295
```

2296 then the enumerateInstances call might return:

```
2297 <INSTANCE>
2298     <PROPERTY NAME="Foo" TYPE="string">
2299         <VALUE>You lose!</VALUE>
2300 </PROPERTY>
```

```
2301     <PROPERTY NAME="Foo" TYPE="string">
2302         <VALUE>Hello, my name is Foo</VALUE>
2303     </PROPERTY>
2304 </INSTANCE>
```

2305
2306 If the client application attempts to retrieve the 'foo' property, which value it obtains (if it doesn't experience an
2307 error) is dependent on the implementation, if not the phase of the moon.

2308
2309 While a class may define a property with the same name as an inherited property, it is not permitted in and of itself
2310 to define two (or more) properties with the same name. This means that the combination of defining class plus
2311 property name uniquely identifies a property. (Most CIM operations that return instances have a flag controlling
2312 whether or not to include the originClass for each property: see e.g. DSP0200, Section 2.3.2.11, enumerateInstances,
2313 and DSP0201, Section 3.1.4, ClassOrigin.) However, the use of class-plus-property-name for identifying properties
2314 makes an application vulnerable to failure if a property is "promoted" up to a superclass in a subsequent schema
2315 release. As a concrete example, consider:

```
2316 class CIM_Top
2317 {
2318 };
2319
2320 class CIM_Middle : CIM_Top
2321 {
2322     uint32    foo;
2323 };
2324
2325 class VENDOR_Bottom : CIM_Middle
2326 {
2327     string    foo;
2328 };
2329
2330
```

2331 An application that identifies the uint32 property as "the property named 'foo' defined by CIM_Middle" will no
2332 longer work if a subsequent release of the CIM schema changes the hierarchy to:

```
2333 class CIM_Top
2334 {
2335     uint32    foo;
2336 };
2337
2338 class CIM_Middle : CIM_Top
2339 {
2340 };
2341
2342 class VENDOR_Bottom : CIM_Middle
2343 {
2344     string    foo;
2345 };
2346
```

2347 Strictly speaking, there is no longer a "property named 'foo' defined by CIM_Middle"; it's now defined by CIM_Top
2348 and merely inherited by CIM_Middle, just as it's inherited by VENDOR_Bottom. An instance of
2349 VENDOR_Bottom returned in XML from a CIMOM might look like this:

```
2350 <INSTANCE CLASSNAME="VENDOR_Bottom">
2351     <PROPERTY NAME="Foo" TYPE="string" CLASSORIGIN="VENDOR_Bottom">
2352         <VALUE>Hello, my name is Foo!</VALUE>
2353     </PROPERTY>
2354     <PROPERTY NAME="Foo" TYPE="uint32" CLASSORIGIN="CIM_Top">
```

```
2356         <VALUE>47</VALUE>
2357     </PROPERTY>
2358 </INSTANCE>
2359
```

2360 A client application looking for a PROPERTY element with NAME="Foo" and CLASSORIGIN="CIM_Middle"
2361 will fail with this XML fragment.

2362 Although CIM_Middle no longer defines a 'foo' property directly in the example above, our intuition says that we
2363 should be able to point to the CIM_Middle class and locate the 'foo' property that's defined in its nearest superclass.
2364 In the general case, the application must be prepared to perform this search, separately obtaining information from
2365 the server when necessary about the (current) class hierarchy and implementing an algorithm to select the
2366 appropriate property information from the instance information returned from a server operation.

2367 Although it's technically allowed, schema writers SHOULD NOT introduce properties that cause name collisions
2368 within the schema, and are strongly discouraged from introducing properties whose names are known to conflict
2369 with property names of any subclass or superclass in another schema.