

Network Working Group
Internet Draft
Intended status: Standard Track
Expires: July 20, 2017

J. Strassner
Huawei Technologies
J. Halpern
S. van der Meer
Ericsson
January 18, 2017

**Generic Policy Information Model for
Simplified Use of Policy Abstractions (SUPA)
draft-ietf-supa-generic-policy-info-model-02**

Abstract

This document defines an information model for representing policies using a common extensible framework that is independent of language, protocol, repository. It is also independent of the level of abstraction of the content and meaning of a policy.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 20, 2017.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided

without warranty as described in the Simplified BSD License.

Strassner, et al.

Expires January 19, 2017

[Page 1]

Table of Contents

*** Note: this is out of date, will be fixed in next version ***

1.	Overview	9
1.1.	Introduction	9
1.2.	Changes Since Version -03	11
2.	Conventions Used in This Document	11
3.	Terminology	12
3.1.	Acronyms	12
3.2.	Definitions	12
3.2.1.	Core Terminology	12
3.2.1.1.	Information Model	12
3.2.1.2.	Data Model	13
3.2.1.3.	Class	13
3.2.1.3.1.	Abstract Class	13
3.2.1.3.2.	Concrete Class	13
3.2.1.4.	Container	13
3.2.1.5.	PolicyContainer	13
3.2.2.	Policy Terminology	14
3.2.2.1.	SUPAPolicyObject	14
3.2.2.2.	SUPAPolicy	14
3.2.2.3.	SUPAPolicyClause	14
3.2.2.4.	SUPAECAPolicyRule	15
3.2.2.5.	SUPAMetadata	15
3.2.2.6.	SUPAPolicyTarget	15
3.2.2.7.	SUPAPolicySource	15
3.2.3.	Modeling Terminology	16
3.2.3.1.	Inheritance	16
3.2.3.2.	Relationship	16
3.2.3.3.	Association	17
3.2.3.4.	Aggregation	17
3.2.3.5.	Composition	17
3.2.3.6.	Association Class	17
3.2.3.7.	Multiplicity	18
3.2.3.8.	Navigability	18
3.3.	Symbology	18
3.3.1.	Inheritance	18
3.3.2.	Association	19
3.3.3.	Aggregation	19
3.3.4.	Composition	19
3.3.5.	Association Class	19
3.3.6.	Abstract vs. Concrete Classes	20
4.	Policy Abstraction Architecture	21
4.1.	Motivation	22
4.2.	SUPA Approach	23

Table of Contents (continued)

4.3.	SUPA Generic Policy Information Model Overview.....	23
4.3.1.	SUPAPolicyObject	25
4.3.2.	SUPAPolicyStructure	26
4.3.3.	SUPAPolicyComponentStructure	26
4.3.4.	SUPAPolicyClause	27
4.3.5.	SUPAPolicyComponentDecorator	27
4.3.6.	SUPAPolicyTarget	28
4.3.7.	SUPAPolicySource	28
4.4.	The Design of the GPIM	28
4.4.1.	Structure of Policies	29
4.4.2.	Representing an ECA Policy Rule	30
4.4.3.	Creating SUPA Policy Clauses	33
4.4.4.	Creating SUPAPolicyClauses	36
4.4.5.	SUPAPolicySources	37
4.4.6.	SUPAPolicyTargets	39
4.4.7.	SUPAPolicyMetadata	39
4.4.7.1.	Motivation	39
4.4.7.2.	Design Approach	40
4.4.7.2.1.	Policies and Actors	42
4.4.7.2.2.	Deployment vs. Execution of Policies	43
4.4.7.2.3.	Using SUPAMetadata for Policy Deployment and Execution	43
4.4.7.3.	Structure of SUPAPolicyMetadata	44
4.5.	Advanced Features	47
4.5.1.	Policy Grouping	47
4.5.2.	Policy Rule Nesting	47
5.	GPIM Model	48
5.1.	Overview	48
5.2.	The Abstract Class "SUPAPolicyObject"	49
5.2.1.	SUPAPolicyObject Attributes	50
5.2.1.1.	Object Identifiers	50
5.2.1.2.	The Attribute "supaPolObjIDContent"	51
5.2.1.3.	The Attribute "supaPolObjIDEncoding"	51
5.2.1.4.	The Attribute "supaPolicyDescription"	51
5.2.1.5.	The Attribute "supaPolicyName"	51
5.2.2.	SUPAPolicy Relationships	52
5.2.2.1.	The Relationship "SUPAHasPolicyMetadata"	52
5.2.2.2.	The Association Class "SUPAHasPolicyMetadataDetail"	52
5.3.	The Abstract Class "SUPAPolicyStructure"	52
5.3.1.	SUPAPolicyStructure Attributes	53
5.3.1.1.	The Attribute "supaPolAdminStatus"	53
5.3.1.2.	The Attribute "supaPolContinuumLevel"	53
5.3.1.3.	The Attribute "supaPolDeployStatus"	54
5.3.1.4.	The Attribute "supaPolExecFailStrategy"	54

Table of Contents (continued)

5.3.2.	SUPAPolicyStructure Relationships	55
5.3.2.1.	The Aggregation "SUPAHasPolicySource"	55
5.3.2.2.	The Association Class	
	"SUPAHasPolicySourceDetail"	55
5.3.2.2.1.	The Attribute "supaPolSrcIsAuthenticated"	55
5.3.2.2.2.	The Attribute "supaPolSrcIsTrusted"	56
5.3.2.3.	The Aggregation "SUPAHasPolicyTarget"	56
5.3.2.4.	The Association Class	
	"SUPAHasPolicyTargetDetail"	56
5.3.2.4.1.	The Attribute "supaPolTgtIsAuthenticated"	56
5.3.2.4.2.	The Attribute "supaPolTgtIsEnabled"	56
5.3.2.5.	The Association "SUPAHasPolExecFailTakeAction"	57
5.3.2.6.	The Association Class	
	"SUPAHasPolExecFailTakeActionDetail"	57
5.3.2.6.1.	The Attribute	
	"supaPolExecFailActionEncoding"	57
5.3.2.6.2.	The Attribute	
	"supaPolExecFailActionName[1..n]"	58
5.3.2.7.	The Aggregation "SUPAHasPolicyClause"	58
5.3.2.8.	The Association Class	
	"SUPAHasPolicyClauseDetail"	58
5.4.	The Abstract Class "SUPAPolicyComponentStructure"	59
5.4.1.	SUPAPolicyComponentStructure Attributes	59
5.4.2.	SUPAPolicyComponentStructure Relationships	59
5.5.	The Abstract Class "SUPAPolicyClause"	59
5.5.1.	SUPAPolicyClause Attributes	60
5.5.1.1.	The Attribute "supaPolClauseDeployStatus"	60
5.5.2.	SUPAPolicyClause Relationships	61
5.6.	The Concrete Class "SUPAEncodedClause"	61
5.6.1.	SUPAEncodedClause Attributes	61
5.6.1.1.	The Attribute "supaEncodedClauseContent"	61
5.6.1.2.	The Attribute "supaEncodedClauseEncoding"	61
5.6.1.3.	The Attribute "supaEncodedClauseLanguage"	62
5.6.1.4.	The Attribute "supaEncodedClauseResponse"	62
5.6.2.	SUPAEncodedClause Relationships	62
5.7.	The Abstract Class "SUPAPolicyComponentDecorator"	62
5.7.1.	The Decorator Pattern	63
5.7.2.	SUPAPolicyComponentDecorator Attributes	64
5.7.2.1.	The Attribute "supaPolCompConstraintEncoding"	65
5.7.2.2.	The Attribute "supaAPolCompConstraint[0..n]"	65
5.7.3.	SUPAPolicyComponentDecorator Relationships	65
5.7.3.1.	The Aggregation	
	"SUPAHasDecoratedPolicyComponent"	66
5.7.3.2.	The Association Class	
	"SUPAHasDecoratedPolicyComponentDetail"	66
5.7.3.2.1.	The Attribute	
	"supaDecoratedConstraintEncoding"	66

5.7.3.2.2. The Attribute	
"supaDecoratedConstraint[0..n]"	67

Table of Contents (continued)

5.7.4. Illustration of Constraints in the Decorator Pattern	67
5.8. The Abstract Class "SUPAPolicyTerm"	68
5.8.1. SUPAPolicyTerm Attributes	69
5.8.1.1. The Attribute "supaPolTermIsNegated"	69
5.8.2. SUPAPolicyTerm Relationships	69
5.9. The Concrete Class "SUPAPolicyVariable"	69
5.9.1. Problems with the RFC3460 Version of PolicyVariable	70
5.9.2. SUPAPolicyVariable Attributes	70
5.9.2.1. The Attribute "supaPolVarName"	70
5.9.3. SUPAPolicyVariable Relationships	70
5.10. The Concrete Class "SUPAPolicyOperator"	70
5.10.1. Problems with the RFC3460 Version	71
5.10.2. SUPAPolicyOperator Attributes	71
5.10.2.1. The Attribute "supaPolOpType"	71
5.10.3. SUPAPolicyOperator Relationships	71
5.11. The Concrete Class "SUPAPolicyValue"	72
5.11.1. Problems with the RFC3460 Version of PolicyValue	72
5.11.2. SUPAPolicyValue Attributes	72
5.11.2.1. The Attribute "supaPolValContent[0..n]"	72
5.11.2.2. The Attribute "supaPolValEncoding"	73
5.11.3. SUPAPolicyValue Relationships	73
5.12. The Concrete Class "SUPAGenericDecoratedComponent"	73
5.12.1. SUPAGenericDecoratedComponent Attributes	74
5.12.1.1. The Attribute	
"supaVendorDecoratedCompContent[0..n]"	74
5.12.1.2. The Attribute "supaVendorDecoratedCompEncoding"	74
5.12.2. SUPAGenericDecoratedComponent Relationships	74
5.13. The Concrete Class "SUPAPolicyCollection"	75
5.13.1. Motivation	75
5.13.2. Solution	75
5.13.3. SUPAPolicyCollection Attributes	76
5.13.3.1. The Attribute "supaPolCollectionContent[0..n]"	76
5.13.3.2. The Attribute "supaPolCollectionEncoding"	76
5.13.3.3. The Attribute "supaPolCollectionFunction"	76
5.13.3.4. The Attribute "supaPolCollectionIsOrdered"	76
5.13.3.5. The Attribute "supaPolCollectionType"	77
5.13.4. SUPAPolicyCollection Relationships	78
5.14. The Concrete Class "SUPAPolicySource"	78
5.14.1. SUPAPolicySource Attributes	78
5.14.2. SUPAPolicySource Relationships	78
5.15. The Concrete Class "SUPAPolicyTarget"	78
5.15.1. SUPAPolicyTarget Attributes	79
5.15.2. SUPAPolicyTarget Relationships	79

Table of Contents (continued)

5.16.	The Abstract Class "SUPAPolicyMetadata"	79
5.16.1.	SUPAPolicyMetadata Attributes	80
5.16.1.1.	The Attribute "supaPolMetadataDescription"	80
5.16.1.2.	The Attribute "supaPolMetadataIDContent"	80
5.16.1.3.	The Attribute "supaPolMetadataIDEncoding"	80
5.16.1.4.	The Attribute "supaPolMetadataName"	81
5.16.2.	SUPAPolicyMetadata Relationships	81
5.16.2.1.	The Aggregation "SUPAHasPolicyMetadata"	81
5.16.2.2.	The Association Class "SUPAHasPolicyMetadataDetail"	81
5.16.2.2.1.	The Attribute "supaPolMetadataIsApplicable"	81
5.16.2.2.2.	The Attribute "supaPolMetadataConstraintEncoding"	82
5.16.2.2.3.	The Attribute "supaPolMetadataConstraint[0..n]"	82
5.17.	The Concrete Class "SUPAPolicyConcreteMetadata"	82
5.17.1.	SUPAPolicyConcreteMetadata Attributes	83
5.17.1.1.	The Attribute "supaPolMDValidPeriodEnd"	83
5.17.1.2.	The Attribute "supaPolMDValidPeriodStart"	83
5.17.2.	SUPAPolicyConcreteMetadata Relationships	83
5.18.	The Abstract Class "SUPAPolicyMetadataDecorator"	83
5.18.1.	SUPAPolicyMetadataDecorator Attributes	83
5.18.2.	SUPAPolicyMetadataDecorator Relationships	83
5.18.2.1.	The Aggregation "SUPAHasMetadataDecorator"	84
5.18.2.2.	The Association Class "SUPAHasMetadataDecoratorDetail"	84
5.19.	The Concrete Class "SUPAPolicyAccessMetadataDef"	84
5.19.1.	SUPAPolicyAccessMetadataDef Attributes	85
5.19.1.1.	The Attribute "supaAccessPrivilegeDef"	85
5.19.1.2.	The Attribute "supaAccessPrivilegeModelName" ..	85
5.19.1.3.	The Attribute "supaAccessPrivilegeModelRef" ...	86
5.20.	The Concrete Class "SUPAPolicyVersionMetadataDef"	86
5.20.1.	SUPAPolicyVersionMetadataDef Attributes	86
5.20.1.1.	The Attribute "supaVersionMajor"	87
5.20.1.2.	The Attribute "supaVersionMinor"	88
5.20.1.3.	The Attribute "supaVersionPatch"	88
5.20.1.4.	The Attribute "supaVersionPreRelease"	88
5.20.1.5.	The Attribute "supaVersionBuildMetadata"	89
6.	SUPA ECAPolicyRule Information Model	89
6.1.	Overview	89
6.2.	Constructing a SUPAECAPolicyRule	91
6.3.	Working With SUPAECAPolicyRules	92
6.4.	The Abstract Class "SUPAECAPolicyRule"	93
6.4.1.	SUPAECAPolicyRule Attributes	95
6.4.1.1.	The Attribute "supaECAPolicyRulePriority"	95

6.4.1.2.	The Attribute "supaECAPolicyRuleStatus"	95
6.4.2.	SUPAECAPolicyRule Relationships	96

Table of Contents (continued)

6.5.	The Concrete Class "SUPAECAPolicyRuleAtomic"	96
6.5.1.	SUPAECAPolicyRuleAtomic Attributes	96
6.5.2.	SUPAECAPolicyRuleAtomic Relationships	96
6.6.	The Concrete Class "SUPAECAPolicyRuleComposite"	96
6.6.1.	SUPAECAPolicyRuleComposite Attributes	96
6.6.1.1.	The Attribute "supaECAEvalStrategy"	97
6.6.2.	SUPAECAPolicyRuleComposite Relationships	97
6.6.2.1.	The Aggregation "SUPAHasECAPolicyRule"	98
6.6.3.	The Association Class "SUPAHasECAPolicyRuleDetail" ..	98
6.6.3.1.	The Attribute "supaECAPolicyIsDefault"	98
6.7.	The Abstract Class "SUPABooleanClause"	98
6.7.1.	SUPABooleanClause Attributes	99
6.7.1.1.	The Attribute "supaBoolClauseIsNegated"	99
6.7.2.	SUPABooleanClause Relationships	99
6.8.	The Concrete Class "SUPABooleanClauseAtomic"	100
6.8.1.	SUPABooleanClauseAtomic Attributes	100
6.8.2.	SUPABooleanClauseAtomic Relationships	100
6.9.	The Concrete Class "SUPABooleanClauseComposite"	100
6.9.1.	SUPABooleanClauseComposite Attributes	100
6.9.1.1.	The Attribute "supaBoolClauseBindValue"	101
6.9.1.2.	The Attribute "supaBoolClauseIsCNF"	101
6.9.2.	SUPABooleanClauseComposite Relationships	101
6.9.2.1.	The Aggregation "SUPAHasBooleanClause"	101
6.9.3.	The Association Class "SUPAHasBooleanClauseDetail" ..	101
6.9.3.1.	SUPAHasBooleanClauseDetail Attributes	101
6.10.	The Abstract Class "SUPAECAComponent"	102
6.10.1.	SUPAECAComponent Attributes	102
6.10.1.1.	The Attribute supaECACompIsTerm	102
6.10.2.	SUPAECAComponent Relationships	102
6.11.	The Concrete Class "SUPAPolicyEvent"	103
6.11.1.	SUPAPolicyEvent Attributes	103
6.11.1.1.	The Attribute "supaPolicyEventIsPreProcessed" ..	103
6.11.1.2.	The Attribute "supaPolicyEventIsSynthetic" ...	103
6.11.1.3.	The Attribute "supaPolicyEventTopic[0..n]" ...	103
6.11.1.4.	The Attribute "supaPolicyEventEncoding[1..n]" ..	103
6.11.1.5.	The Attribute "supaPolicyEventData[1..n]"	104
6.11.2.	SUPAPolicyEvent Relationships	104
6.12.	The Concrete Class "SUPAPolicyCondition"	104
6.12.1.	SUPAPolicyCondition Attributes	105
6.12.1.1.	The Attribute "supaPolicyConditionData[1..n]" ..	105
6.12.1.2.	The Attribute "supaPolicyConditionEncoding" ..	105
6.12.2.	SUPAPolicyEvent Relationships	105
6.13.	The Concrete Class "SUPAPolicyAction"	106
6.13.1.	SUPAPolicyAction Attributes	106
6.13.1.1.	The Attribute "supaPolicyActionData[1..n]" ...	106
6.13.1.2.	The Attribute "supaPolicyActionEncoding"	107

6.13.2.	SUPAPolicyAction Relationships	107
-------------------------	--------------------------------------	---------------------

Strassner, et al.

Expires January 19, 2017

[Page 7]

7.	Examples	107
8.	Security Considerations	107
9.	IANA Considerations	107
10.	Contributors	108
11.	Acknowledgments	108
12.	References	108
12.1.	Normative References	108
12.2.	Informative References	108
	Authors' Addresses	109
	Appendix A. Brief Analyses of Previous Policy Work	110

1. Overview

This document defines an information model for representing policies using a common extensible framework that is independent of language, protocol, repository, and the level of abstraction of the content and meaning of a policy. This enables a common set of concepts defined in this information model to be mapped into different representations of policy (e.g., procedural, imperative, and declarative). It also enables different data models that use different languages, protocols, and repositories to optimize their usage. The definition of common policy concepts also provides better interoperability by ensuring that each data model can share a set of common concepts, independent of its level of detail or the language, protocol, and/or repository that it is using. It is also independent of the target data model that will be generated.

This version of the information model focuses on defining one type of policy rule: the event-condition-action (ECA) policy rule. Accordingly, this document defines two sets of model elements:

1. A framework for defining the concept of policy, independent of how policy is represented or used; this is called the SUPA Generic Policy Information Model (GPIM)
2. A framework for defining a policy model that uses the event-condition-action (ECA) paradigm; this is called the SUPA ECA Policy Rule Information Model (EPRIM), and extends concepts from the GPIM.

The combination of the GPIM and the EPRIM provides an extensible framework for defining policy that uses an event-condition-action representation that is independent of data repository, data definition language, query language, implementation language, and protocol.

The Appendices describe how the structure of the GPIM defines a set of generic concepts that enables other types of policies, such as declarative (or "intent-based") policies, to be added later.

1.1. Introduction

Simplified Use of Policy Abstractions (SUPA) defines a technology-independent neutral information model for creating high-level, possibly network-wide policies as input and producing element configurations (either whole or snippets) as output. SUPA addresses the needs of operators, end-users, and application developers to represent multiple types of ECA policy rules, such as for traffic selection and configuration or security. These ECA policy rules may

vary in the level of abstraction to suit the needs of different actors (e.g., end-users vs. administrators) [[1](#)], [[10](#)].

Different constituencies of users would like to use terminology and concepts that are familiar to each constituency. Rather than require multiple software systems to be used for each constituency, a common information model enables these different concepts and terms to be mapped to elements in the information model. This facilitates the use of a single software system to generate data models for each language. In the example shown in Figure 1 (which is a simplified Policy Continuum [10]), each constituency uses different concepts and terms (according to their skill sets and roles) to formulate (ECA) policy rules that are useful for their job functions. A unified information model is one way to build a consensual lexicon that enables terms from one language to be mapped to terms of another language. This approach enables the syntax of each language to be modified appropriate to its user while keeping a common set of semantics for all languages. This is shown in Figure 1.

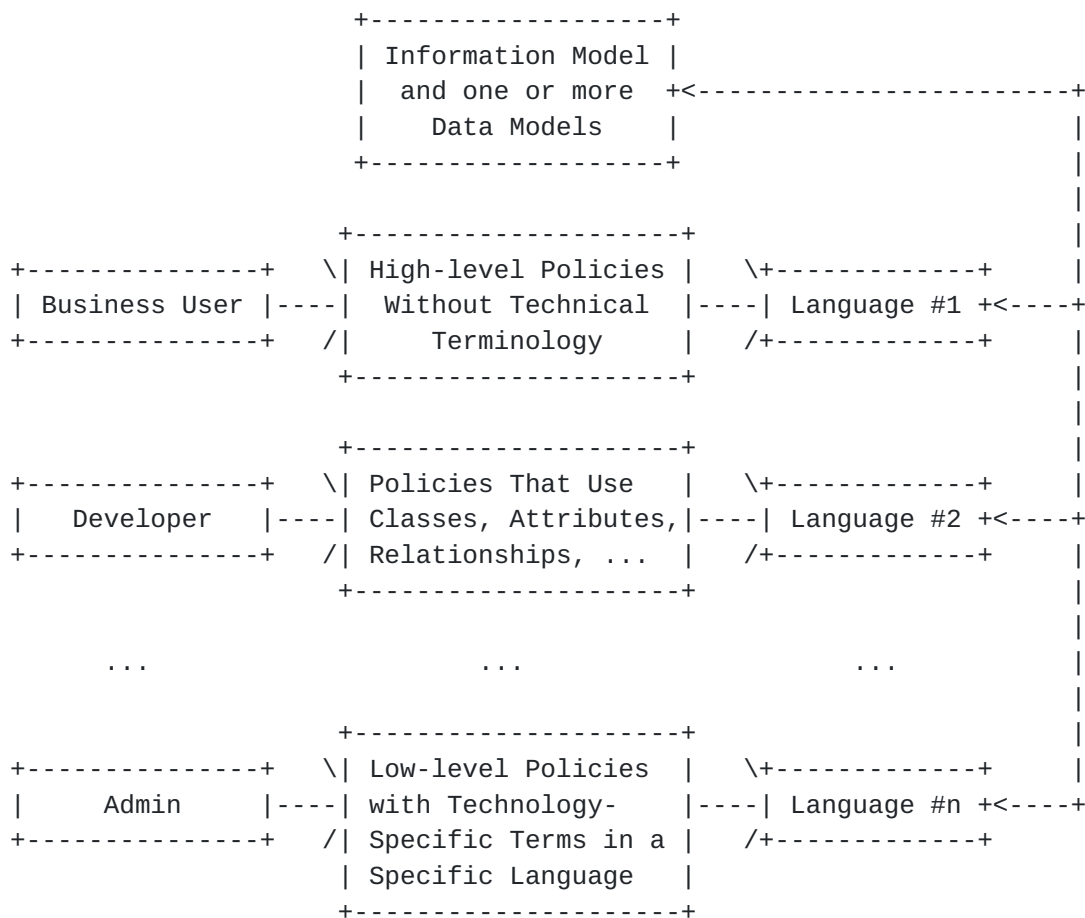


Figure 1. Different Constituencies Need Different Policies

More importantly, an information model defines concepts in a uniform way, enabling formal mapping processes to be developed to translate the information model to a set of data models. This simplifies the process of constructing software to automate the

policy management process. It also simplifies the language generation process, though that is beyond the scope of this document.

Strassner, et al.

Expires January 19, 2017

[Page 10]

This common framework takes the form of an information model that is divided into one high-level module and one or more number of lower-level modules. A lower-level module extends the higher-level module into a new domain; each lower-level domain module can itself be extended to model more granular domain-specific (but still technology- and vendor-independent) concepts as necessary.

Conceptually, a set of model elements (e.g., classes, attributes, constraints, and relationships) are used to define the Generic Policy Information Model (GPIM); this module defines a common set of policy concepts that are independent of the type of policy (e.g., imperative, procedural, declarative, or otherwise). Then, any number of additional modules can be derived from the GPIM; each additional module **MUST** extend the GPIM to define a new type of policy rule by adding to the GPIM. Each additional module **MUST NOT** alter any of the model elements of the GPIM. The use of extensions preserves the interoperability of this approach; if the base GPIM was modified, then this would adversely compromise interoperability.

The SUPA ECA Policy Rule Information Model (EPRIM) extends the GPIM to represent policy rules that use the Event-Condition-Action (ECA) paradigm.

1.2. Changes Since Version -01

There are several changes in this version of this document compared to the previous versions of this document. They are:

- 1) Clarified figure 7 and figure 17
- 2) Aligned enumerations in IM with those in DM and standardized values in the IM version and their explanations
- 3) Removed supaPolExecStatus and supaPolClauseExecStatus
- 4) Added supaPolClauseDeployStatus attribute
- 5) Rewrote SUPAPolicyComponentStructure definition
- 6) Rewrote SUPAPolicyClause definition
- 7) Synchronized information and data models.
- 8) Deleted the attribute "supaEncodedClauseLang[0..n]"
- 9) Fixed typos

2. Conventions Used in This Document

The key 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 [[RFC2119](#)]. In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying [[RFC2119](#)] significance.

3. Terminology

This section defines acronyms, terms, and symbology used in the rest of this document.

3.1. Acronyms

CLI	Command Line Interface
CRUD	Create, Read, Update, Delete
CNF	Conjunctive Normal Form
DNF	Disjunctive Normal Form
ECA	Event-Condition-Action
EPRIM	(SUPA) ECA Policy Rule Information Model
GPIM	(SUPA) Generic Policy Information Model
OAM&P	Operations, Administration, Management, and Provisioning
OID	Object IDentifier
SUPA	Simplified Use of Policy Abstractions
TMF	TeleManagent Forum (TM Forum)
UML	Unified Modeling Language
URI	Uniform Resource Identifier
YANG	A data definition language for use with NETCONF
ZOOM	Zero-touch Orchestration, Operations, and Management (a TMF project that also works on information models)

3.2. Definitions

This section defines the terminology that is used in this document.

3.2.1. Core Terminology

The following subsections define the terms "information model" and "data model", as well as "container" and "policy container".

3.2.1.1. Information Model

An information model is a representation of concepts of interest to an environment in a form that is independent of data repository, data definition language, query language, implementation language, and protocol.

Note: this definition is more specific than that of [[RFC3198](#)], so as to focus on the properties of information models. That definition was: "An abstraction and representation of the entities in a managed environment, their properties, attributes and operations, and the way that they relate to each other. It is independent of any specific repository, software usage, protocol, or platform."

3.2.1.2. Data Model

A data model is a representation of concepts of interest to an environment in a form that is dependent on data repository, data definition language, query language, implementation language, and/or protocol (typically, but not necessarily, all five).

Note: this definition is more specific than that of [[RFC3198](#)], so as to focus on the properties of data models that are generated from information models. That definition was: "A mapping of the contents of an information model into a form that is specific to a particular type of data store or repository."

3.2.1.3. Class

A class is a set of objects that exhibit a common set of characteristics and behavior.

3.2.1.3.1. Abstract Class

An abstract class is a class that cannot be directly instantiated. It MAY have abstract or concrete subclasses. It is denoted with a capital A (for abstract) near the top-left side of the class.

3.2.1.3.2. Concrete Class

A concrete class is a class that can be directly instantiated. Note that classes are either abstract or concrete. In addition, once a class has been defined as concrete in the hierarchy, all of its subclasses MUST also be concrete. It is denoted with a capital C (for concrete) near the top-left side of the class.

3.2.1.4. Container

A container is an object whose instances may contain zero or more additional objects, including container objects. A container provides storage, query, and retrieval of its contained objects in a well-known, organized way.

3.2.1.5. PolicyContainer

In this document, a PolicyContainer is a special type of container that provides at least the following three functions:

1. It uses metadata to define how its content is interpreted
2. It separates the content of the policy from the representation of the policy
3. It provides a convenient control point for OAM&P operations

The combination of these three functions enables a PolicyContainer

to define the behavior of how its constituent components will be accessed, queried, stored, retrieved, and how they operate.

Strassner, et al.

Expires January 19, 2017

[Page 13]

This document does NOT define a specific data type to implementation a PolicyContainer, as many different types of data types can be used. However, the data type chosen SHOULD NOT allow duplicate members in the PolicyContainer. In addition, order is irrelevant, since priority will override any initial order of the members of this PolicyContainer.

3.2.2. Policy Terminology

The following terms define different policy concepts used in the SUPA Generic Policy Information Model (GPIM). Note that the prefix "SUPA" is used for all classes and relationships defined in this model to ensure name uniqueness. Similarly, the prefix "supa" is defined for all SUPA class attributes.

3.2.2.1. SUPAPolicyObject

A SUPAPolicyObject is the root of the GPIM class hierarchy. It is an abstract class that all classes inherit from, except the SUPAPolicyMetadata class and its subclasses.

3.2.2.2. SUPAPolicy

A SUPAPolicy is, in this version of this document, an ECA policy rule that is a type of PolicyContainer. The PolicyContainer MUST contain an ECA policy rule, SHOULD contain one or more SUPAPolicyMetadata objects, and MAY contain other elements that define the semantics of the policy rule. Policies are generically defined as a means to monitor and control the changing and/or maintaining of the state of one or more managed objects [1]. In this context, "manage" means that one or more of the following six fundamental operations are supported: create, read, write, delete, start, and stop) [16].

3.2.2.3. SUPAPolicyClause

A SUPAPolicyClause is an abstract class. Its subclasses define different types of clauses that are used to create the content for different types of SUPAPolicies.

For example, the SUPABooleanClause subclass models the content of a SUPAPolicy as a Boolean clause, where each Boolean clause is made up of a set of reusable objects. In contrast, a SUPAEncodedClause encodes the entire clause as a set of attributes. All types of SUPAPolicies MUST use one or more SUPAPolicyClauses to construct a SUPAPolicy.

3.2.2.4. SUPAECAPolicyRule

An Event-Condition-Action (ECA) Policy (SUPAECAPolicyRule) is an abstract class that is a type of PolicyContainer. It represents a policy rule as a three-tuple, consisting of an event, a condition, and an action clause. In an information model, this takes the form of three different aggregations, one for each clause. Each clause **MUST** be represented by at least one SUPAPolicyClause. Optionally, the SUPAECAPolicyRule **MAY** contain zero or more SUPAPolicySources, zero or more SUPAPolicyTargets, and zero or more SUPAPolicyMetadata objects. Note that for this version of this document, ECA Policy Rules are the ****only**** types of Policies that are defined.

3.2.2.5. SUPAMetadata

Metadata is, literally, data about data. SUPAMetadata is an abstract class that contains prescriptive and/or descriptive information about the object(s) to which it is attached. While metadata can be attached to any information model element, this document only considers metadata attached to classes and relationships.

When defined in an information model, each instance of the SUPAMetadata class **MUST** have its own aggregation relationship with the set of objects that it applies to. However, a data model **MAY** map these definitions to a more efficient form (e.g., flattening the object instances into a single object instance).

3.2.2.6. SUPAPolicyTarget

SUPAPolicyTarget is an abstract class that defines a set of managed objects that may be affected by the actions of a SUPAPolicyClause. A SUPAPolicyTarget may use one or more mechanisms to identify the set of managed objects that it affects; examples include OIDs and URIs.

When defined in an information model, each instance of the SUPAPolicyTarget class **MUST** have its own aggregation relationship with each SUPAPolicy that uses it. However, a data model **MAY** map these definitions to a more efficient form (e.g., flattening the SUPAPolicyTarget, SUPAMetadata, and SUPAPolicy object instances into a single object instance).

3.2.2.7. SUPAPolicySource

SUPAPolicySource is an abstract class that defines a set of managed objects that authored this SUPAPolicyClause. This is required for auditability and authorization policies, as well

as some forms of deontic and alethic logic.

Strassner, et al.

Expires January 19, 2017

[Page 15]

A SUPAPolicySource may use one or more mechanisms to identify the set of managed objects that authored it; examples include OIDs and URIs. Specifically, policy CRUD MUST be subject to authentication and authorization, and MUST be auditable. Note that the mechanisms for doing these three operations are currently not included, and are for further discussion.

When defined in an information model, each instance of the SUPAPolicySource class MUST have its own aggregation relationship with each SUPAPolicy that uses it. However, a data model MAY map these definitions to a more efficient form (e.g., flattening the SUPAPolicySource, SUPAMetadata, and SUPAPolicy object instances into a single object instance).

3.2.3. Modeling Terminology

The following terms define different types of relationships used in the information models of the SUPA Generic Policy Information Model (GPIM).

3.2.3.1. Inheritance

Inheritance makes an entity at a lower level of abstraction (e.g., the subclass) a type of an entity at a higher level of abstraction (e.g., the superclass). Any attributes and relationships that are defined for the superclass are also defined for the subclass. However, a subclass does NOT change the characteristics or behavior of the attributes or relationships of the superclass that it inherits from. Formally, this is called the Liskov Substitution Principle [7]. This principle is one of the key characteristics that is NOT followed in [4], [6], [RFC3060], and [RFC3460].

A subclass MAY add new attributes and relationships that refine the characteristics and/or behavior of it compared to its superclass. A subclass MUST NOT change inherited attributes or relationships.

3.2.3.2. Relationship

A relationship is a generic term that represents how a first set of entities interact with a second set of entities. A recursive relationship sets the first and second entity to the same entity. There are three basic types of relationships, as defined in the subsections below: associations, aggregations, and compositions.

A subclass MUST NOT change the multiplicity (see [section 3.2.3.7](#)) of a relationship that it inherits. A subclass MUST NOT change any attributes of a relation that it inherits that is realized using an association class (see [section 3.2.3.6](#)).

3.2.3.3. Association

An association represents a generic dependency between a first and a second set of entities. In an information model, an association MAY be represented as a class.

3.2.3.4. Aggregation

An aggregation is a stronger type (i.e., more restricted semantically) of association, and represents a whole-part dependency between a first and a second set of entities. Three objects are defined by an aggregation: the first entity, the second entity, and a new third entity that represents the combination of the first and second entities.

The entity owning the aggregation is referred to as the "aggregate", and the entity that is aggregated is referred to as the "part". In an information model, an aggregation MAY be represented as a class.

3.2.3.5. Composition

A composition is a stronger type (i.e., more restricted semantically) of aggregation, and represents a whole-part dependency with two important behaviors. First, an instance of the part is included in at most one instance of the aggregate at a time. Second, any action performed on the composite entity (i.e., the aggregate) is propagated to its constituent part objects. For example, if the composite entity is deleted, then all of its constituent part entities are also deleted. This is not true of aggregations or associations - in both, only the entity being deleted is actually removed, and the other entities are unaffected. In an information model, a composition MAY be represented as a class.

3.2.3.6. Association Class

A relationship may be implemented as an association class. This is used to define the relationship as having its own set of features. (Note: in this document, all relationships are implemented as association classes for consistency and to simplify implementation.) More specifically, if the relationship is implemented as an association class, then the attributes of the association class, as well as other relationships that the association class participates in, may be used to define the semantics of the relationship. If the relationship is not implemented as an association class, then no additional semantics (beyond those defined by the type of the relationship) are expressed by the relationship.


3.2.3.7. Multiplicity

A specification of the range of allowable cardinalities that a set of entities may assume. This is always a pair of ranges, such as 1 - 1 or 0..n - 2..5.

3.2.3.8. Navigability

A relationship may restrict one object from accessing the other object. This document defines two choices:

1. Each object is navigable by the other, which is indicated by NOT providing any additional symbology, or
2. An object A can navigate to object B, but object B cannot navigate to object A. This is indicated by an open-headed arrow pointing to the object that cannot navigate to the other object. An example is shown below:



The above figure shows a navigability restriction. Class A can navigate to Class B, but Class B cannot navigate to Class A. This is a mandatory association, since none of the multiplicities contain a '0'. This association reads as follows:


Class A depends on 3 to 4 instances of Class B, and
Class B depends on 1 to 2 instances of Class A.

3.3. Symbology

The following symbology is used in this document.

3.3.1. Inheritance

Inheritance: a subclass inherits the attributes and relationships of its superclass, as shown below:



+-----+

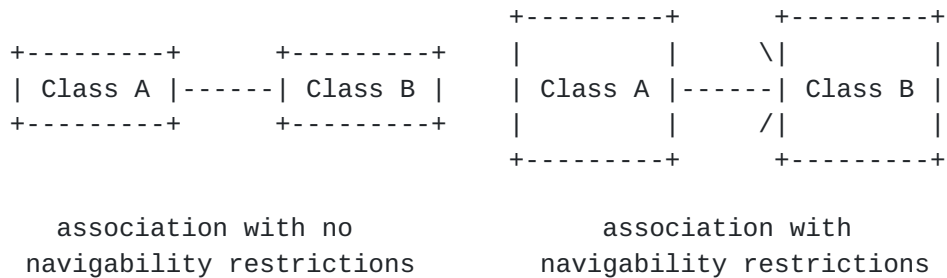
Strassner, et al.

Expires January 19, 2017

[Page 18]

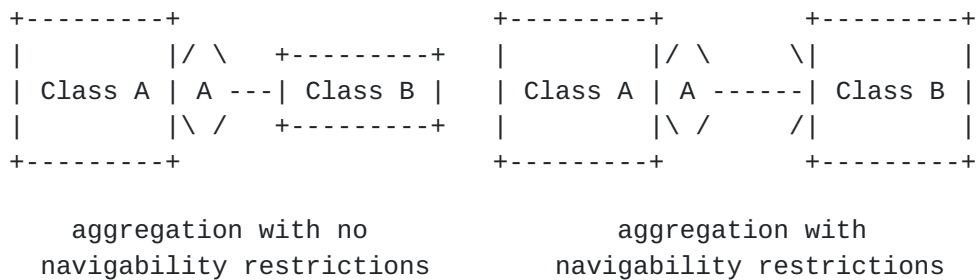
3.3.2. Association

Association: Class B depends on Class A, as shown below:



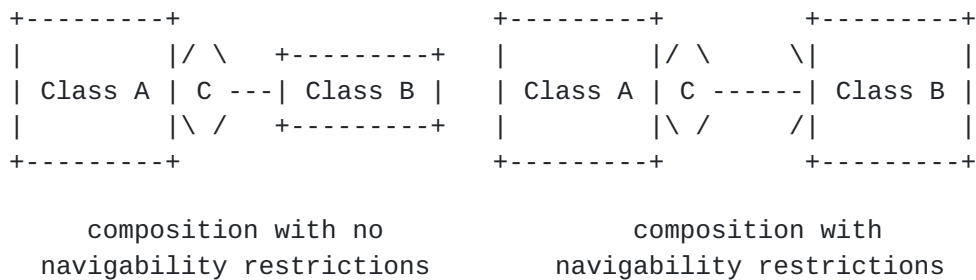
3.3.3. Aggregation

Aggregation: Class B is the part, Class A is the aggregate, as shown below:



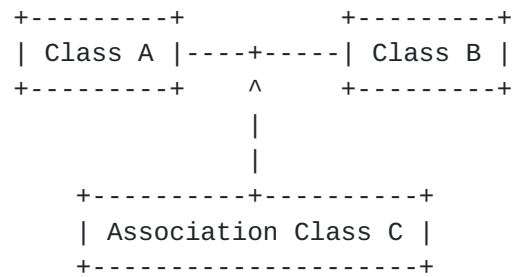
3.3.4. Composition

Composition: Class B is the part, Class A is the composite,
as shown below:



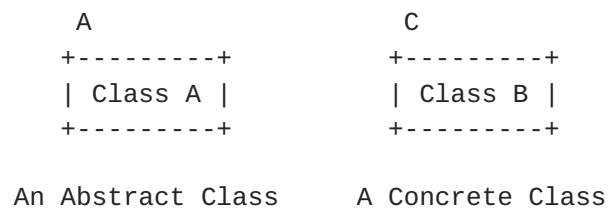
3.3.5. Association Class

Association Class: Class C is the association class implementing the relationship D between classes A and B



[3.3.6.](#) Abstract vs. Concrete Classes

In UML, abstract classes are denoted with their name in italics. For this draft, a capital 'A' will be placed at either the top left or right corner of the class to signify that the class is abstract. Similarly, a capital 'C' will be placed in the same location to represent a concrete class. This is shown below.



4. Policy Abstraction Architecture

This section describes the motivation for the policy abstractions that are used in SUPA. The following abstractions are provided:

- o The GPIM defines a technology-neutral information model that can express the concept of Policy.
 - o All classes, except for SUPAPolicyMetadata, inherit from SUPAPolicyObject, or one of its subclasses.
 - o SUPAPolicyObject and SUPAPolicyMetadata are designed to inherit from classes in another model; the GPIM does not define an "all-encompassing" model.
- o This version of this document restricts the expression of Policy to a set of event-condition-action clauses.
 - o Each clause is defined as a Boolean expression, and MAY also be defined as a reusable object.
 - o Clauses may be combined to form more complex Boolean expressions.
- o The purpose of the GPIM is to enable different policies that have fundamentally different representations to share common model elements. Policy statements, which are implemented as instances of the SUPAPolicyClause class, separates the content of a Policy from its representation. This is supported by:
 - o All policy rules (of which SUPAECAPolicyRule is the first example of a concrete class) are derived from the SUPAPolicyStructure class.
 - o All objects that are components of policy rules are derived from the SUPAPolicyComponentStructure class.
 - o A SUPAPolicy MUST contain at least one SUPAPolicyClause.
 - o A SUPAPolicy MAY specify one or more SUPAPolicyTarget, SUPAPolicySource, and SUPAPolicyMetadata objects to augment the semantics of the SUPAPolicy
- o A SUPAPolicyClause has two subclasses:
 - o A SUPABooleanClause, which is used to build SUPAECAPolicyRules from reusable objects.
 - o A SUPAEncodedClause, which is used for using attributes instead of objects to construct a SUPAECAPolicyRule.
- o A SUPAECAPolicyRule defines the set of events and conditions that are responsible for executing its actions; it MUST have at least one event clause, at least one condition clause, and at least one action clause.
 - o The action(s) of a SUPAECAPolicyRule are ONLY executed if both the event and condition clauses evaluate to TRUE
 - o A SUPAPolicyAction MAY invoke another SUPAECAPolicyRule (see [section 6.13](#)).
- o SUPAMetadata MAY be defined for any SUPAPolicyObject class.
- o SUPAMetadata MAY be prescriptive and/or descriptive in nature.

4.1. Motivation

The power of policy management is its applicability to many different types of systems. There are many different actors that can use a policy management system, including end-users, operators, application developers, and administrators. Each of these constituencies have different concepts and skills, and use different terminology. For example, an operator may want to express an operational rule that states that only Platinum and Gold users can use streaming multimedia applications. As a second example, a network administrator may want to define a more concrete policy rule that looks at the number of dropped packets and, if that number exceeds a programmable threshold, changes the queuing and dropping algorithms used.

SUPA may be used to define other types of policies, such as for systems and operations management; an example is: "All routers and switches must have password login disabled". See section 3 of [\[8\]](#) for additional declarative and ECA policy examples.

All of the above examples are commonly referred to as "policy rules", but they take very different forms, since they are at very different levels of abstraction and typically authored by different actors. The first was very abstract, and did not contain any technology-specific terms, while the second was more concrete, and likely used technical terms of a general (e.g., IP address range, port numbers) as well as a vendor-specific nature (e.g., specific queuing, dropping, and/or scheduling algorithms implemented in a particular device). The third restricted the type of login that was permissible for certain types of devices in the environment.

Note that the first two policy rules could directly affect each other. For example, Gold and Platinum users might need different device configurations to give the proper QoS markings to their streaming multimedia traffic. This is very difficult to do if a common policy model does not exist, especially if the two policies are authored by different actors that use different terminology and have different skill sets. More importantly, the users of these two policies likely have different job responsibilities. They may have no idea of the concepts used in each policy. Yet, their policies need to interact in order for the business to provide the desired service. This again underscores the need for a common policy framework.

Certain types of policy rules (e.g., ECA) may express actions, or other types of operations, that contradict each other. SUPA provides a rich object model that can be used to support language

definitions that can find and resolve such problems.

Strassner, et al.

Expires January 19, 2017

[Page 22]

4.2. SUPA Approach

The purpose of the SUPA Generic Policy Information Model (GPIM) is to define a common framework for expressing policies at different levels of abstraction. SUPA uses the GPIM as a common vocabulary for representing policy concepts that are independent of language, protocol, repository, and level of abstraction. This enables different actors to author and use policies at different levels of abstraction. This forms a policy continuum [1] [2], where more abstract policies can be translated into more concrete policies, and vice-versa.

Most systems define the notion of a policy as a single entity. This assumes that all users of policy have the same terminology, and use policy at the same level of abstraction. This is rarely, if ever, true in modern systems. The policy continuum defines a set of views (much like RM-ODP's viewpoints [9]) that are each optimized for a user playing a specific role. SUPA defines the GPIM as a standard vocabulary and set of concepts that enable different actors to use different formulations of policy. This corresponds to the different levels in the policy continuum, and as such, can make use of previous experience in this area.

It may be necessary to translate a Policy from a general to a more specific form (while keeping the abstraction level the same). For example, the declarative policy "Every network attached to a VM must be a private network owned by someone in the same group as the owner of the VM" may be translated to more formal form (e.g., Datalog (as in OpenStack Congress)). It may also be necessary to translate a Policy to a different level of abstraction. For example, the previous Policy may need to be translated to a form that network devices can process directly. This requires a common framework for expressing policies that is independent of the level of abstraction that a Policy uses.

4.3. SUPA Generic Policy Information Model Overview

Figure 2 illustrates the approach for representing policy rules in SUPA. The top two layers are defined in this document; the bottom layer (Data Models) are defined in separate documents. Conceptually, the GPIM defines a set of objects that define the key elements of a Policy independent of how it is represented or its content. As will be shown, there is a significant difference between SUPAECAPolicyRules (see [Section 6](#)) and other types of policies (see [Section 7](#)). In principle, other types of SUPAPolicies could be defined, but the current charter is restricted to using only event-condition-action SUPAPolicies as exemplars.

Note: the GPIM MAY be used without the EPRIM. However, in order to use the EPRIM, the GPIM MUST also be used.

Strassner, et al.

Expires January 19, 2017

[Page 23]

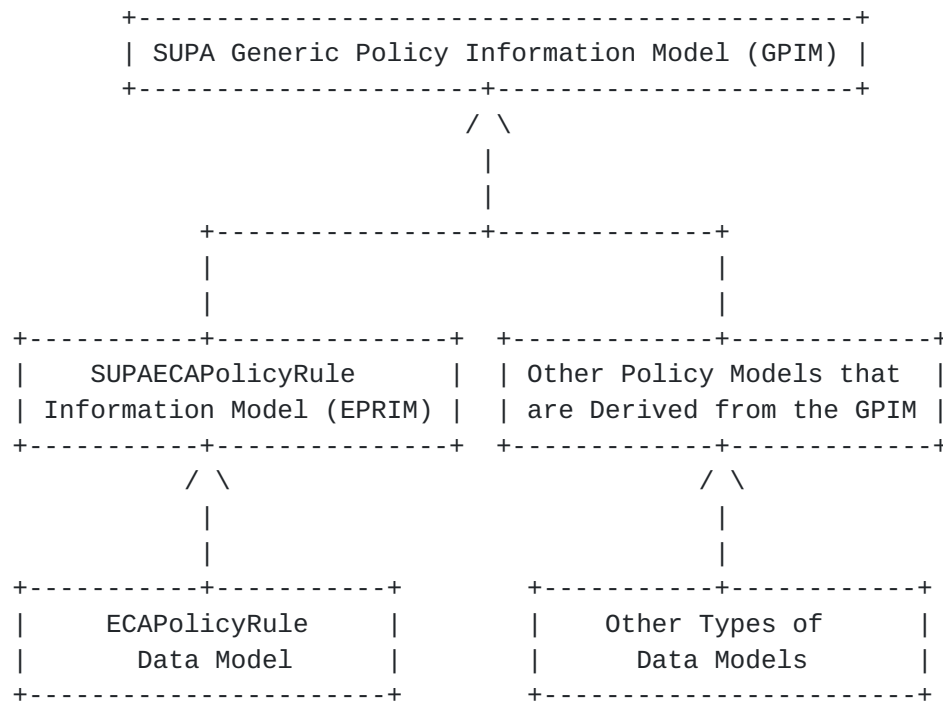


Figure 2. Overview of SUPA Policy Rule Abstractions

This draft defines the GPIM and EPRIM. This draft further assumes that the SUPA Information Model is made up of either the GPIM or the combination of the GPIM and the EPRIM. Extensions to both the GPIM and the EPRIM can be made as long as these extensions do not conflict with the content and structure defined in the GPIM and EPRIM. If the GPIM and EPRIM are part of another information model, then they should collectively still define a single information model. The GPIM defines the following concepts:

- o A class defining the top of the GPIM class hierarchy, called SUPAPolicyObject
- o Four subclasses of SUPAPolicyObject, representing:
 - o the top of the Policy hierarchy, called SUPAPolicyStructure
 - o the top of the Policy component hierarchy, called SUPAPolicyComponentStructure
 - o PolicySource
 - o PolicyTarget

The SUPAPolicyStructure class is the superclass for all types of Policies (e.g., imperative, declarative, and others). This document is currently limited to imperative (e.g., ECA) policies. However, care has been taken to ensure that the attributes and relationships of the SUPAPolicyStructure class are extensible, and can be used for more types of policies than just ECA policies.

This yields the following high-level structure:

Strassner, et al.

Expires January 19, 2017

[Page 24]

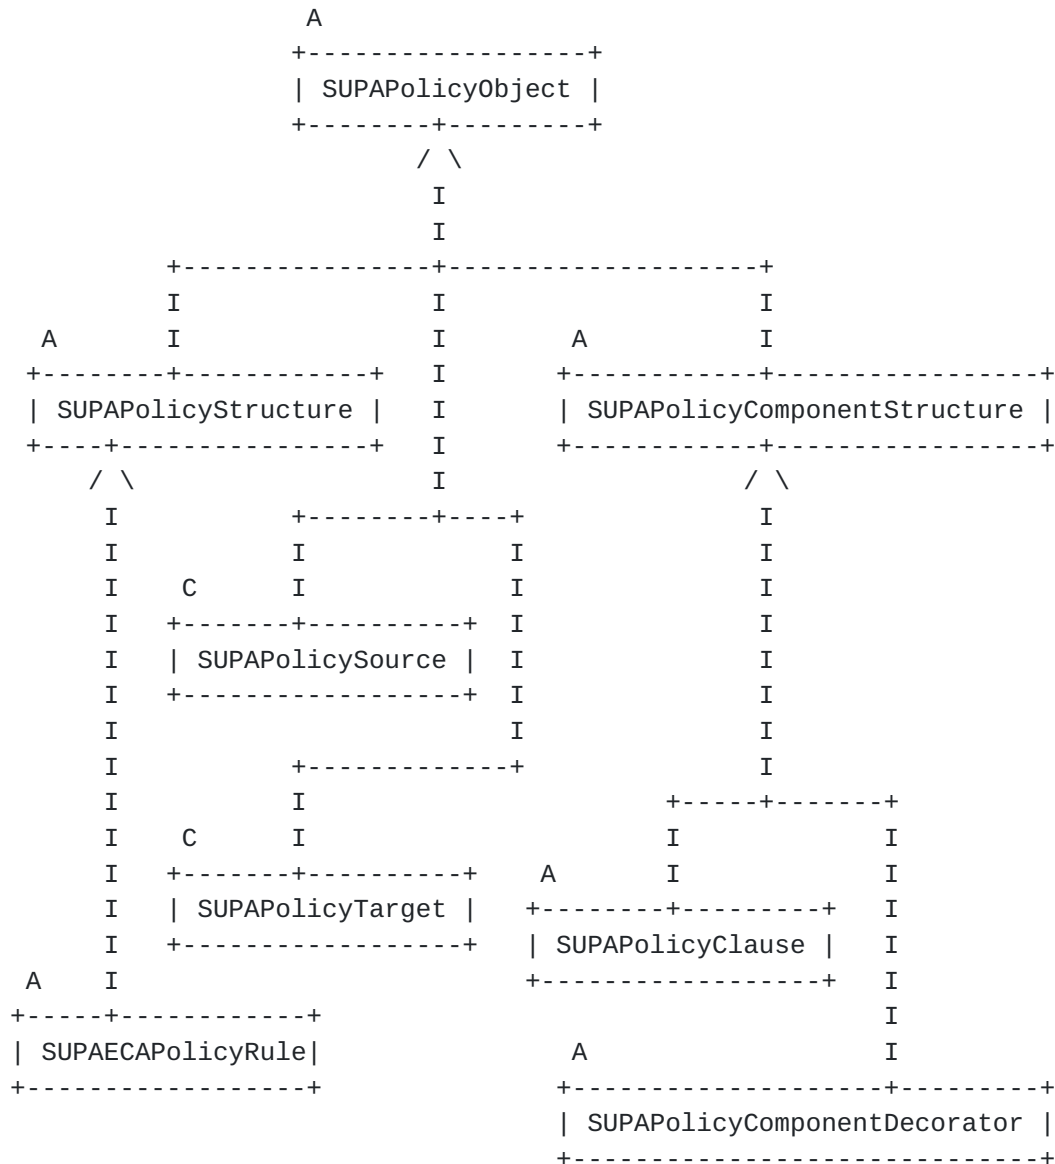


Figure 3. Functional View of the Top-Level GPIM

Note that all classes except the SUPAPolicySource and the SUPAPolicyTarget classes are defined as abstract. This provides more freedom for the data modeler in implementing the data model. For example, if the data model uses an object-oriented language, such as Java, then the above structure enables all of the abstract classes to be collapsed into a single concrete class. If this is done, attributes as well as relationships are inherited.

4.3.1. SUPAPolicyObject

A SUPAPolicyObject serves as a single root of the SUPA system (i.e., all other classes in the model are subclasses of the SUPAPolicyObject class) except for the Metadata objects, which are

in a separate class hierarchy. This simplifies code generation and reusability. It also enables SUPAPolicyMetadata objects to be attached to any appropriate subclass of SUPAPolicyObject.

4.3.2. SUPAPolicyStructure

SUPAPolicyStructure is an abstract superclass that is the base class for defining different types of policies (however, in this version of this document, only ECA policy rules are modeled). It serves as a convenient aggregation point to define atomic (i.e., individual policies that can be used independently) and composite (i.e., hierarchies of policies) SUPAPolicies; it also enables PolicySources and/or PolicyTargets to be associated with a given set of Policies.

SUPAPolicies are defined as either a stand-alone PolicyContainer or a hierarchy of PolicyContainers. A PolicyContainer specifies the structure, content, and optionally, source, target, and metadata information for a SUPAPolicy. This is implemented by the subclasses of SUPAPolicyStructure. For example, the composite pattern is used to create two subclasses of the SUPAECAPolicyRule class; SUPAECAPolicyRuleAtomic is used for stand-alone policies, and SUPAECAPolicyRuleComposite is used to build hierarchies of policies.

This document defines a SUPAPolicy as an ECA Policy Rule, though the GPIM enables other types of policies to be defined and used with an ECA policy rule. The GPIM model is used in [2] and [5], along with extensions that allow [2] and [5] to define multiple types of policies that are derived from the GPIM. They also allow different combinations of different types of policy rules to be used with each other. Most previous work cannot define different types of policy rules; please see [Appendix A](#) for a comparison to previous work.

4.3.3. SUPAPolicyComponentStructure

SUPAPolicyComponentStructure is an abstract superclass that is the base class for defining components of different types of policies. SUPAPolicyStructure subclasses define the structure of a policy, while SUPAPolicyComponentStructure subclasses define the content that is contained in the structure of a policy. For example, a SUPAECAPolicyRule is an imperative policy rule, and defines its structure; its event, condition, and action clauses are populated by SUPAPolicyComponentStructure subclasses. The strength of this design is that different types of policies (e.g., imperative and declarative policies) can be represented using a common set of policy components.

Please see Appendix for a comparison to previous work.

4.3.4. SUPAPolicyClause

All policies derived from the GPIM are made up of one or more SUPAPolicyClauses, which define the content of the Policy. This enables a Policy of one type (e.g., ECA) to invoke Policies of the same or different types. SUPAPolicyClause is an abstract class, and serves as a convenient aggregation point for assembling other objects that make up a SUPAPolicyClause.

The GPIM defines a single concrete subclass of SUPAPolicyClause, called SUPAEncodedClause. This is a generic clause, and can be used by any type of Policy in a stand-alone fashion. It can also be used in conjunction with other SUPAPolicyClauses. The EPRIM also defines a subclass of SUPAPolicyClause; see [section 6.7](#)).

The structure of the GPIM is meant to provide an extensible framework for defining different types of policies. This is demonstrated by the EPRIM (see [section 6](#)) and the LSIM (see the Appendices) that each define new subclasses of SUPAPolicyClause (i.e., SUPABooleanClause and SUPALogicClause, respectively) without defining new classes that have no GPIM superclass.

A SUPAPolicyClause is defined as an object. Therefore, clauses and sets of clauses are objects, which promotes reusability.

4.3.5. SUPAPolicyComponentDecorator

One of the problems in building a policy model is the tendency to have a multitude of classes, and hence object instances, to represent different combinations of policy events, conditions, and actions. This can lead to class and/or relationship explosion. Please see [Appendix A](#) for a comparison to previous work.

SUPAPolicyClauses are constructed using the Decorator Pattern [11]. This is a design pattern that enables behavior to be selectively added to an individual object, either statically or dynamically, without affecting the behavior of other objects from the same class. The decorator pattern uses composition, instead of inheritance, to avoid class and relationship explosion. The decorator pattern also enable new objects to be composed from parts or all of existing objects without affecting the existing objects.

This enables the resulting SUPAPolicyClause to be constructed completely from objects in the SUPA information model. This facilitates the construction of policies at runtime by a machine. This is also true of [2] and [5]; however, this is NOT true of most other models. Please see [Appendix A](#) for a comparison to previous work.

SUPAPolicyComponentDecorator defines four types of objects that can be used to form a SUPAPolicyClause. Each object may be used with all other objects, if desired. The first three are defined in the GPIM, with the last defined in the EPRIM. The objects are:

- o SUPAPolicyTerm, which enables a clause to be defined in a canonical {variable, operator, value} form
- o SUPAGenericDecoratedComponent, which enabled a custom object to be defined and then used in a SUPAPolicyClause
- o SUPAPolicyCollection, which enables a collection of objects to be gathered together and associated with all or a portion of a SUPAPolicyClause
- o SUPAECAComponent, which defines Events, Conditions, and Actions as reusable objects

This approach facilitates the machine-driven construction of policies. Note that this is completely optional; policies do not have to use these constructs.

4.3.6. SUPAPolicyTarget

A SUPAPolicyTarget is a set of managed entities that a SUPAPolicy is applied to. A managed entity can only be designated a SUPAPolicyTarget if it can process actions from a SUPAPolicy.

A managed object may not be in a state that enables management operations to be performed on it. Furthermore, the policy-based management system SHOULD ensure that the management entity performing the management operations has the proper permissions to perform the management operations. The design of the SUPAPolicyTarget addresses both of these criteria.

4.3.7. SUPAPolicySource

A SUPAPolicySource is a set of managed entities that authored, or are otherwise responsible for, this SUPAPolicy. Note that a SUPAPolicySource does NOT evaluate or execute SUPAPolicies. Its primary use is for auditability and the implementation of deontic and/or alethic logic.

4.4. The Design of the GPIM

This section describes the overall design of the GPIM.

The GPIM defines a policy as a type of PolicyContainer. For this version, only ECA Policy Rules will be described. However, it should be noted that the mechanism described is applicable to other types of policies (e.g., declarative) as well.

4.4.1. Structure of Policies

Recall that a PolicyContainer was defined as a special type of container that provides at least the following three functions:

1. It uses metadata to define how its content is described and/or prescribed
2. It separates the content of the policy from the representation of the policy
3. It provides a convenient control point for OAMP operations.

The first requirement is provided by the ability for any subclass of Policy (the root of the information model) to aggregate one or more concrete instances of a SUPAPolicyMetadata class. This is explained in detail in [section 5.2.2](#).

The second requirement is met by representing an ECA Policy as having two parts: (1) a rule part and (2) components that make up the rule. Since functional and declarative policies are not, strictly speaking, "rules", the former is named PolicyStructure, while the latter is named PolicyComponentStructure.

The third requirement is met by the concrete subclasses of PolicyStructure. Since they are PolicyContainers, they are made up of the SUPAECAPolicyRule, its components, and any metadata that applies to the PolicyContainer, the SUPAECAPolicyRule, and/or any components of the SUPAECAPolicyRule. This provides optional low-level control over any part of the SUPAECAPolicyRule.

The above requirements result in the design shown in Figure 4.

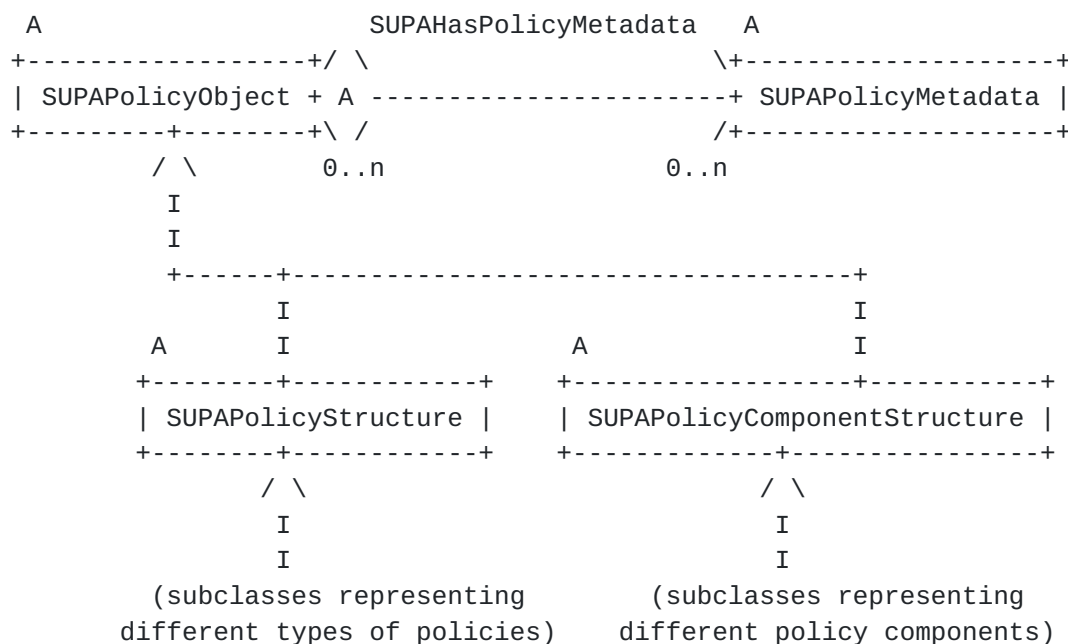


Figure 4. Structure of a Policy

Strassner, et al.

Expires January 19, 2017

[Page 29]

Note that aggregation in Figure 4 (named SUPAHasPolicyMetadata) is realized as an association class, in order to manage which set of Metadata can be aggregated by which SUPAPolicyObject. The combination of these three functions enables a PolicyContainer to define the behavior of how its constituent components will be accessed, queried, stored, retrieved, and how they operate.

It is often necessary to construct groups of policies. The GPIM follows [2] and [5], and uses the composite pattern [11] to implement this functionality, as shown in Figure 5 below. There are a number of advantages to using the composite pattern over a simple relationship, as detailed in [11].

Figure 5 shows that SUPAPolicyStructure has a single subclass, called SUPAECAPolicyRule. Note, however, that other types of policies, such as declarative policies, can be defined as subclasses of SUPAPolicyStructure in the future.

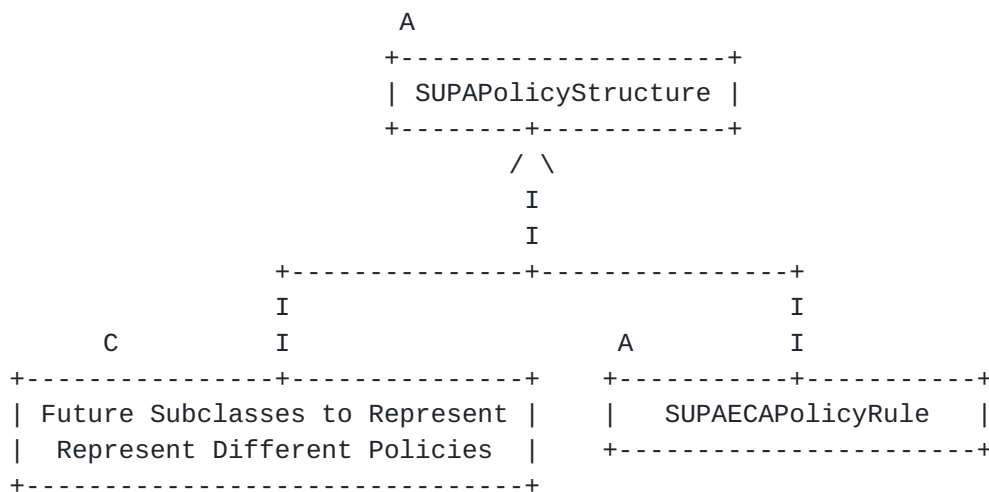


Figure 5. The Composite Pattern Applied to SUPAPolicyStructure

4.4.2. Representing an ECA Policy Rule

An ECA policy rule is a 3-tuple, which is made up of an event clause, a condition clause, and an action clause. Each of these three types of clauses may in turn be made up of a Boolean combination of clauses of that type. Each clause may be viewed as a predicate, as it provides a TRUE or FALSE output. The canonical form of a clause is a 3-tuple of the form "variable operator value", and can be made into more complex Boolean expressions. For example, the SUPAPolicyClause: "((A AND B) OR NOT (C AND D))" consists of two clauses, "(A AND B)" and "(C OR D)", that are combined together using the operators OR and NOT.

A SUPAECAPolicyRule is defined (in the EPRIM) as an abstract subclass of SUPAPolicyStructure.

Strassner, et al.

Expires January 19, 2017

[Page 30]

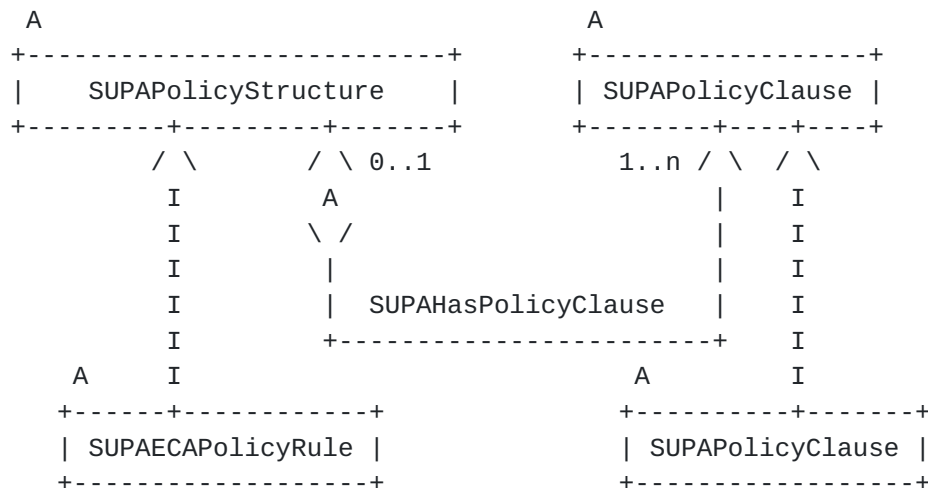
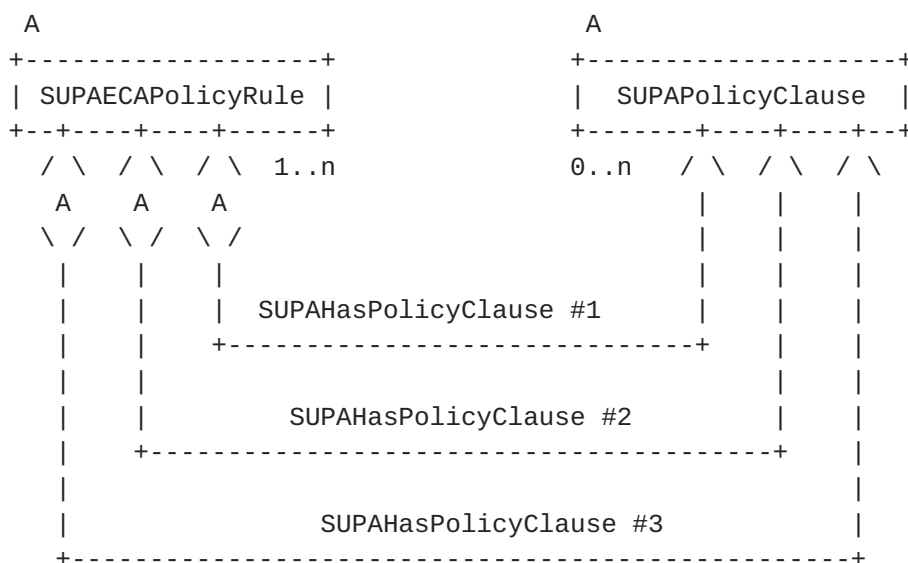


Figure 6. SUPAECAPolicyRule Aggregating SUPAPolicyClauses

Note that the aggregation SUPAHasPolicyClause in Figure 6 is realized as an association class, in order to manage which set of SUPAPolicyClauses can be aggregated by which set of SUPAECAPolicyRules. This aggregation is defined at the SUPAPolicyStructure level, and not at the lower level of SUPAECAPolicyRule, so that non-ECA policies can also use this aggregation.

Since a SUPAECAPolicyRule consists of three SUPAPolicyClauses, at least three separate instances of the SUPAHasPolicyClause aggregation are instantiated in order to make a complete SUPAECAPolicyRule, as shown in Figure 7.



note: all 3 aggregations have a multiplicity of 1..n - 0..n

Figure 7. Instantiating a SUPAECAPolicyRule, part 1

Strassner, et al.

Expires January 19, 2017

[Page 31]

In figure 7, SUPAECAPolicyRule is shown as "owning" these three aggregations, since it inherits them from its superclass (SUPAPolicyStructure). The three aggregations represent the event, condition, and action clauses of a SUPAECAPolicyRule. Note that each of these clauses MAY consist of one or more SUPAPolicyClauses. Similarly, each SUPAPolicyClause MAY consist of one or more predicates. In this way, complex event, condition, and action clauses, which are combinations of Boolean expressions that form a logical predicate) are supported, without having to define additional objects (as is done in previous work; please see [Appendix A](#) for a comparison to previous work).

The multiplicity of the SUPAHasPolicyClause aggregation is 0..n on the aggregate side and 1..n on the part side. This means that a particular SUPAECAPolicyRule MUST aggregate at least one SUPAPolicyClause, and that a given SUPAPolicyClause MAY be aggregated by zero or more SUPAECAPolicyRule objects.

This cardinality MAY be refined to 3..n for SUPAECAPolicyRules, since a SUPAECAPolicyRule MUST have at least three separate clauses. However, since a SUPAPolicyStructure is the owner of this aggregation (which is inherited by SUPAECAPolicyRule), the cardinality is defined to be 1..n on the part side because other types of Policies have different needs. The 0..n cardinality means that a SUPAPolicyClause may be aggregated by zero or more SUPAECAPolicyRules. The zero is provided so that SUPAPolicyClauses can be stored in (for example) a repository before the SUPAECAPolicyRule is created; the "or more" recognizes the fact that multiple SUPAECAPolicyRules could aggregate the same SUPAPolicyClause.

In Figure 7, suppose that SUPAHasPolicyClause#1, #2, and #3 represent the aggregations for the event, condition, and action clauses, respectively. This means that each of these SUPAHasPolicyClause aggregations must explicitly identify the type of clause that it represents.

In looking at Figure 7, there is no difference between any of the three aggregations, except for the type of clause that the aggregation represents (i.e., event, condition, or action clause).

Therefore, three different aggregations, each with their own association class, is not needed. Instead, the GPIM defines a single aggregation (SUPAHasPolicyClause) that is realized using a (single) abstract association class (SUPAHasPolicyClauseDetail); this association class is then subclassed into three concrete subclasses, one each to represent the semantics for an event, condition, and action clause.

The policy management system may use any number of different software mechanisms, such as introspection or reflection, to determine the nature of the aggregation (i.e., what object types are being aggregated) in order to select the appropriate subclass of SUPAHasPolicyClauseDetail. The three subclasses of SUPAHasPolicyClauseDetail are named SUPAHasPolicyEventDetail, SUPAHasPolicyConditionDetail, and SUPAHasPolicyActionDetail, respectively. While Event, Condition, and Action objects are typically used in ECA policy rules, the design in this document enables them to be used as policy components of other types of policies as well. This is shown in Figure 8.

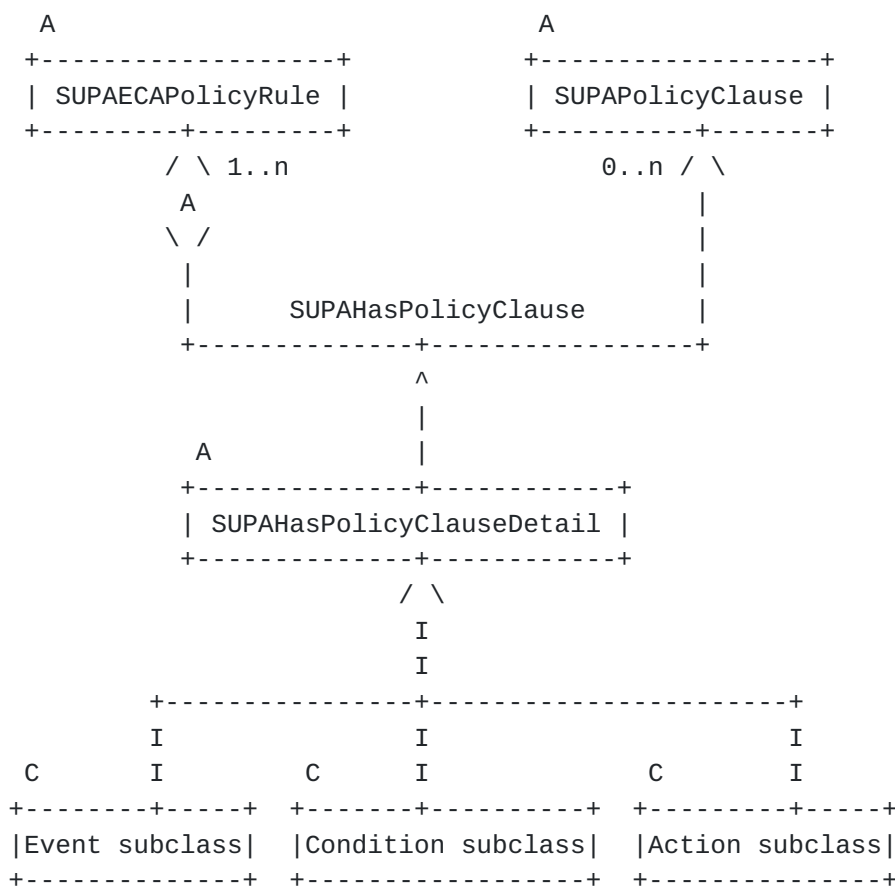


Figure 8. Instantiating a SUPAECAPolicyRule, part 2

4.4.3. Creating SUPA Policy Clauses

There are two different types of Policy Components. They are a SUPAPolicyClause and a SUPAPolicyComponentDecorator. The former is used to construct SUPAECAPolicyRules, while the latter is used to add behavior to a SUPAPolicyClause. This enables the structure

and capabilities of the SUPAPolicyClause to be adjusted dynamically at runtime.

Strassner, et al.

Expires January 19, 2017

[Page 33]

However, since each SUPAECAPolicyRule can be made up of a variable number of SUPAPolicyComponents, the decorator pattern is used to "wrap" any concrete subclass of SUPAPolicyClause with zero or more concrete subclasses of the PolicyComponentDecorator object. This avoids problems of earlier models that resulted in a proliferation of classes and relationships.

Figure 9 shows these two class subclasses. Note that the decorator pattern [11] is used to enable subclasses of the SUPAPolicyComponentDecorator class to add their attributes and/or behavior to a SUPAPolicyClause (as stated in [section 4.3](#)) without affecting the behavior of other objects from the same class. More specifically, concrete subclasses of the (abstract) SUPAPolicyComponentDecorator class can be used to decorate, or "wrap", any of the concrete subclasses of the (abstract) SUPAPolicyClause class.

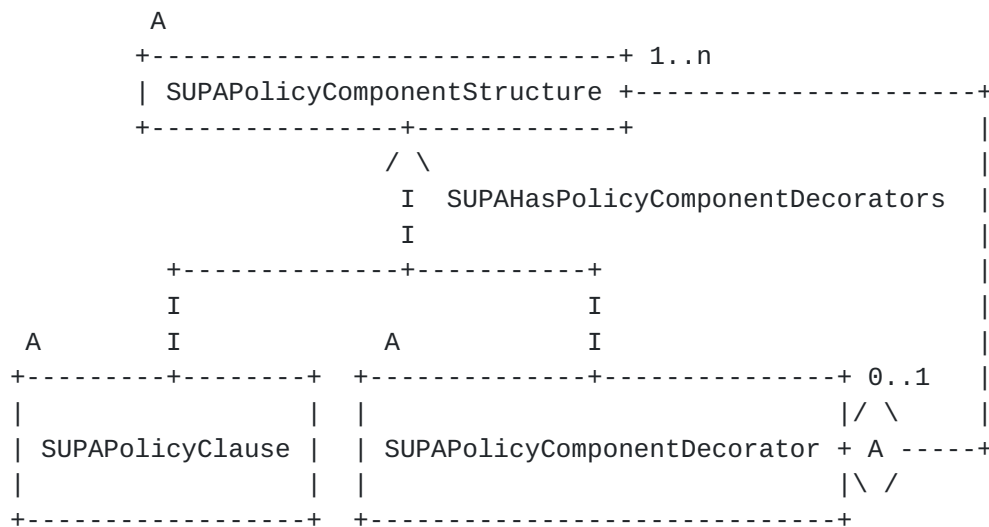


Figure 9. Subclasses of SUPAPolicyComponentStructure

Instead of using inheritance to statically create new classes to represent new types of objects, the decorator pattern uses composition to dynamically combine attributes and behavior from existing objects into new objects. This is done by defining an interface in SUPAPolicyComponent that all of the subclasses of SUPAPolicyComponent conform to. Since the subclasses are of the same type as SUPAPolicyComponent, they all have the same interface. This allows each concrete SUPAPolicyComponentDecorator subclass to add its attributes and/or behavior to the concrete subclass of SUPAPolicyClause that it is decorating (or "wrapping").

This represents an important design optimization for data models. Note that a single SUPAECAPolicyRule can consist of any number of SUPAPolicyClauses, each of very different types. If inheritance

was used, then a subclass AND an aggregation would be required for each separate clause that makes up the policy rule.

Strassner, et al.

Expires January 19, 2017

[Page 34]

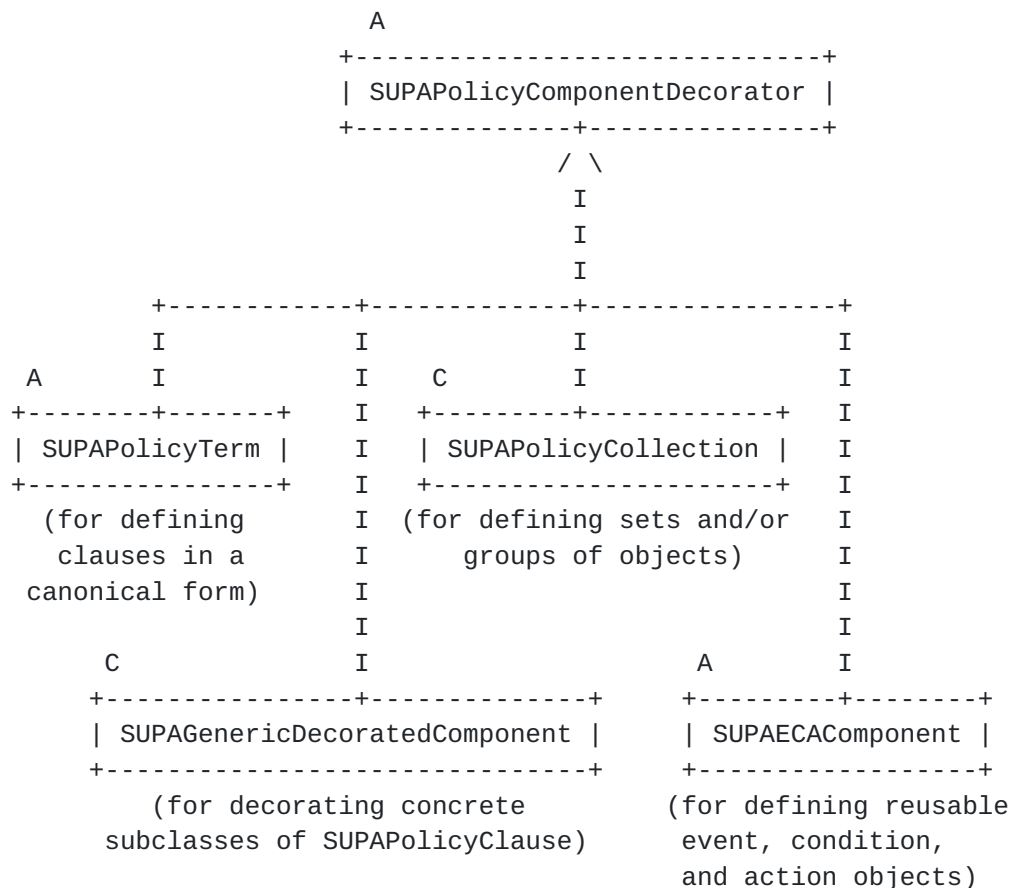


Figure 10. Subclasses of SUPAPolicyComponentDecorator

Strassner, et al.

Expires January 19, 2017

[Page 35]

If a `SUPAEncodedClause` is being used, then there is no need to use any of the `SUPAPolicyComponentDecorator` subclasses, since the `SUPAEncodedClause` already completely defines the content of the `SUPAPolicyClause`.

However, if a `SUPAEncodedClause` is NOT being used, then a `SUPAPolicyClause` will be constructed using one or more types of objects that are each subclasses of `SUPAPolicyComponentDecorator`.

These four subclasses provide four different ways to construct a `SUPAPolicyClause`:

- 1) `SUPAPolicyTerm`: as a {variable, operator, value} clause
- 2) `SUPAEncodedClause`: as an encoded object (e.g., to pass YANG or CLI code)
- 3) `SUPAPolicyCollection`: as a collection of objects that requires further processing in order to be made into a `SUPAPolicyClause`
- 4) `SUPAECAComponent`: subclasses define reusable Event, Condition, or Action objects

These four different types of objects can be intermixed. For example, the first and last types can be combined as follows:

```
Variable == Event.baz                                (A)
Condition BETWEEN VALUE1 and VALUE2                  (B)
(Event.severity == 'Critical' AND
 (SLA.violation == TRUE OR User.class == 'Gold'))      (C)
```

In the above rules, (A) uses `Event.baz` to refer to an attribute of the `Event` class; (B) defines two different instances of a `Value` class, denoted as `Value1` and `Value2`; (C) uses the nomenclature `foo.bar`, where `foo` is the name of a class, and `bar` is the name of an attribute of that class.

4.4.4. Creating SUPAPolicyClauses

The GPIM defines a single subclass of `SUPAPolicyClause`, called `SUPAEncodedClause`. This clause is generic in nature, and MAY be used with any type of policy (ECA or otherwise). The EPRIM defines an ECA-specific subclass of the GPIM, called a `SUPABooleanClause`, which is intended to be used with just ECA policy rules; however, other uses are also possible.

Together, the GPIM and EPRIM provide several alternatives to implement a `SUPAPolicyClause`, enabling the developer to optimize the solution for different constraints:

- 1) The SUPAPolicyClause can be encoded using one or more SUPAEncodedClauses; a SUPAEncodedClause encodes the entire content of its respective event, condition, or action clause.
- 2) The SUPAPolicyClause can be defined using one or more SUPABooleanClauses; each of the three clauses can be defined as either a single SUPABooleanClause, or a combination of SUPABooleanClauses that are logically ANDed, ORed, and/or NOTed.
- 3) The above two mechanisms can be combined (e.g., the first used to define the event clause, and the second used to define the condition and action clauses).

Figure 11 shows the subclasses of SUPAPolicyClause.

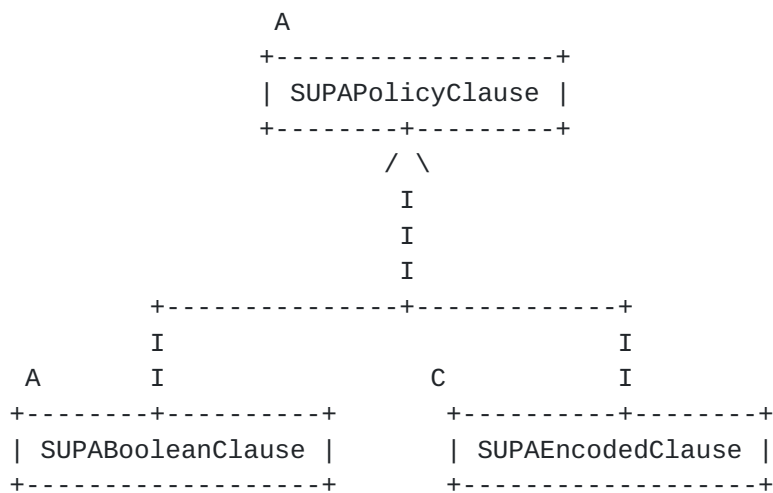


Figure 11. Subclasses of SUPAPolicyClause

SUPABooleanClause is defined in the PRIM, and is used to construct Boolean clauses that collectively make up a SUPAPolicyClause. It is abstract, so that the composite pattern can be applied to it, which enables hierarchies of Boolean clauses to be created. SUPAEncodedClause (see [section 6.7](#)) is used to encode the content of a SUPAPolicyClause as an attribute (instead of reusable objects).

4.4.5. SUPAPolicySources

A SUPAPolicySource is a set of managed entities that authored, or are otherwise responsible for, this SUPAPolicy. Note that a SUPAPolicySource does NOT evaluate or execute SUPAPolicies. Its primary use is for auditability, authorization policies, and other applications of deontic and/or alethic logic.

SUPAPolicyStructure defines four relationships. Two of these (SUPAHasPolicySource and SUPAHasPolicyTarget), which are both aggregations, relate a SUPAPolicyStructure to a SUPAPolicySource and a SUPAPolicyTarget, respectively. Since SUPAECAPolicyRule is a subclass of SUPAPolicyStructure, it (and its subclasses) inherit both of these aggregations. This enables SUPAPolicySources and/or SUPAPolicyTargets to be attached to SUPAECAPolicyRules (but NOT to components of a SUPAPolicy).

Figure 12 shows how SUPAPolicySources and SUPAPolicyTargets are attached to a SUPAPolicy. Note that both of these aggregations are defined as optional, since their multiplicity is 0..n - 0..n. In addition, both of these aggregations are realized as association classes, in order to be able to control which SUPAPolicySources and SUPAPolicyTargets are attached to a given SUPAECAPolicyRule.

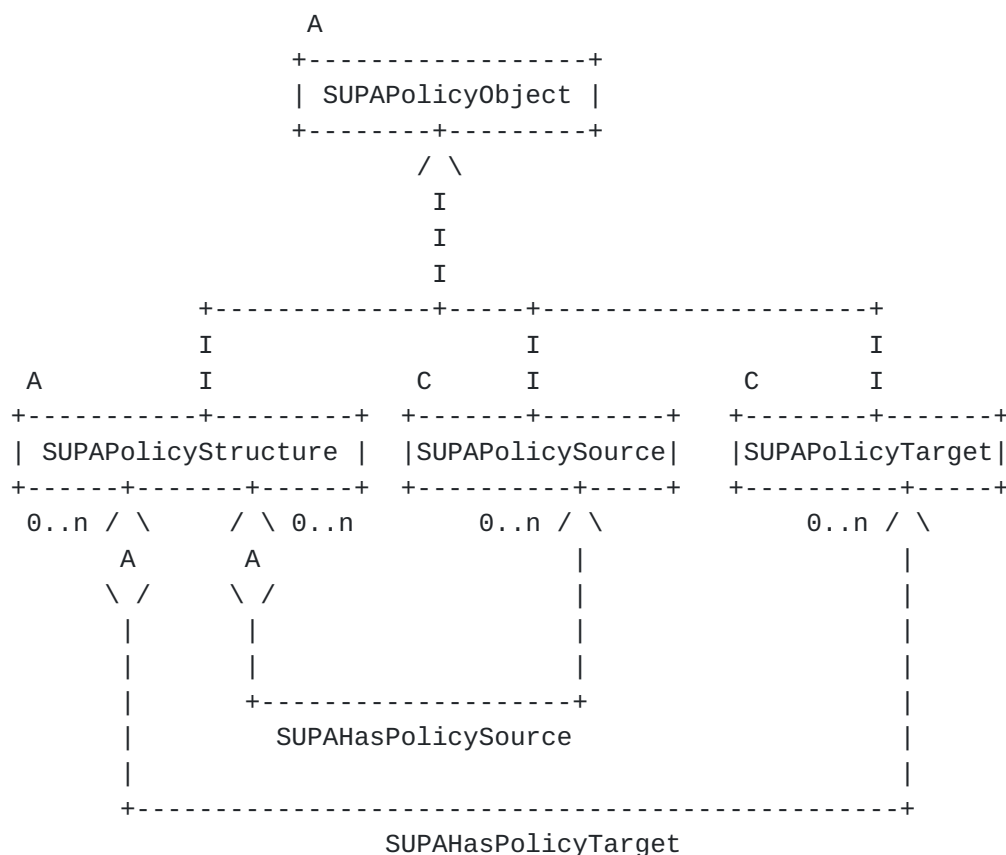


Figure 12. ECAPolicyRules, SUPAPolicySources, and PolicyTargets

A SUPAPolicySource MAY be mapped to a role (e.g., using the role-object pattern [11]); this indirection makes the system less fragile, as entities can be transparently added or removed from the role definition without adversely affecting the definition of

the SUPAPolicy. Note that SUPAPolicyRole is a subclass of SUPAPolicyMetadata.

Strassner, et al.

Expires January 19, 2017

[Page 38]

4.4.6. SUPAPolicyTargets

A SUPAPolicyTarget defines the set of managed entities that a SUPAPolicy is applied to. This is useful for debugging, as well as when the nature of the application requires the set of managed entities affected by a Policy to be explicitly identified. This is determined by two conditions:

- 1) The set of managed entities that are to be affected by the SUPAPolicy must all agree to play the role of a SUPAPolicyTarget. For example, a managed entity may not be in a state that enables SUPAPolicies to be applied to it; hence, in this case, it MUST NOT assume the role of a SUPAPolicyTarget
- 2) A SUPAPolicyTarget must be able to:
 - a) process (either directly or with the aid of a proxy) SUPAPolicies, or
 - b) receive the results of a processed SUPAPolicy and apply those results to itself.

Figure 12 showed how SUPAPolicyTargets are attached to SUPAECAPolicyRules.

A SUPAPolicyTarget MAY be mapped to a role (e.g., using the role-object pattern [11]); this indirection makes the system less fragile, as entities can be transparently added or removed from the role definition without adversely affecting the definition of the SUPAPolicy. Note that SUPAPolicyRole is a subclass of SUPAPolicyMetadata.

4.4.7. Policy Metadata

Metadata is, literally, data about data. As such, it can be descriptive or prescriptive in nature.

4.4.7.1. Motivation

There is a tendency in class design to make certain attributes, such as description, status, validFor, and so forth, bound to a specific class (e.g., [6]). This is bad practice in an information model. For example, different classes in different parts of the class hierarchy could require the use of any of these attributes; if one class is not a subclass of the other, then they must each define the same attribute as part of their class structure. This makes it difficult to find all instances of the attribute and ensure that they are synchronized. Furthermore, context can dynamically change the status of an object, so an easy way to update the status of one object instance without affecting other instances of the same object is required.

Many models, such as [4] and [6], take a simplistic approach of defining a common attribute high in the hierarchy, and making it optional. This violates classification theory, and defeats the purpose of an information model, which is to specify the differences in characteristics and behavior between classes (as well as define how different classes are related to each other). Note that this also violates a number of well-known software architecture principles, including:

- o the Liskov Substitution Principle [13]
(if A is a subclass of B, then objects instantiated from class B may be replaced with objects instantiated from class A WITHOUT ALTERING ANY OF THE PROGRAM SEMANTICS)
- o the Single Responsibility Principle [14]
(every class should have responsibility over one, and only one, part of the functionality provided by the program)
- o the Open/Closed Principle (software should be open for extension, but closed for modification) [17]
- o the Interface-Segregation Principle (clients should not be forced to depend on methods that they do not use) [14]
- o the Dependency Inversion Principle (high-level modules should not depend on low-level modules; both should depend on abstractions) [14]

Most models use inheritance, not composition. The former is simpler, but has some well-known problems. One is called "weak encapsulation", meaning that a subclass can use attributes and methods of a superclass, but if the superclass changes, the subclass may break. Another is that each time a new object is required, a new subclass must be created. These problems are present in [RFC3460], [4], and [6].

Composition is an alternative that provides code that is easier to use. This means that composition can provide data models that are more resistant to change and easier to use. By using composition, we can select just the metadata objects that are needed, instead of having to rely on statically defined objects. We can even create new objects from a set of existing objects through composition. Finally, we can use the decorator pattern to select just the attributes and behaviors that are required for a given instance.

In [2] and [5], a separate metadata class hierarchy is defined to address this problem. This document follows this approach.

4.4.7.2. Design Approach

The goal of the GPIM is to enable metadata to be attached to any subclass of SUPAPolicyObject that requires it. Since this is a

system intended for policy-based management, it therefore makes sense to be able to control which metadata is attached to which policies dynamically (i.e., at runtime).

Strassner, et al.

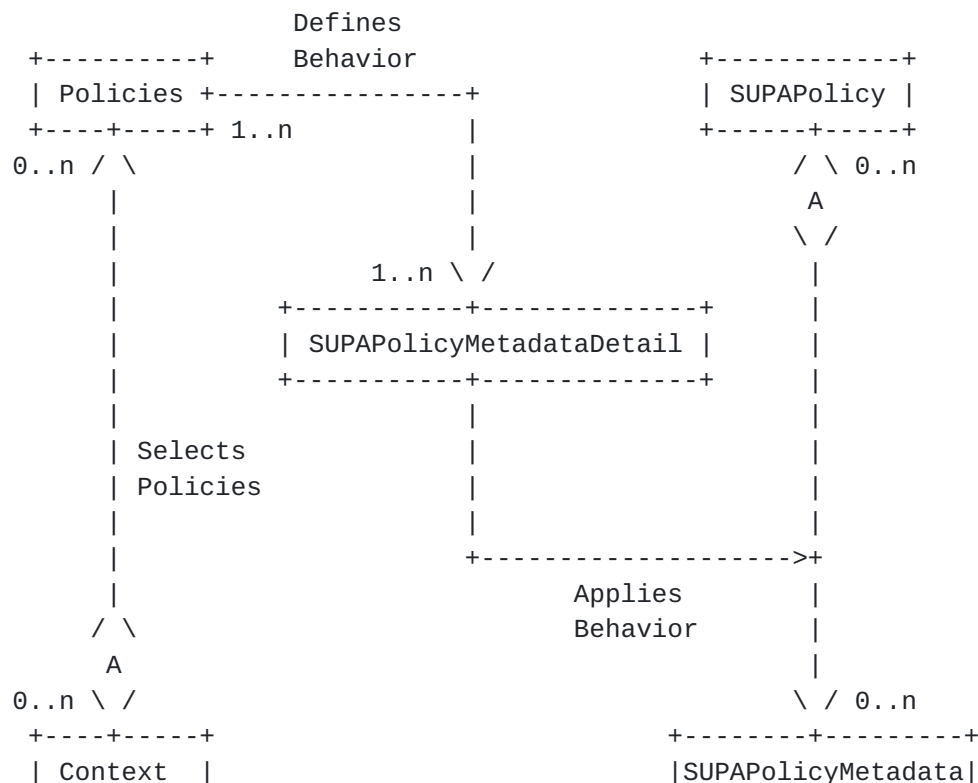
Expires January 19, 2017

[Page 40]

One solution is to use the Policy Pattern [1], [2], [6], [12]. This pattern was built to work with management systems whose actions were dependent upon context. The Policy Pattern works as follows:

- o Context is derived from all applicable system inputs (e.g., OAMP data from network elements, business goals, time of day, geo-location, etc.).
- o Context is then used to select a working set of Policies.
- o Policies are then used to define behavior at various control points in the system.
- o One simple type of control point is an association class. Since the association class represents the semantics of how two classes are related to each other, then
 - o ECAPolicyRule actions can be used to change the attribute values, methods, and relationships of the association class
 - o This has the affect of changing how the two classes are related to each other
- o Finally, as context changes, the working set of policies change, enabling the behavior to be adjusted to follow changes in context (according to appropriate business goals and other factors, of course) in a closed loop manner.

Conceptually, this is accomplished as shown in Figure 13 below.



+-----+

+-----+

Figure 13. Context-Aware Policy Rules

Strassner, et al.

Expires January 19, 2017

[Page 41]

4.4.7.2.1. Policies and Actors

The Policy Continuum ([1] [5] [10] [12]) was defined to associate different actors with different policies at different levels of business and/or technical specificity. Context-aware policy rules, and the Policy Pattern, were defined to realize this association.

Four important functions related to the lifecycle of policies are design, implementation, deployment, and execution. There are many different possible definitions of these functions (even for policy lifecycle management); however, for the purposes of this document, they are defined as follows:

- o Design: The process of defining a software architecture to satisfy user requirements.
- o Development: the process of documenting, programming, testing, and maintaining code and applications as part of a software product
- o Deployment: the process that assembles and transfers completed software artifacts to a state that enables their execution
- o Execution: the process of installing, activating, running, and subsequently deactivating executable software products

The design process is responsible for producing a software architecture. This emphasizes the design, as opposed to the programming, of software systems. In contrast to design, development emphasizes constructing software artifacts via coding and documentation.

Deployment may be described as the process of releasing software. It includes all of the operations required to assemble a completed software product. It typically also includes the process of preparing a software product for execution (e.g., assembling a set of software products into a larger product, determining if the consumer site has appropriate resources to install and execute the software product, and collecting information on the feasibility of using the software product). This contrasts with the execution process, which is the set of processes that follow deployment.

In summary, exemplar states in the policy lifecycle process include:

- o Design: determining how the policy-based management system will operate
- o Development: documenting, programming, testing, and maintaining policies and policy components
- o Deployment: assembling the components of a policy-based

management system
o Execution: installing, enabling, running, disabling,
and uninstalling policies and policy components

Strassner, et al.

Expires January 19, 2017

[Page 42]

4.4.7.2.2. Deployment vs. Execution of Policies

One of the primary reasons for separating the deployment and execution processes is to differentiate between environments that are not ready to execute policies (i.e., deployment) and environments that are ready to execute policies (i.e., execution). This is an important consideration, since policies that are related to the same set of tasks may be deployed in many different places (e.g., in a policy system vs. in a network device). In addition, each managed entity in the set of SUPAPolicyTargets may or may not be in a state that allows SUPAPolicies to be applied to it (see [section 4.4.6.](#)).

Hence, this design includes dedicated class attributes for getting and setting the deployment and execution status, as well as enabling and disabling, SUPAPolicies (see [section 5.3.1.](#)).

4.4.7.2.3. Using SUPAMetadata for Policy Deployment and Execution

One way of encoding deployment and execution status for policies and policy components is to attach Metadata objects to affected SUPAPolicyStructure and SUPAPolicyComponentStructure objects. This provides an extensible and efficient means to describe and/or prescribe deployment and/or execution status of a policy or a policy component. It is extensible, since classes and relationships can be used, as opposed to a set of attributes. It is efficient, because the decorator pattern (see [section 5.7](#)) is used (this enables attributes and/or methods of objects, or the entire object, to be used to add characteristics and/or behavior to a given object).

SUPAPolicyMetadata objects (see sections [5.16](#) - [5.20](#)) may be attached to the SUPAECAPolicyRule and/or any of its components to define additional semantics of the SUPAECAPolicyRule. For example, SUPAAccessMetadataDef (see [section 5.19](#)) and/or SUPAVersionMetadataDef (see [section 5.20](#)) may be attached to define the access privileges and version information, respectively, of a policy rule and/or its components.

The SUPAPolicyStructure defines an attribute, supaPolDeployStatus, (see [section 5.3.1.3.](#)) that SUPAPolicyMetadata objects can use to get and set the deployment and execution status of a SUPAPolicy. This allows metadata to be used to alter the deployment and/or execution state of a policy (or a set of policy components) without having to affect other parts of the policy-based management system. The supaPolDeployStatus attribute indicates that this SUPAPolicy can or cannot be deployed. If it cannot be deployed. Similarly, the supaPolExecStatus attribute is used to indicate if a particular SUPAPolicy has executed, is

currently executing, or is ready to execute, and whether or not the execution of that SUPAPolicy had any failures.

Strassner, et al.

Expires January 19, 2017

[Page 43]

The reverse is also true (and hence, forms a closed-loop system controlled by metadata). For example, if the set of deployed SUPAPolicies are SUPAECAPolicyRules, then when the actions of these SUPAECAPolicyRules are executed, the overall context has changed (see [section 4.4.7.2](#)). The context manager could then change attribute values (directly or indirectly) in the SUPAPolicyMetadataDetail association class. This class represents the behavior of the SUPAHasPolicyMetadata aggregation, which is used to define which SUPAPolicyMetadata can be attached to which SUPAPolicy object in this particular context. For example, the access privileges of a policy and/or policy component could be changed dynamically, according to changes in context.

By using the decorator pattern on SUPAPolicyMetadata, any number of SUPAPolicyMetadata objects (or their attributes, etc.) can be wrapped around a concrete subclass of SUPAPolicyMetadata. This is shown in Figure 14 below.

[4.4.7.3](#). Structure of SUPAPolicyMetadata

SUPAPolicyMetadata also uses the decorator pattern to provide an extensible framework for defining metadata to attach to SUPAPolicy subclasses. Its two principal subclasses are SUPAPolicyConcreteMetadata and SUPAPolicyMetadataDecorator. The former is used to define concrete subclasses of SUPAPolicyMetadata that are attached at runtime to SUPAPolicy subclasses, while the latter is used to define concrete objects that represent reusable attributes, methods, and relationships that can be added to subclasses of SUPAPolicyConcreteMetadata.

For example, concepts like identification, access control, and version information are too complex to represent as a single attribute, or even a couple of attributes - they require the generic power of objects to represent their characteristics and behavior. Furthermore, defining concrete classes to represent these concepts in the policy hierarchy is fragile, because:

1. not all objects that use these concepts need all of the information represented by them (e.g., two subclasses of an Identification Object may be Passport and Certificate, but these two objects are rarely used together, and even those contexts that use one of these classes may not need all of the data in that class)
2. defining a class means defining its attributes, methods, and relationships at a particular place in the hierarchy; this means that defining a relationship between a class A and another class B SHOULD only be done if all of the subclasses of B can use the attributes, methods, and relationships of A (e.g., in the above example, defining a relationship between

an Identification Object and a superclass of a router class
is not appropriate, since routers do not use Passports)

Strassner, et al.

Expires January 19, 2017

[Page 44]



Figure 14. SUPAPolicyMetadata Subclasses and Relationships

Since a class encapsulates attributes, methods, and behavior, defining the Identification Object in the above example as a type of SUPAPolicyMetadata object enables the decorator pattern to be used to attach all or part of that object to other objects that need it.

Figure 15 shows a portion of the SUPAPolicyMetadata hierarchy.



Figure 15. SUPAPolicyMetadata Subclasses and Relationships

Figure 15 shows a relevant portion of the SUPAPolicyMetadata hierarchy. SUPAPolicyConcreteMetadata is a concrete class that subclasses of the SUPAPolicyMetadataDecorator class can wrap. Two such subclasses, SUPAPolicyAccessMetadataDef and SUPAPolicyVersionMetadataDef, are shown in Figure 15. This enables access control and version information to be added statically (at design time) or dynamically (at runtime) to SUPAPolicyConcreteMetadata; this enables metadata-driven systems to adjust the behavior of the management system to changes in context, business rules, services given to end-users, and other similar factors. This is discussed more in sections [5.18](#) - [5.20](#).

4.5. Advanced Features

This section will be completed in the next revision of this document.

4.5.1. Policy Grouping

This section will be completed in the next revision of this document.

4.5.2. Policy Rule Nesting

This section will be completed in the next revision of this document.

5. GPIM Model

This section defines the classes, attributes, and relationships of the GPIM.

5.1. Overview

The overall class hierarchy is shown in Figure 16; section numbers are appended after each class.

(Class of another model that SUPA is integrating into)

```

|
+---SUPAPolicyObject (5.2)
|   |
|   +---SUPAPolicyStructure (5.3)
|   |
|   +---SUPAPolicyComponentStructure (5.4)
|   |   |
|   |   +---SUPAPolicyClause (5.5)
|   |   |   |
|   |   |   +---SUPAEncodedClause (5.6)
|   |   |   |
|   |   +---SUPAPolicyComponentDecorator (5.7)
|   |       |
|   |       +---SUPAPolicyTerm (5.8)
|   |           |
|   |           +---SUPAPolicyVariable (5.9)
|   |           |
|   |           +---SUPAPolicyOperator (5.10)
|   |           |
|   |           +---SUPAPolicyValue (5.11)
|   |               |
|   |               +---SUPAGenericDecoratedComponent (5.12)
|   |                   |
|   |                   +---SUPAPolicyCollection (5.13)
|   |
|   +---SUPAPolicySource (5.14)
|   |
|   +---SUPAPolicyTarget (5.15)

```

(Class of another model that SUPAPolicyMetadata is integrating into)

```

|
+---SUPAPolicyMetadata (5.16)
|   |
|   +---SUPAPolicyConcreteMetadata (5.17)
|   |
|   +---SUPAPolicyMetadataDecorator (5.18)
|       |
|       +---SUPAPolicyAccessMetadataDef (5.19)
|       |
|       +---SUPAPolicyVersionMetadataDef (5.20)

```

Figure 16: Main Classes of the GPIM

Strassner, et al.

Expires January 19, 2017

[Page 48]

SUPAPolicy is the root of the SUPA class hierarchy. For implementations, it is assumed that SUPAPolicy is subclassed from a class from another model. Note that SUPAPolicyMetadata MAY be subclassed from the same or (preferably) a different class in the external model.

Classes, attributes, and relationships that are marked as "mandatory" MUST be part of a conformant implementation (i.e., a schema MUST contain these entities). This does not mean that these entities must be instantiated; rather it means that they must be able to be instantiated. Classes, attributes, and relationships that are marked as "optional" MAY be part of a conformant implementation.

Unless otherwise stated, all classes (and attributes) defined in this section were abstracted from DEN-ng [2], and a version of them are in the process of being added to [5]. However, the work in [5] has been put on hold, and the names of many of the classes, attributes, and relationships are slightly different.

5.2. The Abstract Class "SUPAPolicyObject"

This is a mandatory abstract class. Figure 17 shows the SUPAPolicyObject class, and its four subclasses.



+-----+

Figure 17. SUPAPolicyObject and Its Subclasses

Strassner, et al.

Expires January 19, 2017

[Page 49]

This class is the root of the SUPA class hierarchy. It defines the common attributes and relationships that all SUPA subclasses inherit.

A SUPAPolicyObject MAY be qualified by a set of zero or more SUPAPolicyMetadata objects. This is provided by the SUPAHasPolicyMetadata aggregation (see [Section 5.2.2](#)). This enables the semantics of the SUPAPolicyObject to be more completely specified.

5.2.1. SUPAPolicyObject Attributes

This section defines the attributes of the SUPAPolicyObject class. These attributes are inherited by all subclasses of the GPIM except for the SUPAPolicyMetadata class, which is a sibling class.

5.2.1.1. Object Identifiers

This document defines two class attributes in SUPAPolicyObject, called supaPolObjIDContent and supaPolObjIDEncoding, that together define a unique object ID. This enables all class instances to be uniquely identified.

One of the goals of SUPA is to be able to generate different data models that support different types of protocols and repositories. This means that the notion of an object ID must be generic. It is inappropriate to use data modeling concepts, such as keys, Globally Unique IDentifiers (GUIDs), Universally Unique IDentifiers (UUIDs), Fully Qualified Domain Names (FQDNs), Fully Qualified Path Names (FQPNs), Uniform Resource Identifiers (URIs), and other similar mechanisms, to define the structure of an information model. Therefore, a synthetic object ID is defined using these two class attributes. This can be used to facilitate mapping to different data model object schemes.

The two attributes work together, with the supaPolObjIDContent attribute defining the content of the object ID and the supaPolObjIDEncoding attribute defining how to interpret the content. These two attributes form a tuple, and together enable a machine to understand the syntax and value of an object identifier for the object instance of this class.

Similarly, all SUPA classes and attributes are both uniquely named as well as prepended with the prefixes "SUPA" and "supa", respectively, to facilitate model integration.

5.2.1.2. The Attribute "supaPolObjIDContent"

This is a mandatory string attribute that represents part of the object identifier of an instance of this class. It defines the content of the object identifier. It works with another class attribute, called `supaPolObjIDEncoding`, which defines how to interpret this attribute. These two attributes form a tuple, and together enable a machine to understand the syntax and value of an object identifier for the object instance of this class. This is based on the DEN-ng class design [2].

5.2.1.3. The Attribute "supaPolObjIDEncoding"

This is a mandatory non-zero enumerated integer attribute that represents part of the object identifier of an instance of this class. It defines the format of the object identifier. It works with another class attribute, called `supaPolObjIDContent`, which defines the content of the object ID. These two attributes form a tuple, and together enable a machine to understand the syntax and value of an object identifier for the object instance of this class. The `supaPolObjIDEncoding` attribute is mapped to the following values:

- 0: error
- 1: init
- 2: primary_key
- 3: foreign_key
- 4: GUID
- 5: UUID
- 6: URI
- 7: FQDN
- 8: FQPN
- 9: string_instance_id

The values 0 and 1 represent an error state and an initialization state, respectively. The value 9 defines the canonical representation, in ASCII, of an instance ID of this object.

5.2.1.4. The Attribute "supaPolicyDescription"

This is an optional string attribute that defines a free-form textual description of this object.

5.2.1.5. The Attribute "supaPolicyName"

This is an optional string attribute that defines the name of this Policy. This enables any existing generic naming attribute to be used for generic naming, while allowing this attribute to be used to name Policy entities in a common manner. Note that this is NOT the same as the `commonName` attribute of the Policy class defined

in [[RFC3060](#)], as that attribute is intended to be used with just X.500 cn attributes.

Strassner, et al.

Expires January 19, 2017

[Page 51]

5.2.2. SUPAPolicyObject Relationships

The SUPAPolicyObject class currently defines a single relationship, as defined in the subsection below.

5.2.2.1. The Aggregation "SUPAHasPolicyMetadata"

This is a mandatory aggregation that defines the set of SUPAPolicyMetadata that are aggregated by this particular SUPAPolicyObject. This aggregation is defined in [section 5.16.2](#).

5.2.2.2. The Association Class "SUPAHasPolicyMetadataDetail"

This is a mandatory concrete association class that defines the semantics of the SUPAPolicyMetadata aggregation. This enables the attributes and relationships of the SUPAPolicyMetadataDetail class to be used to constrain which SUPAPolicyMetadata objects can be aggregated by this particular SUPAPolicyObject instance. This association class is defined in [Section 5.16.2.2](#).

5.3. The Abstract Class "SUPAPolicyStructure"

This is a mandatory abstract class that is used to represent the structure of a SUPAPolicy. This class (and all of its subclasses) is a type of PolicyContainer. SUPAPolicyStructure was abstracted from DEN-ng [2], and a version of this class is in the process of being added to [5]. However, the version in [5] differs significantly. First, the class and relationship definitions are different. Second, [5] uses the composite pattern. Neither of these are implemented in this document because of optimizations done to the SUPA class hierarchy that are NOT present in [5].

For this release, the only official type of policy that is supported is the event-condition-action (ECA) type of policy rule. However, the structure of the SUPA hierarchy is defined to facilitate adding new types of rules later.

A SUPAPolicy may take the form of an individual policy or a set of policies. This requirement is supported by applying the composite pattern to subclasses of the SUPAPolicyStructure class, as shown in Figure 5. In this document, this is done for the SUPAECAPolicyRule subclass, and results in two subclasses: SUPAECAPolicyRuleAtomic (for defining stand-alone policies) and SUPAECAPolicyRuleComposite (for defining hierarchies of policies).

Note that there is no need for a "match strategy attribute" that some models [RFC3460], [4], [6] have; this is because the SUPAPolicyStructure class is used just for containment. Hence, the

containers themselves serve as the scoping component for nested policies.

Strassner, et al.

Expires January 19, 2017

[Page 52]

5.3.1. SUPAPolicyStructure Attributes

The following subsections define the attributes of the SUPAPolicyStructure class.

The SUPAPolicyStructure class has a number of attributes that have no counterpart in the SUPAPolicyComponentStructure class. This is because these attributes are only appropriate at the level of a policy rule, not at the level of a policy component.

Care must be taken in adding attributes to this class, because the behavior of future subclasses of this class (e.g., declarative and functional policies) is very different than the behavior of SUPAECAPolicyRules.

5.3.1.1. The Attribute "supaPolAdminStatus"

This is an optional attribute, which is an enumerated non-negative integer. It defines the current administrative status of this SUPAPolicyClause. Values include:

- 0: error
- 1: init
- 2: enabled
- 3: disabled
- 4: in test (i.e., no operational traffic can be passed)

The values 0 and 1 represent an error state and an initialization state, respectively. Values 2 and 3 mean that this SUPAPolicyStructure is administratively enabled or disabled, respectively. A value of 4 means that this SUPAPolicyStructure is currently in a special test mode and SHOULD NOT be used as part of an OAM&P policy.

5.3.1.2. The Attribute "supaPolContinuumLevel"

This is an optional non-negative integer attribute. It defines the level of abstraction, or policy continuum level [10], of this particular SUPAPolicy. The value assignment of this class is dependent on the application; however, it is recommended that for consistency with other SUPA attributes, the values of 0 and 1 are reserved for error and initialization states.

By convention, lower values represent more abstract levels of the policy continuum. For example, a value of 1 could represent business policy, a value of 2 could represent application-specific policies, and a value of 3 could represent low-level policies for network administrators.

5.3.1.3. The Attribute "supaPolDeployStatus"

This is an optional enumerated, non-negative integer attribute. The purpose of this attribute is to indicate that this SUPAPolicy can or cannot be deployed by the policy management system. This attribute enables the policy manager to know which SUPAPolicies to retrieve, and may be useful for the policy execution system for planning the staging of SUPAPolicies. Values include:

- 0: error
- 1: init
- 2: deployed and enabled
- 3: deployed and in test
- 4: deployed but not enabled
- 5: ready to be deployed
- 6: cannot be deployed

The values 0 and 1 represent an error state and an initialization state, respectively. A value of 2 means that the policy management system MAY use this SUPAPolicy. A value of 3-5 means that the policy management system SHOULD NOT use this SUPAPolicy until it is put into an enabled state.

5.3.1.4. The Attribute "supaPolExecFailStrategy"

This is an optional non-negative, enumerated integer that defines what actions, if any, should be taken by this SUPAPolicyStructure object if it fails to execute correctly.

Note that some systems may not be able to support all options specified in this enumeration. If rollback is supported by the system, then option 2 may be skipped. Options 3 and 4 can be used by systems that do and do not support rollback. Values include:

- 0: error
- 1: init
- 2: attempt rollback of all actions taken and stop execution
- 3: attempt rollback of only the action that failed and stop execution
- 4: stop execution but do not rollback any actions
- 5: ignore failure and continue execution

The values 0 and 1 represent an error state and an initialization state, respectively. A value of 2 means that ALL execution is stopped, rollback of all actions (whether successful or not) is attempted, and that SUPAPolicies that otherwise would have been executed are ignored. A value of 3 means that execution is stopped, and rollback is attempted for ONLY the SUPAPolicy that failed to execute correctly. A value of 4 means that execution is stopped, but

no actions are rolled back. A value of 5 means that the failure is ignored, and execution continues.

Strassner, et al.

Expires January 19, 2017

[Page 54]

5.3.2. SUPAPolicyStructure Relationships

The SUPAPolicyStructure class owns four relationships, which are defined in the following subsections.

5.3.2.1. The Aggregation "SUPAHasPolicySource"

This is an optional aggregation, and defines the set of SUPAPolicySource objects that are attached to this particular SUPAPolicyStructure object. The semantics of this aggregation are defined by the SUPAHasPolicySourceDetail association class. PolicySource objects are used for authorization policies, as well as to enforce deontic and alethic logic.

The multiplicity of this aggregation is 0..n - 0..n. This means that it is an optional aggregation; zero or more SUPAPolicySource objects may be aggregated by this SUPAPolicyStructure object, and zero or more SUPAPolicyStructure objects may aggregate this particular SUPAPolicySource object.

5.3.2.2. The Association Class "SUPAHasPolicySourceDetail"

This is an optional concrete association class, and defines the semantics of the SUPAHasPolicySource aggregation. The attributes and relationships of this class can be used to define which SUPAPolicySource objects can be attached to which particular set of SUPAPolicyStructure objects.

5.3.2.2.1. The Attribute "supaPolSrcIsAuthenticated"

This is an optional Boolean attribute. If the value of this attribute is true, then this SUPAPolicySource object has been authenticated by this particular SUPAPolicyStructure object.

5.3.2.2.2. The Attribute "supaPolSrcIsTrusted"

This is an optional Boolean attribute. If the value of this attribute is TRUE, then this particular SUPAPolicySource object has been verified to be trusted by this particular SUPAPolicyStructure object.

5.3.2.3. The Aggregation "SUPAHasPolicyTarget"

This is an optional aggregation, and defines the set of SUPAPolicyTargets that are attached to this particular SUPAPolicyStructure. The semantics of this aggregation is defined by the SUPAHasPolicyTargetDetail association class. The purpose of this class is to explicitly identify managed objects

that will be affected by the execution of one or more SUPAPolicies.

Strassner, et al.

Expires January 19, 2017

[Page 55]

The multiplicity of this aggregation is 0..n - 0..n. This means that it is an optional aggregation; zero or more SUPAPolicyTarget objects may be aggregated by this SUPAPolicyStructure object, and zero or more SUPAPolicyStructure objects may aggregate this particular SUPAPolicyTarget object.

5.3.2.4. The Association Class "SUPAHasPolicyTargetDetail"

This is an optional concrete association class, and defines the semantics of the SUPAPolicyTargetOf aggregation. The attributes and relationships of this class can be used to define which SUPAPolicyTargets can be attached to which particular set of SUPAPolicyStructure objects.

5.3.2.4.1. The Attribute "supaPolTgtIsAuthenticated"

This is an optional Boolean attribute. If the value of this attribute is true, then this SUPAPolicyTarget object has been authenticated by this particular SUPAPolicyStructure object.

5.3.2.4.2. The Attribute "supaPolTgtIsEnabled"

This is an optional Boolean attribute. If its value is TRUE, then this SUPAPolicyTarget is able to be used as a SUPAPolicyTarget. This means that it meets two specific criteria:

1. it has agreed to play the role of a SUPAPolicyTarget (i.e., it is willing to have SUPAPolicies applied to it, and
2. it is able to either process (directly or with the aid of a proxy) SUPAPolicies or receive the results of a processed SUPAPolicy and apply those results to itself.

5.3.2.5. The Association "SUPAHasPolExecFailTakeAction"

This is an optional association that defines which, if any, actions should be taken if this SUPAPolicyStructure object instance fails to execute correctly. The semantics of this association are defined in the SUPAHasPolExecFailTakeActionDetail association class.

For a given SUPAPolicyStructure object A, this association defines a set of policy action objects B to execute if (and only if) the SUPAPolicyStructure object A failed to execute correctly. The multiplicity of this association is defined as 0..n on the owner (A) side and 1..n on the part (B) side. This means that this association is optional; if it is instantiated, then at least one SUPAPolicyStructure MUST be instantiated by this SUPAPolicyStructure object. Similarly, one or more SUPAPolicyStructure objects may be associated with this given

SUPAPolicyStructure object.

Strassner, et al.

Expires January 19, 2017

[Page 56]

5.3.2.6. The Association Class "SUPAHasPolExecFailTakeActionDetail"

This is an optional concrete class that defines the semantics for the SUPAHasPolExecFailTakeAction association. The attributes and/or relationships of this association class can be used to determine which policy action objects are executed in response to a failure of the SUPAPolicyStructure object instance that owns this association. The association relates the policy actions from one SUPAPolicyStructure B to be executed if a SUPAPolicyStructure A fails to execute properly. Figure 18 illustrates this approach.




Figure 18. SUPAHasPolExecFailTakeAction Association

5.3.2.6.1. The Attribute "supaPolExecFailActionEncoding"

This is an optional enumerated, non-negative integer attribute that defines how to find the set of SUPAPolicyActions contained in each element of the supaPolExecFailTakeActionName class attribute. Values include:

- 0: error
- 1: init
- 2: URI
- 3: GUID
- 4: UUID
- 5: FQDN
- 6: FQPN
- 7: string
- 8: string_instance_id

The values 0 and 1 represent an error state and an initialization state, respectively. Values 2-6 define a representation for the SUPAPolicyAction. A value of 7 defines an ASCII string that contains the name of the SUPAPolicyAction to be executed (e.g., to be used in a regex search). A value of 8 defines the canonical representation, in ASCII, of an instance ID of this object.

5.3.2.6.2. The Attribute "supaPolExecFailActionName[1..n]"

This is an optional array of string attributes that identifies the set of SUPAPolicyActions to take if the SUPAPolicyStructure object that owns this association failed to execute properly. The interpretation of this string attribute is defined by the supaPolExecFailTakeActionEncoding class attribute. The association defines the SUPAPolicyStructure that contains the set of policy actions to execute, and this attribute defines which of these actions are to be executed. That there is no need to execute a SUPAPolicy, since the event and failure have already occurred. Note: [1..n] means that this is a multi-valued property that has at least one (and possibly more) attributes.

5.3.2.7. The Aggregation "SUPAHasPolicyClause"

This is an optional aggregation that defines the set of SUPAPolicyClauses that are aggregated by this particular SUPAPolicyStructure instance. The semantics of this aggregation are defined by the SUPAHasPolicyClauseDetail association class.

Every SUPAPolicyStructure object instance MUST aggregate at least one SUPAPolicyClause object instance. However, the converse is NOT true. For example, a SUPAPolicyClause could be instantiated and then stored for later use in a policy repository. Furthermore, the same SUPAPolicyClause could be used by zero or more SUPAPolicyStructure object instances at a given time. Thus, the multiplicity of this aggregation is defined as 0..1 on the aggregate (i.e., the SUPAPolicyStructure side) and 1..n on the part (i.e., the SUPAPolicyClause side). This means that at least one SUPAPolicyClause MUST be aggregated by this SUPAPolicyStructure object. Similarly, a SUPAPolicyClause may be aggregated by this particular SUPAPolicyStructure object.

5.3.2.8. The Association Class "SUPAHasPolicyClauseDetail"

This is an optional abstract association class, and defines the semantics of the SUPAHasPolicyClause aggregation. The attributes and/or relationships of this association class can be used to determine which SUPAPolicyClauses are aggregated by which SUPAPolicyStructure objects.

Attributes will be added to this class at a later time.

5.4. The Abstract Class "SUPAPolicyComponentStructure"

This is a mandatory abstract class that is the superclass of all objects that represent different types of components of a SUPAPolicy. Different types of policies have different types of structural components. This is accommodated by defining two generic subclasses, called SUPAPolicyClause and SUPAPolicyComponentDecorator, which are common to different policy types. These two classes represent convenient control points for defining characteristics and behavior that are common to objects that serve as components of a SUPAPolicy.

SUPAPolicyClause defines a basic building block for writing parts of a SUPAPolicy. It is analogous to a clause in a sentence. For example, in an ECA Policy Rule, the Event, Condition, and Action clauses are each made up of at least one (subclass of a) SUPAPolicyClause. Similarly, declarative Policy Rules can also be defined using (its own subclasses of) SUPAPolicyClauses. This class is defined in [section 5.5](#).

SUPAPolicyComponentDecorator is used to implement the decorator pattern [11]. The decorator pattern enables all or part of one or more objects to "wrap" another concrete object. This enables the definition of an extensible set of subclasses that can augment the definition of a SUPAPolicyClause. This class is defined in [section 5.7](#).

Note that there are significant differences between the definition of the SUPAPolicyComponentStructure class, and its attributes and relationships, and the definition of the corresponding class (and its attributes and relationships) in [5].

5.4.1. SUPAPolicyComponentStructure Attributes

No attributes are currently defined for the SUPAPolicyComponentStructure class.

5.4.2. SUPAPolicyComponentStructure Relationships

SUPAPolicyComponentStructure participates in a single relationship,
SUPAHasDecoratedPolicyComponent, as defined in [section 5.7.3](#).

Strassner, et al.

Expires January 19, 2017

[Page 59]

5.5. The Abstract Class "SUPAPolicyClause"

This is a mandatory abstract class that separates the representation of a SUPAPolicy from its implementation. SUPAPolicyClause was abstracted from DEN-ng [2]. This abstraction is missing in [RFC3060], [RFC3460], [4], and [6]. This class is called PolicyStatement in [5], but the class and relationship definitions differ significantly from the corresponding designs in this document.

A SUPAPolicy, regardless of its structure and semantics, can be abstracted into a set of sentences. Each sentence can in turn be abstracted into a set of clauses. A SUPAPolicyClause is, as its name implies, a clause (i.e., a part of a statement), and defines the content of a SUPAPolicy. The decorator pattern is used to enable an extensible set of objects to "wrap" the SUPAPolicyClause; this enables the contents of a SUPAPolicyClause to be adjusted dynamically at runtime without affecting other objects.

This document defines two different types of policy clauses: SUPAEncodedClause (which is generic, and can be used by any type of policy), and SUPABooleanClause (which is also generic, but is typically used by SUPAECAPolicyRule objects, since it is used specifically to represent Boolean clauses).

SUPAPolicyClauses are objects in their own right, which facilitates their reuse. SUPAPolicyClauses can aggregate a set of any of the subclasses of SUPAPolicyComponentDecorator, which was shown in Figure 10. These four subclasses (i.e., SUPAPolicyTerm, SUPAGenericDecoratedComponent, SUPAECAComponent, and SUPACollection) provide four different ways to construct a SUPAPolicyClause:

- 1) a SUPAPolicyClause can be made up of a set of three SUPAPolicyTerms, which enables constructing a {variable, operator, value} expression for building SUPAPolicyClauses
- 2) a SUPAPolicyClause can be made up of one or more SUPAEncodedClauses, which enables a SUPAPolicyClause to be formed as an encoded object (e.g., to pass YANG or CLI code)
- 3) a SUPAPolicyClause can be made up of a set of SUPACollections, which define a Collection (e.g., set, bag, associative arrays) of objects that can be assembled into SUPAPolicyClauses
- 4) a SUPAPolicyClause can be made up of one or more SUPAECAComponents, which enables a SUPAPolicyClause to be formed using (reusable) Event, Condition, and/or Action objects

SUPAPolicyClauses are aggregated by a SUPAPolicyStructure object, which enables all types of SUPAPolicies to uniformly be

made up of one or more SUPAPolicyClauses.

Strassner, et al.

Expires January 19, 2017

[Page 60]

5.5.1. SUPAPolicyClause Attributes

This section defines the attributes of the SUPAPolicyClause class, which are inherited by all SUPAPolicyClause subclasses.

5.5.1.1. The Attribute "supaPolClauseDeployStatus"

This is an optional enumerated, non-negative integer attribute. The purpose of this attribute is to indicate that this SUPAPolicyClause can or cannot be deployed by the policy management system. This attribute enables the policy manager to know which SUPAPolicyClauses to retrieve, and may be useful for the policy execution system for planning the staging of SUPAPolicies. Values include:

- 0: error
- 1: init
- 2: deployed and enabled
- 3: deployed and in test
- 4: deployed but not enabled
- 5: ready to be deployed
- 6: cannot be deployed

The values 0 and 1 represent an error state and an initialization state, respectively. If the value of this attribute is 0 or 6, then the policy management system SHOULD ignore this SUPAPolicy. Otherwise, the policy management system MAY use this SUPAPolicyClause (once this SUPAPolicyClause is deployed and enabled). However, a value of 4 means that this policy is not administratively enabled for use and SHOULD NOT be used in OAM&P policies.

5.5.2. SUPAPolicyClause Relationships

SUPAPolicyClause participates in a single relationship, SUPAHasPolicyClause, as defined in [section 5.3.2.7](#). Note that SUPAPolicyClause uses the decorator pattern to "wrap" this object with instances of the (concrete) subclasses of the SUPAPolicyComponentDecorator object.

5.6. The Concrete Class "SUPAEncodedClause"

This is a mandatory concrete class that refines the behavior of a SUPAPolicyClause.

This class defines a generalized extension mechanism for representing SUPAPolicyClauses that have not been modeled with other SUPAPolicy objects. Rather, the contents of the policy clause are directly encoded into the attributes of the SUPAEncodedClause. Hence, SUPAEncodedClause objects are reusable

at the object level, whereas SUPABooleanClause clauses are reusable
at the individual Boolean expression level.

Strassner, et al.

Expires January 19, 2017

[Page 61]

This class uses two of its attributes (`supaEncodedClauseContent` and `supaEncodedClauseEncoding`) for defining the content and type of encoding used in a given `SUPAPolicyClause`. The benefit of a `SUPAEncodedClause` is that it enables direct encoding of the text of the `SUPAPolicyClause`, without having the "overhead" of using other objects. However, note that while this method is efficient, it does not reuse other `SUPAPolicy` objects. Furthermore, its potentiality for reuse is reduced, as only `SUPAPolicies` that can use the exact encoding of this clause can reuse this object.

5.6.1. SUPAEncodedClause Attributes

This section defines the attributes of the `SUPAEncodedClause` class. Prescriptive and/or descriptive information about the usage of this `SUPAEncodedClause` may be provided by one or more `SUPAPolicyMetadata` objects, which are each attached to the object instance of this `SUPAEncodedClause`.

5.6.1.1. The Attribute "supaEncodedClauseContent"

This is a mandatory string attribute, and defines the content of this clause. It works with another class attribute, called `supaEncodedClauseEncoding`, which defines how to interpret the value of this attribute (e.g., as a string or reference). These two attributes form a tuple, and together enable a machine to understand the syntax and value of this object instance.

5.6.1.2. The Attribute "supaEncodedClauseEncoding"

This is a mandatory non-negative integer attribute, and defines how to interpret the value of the `supaEncodedClauseContent`. It works with another class attribute (`supaEncodedClauseContent`), which defines the content of the encoded clause. These two attributes form a tuple, and together enable a machine to understand the syntax and value of the encoded clause for the object instance of this class. This attribute is NOT required in all data model implementations. Values include:

- 0: error (i.e., an error state)
- 1: init (i.e., an initialization state)
- 2: primary_key
- 3: foreign_key
- 4: GUID
- 5: UUID
- 6: URI
- 7: FQDN
- 8: FQPN
- 9: string_instance_id

The values 0 and 1 represent an error state and an initialization state, respectively. The value 9 defines the canonical representation, in ASCII, of an instance ID of this object.

5.6.1.3. The Attribute "supaEncodedClauseLanguage"

This is mandatory non-negative integer attribute, and defines the type of language used in this encoded clause. Values include:

- 0: error
- 1: init
- 2: Text
- 3: YANG
- 4: XML
- 5: TL1

The values 0 and 1 represent an error state and an initialization state, respectively.

5.6.1.4. The Attribute "supaEncodedClauseResponse"

This is an optional Boolean attribute that emulates a Boolean response of this clause, so that it may be combined with other subclasses of the SUPAPolicyClause that provide a status as to their correctness and/or evaluation state. This enables this object to be used to construct more complex Boolean clauses. Note that this attribute does NOT have to be implemented by all data model implementations (e.g., [15]).

5.6.2. SUPAEncodedClause Relationships

SUPAPolicyClause participates in a single inherited relationship, SUPAHasPolicyClause, as defined in [section 5.3.2.7](#).

5.7. The Abstract Class "SUPAPolicyComponentDecorator"

This is a mandatory class, and is used to implement the decorator pattern. The decorator pattern enables all or part of one or more objects to "wrap" another concrete object. This means that any concrete subclass of SUPAPolicyClause is wrapped by any concrete subclass of SUPAPolicyComponentDecorator, as shown in Figure 19 below.

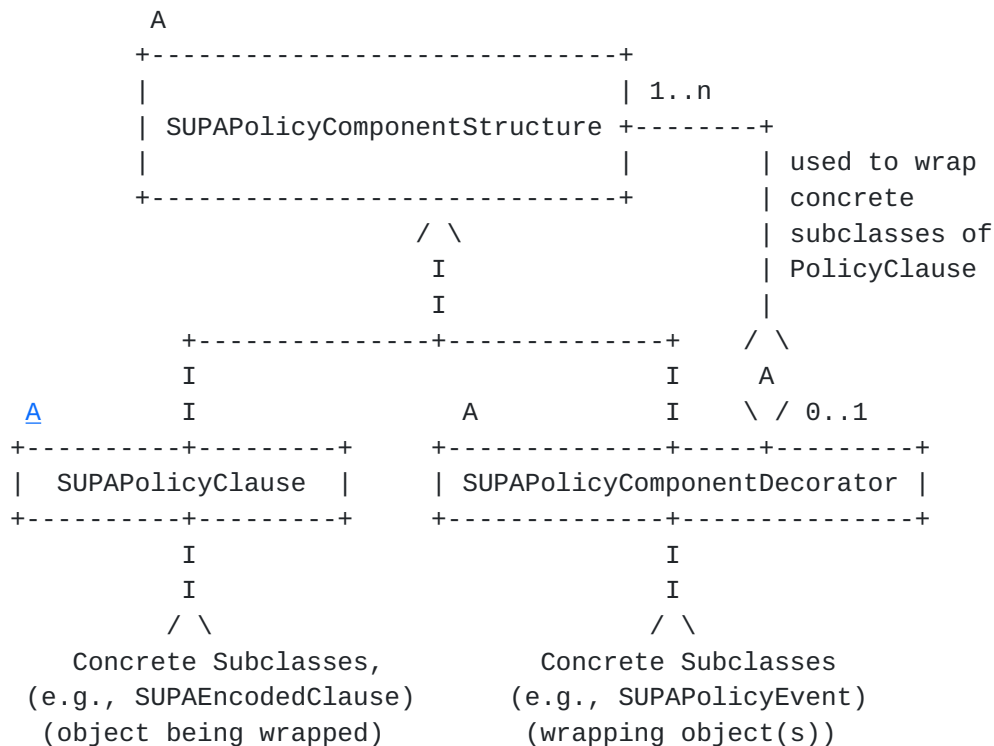


Figure 19. The PolicyComponent Decorator Pattern

5.7.1. The Decorator Pattern

Each **SUPAPolicyComponentDecorator** object **HAS_A** (i.e., wraps) a concrete instance of the **SUPAPolicyClause** object. This means that the **SUPAPolicyComponentDecorator** object has an instance variable that holds a reference to a **SUPAPolicyClause** object. Since the **SUPAPolicyComponentDecorator** object has the same interface as the **SUPAPolicyClause** object, the **SUPAPolicyComponentDecorator** object (and all of its subclasses) are transparent to clients of the **SUPAPolicyClause** object (and its subclasses). This means that **SUPAPolicyComponentDecorator** object instances can add attributes and/or methods to those of the concrete instance of the chosen subclass of **SUPAPolicyClause**.

Figure 20 shows how this is done for methods.

- Figure 20a shows the initial object to be wrapped
- Figure 20b shows the **SUPAPolicyEvent** object wrapping the **SUPAEncodedClause** object
- Figure 20c shows **SUPAGenericDecoratedComponent** object wrapping the **SUPAPolicyEvent** object.


```

+-----+
| SUPAEncodedClause |
|   eval()          |
+-----+

```

(a) Initial Object

==>

```

+-----+
| SUPAPolicyEvent   |
|   eval()          |
| +-----+        |
| | SUPAEncodedClause | |
| |   eval()          | |
| +-----+        |
+-----+

```

(b) SUPAPolicyEvent "wraps" SUPAEncodedClause

==>

```

+-----+
| SUPAGenericDecoratedComponent |
|   eval()                      |
| +-----+                    |
| | SUPAPolicyEvent             | | | |
| |   eval()                    | |
| | +-----+                  | |
| | | SUPAEncodedClause         | | |
| | |   eval()                  | | |
| | +-----+                  | | |
| +-----+                    |
+-----+

```

(c) SUPAGenericDecoratedComponent "wraps" SUPAPolicyEvent

Figure 20. Conceptual Depiction of eval() Decorated Method

When the eval() method is called in the outermost object (SUPAGenericDecoratedComponent), it delegates to the eval() method of SUPAPolicyEvent, which in turn delegates to the eval() method of SUPAEncodedClause. This method executes and returns the results to SUPAPolicyEvent, which executes and returns the results to SUPAGenericDecoratedComponent, which executes and returns the final result.

5.7.2. SUPAPolicyComponentDecorator Attributes

Currently, there are two attributes defined for this class, which are described in the following subsections. Both attributes are used by subclasses to constrain the behavior of that subclass; they do **not** affect the relationship between the concrete

subclass of SUPAPolicyComponentDecorator that is wrapping the
concrete subclass of SUPAPolicyClause.

Strassner, et al.

Expires January 19, 2017

[Page 65]

This is different than the use of similar attributes defined in the SUPAHasDecoratedPolicyComponentDetail association class (which are used to constrain the relationship between the concrete subclass of SUPAPolicyClause and the concrete subclass of the SUPAPolicyComponentDecorator object that is wrapping it).

Note that [2] does not define any attributes for this class.

5.7.2.1. The Attribute "supaPolCompConstraintEncoding"

This is a mandatory non-negative enumerated integer that defines how to interpret each string in the supaPolCompConstraint class attribute. Values include:

- 0: error
- 1: init
- 2: OCL 2.4
- 3: OCL 2.x
- 4: OCL 1.x
- 5: QVT 1.2 - Relations Language
- 6: QVT 1.2 - Operational language
- 7: Alloy
- 8: ASCII Text

Enumerations 1 and 2 signify an error state and an initialization state, respectively. Enumerations 3-5 are dedicated to OCL (with OCL 2.4 being the latest version as of this writing). QVT defines a set of languages [20] (the two most powerful and useful are defined by enumerations 5 and 6). Alloy is a language for describing constraints, and uses a SAT solver to guarantee correctness [21]. Enumeration 8 (ASCII Text) is not recommended (since it is informal, and hence, not verifiable), but is included for completeness.

If this class is instantiated, then this attribute SHOULD also be instantiated, and SHOULD be part of a conformant implementation.

5.7.2.2. The Attribute "supaPolCompConstraint[0..n]"

This is a mandatory array of string attributes. Each attribute specifies a constraint to be applied using the encoding defined in the supaPolCompConstraintEncoding class attribute. This provides a more rigorous and flexible treatment of constraints than is possible in [RFC3460], [4], [5], and [6].

Note: [0..n] means that this is a multi-valued property that may have zero or more attributes.

5.7.3. SUPAPolicyComponentDecorator Relationships

One relationship is currently defined for this class, which is

described in the following subsection.

Strassner, et al.

Expires January 19, 2017

[Page 66]

5.7.3.1. The Aggregation "SUPAHasDecoratedPolicyComponent"

This is a mandatory aggregation, and is part of a decorator pattern. It is used to enable a concrete instance of a SUPAPolicyComponentDecorator to dynamically add behavior to a specific type of SUPAPolicyClause object. The semantics of this aggregation are defined by the SUPAHasDecoratedPolicyComponentDetail association class.

5.7.3.2. The Association Class **"SUPAHasDecoratedPolicyComponentDetail"**

This is a mandatory concrete association class, and defines the semantics of the SUPAHasDecoratedPolicyComponent aggregation. The purpose of this class is to use the Decorator pattern to determine which SUPAPolicyComponentDecorator object instances, if any, are required to augment the functionality of the concrete subclass of SUPAPolicyClause that is being used.

Currently, there are two attributes defined for this class, which are described in the following subsections. Both attributes are used in this association class to constrain the ****relationship**** between the concrete subclass of SUPAPolicyComponentDecorator that is wrapping the concrete subclass of SUPAPolicyClause. Note that class attributes of SUPAPolicyComponentDecorator (see [section 5.9.2](#)) only affect that specific subclass.

5.7.3.2.1. The Attribute "supaDecoratedConstraintEncoding"

This is a mandatory non-negative enumerated integer that defines how to interpret each string in the supaDecoratedConstraint class attribute. Values include:

- 0: error
- 1: init
- 2: OCL 2.4
- 3: OCL 2.x
- 4: OCL 1.x
- 5: QVT 1.2 - Relations Language
- 6: QVT 1.2 - Operational language
- 7: Alloy
- 8: ASCII Text

Enumerations 1 and 2 signify an error state and an initialization state, respectively. Enumerations 3-5 are dedicated to OCL (with OCL 2.4 being the latest version as of this writing). QVT defines a set of languages [\[20\]](#) (the two most powerful and useful are defined by enumerations 5 and 6). Alloy is a language for describing constraints, and uses a SAT solver to guarantee correctness [\[21\]](#).

Enumeration 8 (ASCII Text) is not recommended (since it is informal, and hence, not verifiable), but is included for completeness.

Strassner, et al.

Expires January 19, 2017

[Page 67]

If this class is instantiated, then this attribute SHOULD also be instantiated, and SHOULD be part of a conformant implementation.

5.7.3.2.2. The Attribute "supaDecoratedConstraint[0..n]"

This is a mandatory array of string attributes. Its purpose is to collect a set of constraints to be applied to a decorated object. The interpretation of each constraint in the array is defined in the supaDecoratedConstraintsEncoding class attribute.

Note: [0..n] means that this is a multi-valued property that may have zero or more attributes.

5.7.4. Illustration of Constraints in the Decorator Pattern

Figure 21 builds a simple SUPAPolicyClause that has both types of relationships, and illustrates how the different constraints defined in sections 5.7.2 (class attribute constraints) and [section 5.7.3](#) (relationship constraints) can be used.

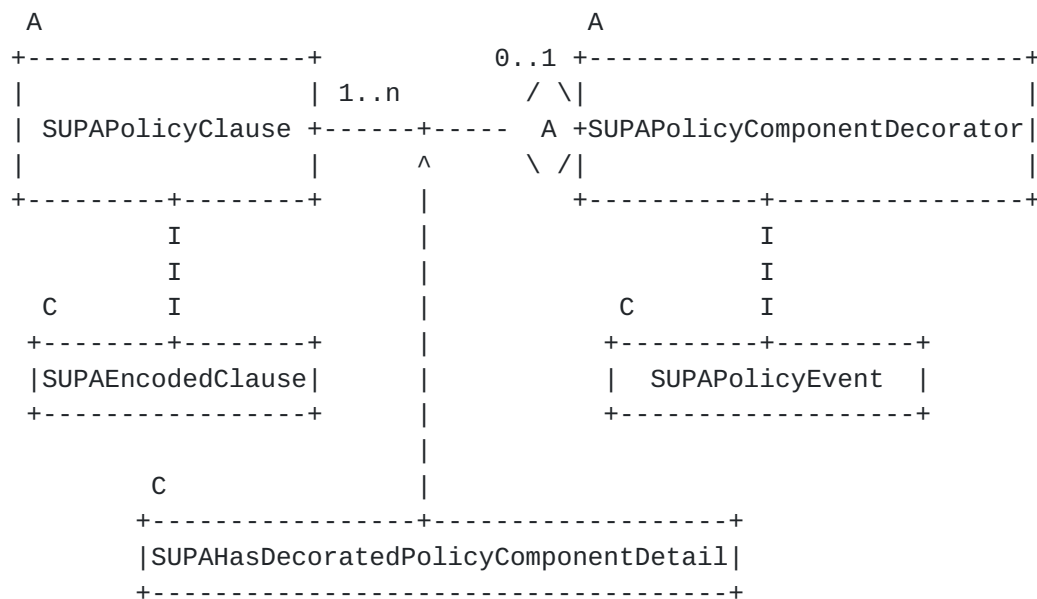


Figure 21. Constraints in the Decorator Pattern

Figure 21 says that a SUPAPolicyClause, realized as a SUPAEncodedClause, is wrapped by a SUPAPolicyEvent object. The attributes in the SUPAPolicyComponentDecorator object are used to constrain the attributes in the SUPAPolicyEvent object, while the attributes in the SUPAHasDecoratedPolicyComponentDetail object are used to constrain the behavior of the aggregation (SUPAHasDecoratedPolicyComponent). For example, the attributes in the SUPAPolicyComponentDecorator object could restrict the data type and range of the components in the SUPAPolicyEvent, while the attributes in the SUPAHasDecoratedPolicyComponentDetail object

could restrict which SUPAPolicyEvent objects are allowed to be used with which SUPAEncodedClauses.

Strassner, et al.

Expires January 19, 2017

[Page 68]

5.8. The Abstract Class "SUPAPolicyTerm"

This is a mandatory abstract class that is the parent of SUPAPolicy objects that can be used to define a standard way to test or set the value of a variable. It does this by defining a 3-tuple, in the form {variable, operator, value}, where each element of the 3-tuple is defined by a concrete subclass of the appropriate type (i.e., SUPAPolicyVariable, SUPAPolicyOperator, and SUPAPolicyValue classes, respectively). For example, a generic test or set of the value of a variable is expressed as:

{variable, operator, value}.

For event and condition clauses, this is typically as written above (e.g., does variable = value); for action clauses, it is typically written as <operator> <variable> <value> (e.g., SET var to 1). A class diagram is shown in Figure 22.

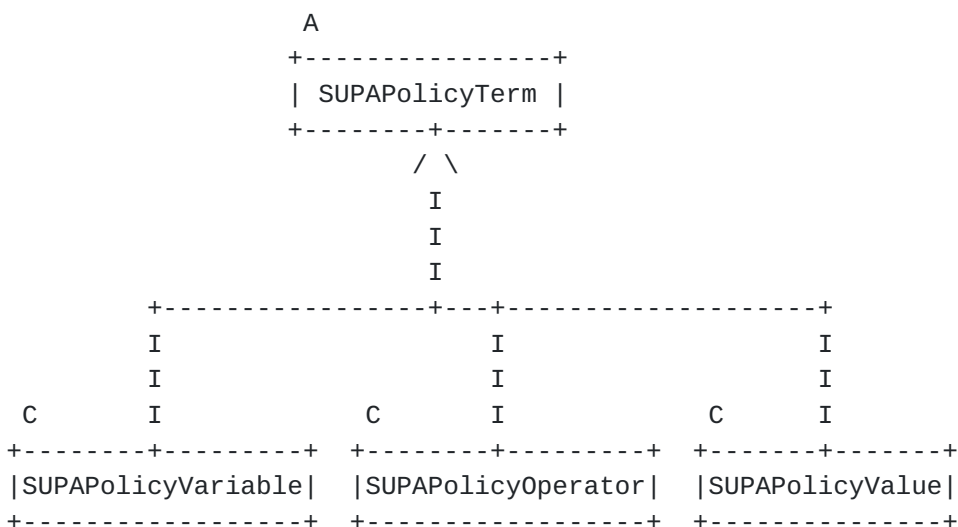


Figure 22. SUPAPolicyTerm Class Hierarchy

Note that generic test and set expressions do not have to only use objects that are subclasses of SUPAPolicyTerm. For example, the polVendorDecoratedContent attribute of the SUPAGenericDecoratedComponent could be used as the variable (or the value) term of a get or set expression.

Hence, the utility of the subclasses of SUPAPolicyTerm is in the ability of its subclasses to define a generic framework for implementing get and set expressions. This is in contrast to previous designs (e.g., [RFC3460] and [6]), which depended on defining a broad set of subclasses of PolicyVariable and PolicyValue. (Note that [4] does not have this generic capability).

5.8.1. SUPAPolicyTerm Attributes

Currently, SUPAPolicyTerm defines a single attribute, as described in the following subsection. Constraints on the subclasses of SUPAPolicyTerm can be applied in two different ways:

1. use SUPAPolicyComponentDecorator attributes to constrain just that individual subclass, and/or
2. use SUPAHasDecoratedPolicyComponentDetail association class attributes to constrain the relationship between the concrete subclass of SUPAPolicyClause and the concrete subclass of the SUPAPolicyTerm class

5.8.1.1. The Attribute "supaPolTermIsNegated"

This is a mandatory Boolean attribute. If the value of this attribute is true, then this particular SUPAPolicyTerm subclass (which represents a term) is negated; otherwise, it is not.

5.8.2. SUPAPolicyTerm Relationships

Currently, no dedicated relationships are defined for the SUPAPolicyTerm class (as there are in [\[RFC3460\]](#) and [\[6\]](#)) that aggregate policy variable and policy value objects into a policy rule). This is:

- 1) to enable the subclasses of SUPAPolicyTerm to be used by other SUPAPolicyComponentDecorator objects, and
- 2) because the decorator pattern replaces how such relationships were used in [\[RFC3460\]](#) and [\[6\]](#).

SUPAPolicyTerm, and its subclasses, inherit the SUPAHasDecoratedPolicyComponent aggregation, which was defined in [section 5.7.3](#).

5.9. The Concrete Class "SUPAPolicyVariable"

This is a mandatory concrete class that defines information that forms a part of a SUPAPolicyClause. It specifies a concept or attribute that represents a variable, which should be compared to a value, as specified in this SUPAPolicyClause. If it is used in a SUPAECAPolicyRule, then its value MAY be able to be changed at any time, including run-time, via use of the decorator pattern. This is not possible in previous designs ([\[RFC3460\]](#), [\[4\]](#), and [\[6\]](#)).

The value of a SUPAPolicyVariable is typically compared to the value of a SUPAPolicyValue using the type of operator defined in a SUPAPolicyOperator. However, other objects may be used instead of a SUPAPolicyValue object, and other operators may be defined

in addition to those defined in the SUPAPolicyOperator class.

Strassner, et al.

Expires January 19, 2017

[Page 70]

SUPAPolicyVariables are used to abstract the representation of a SUPAPolicyRule from its implementation. Some SUPAPolicyVariables are restricted in the values and/or the data type that they may be assigned. For example, port numbers cannot be negative, and they cannot be floating-point numbers. These and other constraints may be defined in two different ways:

1. use SUPAPolicyComponentDecorator attributes to constrain just that individual subclass, and/or
2. use SUPAHasDecoratedPolicyComponentDetail association class attributes to constrain the relationship between the concrete subclass of SUPAPolicyClause and the concrete subclass of the SUPAPolicyVariable class

Please refer to the examples in [section 7](#), which show how to restrict the value, data type, range, and other semantics of the SUPAPolicyVariable when used in a SUPAPolicyClause.

5.9.1. Problems with the [RFC3460](#) Version of PolicyValue

Please see [Appendix A](#) for a detailed comparison.

5.9.2. SUPAPolicyVariable Attributes

SUPAPolicyVariable defines one attribute, as described below.

5.9.2.1. The Attribute "supaPolVarName"

This is an optional string attribute that contains the name of this SUPAPolicyVariable. This variable name forms part of the {variable, operator, value} canonical form of a SUPAPolicyClause.

5.9.3. SUPAPolicyVariable Relationships

Currently, no relationships are defined for the SUPAPolicyVariable class (note that the decorator pattern obviates the need for relationships such as those defined in [[RFC3460](#)] and [[6](#)]). This is because SUPAPolicyVariable, and its subclasses, inherit the SUPAHasDecoratedPolicyComponent aggregation, which was defined in [section 5.7.3](#).

5.10. The Concrete Class "SUPAPolicyOperator"

This is a mandatory concrete class for modeling different types of operators that are used in a SUPAPolicyClause.

The restriction of the type of operator used in a SUPAPolicyClause restricts the semantics that can be expressed in that SUPAPolicyClause. It is typically used with SUPAPolicyVariables

and SUPAPolicyValue to form a SUPAPolicyClause.

Strassner, et al.

Expires January 19, 2017

[Page 71]

[5.10.1.](#) Problems with the [RFC3460](#) Version

Please see [Appendix A](#) for a detailed comparison.

[5.10.2.](#) SUPAPolicyOperator Attributes

Currently, SUPAPolicyOperator defines a single generic attribute, as described below.

[5.10.2.1.](#) The Attribute "supaPolOpType"

This is a mandatory non-negative enumerated integer that specifies the various types of operators that are allowed to be used in this particular SUPAPolicyClause. Values include:

- 0: error
- 1: init
- 2: Greater than
- 3: Greater than or equal to
- 4: Less than
- 5: Less than or equal to
- 6: Equal to
- 7: Not equal to
- 8: IN
- 9: NOT IN
- 10: SET
- 11: CLEAR (0 for integers, "" for strings, FALSE for Booleans)
- 12: BETWEEN (inclusive)

Note that 0 and 1 represent error and initialization states, respectively. Their purpose is to support dynamically building a SUPAPolicyClause by enabling the application to set the value of this attribute to a standard value.

Additional operators may be defined in future work. For example, if SUPAPolicyVariables and SUPAPolicyValues are expanded to/from include structured objects, then "deep" versions of operators 1-6 could also be defined. In this case, values 1-6 will be edited to explicitly indicate that they perform "shallow" comparison operations.

[5.10.3.](#) SUPAPolicyOperator Relationships

Currently, no relationships are defined for the SUPAPolicyOperator class (note that the decorator pattern obviates the need for relationships such as those in [\[6\]](#)). This is because SUPAPolicyOperator, and its subclasses, inherit the SUPAHasDecoratedPolicyComponent aggregation, which was defined in [section 5.7.3](#). Please refer to the examples in [section 7](#), which show how to restrict the value, data type, range, and other

semantics of the SUPAPolicyOperator when used in a SUPAPolicyClause.

Strassner, et al.

Expires January 19, 2017

[Page 72]

5.11. The Concrete Class "SUPAPolicyValue"

The SUPAPolicyValue class is a mandatory concrete class for modeling different types of values and constants that occur in a SUPAPolicyClause.

SUPAPolicyValues are used to abstract the representation of a SUPAPolicyRule from its implementation. Therefore, the design of SUPAPolicyValues depends on two important factors. First, just as with SUPAPolicyVariables (see [Section 5.11](#)), some types of SUPAPolicyValues are restricted in the values and/or the data type that they may be assigned. Second, there is a high likelihood that specific applications will need to use their own variables that have specific meaning to a particular application.

In general, there are two ways to apply constraints to an object instance of a SUPAPolicyValue:

1. use SUPAPolicyComponentDecorator attributes to constrain just that individual subclass, and/or
2. use SUPAHasDecoratedPolicyComponentDetail association class attributes to constrain the relationship between the concrete subclass of SUPAPolicyClause and the concrete subclass of the SUPAPolicyValue class

The value of a SUPAPolicyValue is typically compared to the value of a SUPAPolicyVariable using the type of operator defined in a SUPAPolicyOperator. However, other objects may be used instead of a SUPAPolicyVariable object, and other operators may be defined in addition to those defined in the SUPAPolicyOperator class.

5.11.1. Problems with the [RFC3460](#) Version of PolicyValue

Please see [Appendix A](#) for a detailed comparison.

5.11.2. SUPAPolicyValue Attributes

Currently, SUPAPolicyValue defines two generic attributes, as described below.

5.11.2.1. The Attribute "supaPolValContent[0..n]"

This is a mandatory attribute that defines an array of strings. The array contains the value(s) of this SUPAPolicyValue object instance. Its data type is defined by the supaPolValEncoding class attribute.

Note: [0..n] means that this is a multi-valued property that has zero or more attributes.

5.11.2.2. The Attribute "supaPolValEncoding"

This is a mandatory string attribute that contains the data type of the SUPAPolicyValue object instance. Its value is defined by the supaPolValContent class attribute. Values include:

```
0: error
1: init
2: String
3: Integer
4: Boolean
5: Floating Point
6: DateTime
7: GUID
8: UUID
9: URI
10: DN
11: FQDN
12: FQPN
13: NULL
```

Note that 0 and 1 represent error and initialization states, respectively. A string is a sequence of zero or more characters. An Integer is a whole number, and has no fractional part. A Boolean may take the values TRUE and FALSE. A floating point number may contain fractional values, as well as an exponent. A DateTime represents a value that has a date and/or a time component (as in the Java or Python libraries). A NULL explicitly models the lack of a value.

5.11.3. SUPAPolicyValue Relationships

Currently, no relationships are defined for the SUPAPolicyValue class (note that the decorator pattern obviates the need for relationships such as those in [6]). SUPAPolicyValue, and its subclasses, inherit the SUPAHasDecoratedPolicyComponent aggregation, which was defined in [section 5.7.3](#). Please refer to the examples in [section 7](#), which show how to restrict the value, data type, range, and other semantics of the SUPAPolicyValue when used in a SUPAPolicyClause.

5.12. The Concrete Class "SUPAGenericDecoratedComponent"

A SUPAGenericDecoratedComponent enables a generic object to be defined and used in a SUPAPolicyClause. This class was derived from [2], but is not present in [RFC3460], [4], [5], or [6].

This class should not be confused with the SUPAEncodedClause class. The SUPAGenericDecoratedComponent class represents a single, atomic object that defines a **portion** of a SUPAPolicyClause, whereas a

SUPAEncodedClause represents an ****entire**** SUPAPolicyClause.

Strassner, et al.

Expires January 19, 2017

[Page 74]

5.12.1. SUPAGenericDecoratedComponent Attributes

Currently, SUPAGenericDecoratedComponent defines two generic attributes, as described below.

5.12.1.1. The Attribute "supaVendorDecoratedCompContent[0..n]"

This is a mandatory attribute that defines an array of strings. This array contains the value(s) of the SUPAGenericDecoratedComponent object instance. Its data type is defined by the supaVendorDecoratedEncoding class attribute. Note: [0..n] means that this is a multi-valued property that has zero or more attributes.

5.12.1.2. The Attribute "supaVendorDecoratedCompEncoding"

This is a mandatory integer attribute that defines the format of the supaVendorDecoratedContent class attribute. Values include:

- 0: undefined
- 1: String
- 2: Integer
- 3: Boolean
- 4: Floating Point
- 5: DateTime
- 6: GUID
- 7: UUID
- 8: URI
- 9: FQDN
- 10: FQPN
- 11: NULL

Note that 0 and 1 represent error and initialization states, respectively. A string is a sequence of zero or more characters. An Integer is a whole number (e.g., it has no fractional part). A Boolean represents the values TRUE and FALSE. A floating point number may contain fractional values, as well as an exponent. A DateTime represents a value that has a date and/or a time component (as in the Java or Python libraries). A NULL explicitly models the lack of a value.

5.12.2. SUPAGenericDecoratedComponent Relationships

Currently, no relationships are defined for the SUPAGenericDecoratedComponent class (note that the decorator pattern obviates the need for relationships such as those in [6]). SUPAGenericDecoratedComponent participates in a single relationship, SUPAHasDecoratedPolicyComponent, as defined in [section 5.7.3](#).

**** This class is under discussion - please comment as to its utility ****

5.13. The Concrete Class "SUPAPolicyCollection"

A SUPAPolicyCollection is an optional concrete class that enables a collection (e.g., set, bag, or other, more complex, collections of elements) of ****arbitrary objects**** to be defined and used as part of a SUPAPolicyClause. This class was derived from [2], but is not present in [RFC3460], [4], [5], or [6].

5.13.1. Motivation

One of the problems with ECA policy rules is when a set of events or conditions needs to be tested. For example, if a set of events is received, the policy system may need to wait for patterns of events to emerge (e.g., any number of Events of type A, followed by either one event of type B or two events of type Event C). Similarly, for conditions, testing the value of a set of attributes may need to be performed. Both of these represent behavior similar to a set of if-then-else statements or a switch statement.

It is typically not desirable for the policy system to represent each choice in such clauses as its own policy clause (i.e., a 3-tuple), as this creates object explosion and poor performance. Furthermore, in these cases, it is often required to have a set of complex logic to be executed, where the logic varies according to the particular event or condition that was selected. It is much too complex to represent this using separate objects, especially when the logic is application- and/or vendor-specific. However, recall that one of the goals of this document was to facilitate the machine-driven construction of policies. Therefore, a solution to this problem is needed.

5.13.2. Solution

Therefore, this document defines the concept of a collection of entities, called a SUPAPolicyCollection. Conceptually, the items to be collected (e.g., events or conditions) are aggregated in one or more SUPAPolicyCollection objects of the appropriate type. Another optional SUPAPolicyCollection object could be used to aggregate logic blocks (including SUPAPolicies) to execute. Once finished, all appropriate SUPAPolicyCollection objects are sent to an external system for evaluation.

The computation(s) represented by the SUPAPolicyCollection may be part of a larger SUPAPolicyClause, since SUPAPolicyCollection is a subclass of SUPAPolicyComponentDecorator, and can be used to decorate a SUPAPolicyClause. Therefore, the external system is responsible for providing a Boolean TRUE or FALSE return value, so

that the policy system can use that value to represent the computation of the function(s) performed in the SUPAPolicyCollection in a Boolean clause.

Strassner, et al.

Expires January 19, 2017

[Page 76]

5.13.3. SUPAPolicyCollection Attributes

Currently, SUPAGenericDecoratedComponent defines five attributes, as described below.

5.13.3.1. The Attribute "supaPolCollectionContent[0..n]"

This is an optional attribute that defines an array of strings. Each string in the array identifies a domain-suitable identifier of an object that is collected by this SUPAPolicyCollection instance. Note: [0..n] means that this is a multi-valued property that has zero or more attributes.

5.13.3.2. The Attribute "supaPolCollectionEncoding"

This is an optional non-negative enumerated integer that defines the data type of the content of this collection instance. Values include:

- 0: error
- 1: init
- 2: by regex (regular expression)
- 3: by URI

Note that 0 and 1 represent error and initialization states, respectively. If the value of this attribute is 2, then each of the strings in the supaPolCollectionContent attribute represent a regex that contains all or part of a string to match the class name of the object that is to be collected by this instance of a SUPAPolicyCollection class. If the value of this attribute is 3, then each of the strings in the supaPolCollectionContent attribute represent a URI.

5.13.3.3. The Attribute "supaPolCollectionFunction"

This is an optional non-negative enumerated integer that defines the function of this collection instance. Values include:

- 0: error
- 1: init
- 2: event collection
- 3: condition collection
- 4: action collection
- 5: logic collection

Note that 0 and 1 represent error and initialization states, respectively. Values 2-4 define a collection of objects that are to be used to populate the event, condition, or action clauses, respectively, of a SUPAECAPolicyRule. A value of 4 indicates that

this collection contains objects that define logic for processing
a SUPAPolicy.

Strassner, et al.

Expires January 19, 2017

[Page 77]

5.13.3.4. The Attribute "supaPolCollectionIsOrdered"

This is an optional Boolean attribute. If the value of this attribute is TRUE, then all elements in this instance of this SUPAPolicyCollection are ordered.

5.13.3.5. The Attribute "supaPolCollectionType"

This is an optional non-negative enumerated integer that defines the type of collection that this instance is. Values include:

- 0: error
- 1: init
- 2: set
- 3: bag (e.g., multi-set)
- 4: dictionary (e.g., associative array)

Note that 0 and 1 represent error and initialization states, respectively. A set is an unordered collection of elements that MUST NOT have duplicates. A bag is an unordered collection of elements; it MAY have duplicates. A dictionary is a table that associates a key with a value.

Sets have a number of important functions, including:

- o membership: returns TRUE if the element being tested is in the set, and FALSE otherwise
- o subset: returns TRUE if all elements in the first set are also in the second set
- o union: returns all elements from both sets with no duplicates
- o intersection: returns all elements that are in both sets with no duplicates
- o difference: returns all elements in the first set that are not in the second set

Bags have a number of important functions in addition to the functions defined for sets (note that while the above set of functions for a set and a bag are the same, a bag is a different data type than a set):

- o multiplicity: returns the number of occurrences of an element in the bag
- o count: returns the number of all items, including duplicates
- o countDistinct: returns the number of items, where all duplicates are ignored

A dictionary is an unordered set of key:value pairs, where each key is unique within a given dictionary. The combination of a key and a value is called an item. The format of an item is defined as one element (the key) followed by a colon followed by a second element (the value). Each item in a set of items is separated by a comma. Keys MUST NOT be NULL; values MAY be NULL.

An example of a dictionary is {20:"FTP", 21:"FTP", 22: "SSH"}.
An example of a null dictionary is simply {}.

[5.13.4.](#) SUPAPolicyCollection Relationships

Currently, no relationships are defined for the SUPAGenericDecoratedComponent class (note that the decorator pattern obviates the need for relationships such as those in [\[6\]](#)). SUPAPolicyCollection participates in a single relationship, SUPAHasDecoratedPolicyComponent, as defined in [section 5.7.3](#).

[5.14.](#) The Concrete Class "SUPAPolicySource"

This is an optional class that defines a set of managed entities that authored, or are otherwise responsible for, this SUPAPolicyRule. Note that a SUPAPolicySource does NOT evaluate or execute SUPAPolicies. Its primary use is for auditability and the implementation of deontic and/or alethic logic. A class diagram is shown in Figure 12.

A SUPAPolicySource SHOULD be mapped to a role or set of roles (e.g., using the role-object pattern [\[11\]](#)). This enables role-based access control to be used to restrict which entities can author a given policy. Note that Role is a type of SUPAPolicyMetadata.

[5.14.1.](#) SUPAPolicySource Attributes

Currently, no attributes are defined for this class.

[5.14.2.](#) SUPAPolicySource Relationships

SUPAPolicySource participates in a single relationship, SUPAHasPolicySource, as defined in [section 5.3.2.1](#). SUPAPolicySource, and its subclasses, inherit the SUPAHasDecoratedPolicyComponent aggregation, which was defined in [section 5.7.3](#).

[5.15.](#) The Concrete Class "SUPAPolicyTarget"

This is an optional class that defines a set of managed entities

that a SUPAPolicy is applied to. Figure 12 shows a class diagram of the SUPAPolicyTarget.

A managed object must satisfy two conditions in order to be defined as a SUPAPolicyTarget. First, the set of managed entities that are to be affected by the SUPAPolicy must all agree to play the role of a SUPAPolicyTarget. In general, a managed entity may or may not be in a state that enables SUPAPolicies to be applied to it to change its state; hence, a negotiation process may need to occur to enable the SUPAPolicyTarget to signal when it is willing to have SUPAPolicies applied to it. Second, a SUPAPolicyTarget must be able to process (directly or with the aid of a proxy) SUPAPolicies.

If a proposed SUPAPolicyTarget meets both of these conditions, it SHOULD set its supaPolicyTargetEnabled Boolean attribute to a value of TRUE.

A SUPAPolicyTarget SHOULD be mapped to a role (e.g., using the role-object pattern). This enables role-based access control to be used to restrict which entities can author a given policy. Note that Role is a type of SUPAPolicyMetadata.

5.15.1. SUPAPolicyTarget Attributes

Currently, no attributes are defined for the SUPAPolicyTarget class.

5.15.2. SUPAPolicyTarget Relationships

SUPAPolicyTarget participates in a single relationship, SUPAHasPolicyTarget, as defined in [section 5.3.2.3](#).

5.16. The Abstract Class "SUPAPolicyMetadata"

Metadata is information that describes and/or prescribes characteristics and behavior of another object that is ****not**** an inherent, distinguishing characteristic or behavior of that object (otherwise, it would be an integral part of that object).

For example, a socialSecurityNumber attribute should not be part of a generic Person class. First, most countries in the world do not know what a social security number is, much less use them. Second, a person is not created with a social security number; rather, a social security number is used to track people for administering social benefits, though it is also used as a form of identification.

Continuing the example, a better way to add this capability to a model would be to have a generic Identification class, then define a SocialSecurityNumber subclass, populate it as necessary, and then define a composition between a Person and it (this is a

composition because social security numbers are not reused).

Strassner, et al.

Expires January 19, 2017

[Page 80]

Since social security numbers are given to US citizens, permanent residents, and temporary working residents, and because it is also used to administer benefits, the composition is realized as an association class to define how it is being used.

An example of descriptive metadata for network elements would be documentation about best current usage practices (this could also be in the form of a reference). An example of prescriptive metadata for network elements would be the definition of a time period during which specific types of operations are allowable.

This is an optional class that defines the top of a hierarchy of model elements that are used to define different types of metadata that can be applied to policy and policy component objects. This enables common metadata to be defined as objects and then reused when the metadata are applicable. One way to control whether SUPAPolicyMetadata objects are reused is by using the attributes of the SUPAHasPolicyMetadataDetail association class.

It is recommended that this class, along with its SUPAPolicyConcreteMetadata and SUPAPolicyMetadataDecorator subclasses, be used as part of a conformant implementation. It is defined to be optional, since metadata is not strictly required. However, metadata can help specify and describe SUPAPolicyObject entities, and can also be used to drive dynamic behavior.

5.16.1. SUPAPolicyMetadata Attributes

This section defines the attributes of the SUPAPolicyMetadata class.

5.16.1.1. The Attribute "supaPolMetadataDescription"

This is an optional string attribute that defines a free-form textual description of this metadata object.

5.16.1.2. The Attribute "supaPolMetadataIDContent"

This is a mandatory string attribute that represents part of the object identifier of an instance of this class. It defines the content of the object identifier. It works with another class attribute, called supaPolMetadataIDEncoding, which defines how to interpret this attribute. These two attributes form a tuple, and together enable a machine to understand the syntax and value of an object identifier for the object instance of this class.

5.16.1.3. The Attribute "supaPolMetadataIDEncoding"

This is an optional non-zero enumerated integer attribute that represents part of the object identifier of an instance of this class. It defines the format of the object identifier. It works with another class attribute, called `supaPolMetadataIDContent`, which defines the content of the object ID.

These two attributes form a tuple, and together enable a machine to understand the syntax and value of an object identifier for the object instance of this class. The `supaPolMetadataIDEncoding` attribute is mapped to the following values:

```
0: error
1: init
2: GUID
3: UUID
4: URI
5: FQDN
6: FQPN
7: string_instance_id
```

Note that 0 and 1 represent error and initialization states, respectively. Values 2-6 define the content as a reference. Value 7 defines the content as a string that is the canonical representation, in ASCII, of an instance ID of this object.

5.16.1.4. The Attribute "supaPolMetadataName"

This is an optional string attribute that defines the name of this `SUPAPolicyMetadata` object.

5.16.2. SUPAPolicyMetadata Relationships

`SUPAPolicyMetadata` participates in a single aggregation, which is defined in the following subsections.

5.16.2.1. The Aggregation "SUPAHasPolicyMetadata"

This is a mandatory aggregation that defines the set of `SUPAPolicyMetadata` that are aggregated by this particular `SUPAPolicyObject`. It is recommended that this aggregation be used as part of a conformant implementation.

The multiplicity of this relationship is defined as 0..n on the aggregate (`SUPAPolicyObject`) side, and 0..n on the part (`SUPAPolicyMetadata`) side. This means that this relationship is optional. The semantics of this aggregation are implemented using the `SUPAHasPolicyMetadataDetail` association class.

5.16.2.2. The Abstract Class "SUPAHasPolicyMetadataDetail"

This is a mandatory concrete association class, and defines the semantics of the SUPAHasPolicyMetadata aggregation. Its purpose is to determine which SUPAPolicyMetadata object instances should be attached to which particular object instances of the SUPAPolicyObject class. This is done by using the attributes and relationships of the SUPAPolicyMetadataDetail class to constrain which SUPAPolicyMetadata objects can be aggregated by which particular SUPAPolicyObject instances. It is recommended that this association class be used as part of a conformant implementation.

5.16.2.2.1. The Attribute "supaPolMetadataIsApplicable"

This is an optional Boolean attribute. If the value of this attribute is TRUE, then the SUPAPolicyMetadata object(s) of this particular SUPAHasPolicyMetadata aggregation SHOULD be aggregated by this particular SUPAPolicyObject.

5.16.2.2.2. The Attribute "supaPolMetadataConstraintEncoding"

This is an optional non-negative enumerated integer that defines how to interpret each string in the supaPolMetadataConstraint class attribute. Values include:

- 0: error
- 1: init
- 2: OCL 2.4
- 3: OCL 2.x
- 4: OCL 1.x
- 5: QVT 1.2 - Relations Language
- 6: QVT 1.2 - Operational language
- 7: Alloy
- 8: ASCII Text

Enumerations 1 and 2 signify an error state and an initialization state, respectively. Enumerations 3-5 are dedicated to OCL (with OCL 2.4 being the latest version as of this writing). QVT defines a set of languages [20] (the two most powerful and useful are defined by enumerations 5 and 6). Alloy is a language for describing constraints, and uses a SAT solver to guarantee correctness [21]. Enumeration 8 (ASCII Text) is not recommended (since it is informal, and hence, not verifiable), but is included for completeness.

If this class is instantiated, then this attribute SHOULD also be instantiated, and SHOULD be part of a conformant implementation.

5.16.2.2.3. The Attribute "supaPolMetadataConstraint[0..n]"

This is an optional array of string attributes. Each attribute specifies a constraint to be applied using the format identified by the value of the supaPolMetadataPolicyConstraintEncoding class attribute. This provides a more rigorous and flexible treatment of constraints than is possible in [\[RFC3460\]](#).

If this class is instantiated, then this attribute SHOULD also be instantiated, and should be part of a conformant implementation. Note: [0..n] means that this is a multi-valued property that has zero or more attributes.

5.17. The Concrete Class "SUPAPolicyConcreteMetadata"

This is an optional concrete class. It defines an object that will be wrapped by concrete instances of the SUPAPolicyMetadataDecorator class. It can be viewed as a "carrier" for metadata that will be attached to a subclass of SUPAPolicyObject. Since the decorator pattern is used, any number of concrete subclasses of the SUPAPolicyMetadataDecorator class can wrap an instance of the SUPAPolicyConcreteMetadata class.

It is recommended that this class be used as part of a conformant implementation.

5.17.1. SUPAPolicyConcreteMetadata Attributes

Currently, two attributes are defined for the SUPAPolicyConcreteMetadata class, and are described in the following subsections.

5.17.1.1. The Attribute "supaPolMDValidPeriodEnd"

This is an optional attribute. Its data type should be able to express a date and a time. This attribute defines the ending date and time that this Metadata object is valid for.

5.17.1.2. The Attribute "supaPolMDValidPeriodStart"

This is an optional attribute. Its data type should be able to express a date and a time. This attribute defines the starting date and time that this Metadata object is valid for.

5.17.2. SUPAPolicyConcreteMetadata Relationships

This class inherits the relationships of the SUPAPolicyMetadata class; see [section 5.16.2](#). It can also be used by subclasses of the SUPAPolicyMetadataDecorator class, and hence, can participate

in the SUPAHasMetadataDecorator aggregation; see [section 5.18.2](#).

Strassner, et al.

Expires January 19, 2017

[Page 84]

5.18. The Abstract Class "SUPAPolicyMetadataDecorator"

This is an optional class, and is used to implement the decorator pattern (see [section 5.7.1.](#)) for metadata objects. This pattern enables all or part of one or more SUPAPolicyMetadataDecorator subclasses to "wrap" a SUPAPolicyConcreteMetadata object instance.

It is recommended that this class be used as part of a conformant implementation.

5.18.1. SUPAPolicyMetadataDecorator Attributes

Currently, no attributes are defined for the SUPAPolicyMetadataDecorator class.

5.18.2. SUPAPolicyMetadataDecorator Relationships

This class inherits the relationships of the SUPAPolicyMetadata class; see [section 5.16.2.](#) It also defines a single aggregation, SUPAHasMetadataDecorator, which is used to implement the decorator pattern, as described in the following subsections.

5.18.2.1. The Aggregation "SUPAHasMetadataDecorator"

This is an optional aggregation, and is part of a decorator pattern. It is used to enable a concrete instance of a SUPAPolicyMetadataDecorator to dynamically add behavior to a SUPAPolicyConcreteMetadata object instance. The semantics of this aggregation are defined by the SUPAHasMetadataDecoratorDetail association class.

It is recommended that this aggregation be part of a conformant implementation.

The multiplicity of this aggregation is 0..1 on the aggregate (SUPAPolicyMetadataDecorator) side and 1..n on the part (SUPAPolicyMetadata) side. This means that if this aggregation is defined, then at least one SUPAPolicyMetadata object (e.g., a concrete subclass of SUPAPolicyMetadataDecorator) must also be instantiated and wrapped by this SUPAPolicyConcreteMetadata object instance. The semantics of this aggregation are defined by the SUPAHasMetadataDecoratorDetail association class.

5.18.2.2. The Association Class "SUPAHasMetadataDecoratorDetail"

This is an optional concrete association class, and defines the semantics of the SUPAHasMetadataDecorator aggregation. The purpose of this class is to use the Decorator pattern to determine which SUPAPolicyMetadataDecorator object instances, if any, are required to augment the functionality of the

SUPAPolicyConcreteMetadata object instance that is being used.

Strassner, et al.

Expires January 19, 2017

[Page 85]

It is recommended that this association class be part of a conformant implementation.

Attributes for this association class will be defined in a future version of this document.

5.19. The Concrete Class "SUPAPolicyAccessMetadataDef"

This is an optional concrete class that defines access control information, in the form of metadata, that can be added to a SUPAPolicyObject. This is done using the SUPAHasPolicyMetadata aggregation (see [section 5.2.2.](#)). This enables all or part of a standardized description and/or specification of access control for a given SUPAPolicyObject to be easily changed at runtime by wrapping an object instance of the SUPAPolicyConcreteMetadata class (or its subclass) with all or part of this object, and then adorning the SUPAPolicyObject with the SUPAPolicyConcreteMetadata object instance.

5.19.1. SUPAPolicyAccessMetadataDef Attributes

Currently, the SUPAPolicyAccessMetadataDef class defines three attributes; these are described in the following subsections.

5.19.1.1. The Attribute "supaPolAccessPrivilegeDef"

This is an optional non-negative enumerated integer attribute. It specifies the access privileges that external Applications have when interacting with a specific SUPAPolicyObject that is adorned with an instance of this SUPAPolicyAccessMetadataDef object. This enables the management system to control, in a consistent manner, the set of operations that external Applications have for SUPAPolicies and components of SUPAPolicies. Values include:

- 0: error
- 1: init
- 2: read only (for all policy components)
- 3: read and write (for all policy components)
- 4: privileges are specified by an external MAC model
- 5: privileges are specified by an external DAC model
- 6: privileges are specified by an external RBAC model
- 7: privileges are specified by an external ABAC model
- 8: privileges are specified by an external custom model

Note that 0 and 1 represent error and initialization states, respectively. Values 4-8 indicate that a formal external access control model is used. The name of this model, and its location, are specified in two other class attributes, called `supaPolAccessPrivilegeModelName` and `supaPolAccessPrivilegeModelRef`.

MAC, DAC, RBAC, and ABAC (values 4-7 stand for Mandatory Access Control, Discretionary Access Control, Role-Based Access Control, and Attribute-Based Access Control, respectively. They are defined in [22]. A value of 8 indicates that a formal external model that is not MAC, DAC, RBAC, or ABAC is used.

5.19.1.2. The Attribute "supaPolAccessPrivilegeModelName"

This is an optional string attribute that contains the name of the access control model being used. If the value of the `supaPolAccessPrivilegeDef` is 0-2, then the value of this attribute is not applicable. Otherwise, the text in this class attribute should be interpreted according to the value of the `supaPolAccessPrivilegeModelRef` class attribute.

5.19.1.3. The Attribute "supaPolAccessPrivilegeModelRef"

This is an optional non-negative enumerated integer attribute that defines the data type of the `supaPolAccessPrivilegeModelName` attribute. If the value of the `supaPolAccessPrivilegeDef` class attribute is 0-2, then the value of this attribute is not applicable. Otherwise, the value of this class attribute defines how to interpret the text in the `supaPolAccessPrivilegeModelRef` class attribute. Values include:

- 0: error
- 1: init
- 2: GUID
- 3: UUID
- 4: URI
- 5: FQDN
- 6: FQPN
- 7: string_instance_id

Note that 0 and 1 represent error and initialization states, respectively. Values 2-6 define the content as a reference. Value 7 defines the content as a string that is the canonical representation, in ASCII, of an instance ID of this object.

5.20. The Concrete Class "SUPAPolicyVersionMetadataDef"

This is an optional concrete class that defines versioning information, in the form of metadata, that can be added to a `SUPAPolicyObject`. This enables all or part of a standardized description and/or specification of version information for a given `SUPAPolicyObject` to be easily changed at runtime by wrapping an object instance of the `SUPAPolicyConcreteMetadata` class (or its subclass) with all or part of this object.

5.20.1. SUPAPolicyVersionMetadataDef Attributes

Version information is defined in a generic format based on the Semantic Versioning Specification [18] as follows:

```
<major>.<minor>.<patch>[<pre-release>][<build-metadata>]
```

where the first three components (major, minor, and patch) MUST be present, and the latter two components (pre-release and build-metadata) MAY be present. A version number MUST take the form <major>.<minor>.<patch>, where <major>, <minor>, and <patch> are each non-negative integers that MUST NOT contain leading zeros. In addition, the value of each of these three elements MUST increase numerically.

In this approach:

- o supaVersionMajor denotes a new release; this number MUST be incremented when either changes are introduced that are not backwards-compatible, and/or new functionality not previously present is introduced
- o supaVersionMinor denotes a minor release; this number MUST be incremented when new features and/or bug fixes to a major release that are backwards-compatible are introduced, and/or if any features are marked as deprecated
- o supaVersionPatch denotes a version that consists ONLY of bug fixes, and MUST be incremented when these bug fixes are Not backwards-compatible

When multiple versions exist, the following rules define their precedence:

1. Precedence MUST be calculated by separating the version into major, minor, patch, and pre-release identifiers, in that order. Note that build-metadata is NOT used to calculate precedence.
2. Precedence is determined by the first difference when comparing each of these identifiers, from left to right, as follows:
 - a. Major, minor, and patch versions are always compared numerically (e.g., 1.0.0 < 2.0.0 < 2.1.0 < 2.1.1)
 - b. When major, minor, and patch are equal, a pre-release version has LOWER precedence than a normal version (e.g., 1.0.0-alpha < 1.0.0)
 - c. Precedence for two pre-release versions with the same major, minor, and patch version MUST be determined by comparing each dot separated identifier from left to right until a difference is found as follows:
 - identifiers consisting only of digits are compared

numerically and identifiers with letters and/or
hyphens are compared lexically in ASCII sort order

Strassner, et al.

Expires January 19, 2017

[Page 88]

- Numeric identifiers always have lower precedence than non-numeric identifiers
 - A larger set of pre-release fields has a higher precedence than a smaller set, if all of the preceding identifiers are equal
3. Example: 1.0.0-alpha < 1.0.0-alpha.1 < 1.0.0-alpha-beta < 1.0.0-beta < 1.0.0-beta.2 < 1.0.0-rc.1 < 1.0.0.

Currently, the SUPAPolicyVersionMetadataDef class defines five attributes; these are described in the following subsections.

5.20.1.1. The Attribute "supaVersionMajor"

This is a mandatory string attribute, and contains a string representation of an integer that is greater than or equal to zero. It indicates that a significant increase in functionality is present in this version. It MAY also indicate that this version has changes that are NOT backwards-compatible; this MAY be denoted in the supaVersionBuildMetadata class attribute.

The special string "0.1.0" is for initial development that MUST NOT be considered stable. Improvements to this initial version, before they are released to the public, are denoted by incrementing the minor and patch version numbers.

The major version X (i.e., X.y.z, where $X > 0$) MUST be incremented if any backwards incompatible changes are introduced. It MAY include minor and patch level changes. The minor and patch version numbers MUST be reset to 0 when the major version number is incremented.

5.20.1.2. The Attribute "supaVersionMinor"

This is a mandatory string attribute, and contains a string representation of an integer that is greater than or equal to zero. It indicates that this release contains a set of features and/or bug fixes that MUST be backwards-compatible.

The minor version Y (i.e., x.Y.z, where $x > 0$) MUST be incremented if new, backwards-compatible changes are introduced. It MUST be incremented if any features are marked as deprecated. It MAY be incremented if new functionality or improvements are introduced. It MAY include patch level changes. The patch version number MUST be reset to 0 when the minor version number is incremented.

5.20.1.3. The Attribute "supaVersionPatch"

This is a mandatory string attribute, and contains a string representation of an integer that is greater than or equal to zero. It indicates that this version contains ONLY bug fixes.

The patch version Z (i.e., x.y.Z, where $x > 0$) MUST be incremented if new, backwards-compatible changes are introduced. A bug fix is defined as an internal change that fixes incorrect behavior.

5.20.1.4. The Attribute "supaVersionPreRelease"

This is an optional string attribute, and contains a string defining the pre-release version.

A pre-release version MAY be denoted by appending a hyphen and a series of dot-separated identifiers immediately following the patch version. Identifiers MUST comprise only ASCII alphanumerics and a hyphen. Identifiers MUST NOT be empty. Numeric identifiers MUST NOT include leading zeroes. Pre-release versions have a lower precedence than the associated normal version. A pre-release version indicates that the version is unstable and might not satisfy the intended compatibility requirements as denoted by its associated normal version. Examples include: 1.0.0-alpha, 1.0.0-alpha.1, 1.0.0-0.3.7, and 1.0.0-x.7.z.92.

5.20.1.5. The Attribute "supaVersionBuildMetadata"

This is an optional string attribute, and contains a string defining the build metadata. Build metadata MAY be denoted by appending a plus sign and a series of dot-separated identifiers immediately following the patch or pre-release version. Identifiers MUST be made up of only ASCII alphanumerics and a hyphen. Identifiers MUST NOT be empty. Build metadata SHOULD be ignored when determining version precedence. Examples include: 1.0.0.-alpha+1, 1.0.0+20130313144700, and 1.0.0-beta+exp.sha.5114f85.

6. SUPA ECAPolicyRule Information Model

This section defines the classes, attributes, and relationships of the SUPA ECAPolicyRule Information Model (EPRIM). Unless otherwise stated, all classes (and attributes) defined in this section were abstracted from DEN-ng [2], and a version of them are in the process of being added to [5].

6.1. Overview

Conceptually, the EPRIM is a set of subclasses that specialize the concepts defined in the GPIM for representing the components of a Policy that uses ECA semantics. This is shown in Figure 23 (only new EPRIM subclasses and their GPIM superclasses are shown; note that the SUPAPolicyMetadata hierarchy is used *as is*).

(Class of another model that SUPA is integrating into)



Figure 23. The EPRIM Class Hierarchy

Specifically, the EPRIM specializes the SUPAPolicyStructure class to create a SUPAECAPolicyRule (see sections [6.4](#) - [6.6](#)); it also specializes two subclasses of the SUPAPolicyComponentStructure class to create two new sets of policy components. These two SUPAPolicyComponentStructure subclasses are:

- o a new subclass of SUPAPolicyClause, called SUPABooleanClause (see sections [6.7](#) - [6.9](#)), is defined for constructing Boolean clauses that are specific to the needs of ECA Policy Rules
- o a new subclass of SUPAPolicyComponentDecorator, called SUPAECAComponent (see sections [6.10](#) - [6.13](#)), is defined for constructing reusable objects that represent Events, Conditions, and Actions

The EPRIM provides new functionality, based on the GPIM, by extending the GPIM to define new classes and relationships. The EPRIM does NOT define new classes that are not inherited from existing GPIM classes. This ensures that the semantics of the GPIM are not changed, even though new functionality (for ECA Policy Rules and components) are being defined.

The overall strategy for refining the GPIM is as follows:

- o SUPAECAPolicyRule is defined as a subclass of the GPIM SUPAPolicyStructure class
- o A SUPAECAPolicyRule has event, condition, and action clauses
 - o Conceptually, this can be viewed as three aggregations between the SUPAECAPolicyRule and each clause
 - o Each aggregation uses an instance of a concrete subclass of SUPAPolicyClause; this can be a SUPABooleanClause (making it ECA-specific), a SUPAEncodedClause (making it generic in nature), or a new subclass of SUPAPolicyClause
 - o Concrete subclasses of SUPAPolicyClause may be decorated with zero or more concrete subclasses of the SUPAPolicyComponentDecorator class
- o An optional set of GPIM SUPAPolicySource objects can be defined to represent the authoring of a SUPAECAPolicyRule
- o An optional set of GPIM SUPAPolicyTarget objects can be defined to represent the set of managed entities that will be affected by this SUPAECAPolicyRule
- o An optional set of SUPAPolicyMetadata can be defined for any of the objects that make up a SUPAECAPolicyRule, including any of its components

6.2. Constructing a SUPAECAPolicyRule

There are several different ways to construct a SUPAECAPolicyRule; they differ in which set of components are used to define the content of the SUPAECAPolicyRule, and whether each component is decorated or not. The following are some examples of creating a SUPAECAPolicyRule:

- o Define three types of SUPABooleanClauses, one each for the event, condition, and action clauses that make up a SUPAECAPolicyRule
- o For one or more of the above clauses, associate an appropriate set of SUPAPolicyEvent, SUPAPolicyCondition, or SUPAPolicyAction objects, and complete the clause using an appropriate SUPAPolicyOperator and a corresponding SUPAPolicyValue or SUPAPolicyVariable
- o Note that compound Boolean clauses may be formed using one or more SUPABooleanClauseComposite objects with one or more SUPABooleanClauseAtomic objects
- o Define a SUPAPolicyCollection component, which is used to aggregate a set of objects appropriate for a clause, and complete the clause using an appropriate SUPAPolicyOperator and a corresponding SUPAPolicyValue or SUPAPolicyVariable
- o Create a new concrete subclass of SUPAPolicyComponentStructure (i.e., a sibling class of SUPAPolicyComponentDecorator and SUPAPolicyClause), and use this new subclass in a concrete subclass of SUPABooleanClause; note that this approach enables the new concrete subclass of SUPAPolicyComponentStructure to optionally be decorated as well
- o Create a new subclass of SUPAPolicyComponentDecorator (e.g., a sibling of SUPAECAComponent) that provides ECA-specific functionality, and use that to decorate a SUPAPolicyClause
- o Create a new concrete subclass of SUPAPolicyStructure that provides ECA-specific functionality, and define all or part of its content by aggregating a set of SUPAPolicyClauses

6.3. Working With SUPAECAPolicyRules

A SUPAECAPolicyRule is a type of SUPAPolicy. It is a tuple that MUST have three clauses, defined as follows:

- o The event clause defines a Boolean expression that, if TRUE, triggers the evaluation of its condition clause (if the event clause is not TRUE, then no further action for this policy rule takes place).
- o The condition clause defines a Boolean expression that, if TRUE, enables the actions in the action clause to be executed (if the condition clause is not TRUE, then no further action for this policy rule takes place).

- o The action clause contains a set of actions (note that an action MAY invoke another SUPAECAPolicyRule; see [section 6.13](#)).

Strassner, et al.

Expires January 19, 2017

[Page 93]

Each of the above clauses can be a simple Boolean expression (of the form {variable operator value}, or a compound Boolean expression consisting of Boolean combinations of clauses. Compound Boolean expressions SHOULD be in CNF or DNF.

Note that each of the above three clauses MAY have a set of SUPAPolicyMetadata objects that can constrain, or otherwise affect, how that clause is treated. For example, a set of SUPAPolicyMetadata MAY affect whether none, some, or all actions are executed, and what happens if an action fails.

Each of the three clauses can be constructed from either a SUPAEncodedClause or a SUPABooleanClause. The advantage of using SUPAEncodedClauses is simplicity, as the content of the clause is encoded directly into the attributes of the SUPAEncodedClause. The advantage of using SUPABooleanClauses is reusability, since each term in each clause is potentially a reusable object.

Since a SUPABooleanClause is a subclass of a SUPAPolicyClause (see [Section 6.7](#)), it can be decorated by one or more concrete subclasses of SUPAPolicyComponentDecorator. Therefore, a SUPAECAPolicyRule can be built entirely from objects defined in the GPIM and EPRIM, which facilitates the construction of SUPAPolicies by a machine.

The relation between a SUPAECAPolicyRule and a SUPAPolicyClause is shown in Figure 24, and is explained in further detail in [Section 6.4](#).

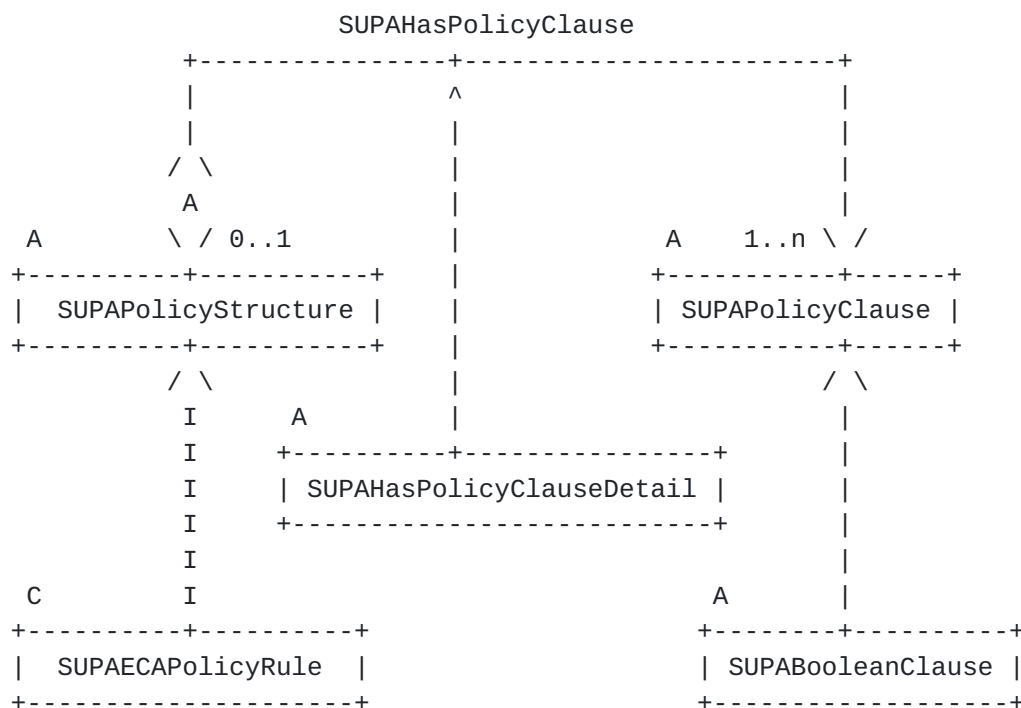


Figure 24. SUPAECAPolicyRule Clauses

Strassner, et al.

Expires January 19, 2017

[Page 94]

The SUPAHasPolicyClause aggregation is implemented using the SUPAHasPolicyClauseDetail association class. These were described in sections [5.4.2.1](#) and [5.4.2.2](#), respectively.

6.4. The Abstract Class "SUPAECAPolicyRule"

This is a mandatory abstract class, which is a PolicyContainer that aggregates PolicyEvents, PolicyConditions, PolicyActions into a type of policy rule known as an Event-Condition-Action (ECA) policy rule. As previously explained, this has the following semantics:

```
IF the event clause evaluates to TRUE
  IF the condition clause evaluates to TRUE
    THEN execute actions in the action clause
  ENDIF
ENDIF
```

The event clause, condition clause, and action clause collectively form a three-tuple. Each clause **MUST** be defined by at least one SUPAPolicyClause (which **MAY** be decorated with other elements, as described in [section 5.7](#)).

Each of the three types of clauses is a 3-tuple of the form:

```
{variable operator value}
```

Each of the three clauses **MAY** be combined with additional clauses using any combination of logical AND, OR, and NOT operators; this forms a "compound" Boolean clause. For example, if A, B, and C are three attributes in an event, then a valid event clause could be:

```
(A AND B) OR C
```

Note that the above expression is in DNF; the equivalent CNF form is ((A OR C) AND (B OR C)). In either case, the output of all three clauses is either TRUE or FALSE; this facilitates combining and chaining SUPAECAPolicyRules.

An action clause **MAY** invoke a new SUPAECAPolicyRule; see [section 6.13](#) for more details.

An ECAPolicyRule **MAY** be optionally augmented with PolicySources and/or PolicyTargets (see sections [5.16](#) and [5.17](#), respectively). In addition, all objects that make up a SUPAECAPolicyRule **MAY** have SUPAPolicyMetadata (see [section 5.16](#)) attached to them to further describe and/or specify behavior.

When defined in an information model, each of the event, condition, and action clauses MUST be represented as an aggregation between a SUPAECAPolicyRule (the aggregate) and a set of event, condition, or action objects (the components). However, a data model MAY map these definitions to a more efficient form (e.g., by flattening these three types of object instances, along with their respective aggregations, into a single object instance).

The composite pattern [3] is applied to the SUPAECAPolicyRule class, enabling its (concrete) subclasses to be used as either a stand-alone policy rule or as a hierarchy of policy rules. SUPAECAPolicyRuleComposite and SUPAECAPolicyRuleAtomic both inherit from SUPAECAPolicyRule. This means that they are both a type of SUPAECAPolicyRule. Hence, the HasSUPAECAPolicyRule aggregation enables a particular SUPAECAPolicyRuleComposite object to aggregate both SUPAECAPolicyRuleComposite as well as SUPAECAPolicyRuleAtomic objects. In contrast, a SUPAECAPolicyRuleAtomic can NOT aggregate either a SUPAECAPolicyRuleComposite or a SUPAECAPolicyRuleAtomic. SUPAECAPolicyRuleAtomic and SUPAECAPolicyRuleComposite are defined in sections 6.5 and 6.6, respectively. This is shown in Figure 25.

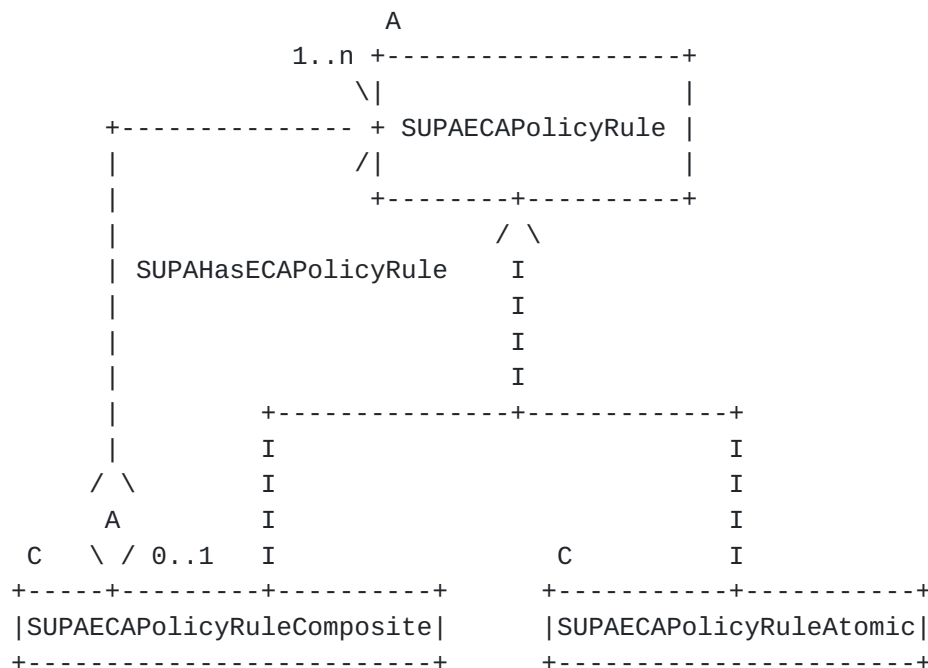


Figure 25. The Composite Pattern Applied to a SUPAECAPolicyRule

Note that the HasSUPAECAPolicyRule aggregation is defined by the

HasSUPAECAPolicyRuleDetail association class; both are defined in sections [6.6.2](#) and [6.6.3](#), respectively.

Strassner, et al.

Expires January 19, 2017

[Page 96]

6.4.1. SUPAECAPolicyRule Attributes

Currently, the SUPAECAPolicyRule defines two attributes, as described in the following subsections.

6.4.1.1. The Attribute "supaECAPolicyRulePriority"

This is a mandatory non-negative integer attribute that defines the priority of this particular SUPAECAPolicyRule. A larger value indicates a higher priority. A default value of 0 MAY be assigned.

Priority is used primarily for 2 reasons: (1) to resolve conflicts among policy actions (e.g., given a set of conflicting actions, which one will execute) and (2) to define the execution order of policy actions (e.g., when one action may depend on the output of one or more previous actions).

6.4.1.2. The Attribute "supaECAPolicyRuleStatus"

This is an optional non-negative enumerated integer whose value defines the current status of this policy rule. Values include:

- 0: error
- 1: init
- 2: In development, not ready to be deployed
- 1: Ready to be deployed
- 2: Deployed but not enabled
- 3: Deployed and enabled, but not executed
- 4: Executed without errors
- 5: Executed with errors
- 6: Aborted during execution

Note that 0 and 1 represent error and initialization states, respectively.

6.4.2. SUPAECAPolicyRule Relationships

Currently, the SUPAECAPolicyRule does not define any relationships. It inherits all four relationships defined by the SUPAPolicyStructure class (see [section 5.3.2.](#)).

6.5. The Concrete Class "SUPAECAPolicyRuleAtomic"

This is a mandatory concrete class. This class is a type of PolicyContainer, and represents a SUPAECAPolicyRule that can operate as a single, stand-alone, manageable object. Put another way, a SUPAECAPolicyRuleAtomic object can NOT be modeled as a set of hierarchical SUPAECAPolicyRule objects; if this is required, then a SUPAECAPolicyRuleComposite object should be used instead.

6.5.1. SUPAECAPolicyRuleAtomic Attributes

Currently, the SUPAECAPolicyRuleAtomic class does not define any attributes.

6.5.2. SUPAECAPolicyRuleAtomic Relationships

Currently, the SUPAECAPolicyRuleAtomic class does not define any relationships. It inherits all four relationships defined by the SUPAPolicyStructure class (see [section 5.3.2.](#)).

6.6. The Concrete Class "SUPAECAPolicyRuleComposite"

This is a mandatory concrete class. This class is a type of PolicyContainer, and represents a SUPAECAPolicyRule as a hierarchy of SUPAPolicy objects, where the hierarchy contains instances of a SUPAECAPolicyRuleAtomic and/or SUPAECAPolicyRuleComposite objects. Each of the SUPAPolicy objects, including the outermost SUPAECAPolicyRuleComposite object, are separately manageable. More importantly, each SUPAECAPolicyRuleComposite object represents an aggregated object that is itself manageable.

6.6.1. SUPAECAPolicyRuleComposite Attributes

Currently, the SUPAECAPolicyRuleComposite defines one attribute, as described in the following subsection.

6.6.1.1. The Attribute "supaECAEvalStrategy"

This is a mandatory, non-zero, integer attribute that enumerates a set of allowable alternatives that define how the set of SUPAECAPolicyRule object instances in a SUPAECAPolicyRuleComposite object are evaluated. It is assumed that the event and condition clauses of the SUPAECAPolicyRules have evaluated to TRUE (e.g., the event has occurred and the conditions were met). Values include:

- 0: error
- 1: init
- 2: execute the first SUPAECAPolicyRule in the SUPAECAPolicyRuleComposite and then terminate
- 3: execute only the highest priority SUPAECAPolicyRule(s) in the SUPAECAPolicyRuleComposite and then terminate
- 4: execute all SUPAECAPolicyRules in prioritized order (if any) regardless of whether their SUPAPolicyActions succeed or fail
- 5: execute all SUPAECAPolicyRules in prioritized order (if any) until at least one SUPAPolicyAction in a SUPAECAPolicyRule fails, and then terminate

Note that 0 and 1 represent error and initialization states,
respectively.

Strassner, et al.

Expires January 19, 2017

[Page 98]

If the value of `supaECAEvalStrategy` is 3 or 4, then all `SUPAECAPolicyRule`s that have a priority will be executed first (starting with the `SUPAECAPolicyRule(s)` that have the highest priority, and descending); all `SUPAECAPolicyRule(s)` that do not have a priority are then executed (in any order).

Assume that the actions in a given `SUPAECAPolicyRuleComposite` are defined as follows

```
SUPAECAPolicyRule A, priority 0
SUPAECAPolicyRule B, priority 10
SUPAECAPolicyRule C, priority 5
SUPAECAPolicyRule D, priority 10
SUPAECAPolicyRule E, priority 2
```

Then, if the `supaECAEvalStrategy` attribute value equals:

- 0: an error is issued
- 1: this `SUPAECAPolicyRule` MUST NOT be used (since it is not yet properly initialized)
- 2: only `SUPAECAPolicyRule A` is executed
- 3: only `SUPAECAPolicyRules B` and `D` are executed
- 4: all `SUPAECAPolicyRules` are executed, regardless of any failures in their `SUPAPolicyActions`
- 5: all `SUPAECAPolicyRules` are executed until a failure is detected, and then execution for all `SUPAECAPolicyRules` terminates

6.6.2. SUPAECAPolicyRuleComposite Relationships

Currently, the `SUPAECAPolicyRuleComposite` defines a single aggregation between it and `SUPAECAPolicyRule`, as described below.

6.6.2.1. The Aggregation "SUPAHasECAPolicyRule"

This is an optional aggregation that implements the composite pattern. The multiplicity of this aggregation is 0..1 on the aggregate (`SUPAECAPolicyRuleComposite`) side and 1..n on the part (`SUPAECAPolicyRule`) side. This means that if this aggregation is defined, then at least one `SUPAECAPolicyRule` object (which may be either an instance of a `SUPAECAPolicyRuleAtomic` or a `SUPAECAPolicyRuleComposite` class) must also be instantiated and aggregated by this particular `SUPAECAPolicyRuleComposite` object. The semantics of this aggregation are defined by the `SUPAHasECAPolicyRuleDetail` association class.

6.6.3. The Association Class "SUPAHasECAPolicyRuleDetail"

This is an optional concrete association class, and defines the semantics of the SUPAHasECAPolicyRule aggregation. This enables the attributes and relationships of the SUPAHasECAPolicyRuleDetail class to be used to constrain which SUPAHasECAPolicyRule objects can be aggregated by this particular SUPAECAPolicyRuleComposite object instance.

6.6.3.1. The Attribute "supaECAPolicyIsDefault"

This is an optional Boolean attribute. If the value of this attribute is true, then this SUPAECAPolicyRule is a default policy, and will be executed if no other SUPAECAPolicyRule in the SUPAECAPolicyRuleComposite container has been executed. This is a convenient way for error handling, though care should be taken to ensure that only one default policy rule is defined per SUPAECAPolicyRuleComposite container.

6.7. The Abstract Class "SUPABooleanClause"

A SUPABooleanClause specializes a SUPAPolicyClause, and defines a Boolean expression consisting of a standard structure in the form of a SUPAPolicyVariable, a SUPAPolicyOperator, and a SUPAPolicyValue. For example, this enables the following Boolean clause to be defined:

Foo >= Baz

where 'Foo' is a PolicyVariable, '>=' is a PolicyOperator, and 'Baz' is a PolicyValue.

Note that in this approach, the SUPAPolicyVariable and SUPAPolicyValue terms are defined as an appropriate subclass of the SUPAPolicyComponentDecorator class; it is assumed that the SUPAPolicyOperator is an instance of the SUPAPolicyOperator class.

This enables the EPRIM, in conjunction with the GPIM, to be used as a reusable class library. This encourages interoperability, since each element of the clause is itself an object defined by the SUPA object hierarchy.

An entire SUPABooleanClause may be negated by setting the supaBoolClauseIsNegated class attribute of the SUPABooleanClause class to TRUE. Individual terms of a Boolean clause can be negated by using the supaTermIsNegated Boolean attribute in the SUPAPolicyTerm class (see [section 5.10](#)).

A PolicyClause is in Conjunctive Normal Form (CNF) if it is a

sequence of logically ANDed terms, where each term is a sequence of logically ORed terms.

Strassner, et al.

Expires January 19, 2017

[Page 100]

A PolicyClause is in Disjunctive Normal Form (DNF) if it is a sequence of logically ORed terms, where each term is a sequence of logically ANDed terms.

The construction of more complex clauses, which consist of a set of simple clauses in CNF or DNF (as shown in the above example), is provided by using the composite pattern [\[3\]](#) to construct two concrete subclasses of the abstract SUPABooleanClause class. These are called SUPABooleanClauseAtomic and SUPABooleanClauseComposite, and are defined in sections [6.8](#) and [6.9](#), respectively. This enables instances of either a SUPABooleanClauseAtomic and/or a SUPABooleanClauseComposite to be aggregated into a SUPABooleanClauseComposite object.

[6.7.1.](#) SUPABooleanClause Attributes

The SUPABooleanClause class currently defines one attribute, which are defined in the following subsections.

[6.7.1.1.](#) The Attribute "supaBoolClauseIsNegated"

This is a mandatory Boolean attribute. If the value of this attribute is TRUE, then this (entire) SUPABooleanClause is negated. Note that the supaPolTermIsNegated class attribute of the SUPAPolicyTerm class is used to negate a single term.

[6.7.2.](#) SUPABooleanClause Relationships

Currently, no relationships are defined for the SUPABooleanClause class. It inherits the relationships of SUPAPolicyClause (see [section 5.5.](#)).

[6.8.](#) The Concrete Class "SUPABooleanClauseAtomic"

This is a mandatory concrete class that represents a SUPABooleanClause that can operate as a single, stand-alone, manageable object. A SUPABooleanClauseAtomic object can NOT be modeled as a set of hierarchical clauses; if this functionality is required, then a SUPABooleanClauseComposite object must be used. Examples of Boolean clauses that could be contained in a SUPABooleanClauseAtomic include P, NOT P, and (P OR Q), where P and Q are literals (e.g., a variable name that can be either true or false, or a formula that evaluates to a literal). Examples of Boolean clauses that are NOT in CNF are NOT(P AND Q), (P AND Q) OR R, and P AND (Q OR (R AND S)); their CNF equivalent forms are NOT P AND NOT Q, (P AND R) OR (Q AND R), and P AND (Q OR S) AND (Q OR S), respectively.

6.8.1. SUPABooleanClauseAtomic Attributes

No attributes are currently defined for the SUPABooleanClauseAtomic class.

6.8.2. SUPABooleanClauseAtomic Relationships

Currently, no relationships are defined for the SUPABooleanClauseAtomic class. It inherits the relationships of SUPAPolicyClause (see [section 5.5.](#)).

6.9. The Concrete Class "SUPABooleanClauseComposite"

This is a mandatory concrete class that represents a SUPABooleanClause that can operate as a hierarchy of PolicyClause objects, where the hierarchy contains instances of SUPABooleanClauseAtomic and/or SUPABooleanClauseComposite objects. Each of the SUPABooleanClauseAtomic and SUPABooleanClauseComposite objects, including the outermost SUPABooleanClauseComposite object, are separately manageable.

More importantly, each SUPAECAPolicyRuleComposite object represents an aggregated object that is itself manageable. Examples of Boolean clauses that could be contained in a SUPABooleanClauseAtomic include $((P \text{ OR } Q) \text{ AND } R)$, and $((\text{NOT } P \text{ OR } Q) \text{ AND } (R \text{ OR } \text{NOT } S) \text{ AND } T)$, where P, Q, R, S, and T are literals.

6.9.1. SUPABooleanClauseComposite Attributes

Two attributes are currently defined for the SUPABooleanClauseComposite class, and are described in the following subsections.

6.9.1.1. The Attribute "supaBoolClauseBindValue"

This is a mandatory non-zero integer attribute, and defines the order in which terms bind to a clause. For example, the Boolean expression $((A \text{ AND } B) \text{ OR } (C \text{ AND } \text{NOT } (D \text{ OR } E)))$ has the following binding order: terms A and B have a bind value of 1; term C has a binding value of 2, and terms D and E have a binding value of 3.

6.9.1.2. The Attribute "supaBoolClauseIsCNF"

This is an optional Boolean attribute. If the value of this attribute is TRUE, then this SUPABooleanClauseComposite is in CNF form. Otherwise, it is in DNF form.

6.9.2. SUPABooleanClauseComposite Relationships

Currently, the SUPABooleanClauseComposite class defined a single aggregation, which is described in the subsections below.

Strassner, et al.

Expires January 19, 2017

[Page 102]

6.9.2.1. The Aggregation "SUPAHasBooleanClause"

This is a mandatory aggregation that defines the set of SUPABooleanClause objects that are aggregated by this SUPABooleanClauseComposite object.

The multiplicity of this relationship is 0..1 on the aggregate (SUPABooleanClauseComposite) side, and 1..n on the part (SUPABooleanClause) side. This means that one or more SUPABooleanClauses are aggregated and used to define this SUPABooleanClauseComposite object. The 0..1 cardinality on the SUPABooleanClauseComposite side is necessary to enable SUPABooleanClauses to exist (e.g., in a PolicyRepository) before they are used by a SUPABooleanClauseComposite. The semantics of this aggregation is defined by the SUPAHasBooleanClauseDetail association class.

6.9.3. The Association Class "SUPAHasBooleanClauseDetail"

This is a mandatory concrete association class that defines the semantics of the SUPAHasBooleanClause aggregation. This enables the attributes and relationships of the SUPAHasBooleanClauseDetail class to be used to constrain which SUPABooleanClause objects can be aggregated by this particular SUPABooleanClauseComposite object instance.

6.9.3.1. SUPAHasBooleanClauseDetail Attributes

The SUPAHasBooleanClauseDetail class currently does not define any attributes.

6.10. The Abstract Class "SUPAECAComponent"

This is a mandatory abstract class that defines three concrete subclasses, one each to represent the concepts of reusable events, conditions, and actions. They are called SUPAPolicyEvent, SUPAPolicyCondition, and SUPAPolicyAction, respectively.

SUPAECAComponents provide two different ways to construct SUPAPolicyClauses. The first is for the SUPAECAComponent to be used as either a SUPAPolicyVariable or a SUPAPolicyValue, and the second is for the SUPAECAComponent to contain the entire clause text.

For example, suppose it is desired to define a policy condition clause with the text 'queueDepth > 10'. The two approaches could satisfy this as follows:

Approach #1 (canonical form):

```
SUPAPolicyCondition.supapolicyConditionData contains the text
'queueDepth'
SUPAPolicyOperator.supapolicyOpType is set to '1' (greater than)
SUPAPolicyValue.supapolicyValContent is set to '10'
```

Approach #2 (SUPAECAComponent represents the entire clause):

```
SUPAPolicyCondition.supapolicyConditionData contains the text
'queueDepth > 10'
```

The class attribute `supaECACompIsTerm`, defined in [subsection 6.10.1.1](#), is used to identify which of these two approaches is used by an object instance of this class.

[6.10.1](#). SUPAECAComponent Attributes

A single attribute is currently defined for this class, and is described in the following subsection.

[6.10.1.1](#). The Attribute `supaECACompIsTerm`

This is an optional Boolean attribute. If the value of this attribute is TRUE, then this SUPAECAComponent is used as the value of a SUPAPolicyTerm to construct a SUPAPolicyClause (this is approach #1 in [section 6.10](#) above). If the value of this attribute is FALSE, then this SUPAECAComponent contains the text of the entire corresponding SUPAPolicyClause (this is approach #2 in [section 6.10](#) above).

[6.10.2](#). SUPAECAComponent Relationships

No relationships are currently defined for this class.

[6.11](#). The Concrete Class "SUPAPolicyEvent"

This is a mandatory concrete class that represents the concept of an Event that is applicable to a policy management system. Such an Event is defined as anything of importance to the management system (e.g., a change in the system being managed and/or its environment) occurring on a time-axis (as defined in [\[19\]](#)). SUPAPolicyEvents can be used as part of a SUPAPolicyClause; this is done by specifying the attribute name and value of an Event in the `supaPolicyEventData` attribute of the SUPAPolicyEvent. This enables event attributes to be used as part of a SUPAPolicyClause.

Note: this class does NOT model the "raw" occurrences of Events. Rather, it represents the concept of an Event object whose class attributes describe pertinent attributes that can trigger the evaluation of a SUPAECAPolicyRule.

6.11.1. SUPAPolicyEvent Attributes

Currently, five attributes are defined for the SUPAPolicyEvent class, which are described in the following subsections.

6.11.1.1. The Attribute "supaPolicyEventIsPreProcessed"

This is an optional Boolean attribute. If the value of this attribute is TRUE, then this SUPAPolicyEvent has been pre-processed by an external entity, such as an Event Service Bus, before it was received by the Policy Management System.

6.11.1.2. The Attribute "supaPolicyEventIsSynthetic"

This is an optional Boolean attribute. If the value of this attribute is TRUE, then this SUPAPolicyEvent has been produced by the Policy Management System. If the value of this attribute is FALSE, then this SUPAPolicyEvent has been produced by an entity in the system being managed.

6.11.1.3. The Attribute "supaPolicyEventTopic[0..n]"

This is a mandatory array of string attributes, and contains the subject that this PolicyEvent describes.

Note: [0..n] means that this is a multi-valued property that has zero or more attributes.

6.11.1.4. The Attribute "supaPolicyEventEncoding"

This is a mandatory non-zero enumerated integer attribute, and defines how to interpret the supaPolicyEventData class attribute. These two attributes form a tuple, and together enable a machine to understand the syntax and value of the data carried by the object instance of this class. Values include:

- 0: error
- 1: init
- 2: String
- 3: Integer
- 4: Boolean
- 5: Floating Point
- 6: DateTime

6.11.1.5. The Attribute "supaPolicyEventData[1..n]"

This is a mandatory attribute that defines an array of strings. Each string in the array represents an attribute name and value of an Event object. The format of each string is defined as name:value. The 'name' part is the name of the SUPAPolicyEvent

attribute, and the 'value' part is the value of that attribute.

Strassner, et al.

Expires January 19, 2017

[Page 105]

Note: [1..n] means that this is a multi-valued property that has at least one (and possibly more) attributes. For example, if this value of this attribute is:

```
{(startTime:0800), (endTime:1700), (date:2016-05-11),  
 (timeZone:-08:00)}
```

then this attribute contains four properties, called `startTime`, `endTime`, `date`, and `timeZone` whose values are 0800, 1700, May 11 2016, and Pacific Standard Time, respectively.

Note that the `supaPolicyEventEncoding` class attribute defines how to interpret the value portion of this attribute.

This attribute works with another class attribute, called `supaPolicyEventEncoding`, which defines how to interpret this attribute. These two attributes form a tuple, and together enable a machine to understand the syntax and value of the data carried by the object instance of this class.

6.11.2. SUPAPolicyEvent Relationships

No relationships are currently defined for this class. It inherits the relationships defined by the `SUPAPolicyComponentDecorator` (see [section 5.7.3.](#)).

6.12. The Concrete Class "SUPAPolicyCondition"

This is a mandatory concrete class that represents the concept of an Condition that will determine whether or not the set of Actions in the `SUPAECAPolicyRule` to which it belongs are executed or not. `SUPAPolicyConditions` can be used as part of a `SUPAPolicyClause` (e.g., `var = SUPAPolicyCondition.supapolicyconditiondata`) or as a stand-alone `SUPAPolicyClause` (e.g., the `supapolicyconditiondata` attribute contains text that defines the entire condition clause).

6.12.1. SUPAPolicyCondition Attributes

Currently, two attributes are defined for the `SUPAPolicyCondition` class, which are described in the following subsections.

6.12.1.1. The Attribute "supapolicyconditiondata[1..n]"

This is a mandatory array of string attributes that contains the content of this `SUPAPolicyCondition` object.

Note: [1..n] means that this is a multi-valued property that has at least one (and possibly more) attributes.

This attribute works with another class attribute, called `supaPolicyConditionEncoding`, which defines how to interpret this attribute. These two attributes form a tuple, and together enable a machine to understand the syntax and value of the data carried by the object instance of this class.

6.12.1.2. The Attribute "supaPolicyConditionEncoding"

This is a mandatory non-zero enumerated integer attribute, and defines the data type of the `supaPolicyConditionData` attribute. These two attributes form a tuple, and together enable a machine to understand the syntax and value of the content of this `SUPAPolicyCondition` object. Values include:

- 0: error
- 1: init
- 2: OCL 2.4
- 3: OCL 2.x
- 4: OCL 1.x
- 5: QVT 1.2 - Relations Language
- 6: QVT 1.2 - Operational language
- 7: Alloy
- 8: ASCII Text

Enumerations 1 and 2 signify an error state and an initialization state, respectively. Enumerations 3-5 are dedicated to OCL (with OCL 2.4 being the latest version as of this writing). QVT defines a set of languages [\[20\]](#) (the two most powerful and useful are defined by enumerations 5 and 6). Alloy is a language for describing constraints, and uses a SAT solver to guarantee correctness [\[21\]](#). Enumeration 8 (ASCII Text) is not recommended (since it is informal, and hence, not verifiable), but is included for completeness.

If this class is instantiated, then this attribute SHOULD also be instantiated, and SHOULD be part of a conformant implementation.

6.12.2. SUPAPolicyEvent Relationships

No relationships are currently defined for this class. It inherits the relationships defined by the `SUPAPolicyComponentDecorator` (see [section 5.7.3.](#)).

6.13. The Concrete Class "SUPAPolicyAction"

This is a mandatory concrete class that represents the concept of an Action, which is a part of a `SUPAECAPolicyRule`, which may be executed when both the event and the condition clauses of its owning `SUPAECAPolicyRule` evaluate to true. The execution of this action is determined by its `SUPAECAPolicyRule` container, and any applicable `SUPAPolicyMetadata` objects. `SUPAPolicyActions` can be

used in three different ways:

Strassner, et al.

Expires January 19, 2017

[Page 107]

- o as part of a SUPAPolicyClause (e.g., var = SUPAPolicyAction.supapolicyActionData)
- o as a stand-alone SUPAPolicyClause (e.g., the supapolicyActionData attribute contains text that defines the entire action clause)
- o to invoke a SUPAECAPolicyRule

In the third case, the execution semantics SHOULD be to suspend the current execution of the set of SUPAPolicyActions that are executing, transfer execution control to the invoked SUPAECAPolicyRule, and resume the execution of the original set of SUPAPolicyActions when the invoked SUPAECAPolicyRule has finished execution.

6.13.1. SUPAPolicyAction Attributes

Currently, two attributes are defined for the SUPAPolicyCondition class, which are described in the following subsections.

6.13.1.1. The Attribute "supapolicyActionData[1..n]"

This is a mandatory array of string attributes that contains the content of this SUPAPolicyAction object. This attribute works with another class attribute, called supapolicyActionEncoding, which defines how to interpret this attribute. These two attributes form a tuple, and together enable a machine to understand the syntax and value of the data carried by the object instance of this class. Note: [1..n] means that this is a multi-valued property that has at least one (and possibly more) attributes.

Since this attribute could represent a term in a SUPAPolicyClause (e.g., var = SUPAPolicyAction.supapolicyActionData), a complete SUPAPolicyClause (e.g., the supapolicyActionData attribute contains text that defines the entire action clause), or the name of a SUPAECAPolicyRule to invoke, each element in the string array is prepended with one of the following strings:

- o 't:' (or 'term:'), to denote a term in a SUPAPolicyClause
- o 'c:' (or 'clause:'), to denote an entire SUPAPolicyClause
- o 'r:' (or 'rule:'), to invoke a SUPAECAPolicyRule

6.13.1.2. The Attribute "supapolicyActionEncoding"

This is a mandatory non-zero enumerated integer attribute, and defines the data type of the supapolicyActionData attribute. This attribute works with another class attribute, called supapolicyActionData, which contains the content of the action. These two attributes form a tuple, and together enable a machine to understand the syntax and value of the content of this SUPAPolicyAction object. Values include:

0: error
1: init
2: OCL 2.4
3: OCL 2.x
4: OCL 1.x
5: QVT 1.2 - Relations Language
6: QVT 1.2 - Operational language
7: Alloy
8: ASCII Text
9: GUID
10: UUID
11: URI
12: FQDN

Enumerations 1 and 2 signify an error state and an initialization state, respectively. Enumerations 3-5 are dedicated to OCL (with OCL 2.4 being the latest version as of this writing). QVT defines a set of languages [\[20\]](#) (the two most powerful and useful are defined by enumerations 5 and 6). Alloy is a language for describing constraints, and uses a SAT solver to guarantee correctness [\[21\]](#). Enumeration 8 (ASCII Text) is not recommended (since it is informal, and hence, not verifiable), but is included for completeness. Enumerations 9-12 define a reference to the SUPAPolicyAction.

If this class is instantiated, then this attribute SHOULD also be instantiated, and SHOULD be part of a conformant implementation.

[6.13.2.](#) SUPAPolicyAction Relationships

No relationships are currently defined for this class. It inherits the relationships defined by the SUPAPolicyComponentDecorator (see [section 5.7.3.](#)).

Enumerations 1-4 are used to provide a reference to an action object. Enumerations 5-10 are used to express the action to perform as a string.

[7.](#) Examples

This will be defined in the next version of this document.

[8.](#) Security Considerations

This will be defined in the next version of this document.

9. IANA Considerations

This document has no actions for IANA.

10. Contributors

The following people contributed to creating this document, and are listed in alphabetical order:

Jason Coleman

11. Acknowledgments

This document has benefited from reviews, suggestions, comments and proposed text provided by the following members, listed in alphabetical order: Andy Bierman, Bert Wijnen, Bob Natale, Dave Hood, Fred Feisullin, Georgios Karagiannis, Liu (Will) Shucheng, Marie-Jose Montpetit.

12. References

This section defines normative and informative references for this document.

12.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

12.2. Informative References

[RFC3060] Moore, B., Ellessen, E., Strassner, J., Westerinen, A., "Policy Core Information Model -- Version 1 Specification", [RFC 3060](#), February 2001

[RFC3198] Westerinen, A., Schnizlein, J., Strassner, J., Scherling, M., Quinn, B., Herzog, S., Huynh, A., Carlson, M., Perry, J., Waldbusser, S., "Terminology for Policy-Based Management", [RFC 3198](#), November, 2001

- [RFC3460] Moore, B., ed., "Policy Core Information Model (PCIM) Extensions, [RFC 3460](#), January 2003
- [1] Strassner, J., "Policy-Based Network Management", Morgan Kaufman, ISBN 978-1558608597, Sep 2003
- [2] Strassner, J., ed., "The DEN-ng Information Model", add stable URI
- [3] Riehle, D., "Composite Design Patterns", Proceedings of the 1997 Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA '97). ACM Press, 1997, Page 218-228
- [4] DMTF, CIM Schema, v2.44, http://dmtf.org/standards/cim/cim_schema_v2440
- [5] Strassner, J., ed., "ZOOM Policy Architecture and Information Model Snapshot", TR235, part of the TM Forum ZOOM project, October 26, 2014
- [6] TM Forum, "Information Framework (SID), GB922 and associated Addenda, v14.5, <https://www.tmforum.org/information-framework-sid/>
- [7] Liskov, B.H., Wing, J.M., "A Behavioral Notion of subtyping", ACM Transactions on Programming languages and Systems 16 (6): 1811 - 1841, 1994
- [8] Klyus, M., Strassner, J., Liu, W., Karagiannis, G., Bi, J., "SUPA Value Proposition", [draft-klyus-supa-value-proposition-00](#), March 21, 2016
- [9] ISO/IEC 10746-3 (also ITU-T Rec X.903), "Reference Model Open Distributed Processing Architecture", April 20, 2010
- [10] Davy, S., Jennings, B., Strassner, J., "The Policy Continuum - A Formal Model", Proc. of the 2nd Intl. IEEE Workshop on Modeling Autonomic Communication Environments (MACE), Multicon Lecture Notes, No. 6, Multicon, Berlin, 2007, pages 65-78
- [11] Gamma, E., Helm, R., Johnson, R., Vlissides, J., "Design Patterns - Elements of Reusable Object-Oriented Software", Addison-Wesley, 1994, ISBN 0-201-63361-2

- [12] Strassner, J., de Souza, J.N., Raymer, D., Samudrala, S., Davy, S., Barrett, K., "The Design of a Novel Context-Aware Policy Model to Support Machine-Based Learning and Reasoning", Journal of Cluster Computing, Vol 12, Issue 1, pages 17-43, March, 2009
- [13] Liskov, B.H., Wing, J.M., "A Behavioral Notion of subtyping", ACM Transactions on Programming languages and Systems, 16 (6): 1811 - 1841, 1994
- [14] Martin, R.C., "Agile Software Development, Principles, Patterns, and Practices", Prentice-Hall, 2002, ISBN: 0-13-597444-5
- [15] Halpern, J., Strassner, J., "Generic Policy Data Model for Simplified Use of Policy Abstractions (SUPA)" [draft-ietf-supa-generic-policy-data-model-02](#), July 13, 2016
- [16] Wang, Y., Esposito, F., Matta, I., Day, J., "RINA: An Architecture for Policy-based Dynamic Service Management", Tech Report BUCS-TR-2013-014, 2013
- [17] Meyer, B., "Object-Oriented Software Construction", Prentice Hall, second edition, 1997 ISBN 0-13-629155-4
- [18] <http://semver.org/>
- [19] ISO/IEC:2004(E), "Data elements and interchange formats -- Information interchange -- Representation of dates and times", 2004
- [20] <http://www.omg.org/spec/QVT/>
- [21] <http://alloy.mit.edu/alloy/>

Authors' Addresses

John Strassner
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95138 USA
Email: john.sc.strassner@huawei.com

Joel Halpern
Ericsson
P. O. Box 6049
Leesburg, VA 20178
Email: joel.halpern@ericsson.com

Sven van der Meer
LM Ericsson Ltd.
Ericsson Software Campus
Garrycastle
Athlone
N37 PV44
Ireland
Email: sven.van.der.meer@ericsson.com

[Appendix A](#). Brief Analyses of Previous Policy Work

This appendix describes of some of the important problems with previous IETF policy work., and describes the rationale for taking different design decisions in this document.

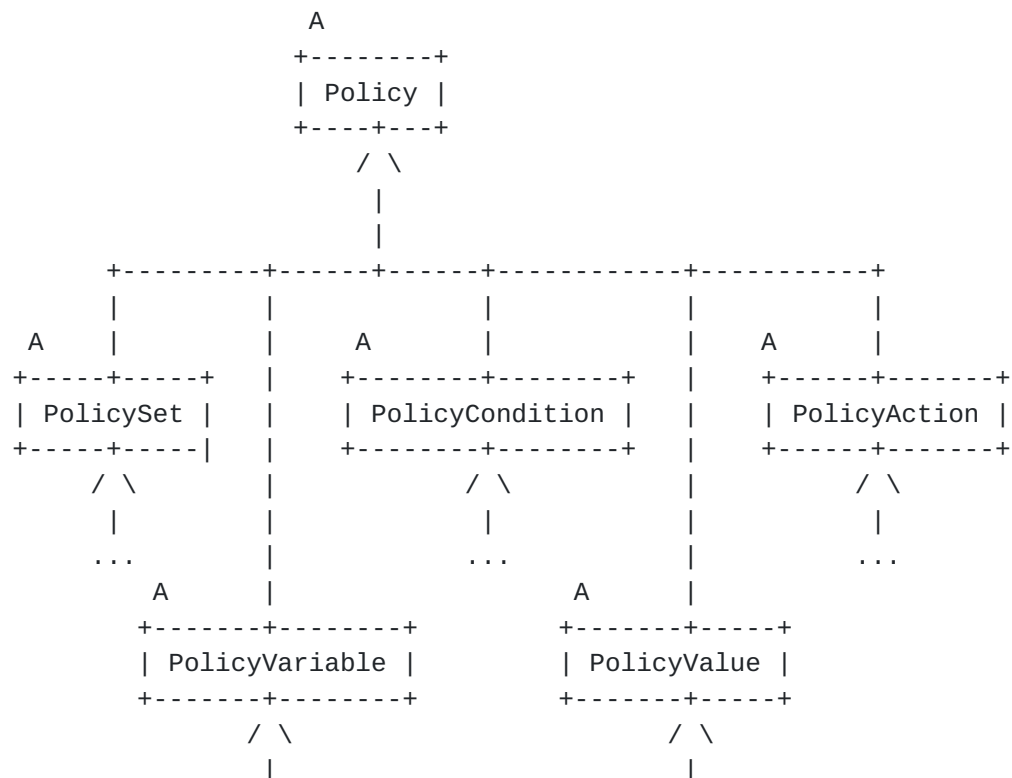
[A.1](#). PolicySetComponent vs. SUPAPolicyStructure

The ability to define different types of policy rules is not present in [\[RFC3060\]](#) and [\[RFC3460\]](#), because both are based on [\[4\]](#), and this ability is not present in [\[4\]](#). [\[RFC3060\]](#), [\[RFC3460\]](#), and [\[4\]](#) are all limited to CA (condition-action) policy rules. In addition, events are NOT defined. These limitations mean that [RFC3060](#), [RFC3460](#), and [\[4\]](#) can only represent CA Policy Rules.

In contrast, the original design goal of SUPA was to define a single class hierarchy that could represent different types of policies (e.g., imperative and declarative). Hence, it was decided to make SUPAPolicyStructure generic in nature, so that different types of policies could be defined as subclasses. This enables a single Policy Framework to support multiple types of policies.

[A.2](#). Flat Hierarchy vs. SUPAPolicyComponentStructure

Figure 26 shows a portion of the class hierarchy of [\[RFC3460\]](#).



...

...

Figure 26. Simplified Class Hierarchy of [[RFC3460](#)]

Strassner, et al.

Expires January 19, 2017

[Page 114]

[RFC3060](#)], [[RFC3460](#)], and [[4](#)] defined PolicyConditions and PolicyActions as subclasses of Policy (along with PolicySet, which is the superclass of PolicyRules and PolicyGroups). This means that there is no commonality between PolicyConditions and PolicyActions, even though they are both PolicyRule components. From an object-oriented point-of-view, this is incorrect, since a PolicyRule aggregates both PolicyConditions and PolicyActions.

In addition, note that both PolicyVariables and PolicyValues are siblings of PolicyRules, PolicyConditions, and PolicyActions. This is incorrect for several reasons:

- o a PolicyRule cannot rectly contain PolicyVariables or PolicyValues, so they shouldn't be at the same level of the class hierarchy
- o both PolicyConditions and PolicyActions can contain PolicyVariables and PolicyValues, which implies that both PolicyVariables and PolicyValues should be lower in the class hierarchy

Note that in the current version of [[4](#)], PolicyVariable and PolicyValue are both deleted. There are other changes as well, but they are beyond the scope of this Appendix.

The original design goal of SUPA was to define a single class hierarchy that could represent different types of policies and policy components. This cannot be accomplished in [[RFC3460](#)], since there is no notion of a policy component (or alternatively, PolicyCondition, PolicyAction, PolicyVariable, and PolicyValue are all components at the same abstraction level, which is clearly not correct). Hence, SUPA defined the SUPAPolicyComponentStructure class to capture the concept of a reusable policy component.

In summary, SUPAPolicyStructure subclasses define the structure of a policy in a common way, while SUPAPolicyComponentStructure subclasses define the content that is contained in the structure of a policy, also in a common way.

[A.3.](#) PolicyRules and PolicyGroups vs. SUPAPolicyRules

A PolicySetComponent is an aggregation, implemented as an association class, that "collects instances of PolicySet subclasses into coherent sets of Policies". This is a recursive aggregation, with multiplicity 0..n - 0..n, on the PolicySet class.

Since this is a recursive aggregation, it means that a PolicySet can aggregate zero or more PolicySets. This is under-specified, and can be interpreted in one of two ways:

1. A PolicySet subclass can aggregate any PolicySet subclass (PolicyRules can aggregate PolicyRules and PolicyGroups, and vice-versa)
2. PolicyRules can aggregate PolicyRules, and PolicyGroups can aggregate PolicyGroups, but neither class can aggregate the other type of class

Both interpretations are ill-suited for policy-based management. The problem with the first is that if PolicyGroup is the mechanism for grouping, why can a PolicyRule aggregate a PolicyGroup? This implies that PolicyGroups are not needed. The problem with the second is that PolicyGroups cannot aggregate PolicyRules (which again implies that PolicyGroups are not needed).

Furthermore, there are no mechanisms defined in the [\[RFC3460\]](#) model to prevent loops of PolicyRules. This is a problem, because EVERY PolicyRule and PolicyGroup inherits this recursive aggregation.

This is why this document uses the composite pattern. First, this pattern clearly shows what object is aggregating what other object (i.e., a SUPAECAPolicyRuleAtomic cannot aggregate a SUPAECAPolicyRuleComposite). Second, it does not allow a SUPAECAPolicyRule to be aggregated by another SUPAECAPolicyRule (this is discussed more in the following subsection).

[A.3.1.](#) Sub-rules

Sub-rules (also called nested policy rules) enable a policy rule to be contained within another policy rule. These have very complex semantics, are very hard to debug, and provide limited value. They also require a complex set of aggregations (see section A.4.).

The main reason for defining sub-rules in [\[RFC3460\]](#) is to enable "complex policy rules to be constructed from multiple simpler policy rules". However, the composite pattern does this much more efficiently than a simple recursive aggregation, and avoids the ambiguous semantics of a recursive aggregation. This latter point is important, because if PolicyRule and/or PolicyGroup is subclassed, then all subclasses still inherit this recursive aggregation, along with its ambiguous semantics.

[A.4.](#) PolicyConditions and PolicyActions vs. SUPAECAComponent

There is no need to use the SimplePolicyCondition and ComplexPolicyCondition objects defined in [\[RFC3460\]](#), since the SUPAPolicyComponentStructure uses the decorator pattern (see [section 5.7](#)) to provide more extensible types of conditions than is possible with those classes. This also applies for the

SimplePolicyAction and the ComplexPolicyAction classes defined in [\[RFC3460\]](#).

Strassner, et al.

Expires January 19, 2017

[Page 116]

More importantly, this removes the need for a complex set of aggregations (i.e., PolicyComponent, PolicySetComponent, PolicyConditionStructure, PolicyConditionInPolicyRule, PolicyConditionInPolicyCondition, PolicyActionStructure, PolicyActionInPolicyRule, and PolicyActionInPolicyAction). Instead, ANY SUPAECAComponent is defined as a decorator (i.e., a subclass of SUPAPolicyComponentDecorator), and hence, Any SUPAECAComponent is wrapped onto a concrete subclass of SUPAPolicyClause using the SAME aggregation (SUPAHasDecoratedPolicyComponent). This is a significantly simpler design that is also more powerful.

A.5. The SUPAPolicyComponentDecorator Abstraction

One of the problems in building a policy model is the tendency to have a multitude of classes, and hence object instances, to represent different combinations of policy events, conditions, and actions. This can lead to class and/or relationship explosion, as is the case in [\[RFC3460\]](#), [\[4\]](#), and [\[6\]](#).

For example, [\[RFC3460\]](#) defines five subclasses of PolicyCondition: PolicyTimePeriodCondition, VendorPolicyCondition, SimplePolicyCondition, CompoundPolicyCondition, and CompoundFilterCondition. Of these:

- o PolicyTimePeriodCondition is a data structure, not a class
- o VendorPolicyCondition represents a condition using two attributes that represent a multi-valued octet string
- o SimplePolicyCondition, CompoundPolicyCondition, and CompoundFilterCondition all have ambiguous semantics

SimplePolicyCondition represents an ordered 3-tuple, in the form {variable, match, value}. However, the match operator is not formally modeled. Specifically, "the 'match' relationship is to be interpreted by analyzing the variable and value instances associated with the simple condition". This becomes problematic for several cases, such as shallow vs. deep object comparisons. More importantly, this requires two separate aggregations (PolicyVariableInSimplePolicyCondition and PolicyValueInSimplePolicyCondition) to associate variables and values to the SimplePolicyCondition, respectively. Since [\[RFC3460\]](#) defines all relationships as classes, this means that the expression "Foo > Bar" requires a total of FIVE objects (one each for the variable and value, one for the SimplePolicyCondition, and one each to associate the variable and value with the SimplePolicyCondition).

This is exacerbated when SimplePolicyConditions are used to build CompoundPolicyConditions. In addition to the above complexity (which is required for each SimplePolicyCondition), a new aggregation (PolicyConditionInPolicyCondition) is required to aggregation PolicyConditions. Thus, the compound expression: "((Foo > Bar) AND (Foo < Baz))" requires a total of THIRTEEN objects (five for each of the terms being ANDed, plus one for the CompoundPolicyCondition, and two to aggregate each term to the CompoundPolicyCondition).

Note that in the above examples, the superclasses of each of the relationships are omitted for clarity. In addition, [\[RFC3460\]](#) is built using inheritance; this means that if a new function is required, a new class must be built (e.g., CompoundFilterCondition is a subclass, but all it adds is one attribute).

In contrast, the Decorator Pattern enables behavior to be selectively added to an individual object, either statically or dynamically, without having to build association classes. In addition, the decorator pattern uses composition, instead of inheritance, to avoid class explosion. This means that a new variable, value, or even condition class can be defined at runtime, and then all or part of that class can dynamically wrap an existing object without need for recompilation and redeployment.

[A.6.](#) The Abstract Class "SUPAPolicyClause"

This abstraction is missing in [\[RFC3060\]](#), [\[RFC3460\]](#), [\[4\]](#), and [\[6\]](#). SUPAPolicyClause was abstracted from DEN-ng [\[2\]](#), and a version of this class is in the process of being added to [\[5\]](#). However, the class and relationship design in [\[5\]](#) differs significantly from the corresponding designs in this document.

SUPAPolicyClause further reinforces the difference between a policy rule and a component of a policy rule by abstracting the content of a policy rule as a reusable object. This is fundamental for enabling different types of policy rules (e.g., imperative and declarative) to be represented using the same constructs.

[A.7.](#) Problems with the [RFC3460](#) Version of PolicyVariable

The following subsections define a brief, and incomplete, set of problems with the implementation of [\[RFC3460\]](#) (note that [\[RFC3060\]](#) did not define variables, operators, and/or values).

[A.7.1.](#) Object Bloat

[\[RFC3460\]](#) used two different and complex mechanisms for providing generic get and set expressions. PolicyVariables were subclassed

into two subclasses, even though they performed the same semantic function. This causes additional problems:

Strassner, et al.

Expires January 19, 2017

[Page 118]

- o PolicyExplicitVariables are for CIM compatibility; note that the CIM does not contain either PolicyVariables or PolicyValues ([4])
- o PolicyImplicitVariable subclasses do not define attributes; rather, they are bound to an appropriate subclass of PolicyValue using an association

Hence, defining a variable is relatively expensive in [RFC3460], as in general, two objects and an association must be used. The objects themselves do not define content; rather, their names are used as a mechanism to identify an object to match. This means that an entire object must be used (instead of, for example, an attribute), which is wasteful. It also make it difficult to adjust constraints at runtime, since the constraint is defined in a class that is statically defined (and hence, requires recompilation and possibly redeployment if it is changed).

A.7.2. Object Explosion

The above three problems lead to class explosion (recall that in [RFC3060], [RFC3460], and [4], associations are implemented as classes).

In contrast to this approach, the approach in this document keeps the idea of the class hierarchy for backwards compatibility, but streamlines the implementation. Specifically:

1. The decorator pattern is an established and very used software pattern (it dates back to at least 1994 [11]).
2. The use of a single association class (i.e., SUPAHasDecoratedPolicyComponentDetail) can represent more constraints than is possible in the approaches of [RFC3460] and [4] in a much more flexible manner, due to its function as a decorator of other objects.
3. Note that there is no way to enforce the constraint matching in [RFC3460] and [6]; the burden is on the developer to check and see if the constraints specified in one class are honored in the other class.
4. If these constraints are not honored, there is no mechanism specified to define the clause as incorrectly formed.

A.7.3. Specification Ambiguities

There are a number of ambiguities in [RFC3460].

First, [RFC3460] says: "Variables are used for building individual conditions". While this is true, variables can also be used for building individual actions. This is reflected in the definition for SUPAPolicyVariable.

Second, [\[RFC3460\]](#) says: "The variable specifies the property of a flow or an event that should be matched when evaluating the condition." While this is true, variables can be used to test many other things than "just" a flow or an event. This is reflected in the SUPAPolicyVariable definition.

Third, the [\[RFC3460\]](#) definition requires the use of associations in order to properly constrain the variable (e.g., define its data type, the range of its allowed values, etc.). This is both costly and inefficient.

Fourth, [\[RFC3460\]](#) is tightly bound to the DMTF CIM schema [\[4\]](#). The CIM is a data model (despite its name), because:

- o It uses keys and weak relationships, which are both concepts from relational algebra and thus, not technology-independent
- o It has its own proprietary modeling language
- o It contains a number of concepts that are not defined in UML (including overriding keys for subclasses)

Fifth, the class hierarchy has two needless classes, called SUPAImplicitVariable and SUPAExplicitVariable. These classes do not define any attributes or relationships, and hence, do not add any semantics to the model.

Finally, in [\[RFC3460\]](#), defining constraints for a variable is limited to associating the variable with a PolicyValue. This is both cumbersome (because associations are costly; for example, they equate to a join in a relational database management system), and not scalable, because it is prone to proliferating PolicyValue classes for every constraint (or range of constraints) that is possible. Therefore, in SUPA, this mechanism is replaced with using an association to an association class that defines constraints in a much more general and powerful manner (i.e., the SUPAHasDecoratedPolicyComponentDetail class).

[A.8.](#) Problems with the [RFC3460](#) Version of PolicyValue

The following subsections define a brief, and incomplete, set of problems with the implementation of [\[RFC3460\]](#) (note that [\[RFC3060\]](#) did not define variables, operators, and/or values).

[A.8.1.](#) Object Bloat

[\[RFC3460\]](#) defined a set of 7 subclasses; three were specific to networking (i.e., IPv4 Address, IPv6 Address, MAC Address) and 4 (PolicyStringValue, PolicyBitStringValue, PolicyIntegerValue, and PolicyBooleanValue) were generic in nature. However, each of these objects defined a single class attribute. This has the same two

problems as with PolicyVariables (see [section 5.9.1.1](#)):

Strassner, et al.

Expires January 19, 2017

[Page 120]

1. Using an entire object to define a single attribute is very wasteful and expensive
2. It also make it difficult to adjust constraints at runtime, since the constraint is defined in a class that is statically defined (and hence, requires recompilation and possibly redeployment if it is changed).

[A.8.2.](#) Object Explosion

[RFC3460] definition requires the use of associations in order to properly constrain the variable (e.g., define its data type, the range of its allowed values, etc.). This is both costly and inefficient (recall that in [\[RFC3060\]](#), [\[RFC3460\]](#), and [\[4\]](#), associations are implemented as classes).

[A.8.3.](#) Lack of Constraints

There is no generic facility for defining constraints for a PolicyValue. Therefore, there is no facility for being able to change such constraints dynamically at runtime.

[A.8.4.](#) Tightly Bound to the CIM Schema

[RFC3460] is tightly bound to the DMTF CIM schema [\[4\]](#). The CIM is a data model (despite its name), because:

- o It uses keys and weak relationships, which are both concepts from relational algebra and thus, not technology-independent
- o It has its own proprietary modeling language
- o It contains a number of concepts that are not defined in UML (including overriding keys for subclasses)

[A.8.5.](#) Specification Ambiguity

[RFC3460] says: It is used for defining values and constants used in policy conditions". While this is true, variables can also be used for building individual actions. This is reflected in the SUPAPolicyVariable definition.

[A.8.6.](#) Lack of Symmetry

Most good information models show symmetry between like components. [\[RFC3460\]](#) has no symmetry in how it defines variables and values. In contrast, this document recognizes that variables and values are just terms in a clause; hence, the only difference in the definition of the SUPAPolicyVariable and SUPAPolicyValue classes is that the content attribute in the former is a single string, whereas the content attribute in the latter is a string array. In particular, the semantics of both variables and values are defined using the decorator pattern, along with the attributes of

the SUPAPolicyComponentDecorator and the
SUPAHasDecoratedPolicyComponentDetail classes.

Strassner, et al.

Expires January 19, 2017

[Page 121]