

Drools Usage Manual

version 2.0-beta-12
generated January 3, 2004

The Drools Project

drools.org

Preface

Why rules?

Rules are important...

Contents

Preface	i
Why rules?	i
List of Figures	vii
Acknowledgements	ix
About this manual	x
Typographic conventions	x
Production notes	x
About the authors	xi
Bob McWhirter	xi
Robert Searle	xi
 Part I Drools—How and why	 1
1 50,000ft view	2
2 Work Flow Engine	3
 Part II Drools	 4
3 Functional Overview	5
3.1 Rules	5
3.1.1 Declarative Form	5
3.1.2 Applicable Context	6
3.2 Knowledge	6
3.3 Why a rule-engine?	6
3.4 A note about “business rules”	7

4	Drools Rule Language	8
4.1	Introduction	8
4.2	Core DRL namespace	8
4.3	Core DRL tags	9
4.3.1	<rule-set>	9
4.3.2	<rule>	9
4.3.3	<parameter>	10
4.3.4	<declaration>	10
4.3.5	<duration>	11
5	Client API	12
5.1	Introduction	12
5.2	Retrieving a RuleBase	12
5.2.1	Looking up a RuleBase from JNDI	12
5.2.2	Deserializing a RuleBase from a file	13
5.3	Creating a WorkingMemory	13
5.4	Working with knowledge	13
5.4.1	Adding knowledge	13
5.4.2	Removing knowledge	14
5.4.3	Modifying knowledge	14
5.5	Firing rules	14
6	Administrative API	16
6.1	Introduction	16
6.2	Building rules	16
6.3	Building a RuleBase	16
6.3.1	Loading rules	16
6.3.2	Conflict resolution	17
6.4	Deploying a RuleBase to a JNDI data store	17
6.4.1	JNDI deployment utility	18
6.5	Serializing a RuleBase to a file	18
6.5.1	Serialization utility	19
7	Conflict Resolution	20
7.1	Introduction	20
7.2	SalienceConflictResolutionStrategy	20
7.3	ComplexityConflictResolutionStrategy	21
7.4	SimplicityConflictResolutionStrategy	21
	Part III Semantic Modules	22
8	Java Semantic Module	23
8.1	Introduction	23
8.2	Java semantic module namespace URI	23
8.3	Java Semantic Module tags	24

8.3.1	<class>	24
8.3.2	<condition>	24
8.3.3	<extractor>	24
8.3.4	<consequence>	24
9	Python Semantic Module	25
9.1	Introduction	25
9.2	<class>	25
9.3	<condition>	25
9.4	<extractor>	25
9.5	<consequence>	25
10	XML Semantic Module	26
10.1	Introduction	26
10.2	<document>	26
10.3	<expr>	26
10.4	<extractor>	26
10.5	<consequence>	26
11	Dispatcher Semantic Module	27
11.1	Introduction	27
11.2	<rule>	27
11.3	<consequence>	27
Part IV	Drools Internals	28
12	Architecture	29
12.1	Rules, rule-sets, and rule-bases	29
12.1.1	Rule, Condition and Consequence	29
12.1.2	RuleSet	30
12.1.3	RuleBase	30
12.2	Knowledge	30
12.3	Complete Model	31
13	Semantic Module Framework	32
13.1	Introduction	32
13.2	Declaration objects	32
13.3	Tuple objects	32
13.4	Semantic components	33
13.4.1	RuleType	33
13.4.2	ObjectType	33
13.4.3	Condition	34
13.4.4	Extractor	35
13.4.5	Condition	35
13.4.6	Consequence	35

13.5 The Configuration structure	35
13.6 Semantic Module Descriptor	36
13.7 SemanticsRepository	36
13.8 The DefaultSemanticsRepository helper	37
14 Rule Assembly API	38
14.1 Introduction	38
14.2 Requirements	38
 Part V Examples	 39
15 Sisters and Pets	40
16 Fibonacci Calculation	41
17 Trouble-Ticket Escalation	42
18 Manners Benchmark	43
19 Fish Monger	44
19.1 Java Code	44
19.1.1 Creating a Rule Base	44
19.1.2 Loading the Rules—XML file	44
19.1.3 Adding the Rules to the Rules base	44
19.1.4 Feature Request—Two step Rules base	44
19.2 Working Memory 101	45
19.2.1 Telling the Engine about your Objects	45
19.2.2 Telling the Engine to Run!	45
19.3 Drools Rule Logic (.drl) File	46
19.3.1 Explode Cart	46
19.3.2 Free Fish Food Special	46
19.3.3 Suggest a Tank	47
19.3.4 Apply a Discount	47
19.3.5 Source code of Fish Monger.drl	48
19.4 Output	49
19.5 Sample Code	51
19.5.1 Cart Item	51
19.5.2 FishMonger	52
19.5.3 Shopping Cart	53
 20 Family Tree of Baggins	 54
20.1 Java Code	54
20.2 Sample Code	55
20.2.1 Person	55
20.2.2 Main	56
20.2.3 Source code of Family Tree.drl	57

Part VI Appendices	58
A Algorithms	59
A.1 Efficient Matching	59
A.2 Rete	59
A.3 Rete-OO	62
B Frequently Asked Questions	64
B.1 Why are “or” conditions not allowed?	64
B.2 Can I nest rules?	64
B.3 What happened to version 1.0?	64
C Project Information	65
C.1 Web Site	65
C.2 Mailing Lists	65
C.3 Source Repository	65
C.4 Internet Relay Chat	66
C.5 Bug, Issue & Feature Tracking	66
C.6 Project Team	67
C.6.1 Bob McWhirter	67
C.6.2 Thomas Diesler	67
C.6.3 Roger F. Gay	67
C.6.4 Contributors	67
D Licensing	69
D.1 Drools license	69
D.1.1 The license	69
D.1.2 Summary	70
E JSR-94 : JavaTM Rule-Engine API	71
E.1 Introduction	71
E.2 Scope	71
E.3 JSR-94 and Drools	71
Index	72

List of Figures

4.1	Binding the DRL namespace URI to a prefix	8
4.2	Binding the DRL namespace URI as the default namespace	8
5.1	Loading a RuleBase from a JNDI data source	12
5.2	Deserializing a RuleBase from a file	13
5.3	Creation of a WorkingMemory from a RuleBase	13
5.4	Aasserting knowledge into a WorkingMemory	14
5.5	Retracting knowledge from a WorkingMemory	14
5.6	Modifying knowledge within a WorkingMemory	15
5.7	Firing rules	15
6.1	Loading a RuleSet from a URL	17
6.2	Conflict-resolution strategies	17
6.3	Using RuleBaseBuilder to build a RuleBase from a RuleSet	18
6.4	Deploying a RuleBase to a JNDI data store	18
6.5	Serializing a RuleBase to a file	19
7.1	Using a SalienceConflictResolutionStrategy	21
8.1	Binding of Java semantic module namespace URI to a prefix	23
8.2	Example of the <class> tag	24
12.1	Object model for Rule , Condition and Consequence	29
12.2	Object model for RuleSet and Rule	30
12.3	Object model for RuleBase and RuleSet	30
12.4	Object model for WorkingMemory and RuleBase	30
12.5	Complete object model	31
13.1	Method to retrieve an object by its Declaration from a Tuple	33
13.2	RuleType interface	33
13.3	ObjectType interface	34
13.4	ConfigurableObjectType interface	34
13.5	Condition interface	35
13.6	Attribute-related operations of Configuration	35
13.7	Text-related operation of Configuration	36

13.8 Child-related operations of <code>Configuration</code>	36
13.9 Usage of the <code>SemanticsRepository</code>	36
13.10 Retrieving and using the <code>DefaultSemanticsRepository</code> helper	37
20.1 Baggins Family Tree	54
A.1 Rete network	61
A.2 Example tuple sets	61
A.3 Rete-OO network	63

Acknowledgements

About this manual

Typographic conventions

Production notes

This manual was produced without the aid of Microsoft products. It was authored on a Linux laptop using XEmacs and \LaTeX . Drafts were previewed using `xdvi`. Diagrams were prepared using `dia`, exported in EPS format and converted for inclusion in the PDF document with `epstopdf`. The index was prepared using the `makeindex` package.

About the authors

Bob McWhirter

Robert Searle

Robert Searle is a senior consultant at Platinum Architecture Group, a firm specializing in software and business work flow design.

Mr. Searle has managed small teams of software developers, developed software architectures for both JAVA and Microsoft's Windows environments, and improved many different companies reporting systems. His work experience from various sectors including: Banks, Embedded Systems, Financial, Manufacturing, Transportation, and Research & Development.

Mr. Searle graduated from [Carleton University](#), Ottawa, Canada in 1995. He is a member of the Institute of Electrical and Electronics Engineers, the Association for Computing Machinery and the Professional Engineers of Ontario.

Robert Searle spends his free time training the family's German Shepherd, Rudi. Rudi will hopefully become a member of the volunteer search and rescue dog club.

Please contact Mr. Searle through his e-mail account RobertSearle@hotmail.com.

Part I

Drools—How and why

Chapter 1

50,000ft view

Chapter 2

Work Flow Engine

Part II

Drools

Chapter 3

Functional Overview

3.1 Rules

3.1.1 Declarative Form

Drools directly supports *declarative* rules, as opposed to *procedural* logic. Declarative rules typically take the form as follows:

```
if condition  
then consequence
```

By “declarative”, it is meant that the rules *declare*, by way of the *condition*, what should occur but do not specify the procedure for actually testing the conditions. For example, a procedural method for ensuring you have an umbrella if it is raining would be:

```
Step outside and determine if it is raining. If it is raining, then go  
to the closet and get an umbrella.
```

In a declarative form, the above could be represented by two rules:

- **if** *it is raining*
 then *you need an umbrella*
- **if** *you need an umbrella*
 then *get one from the closet*

Given declarative rules, the knowledge that it is raining could produce two courses of action:

1. You already have an umbrella, perhaps because you always carry one, in which case, you're ready.
2. You don't have an umbrella, so you go get one from the closet.

3.1.2 Applicable Context

The set of rules to be considered at any point of time depend upon the current context. As a human, you have certain rules to think about when you dine at a fine restaurant, which are exclusive to the set of rules you consider when spending a sunny day at the swimming pool. Even so, there are some other overriding rules that may be important regardless of the context, such as laws against homicide.

So, in a given context, different sets of rules are pertinent. A single set of rules may be user in one context, while a different set is used in a different context. Within a rule-engine, available sets of rules are called *rule sets*. The set or sets of rules currently applicable given the context is called a *rule base*.

3.2 Knowledge

You become aware of knowledge, in the form of facts, over time. Likewise, over time, facts may change or cease to be true, in which case, they facts as you know them must be altered or purged from your memory. Likewise, within a rule-engine, there are operations for becoming aware of a fact, purging a fact, or modifying a known fact. These operations are as follow:

- **Assert** Add a fact to what is known.
The weather is rainy. Betty is in the room.
- **Retract** Remove a fact from what is known.
Removing the fact *Betty is in the room* when Betty leave the room.
- **Modify** Alter a fact from what is know.
Changing the knowledge about the weather from *The weather is rainy* to *The weather is sunny* when the rain stops and the clouds go away.

Within a rule-engine, the collected knowledge is called the *working memory*. Knowledge is asserted, retracted and modified within the working memory and the rules are evaluated to determine what actions, if any, should be taken.

3.3 Why a rule-engine?

While the logic expressed in a rule can and has often been written within the code a system, a rule-engine offers many benefits. Instead of locking the logic up in code written by developers, the logic can be moved out-board external to the actual application. In this way it is possible for non-developers to change the logic without having to rebuild the system. Additionally, by codifying all of the system rules in a central location, they are no longer scattered throughout the application. This allows for easier validation of the system's requirements and analysis of the logic of the system.

Additionally, a rule-engine such as Drools is built upon an intelligent algorithm that allows for the evaluation of many rules against many facts in an efficient manner. In a procedural system, a change in a single fact might require double-checking *every* rule to determine if any action needs to be taken. A rule-engine which uses the *Rete* algorithm is optimized to minimize the amount of processing effort that is required to evaluate the rules that may have been affected by a change in knowledge.

3.4 A note about “business rules”

A higher-level form of rule is the *business rule*. Business rules do not necessarily follow the *if-then* form, but may be specified in different formats that are not as closely linked to the underlying rule-engine implementation. Business rules tend to use the *must* or *must not* form to expression constraints or inferences.

- An order **must not** be billed before it ships.
- An applicant for store credit **must** be 18 years of age.

Drools does not directly support this level of business rules, but other projects built upon Drools¹ may easily support such notation.

¹The **Fluxtapose** suite of tools from The Werken Company is one such product that supports business rules.

Chapter 4

Drools Rule Language

4.1 Introduction

The Drools Rule Language (DRL) is an XML-based language that uses an extensible tag mechanism. The available tags can be extended through the use of *semantic modules*. By itself the DRL is not fully usable. At least one semantic module must be selected and used. What follows is a reference for the available core DRL tags.

4.2 Core DRL namespace

The core DRL tags belong to the `http://drools.org/rules` namespace. In order to use the tags, the root tag must include an `xmlns` declaration binding the DRL namespace to either a prefix (Figure 4.1) or as the default namespace (Figure 4.2). Common practice is to bind the DRL namespace URI to the default namespace upon the root `<rule-set>` tag of a document.

```
<drl:rule-set xmlns:drl="http://drools.org/rules">
```

Figure 4.1: Binding the DRL namespace URI to a prefix

```
<rule-set xmlns="http://drools.org/rules">
```

Figure 4.2: Binding the DRL namespace URI as the default namespace

4.3 Core DRL tags

4.3.1 <rule-set>

The root tag of a DRL document is the **<rule-set>** tag, which serves to collect a group of rules. A rule-set must have a **name** attribute defining its name.

<rule-set>	
Attribute	Description
name	Name of the rule-set.
Tag	Description
<rule>	One or more rule members of the set.

```
<rule-set xmlns="http://drools.org/rules"
          name="My First Rule Set">
    ...
</rule-set>
```

4.3.2 <rule>

A **<rule>** tag defines a single rule within the rule-set. A rule must have at least one parameter¹. A rule contains conditions, extractors and exactly one consequence.

<rule>	
Attribute	Description
name	Name of the rule.
salience	Optional rule salience/priority.
Tag	Description
<parameter>	Parameter declaration.
<declaration>	Additional local declaration.
<ns:condition>	Semantic condition.
<ns:extractor>	Semantic extractor.
<ns:consequence>	Semantic consequence.
<duration>	Truthness duration.

¹Semantic modules may add implicit parameters to rules, so it's not strictly true that each **<rule>** must have a **<parameter>** or a **<consequence>**

```

<rule-set xmlns="http://drools.org/rules"
          name="My First Rule Set">

    <rule name="My First Rule"
          salience="10">
        ...
    </rule>

</rule-set>

```

4.3.3 <parameter>

A <parameter> defines an input parameter for a rule. Each input parameter will be supplied an object from the working memory in order to attempt satisfying the conditions. Each <parameter> is required to have an **identifier** that is unique within the rule and a semantic object type as its child element.

<parameter>	
Attribute	Description
identifier	Parameter identifier.
Tag	Description
<ns:objectType>	Object type of the parameter.

```

<rule name="My First Rule">

    <parameter identifier="factOne">
        ...
    </parameter>
    <parameter identifier="factTwo">
        ...
    </parameter>

</rule>

```

4.3.4 <declaration>

A <declaration> defines an addition *local* declared fact which may be the target of an *extractor*. Its **identifier** must be unique across all <parameter> and <declaration> tags. As with a <parameter>, it must have a semantic object type as its child element.

<declaration>

Attribute	Description
identifier	Declaration identifier.

Tag	Description
<ns:objectType>	Object type of the local declaration.

```

<rule name="My First Rule">

    <declaration identifier="factOne">
        ...
    </declaration>
    <declaration identifier="factTwo">
        ...
    </declaration>

</rule>

```

4.3.5 <duration>

The <duration> tag is optionally used to specify the truthness duration of the rule. The truthness duration determines how long the rule's conditions must be *continually* true in order to activate the rule.

Implementor's Note

The RuleSetReader currently doesn't allow <duration> tags, which is merely an oversight.

<duration>

Attribute	Description
weeks	Number of weeks.
days	Number of days.
hours	Number of hours.
minutes	Number of minutes.
seconds	Number of seconds.

```

<rule name="My First Rule">

    <duration hours="2"
              minutes="30"/>

</rule>

```

Chapter 5

Client API

5.1 Introduction

The Drools client API is simply a fairly lightweight interface for starting a session, manipulating knowledge and executing rules.

5.2 Retrieving a RuleBase

Before any activities may be performed, the client application must obtain a `RuleBase` through either building one using the administrative API or by performing a lookup within a JNDI directory or other managed object data source. Building a `RuleBase` is discussed in the next chapter.

5.2.1 Looking up a RuleBase from JNDI

A `RuleBase` may be a managed object which is deployed by an administrator within a JNDI directory. This method of deployment allows an application to be isolated from changes in the rules. The rules can be easily altered and redeployed, enabling an on-the-fly modification of the behaviour of an application. See Figure 5.1 for an example of accessing a `RuleBase` deployed within a JNDI data source.

```
1  DirContext ctx = new InitialDirContext( props );  
2  
3  RuleBase ruleBase = ctx.lookup( "cn='MyRuleBase'" );
```

Figure 5.1: Loading a `RuleBase` from a JNDI data source


```
1 File ser = new File(...);
```

Figure 5.2: Deserializing a RuleBase from a file

```
1 RuleBase ruleBase = obtainRuleBase();  
2  
3 WorkingMemory workingMemory = ruleBase.newWorkingMemory();
```

Figure 5.3: Creation of a WorkingMemory from a RuleBase

5.2.2 Deserializing a RuleBase from a file

Fully assembled `RuleBase` objects may be serialized to a file or other bytestream. This allows for the rules to be assembled before deployment and enables easy loading. See Figure 5.2 for an example of loading a serialized `RuleBase` from a file.

5.3 Creating a WorkingMemory

A `WorkingMemory` object represents an isolated rule session where knowledge is manipulated and evaluated to determine which rules should fire. Each `WorkingMemory` is independent from all other instances. Each `WorkingMemory` is backed by the `RuleBase` that created it (Figure 5.3) using `newWorkingMemory()`.

5.4 Working with knowledge

The three knowledge manipulation operations that can be performed are *assertion*, *retraction*, and *modification*. The `WorkingMemory` class models these three operations through three methods:

- `assertObject(...)` Adds a fact to the working knowledge.
- `retractObject(...)` Removes a fact from the working knowledge.
- `modifyObject(...)` Modifies a fact within the working knowledge.

5.4.1 Adding knowledge

To add knowledge to a `WorkingMemory`, a fact, in the form of any Java object is asserted into it (Figure 5.4 on the next page). A `FactHandle` is returned from each assertion and is used when performing further manipulation of the asserted fact. Once asserted, the fact may be used to satisfy some portion of any rule's activation condition.

```
1 WorkingMemory memory = ruleBase.newWorkingMemory();
2
3 FactHandle handle1 = memory.assertObject( factOne );
4 FactHandle handle2 = memory.assertObject( factTwo );
5 FactHandle handle3 = memory.assertObject( factThree );
```

Figure 5.4: Aasserting knowledge into a WorkingMemory

```
1 WorkingMemory memory = ruleBase.newWorkingMemory();
2
3 FactHandle handle = memory.assertObject( factOne );
4
5 memory.retractObject( handle );
```

Figure 5.5: Retracting knowledge from a WorkingMemory

5.4.2 Removing knowledge

To remove knowledge from a `WorkingMemory`, a `FactHandle` obtained from a previous invocation of `assertObject(...)` is used to reference the fact to be removed (Figure 5.5). Once retracted, the fact can not be used to satisfy any portion of any rule’s activation condition.

5.4.3 Modifying knowledge

Over time, knowledge may change. Instead of always retracting old facts and asserting new ones, it is sometimes advisable to view a change in knowledge as the *modification of a single fact*. For example, the fact about the current state of the weather may change as the weather changes. Instead of retracting a “rainy” fact and asserting a “sunny” fact, it is possible to simply change the value of the general “weather” fact (Figure 5.6 on the next page). As with retraction, modification of knowledge depends on a `FactHandle` obtained from a previous `assertObject(...)` call (Figure 5.6 on the following page).

The most important reason to perform fact modification instead of an assertion and a retraction is that rules with truth durations will automatically deactivate if a matching fact is retracted. If a fact is modified but still matches a rule with a truth duration, the clock on the duration timer is *not* reset.

5.5 Firing rules

As knowledge is manipulated within a `WorkingMemory`, no rules are actually triggered to be fired. Instead, rules are *activated*, which makes them candidates for being fired. Rule activations are placed upon the *agenda*.

```
1 WorkingMemory memory = ruleBase.newWorkingMemory();
2
3 FactHandle weatherHandle = memory.assertObject( new Weather( "rainy" ) );
4
5 memory.modifyObject( weatherHandle,
6                     new Weather( "sunny" ) );
```

Figure 5.6: Modifying knowledge within a `WorkingMemory`

```
1 WorkingMemory memory = ruleBase.newWorkingMemory();
2
3 FactHandle weatherHandle = memory.assertObject( new Weather( "rainy" ) );
4
5 memory.fireAllRules();
```

Figure 5.7: Firing rules

When an application determines that the current state of knowledge is consistent and should be acted upon, the application may cause the activations on the agenda to begin firing (Figure 5.7). The `fireAllRules(...)` method blocks until the agenda has been completely cleared of activations.

Note that as an activation fires, its consequence may perform knowledge manipulation by assertion new facts, retracting existing facts, or modifying existing facts. All of these activities may cause activations to be added to or removed from the agenda.

Chapter 6

Administrative API

6.1 Introduction

The administrative functions of drools include the assembly and deployment of rules. These activities may certainly be performed by the same application that acts as the client, The API has been designed to allow for the separation between the client and administrative responsibilities.

6.2 Building rules

Programmatically building rules using the Java API is covered in Chapter 14 on page 38. The next section discusses loading rules that were built using the Drools Rule Language, which is described in depth in Chapter 4 on page 8.

6.3 Building a RuleBase

6.3.1 Loading rules

Rule-sets can be loaded using the `RuleSetReader` from a variety of sources including `URL`, `InputStream` and `Reader` objects. A `RuleSetReader` builds a `RuleSet` from the contents of a DRL document (Chapter 4 on page 8).

The `RuleSetReader` by default may use any and all semantic modules available on the classpath (see Chapter 13 on page 32). It may be optionally parameterized with a `SAXReader` and a `SemanticsRepository`. By default, the `RuleSetReader` uses the Java API for XML Parsing (JAXP) for reading the XML documents. See Figure 6.1 on the following page for an example of loading a `RuleSet` from a `URL`.

```

1 RuleSetReader reader      = new RuleSetReader();
2 URL           ruleSetUrl = new URL( "http://myco.com/theRuleSet.xml" );
3 RuleSet       ruleSet    = reader.read( ruleSetUrl );

```

Figure 6.1: Loading a RuleSet from a URL

6.3.2 Conflict resolution

When constructing a RuleBase, a *conflict resolution strategy* must be selected. Conflict resolution is described fully in Chapter 7 on page 20, but a summary of available strategies appears in Figure 6.2.

Class	Description
SalienceConflictResolutionStrategy	Resolve conflicts based upon the <i>salience</i> of the rules. If two conflicting rules have the same salience, one is selected <i>at random</i> to be fired first.
ComplexityConflictResolutionStrategy	Resolves conflicts based upon the complexity of the rules as measured by the number of conditions. Rules with more conditions have higher priority than those with fewer conditions. If two conflicting rules have the same measure of complexity, conflict is then resolved as with the salience conflict resolution strategy.
SimplicityConflictResolutionStrategy	Resolves conflicts based upon the simplicity of the rules as measured by the number of conditions. Rules with fewer conditions have higher priority than those with more conditions. If two conflicting rules have the same measure of complexity, conflict is then resolved as with the salience conflict resolution strategy.

Figure 6.2: Conflict-resolution strategies

6.4 Deploying a RuleBase to a JNDI data store

A fully-constructed RuleBase may be serialized and stored within a JNDI-accessible directory (Figure 6.4 on the next page). This method allows for an administrator to deploy a RuleBase as a managed object. Multiple applications may use the JNDI-accessible RuleBase without being concerned with parsing DRL documents and building a RuleBase directly. Additionally, the deployed RuleBase may be easily replaced, with the replacement being immediately ac-

```

1 RuleBaseBuilder builder = new RuleBaseBuilder();
2
3 builder.addRuleSet( theRuleSet );
4
5 builder.setConflictResolutionStrategy(
6     SaliencyConflictResolutionStrategy.getInstance() );
7
8 RuleBase ruleBase = builder.buildRuleBase();

```

Figure 6.3: Using RuleBaseBuilder to build a RuleBase from a RuleSet

```

1 DirContext ctx = new InitialDirContext( props );
2 ctx.bind( "cn='MyRuleBase'" );

```

Figure 6.4: Deploying a RuleBase to a JNDI data store

cessible to a running application.

All components of the Rule *must* be serializable in order for the constructed RuleBase to be deployed against a JNDI datastore. The various semantic component interfaces and rule-assembly classes have already been marked as **Serializable**. The class definition files must be in the application's classpath in order to deserialize the object structure once retrieved from the directory. Failure to have the classes available will result in an exception being thrown.

6.4.1 JNDI deployment utility

Implementor's Note

...Not yet done

The command-line utility **name-me-please** is available to assist in the creation of a serialized RuleBase in a JNDI directory from a DRL rule definition file (Chapter 4 on page 8).

```

1 name-me-please -Dprop=value --classpath xx:yy:zz myRules.drl 'cn=MyRules'

```

6.5 Serializing a RuleBase to a file

A RuleBase can be serialized in full to a file on disk (Figure 6.5 on the next page). This method allows for a rule-base to be “frozen” prior to deployment. An application merely has to unserialize the file to use the contained RuleBase.

```
1 File ruleBaseFile = new File( pathToFile );
2
3 FileOutputStream fileOut = new FileOutputStream( ruleBaseFile );
4 ObjectOutputStream objOut = new ObjectOutputStream( fileOut );
5
6 objOut.writeObject( ruleBase );
```

Figure 6.5: Serializing a RuleBase to a file

6.5.1 Serialization utility

Implementor's Note

...Not yet done

The command-line utility **name-me-please** is available to assist in the creation of a serialized RuleBase from a DRL rule definition file (Chapter 4 on page 8).

```
1 name-me-please --classpath xx:yy:zz myRules.drl myRules.ser
```

Chapter 7

Conflict Resolution

7.1 Introduction

When knowledge is manipulated, many rules may become *activated*. The consequences of the rules fire serially and thus must be ordered in the activation queue¹. Since firing one consequence may manipulate knowledge, the order of rule firing is important. Different end-results may occur based upon different firing order. Applying an order to the activations is called a *conflict resolution strategy*. Each `RuleBase` has a conflict resolution strategy that determines how rules are prioritized. Custom conflict resolution strategies may be devised for each rule-base, but several standard ones are available for use.

7.2 SalienceConflictResolutionStrategy

Each `Rule` has a *salience* property attached to it. By default every rule has a salience of zero, but during assembly, either programatically using the Rule Assembly API (Chapter ?? on page ??) or through the DRL (Chapter 4 on page 8) a specific non-default salience value may be applied to each rule.

Salience is a form of priority where rules with higher salience values are given higher priority when ordered in the activation queue. In the event that multiple rules are assigned the same salience value, they are placed upon the queue in an arbitrary order.

To use a salience-based conflict resolution strategy, the `RuleBase` should be built using an instance of the `SalienceConflictResolutionStrategy` class. For ease of use a `getInstance()` method is provided (Figure 7.1 on the next page).

¹The activation queue is formally known as the *agenda*.


```
1 RuleBaseBuilder builder = new RuleBaseBuilder();  
2  
3 builder.setConflictResolutionStrategy(  
4     SalienceConflictResolutionStrategy.getInstance() );
```

Figure 7.1: Using a SalienceConflictResolutionStrategy

7.3 ComplexityConflictResolutionStrategy

Built upon the `SalienceConflictResolutionStrategy`, a strategy that first takes into account the *complexity* of the conflicting rules is available. The `ComplexityConflictResolutionStrategy` initially orders activations by their complexity as measured by the number of conditions in each rule. Rules with more conditions have a higher complexity and thus a higher priority when compared to rules with fewer conditions. When rules have the same complexity, ties are broken using their salience values.

7.4 SimplicityConflictResolutionStrategy

Also built upon the `SalienceConflictResolutionStrategy`, a strategy that first takes into account the *simplicity* of the conflicting rules is available. The `SimplicityConflictResolutionStrategy` may be thought of as the inverse of the complexity-based strategy. Rules with less conditions are given higher priority than those with more. Once again, when conflicting rules have the same measure of simplicity, the strategy falls back to resolving conflicts based upon salience.

Part III

Semantic Modules

Chapter 8

Java Semantic Module

8.1 Introduction

The Java semantic module defines addition tags which may be used within a DRL document (Chapter 4 on page 8). The tags provided by this module allow for semantics based upon the Java programming language. Object types may be determined using Java classes while conditions and extractors are formulated in terms of Java expressions. A consequence may be an arbitrary block of Java statements.

8.2 Java semantic module namespace URI

That tags for the Java semantic module are defined within the XML namespace `http://drools.org/semantics/java`. In order to use the tags of the Java semantic module, this namespace should be bound to a prefix with the DRL document. Typically this is done on the root `rule-set` tag (Figure 8.1).

```
<rule-set xmlns="http://drools.org/rules"
          xmlns:java="http://drools.org/semantics/java">
  ...
</rule-set>
```

Figure 8.1: Binding of Java semantic module namespace URI to a prefix

```
<java:class>
```

Figure 8.2: Example of the `<class>` tag

8.3 Java Semantic Module tags

8.3.1 `<class>`

		<code><class></code>
<i>No attributes</i>		
Content		
Name of the java class.		

8.3.2 `<condition>`

8.3.3 `<extractor>`

8.3.4 `<consequence>`

Chapter 9

Python Semantic Module

9.1 Introduction

9.2 `<class>`

See Section 8.3.1 on the preceding page.

9.3 `<condition>`

9.4 `<extractor>`

9.5 `<consequence>`

Chapter 10

XML Semantic Module

Implementor's Note

Not yet implemented...

10.1 Introduction

10.2 <document>

10.3 <expr>

10.4 <extractor>

10.5 <consequence>

Chapter 11

Dispatcher Semantic Module

Implementor's Note

Not yet implemented...

11.1 Introduction

11.2 <rule>

11.3 <consequence>

Part IV

Drools Internals

Chapter 12

Architecture

12.1 Rules, rule-sets, and rule-bases

Within Drools, the concepts of rules, rule-sets and rule-bases are directly modelled by the classes `Rule`, `RuleSet` and `RuleBase`. A `Rule` may be a member of multiple `RuleSets`, and multiple `RuleSets` may be active within a given `RuleBase`.

12.1.1 Rule, Condition and Consequence

A single `Rule` may have one-or-more `Conditions` associated with it. Each condition must be met before the rule is considered to be activated. Once activated the rule's `Consequence` is a candidate for being fired. In pattern parlance, the `Condition` class is simply a *predicate object* which evaluates itself against the known facts to return a boolean value of either *true* or *false*. The `Consequence` class is likewise simply a *functor* which objectifies a function and performs an arbitrary task when executed.

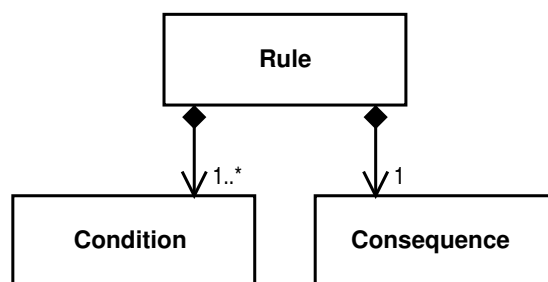


Figure 12.1: Object model for Rule, Condition and Consequence




Figure 12.2: Object model for RuleSet and Rule




Figure 12.3: Object model for RuleBase and RuleSet

12.1.2 RuleSet

A **RuleSet** is simply a collection of **Rules**. It serves only to associate a group of rules with one another so that they may be worked with as a set.

12.1.3 RuleBase

A **RuleBase** is an *active* collection of **RuleSets**. A **RuleBase** contains rules that are all considered to be in effect for a given set of knowledge. Multiple **RuleSets** may be a part of a given **RuleBase**.

12.2 Knowledge

The set of knowledge that is examined is modelled by the class **WorkingMemory**, which is backed by a particular **RuleBase**. It is through the **WorkingMemory** that knowledge is *asserted*, *retracted* and *modified*. Each **WorkingMemory** is backed by exactly one **RuleBase** which determines which rules are evaluated as knowledge is manipulated.




Figure 12.4: Object model for WorkingMemory and RuleBase

12.3 Complete Model




Figure 12.5: Complete object model

Chapter 13

Semantic Module Framework

13.1 Introduction

Implementor's Note

*Wait wait wait...
Reworking the SMF. ConfigurableFoo is just crappy.*

The Semantic Module Framework (SMF) provides an extension point to the Drools Rule Language syntax (see Chapter 4 on page 8). By creating semantic modules, domain-specific rule types, object-types, conditions, extractors and consequences can be added to the core DRL language.

13.2 Declaration objects

Each has a set of parameters that specify what input objects are required through the evaluation of the rule. Each parameter is modelled through a **Declaration** object which specifies the parameter's *identifier* and *object type*. The identifier of a **Declaration** is simply a string name that uniquely distinguishes one parameter from another. The object type is defined through an **ObjectType** component (Section ?? on page ??). Various semantic components, notably **Condition** and **Consequence** implementations must be cognizant of the available declarations of a rule.

13.3 Tuple objects

Object flow through the conditions and consequence of a rule in the form of a **Tuple**. A tuple is merely an associative array, not unlike a **HashMap**, which indexes the value of the objects by its **Declaration**. **Condition** and **Consequence** implementations operate upon **Tuple** objects by retrieving values by their **Declaration** (Figure 13.1 on the next page).

```
1 package org.drools.spi;
2
3 public interface Tuple
4 {
5     ...
6     Object get(Declaration declaration);
7     ...
8 }
```

Figure 13.1: Method to retrieve an object by its Declaration from a Tuple

```
1 package org.drools.smf;
2
3 public interface RuleType
4 {
5     void initializeRule(Rule rule,
6                        Configuration config)
7         throws RuleTypeException;
8 }
```

Figure 13.2: RuleType interface

13.4 Semantic components

13.4.1 RuleType

Implementor's Note

This doesn't exist at the moment, but will soon. RuleType components will allow a module to be responsible for the initialization of rule, possibly adding parameters, condition and consequences implicitly.

Implementations of RuleType allow a module to initialize a new rule, possibly by adding parameters, conditions or consequences. The RuleType is passed the new Rule for initialization.

13.4.2 ObjectType

Implementations of ObjectType allow rules to be defined in terms of semantic types. Instead of requiring all rules to be defined in terms of Java classes, rules may be defined in terms of higher semantics. For example, while all XML documents may be instances of `org.w3c.dom.Document` each document may have a different semantic type based upon the name and namespace of the root tag. Given an Object, an ObjectType implementation must simply determine if it matches its semantic type (Figure 13.3 on the following page).

```
1 package org.drools.spi;
2
3 public interface ObjectType
4     extends SemanticComponent
5 {
6     boolean matches(Object object);
7 }
```

Figure 13.3: ObjectType interface

```
1 package org.drools.smf;
2
3 public interface ConfigurableObjectType
4     extends ObjectType
5 {
6     void configure(Configuration config) throws ConfigurationException;
7 }
```

Figure 13.4: ConfigurableObjectType interface

ConfigurableObjectType

Configuration information may be passed to an `ObjectType` if the implementation is marked with the `ConfigurableObjectType` interface. A `configure(...)` method is added by the `ConfigurableObjectType` interface, through which additional information may be passed (Figure 13.4).

13.4.3 Condition

Implementations of the `Condition` interface allow for custom conditions to be created. Each condition is effectively a predicate object that returns a boolean value given some input data. In order to insert the `Condition` into the appropriate location within the Rete graph, each `Condition` must specify the declarations of the variables with which it analyzes (Section 13.2 on page 32).

Figure 13.5 on the next page shows the `Condition` interface. Any value expected to be used by the `isAllowed(...)` method must be accounted for in the array of `Declaration` objects returned by `getRequiredTupleMembers`. For example, a rule may declare three objects: `a`, `b`, and `c`. A particular `Condition` implementation might only test attributes of the `b` object. In that case, the `Declaration` associated with the `b` object must be returned from `getRequiredTupleMembers(...)`.

The `isAllowed(...)` method is passed a `Tuple` when enough knowledge to satisfy the required members is available. It may then use the operations of `Tuple` to retrieve the objects and perform its test to return a boolean.

```

1 package org.drools.spi;
2
3 public interface Condition
4     extends SemanticComponent
5 {
6     Declaration[] getRequiredTupleMembers();
7
8     boolean isAllowed(Tuple tuple) throws ConditionException;
9 }

```

Figure 13.5: Condition interface

```

1 String[] attrNames = config.getAttributeNames();
2 String    someValue = config.getAttribute( someName );

```

Figure 13.6: Attribute-related operations of Configuration

ConfigurableCondition

13.4.4 Extractor

ConfigurableExtractor

13.4.5 Condition

ConfigurableCondition

13.4.6 Consequence

ConfigurableConsequence

13.5 The Configuration structure

For the `Configurable...` form of the semantic components, configuration information is communicated through a tree of `Configuration` objects. Each `Configuration` object acts as a node in the tree, and may contain the following data:

- **Attributes.** Zero or more name/value pairs of strings (Figure 13.6).
- **Text.** A single string text value (Figure 13.7 on the next page).
- **Child Configuration nodes.** Zero or more named child `Configuration` nodes (Figure 13.8 on the following page).

The tree of `Configuration` nodes may be thought of as a simplified version of an XML structure. For configurable semantic components used through the

```

1 String nodeText = config.getText();

```

Figure 13.7: Text-related operation of Configuration

```

1 Configuration[] allChildren    = config.getChildren();
2 Configuration  someFirstChild = config.getChild( someName );
3 Configuration[] someChildren  = config.getChildren( someName );

```

Figure 13.8: Child-related operations of Configuration

DRL (Chapter 4 on page 8), the root `Configuration` is based upon the component's own tag, and children tags are represented by children `Configuration` nodes.

13.6 Semantic Module Descriptor

Conforming semantic modules are packaged as individual JAR files which can be added to the application's classpath. Each JAR should contain, in the `META-INF` directory a file named `drools-semantic.properties` which provides meta-information about the module and its available semantic components.

Implementor's Note

Still working on the module descriptor format.

13.7 SemanticRepository

A `SemanticRepository` manages a set of `SemanticModule` objects and allows each to be looked-up by its URI (Figure 13.9). Primary a `SemanticRepository` is used by a `RuleSetReader` (Section 6.3.1 on page 16) in order to extend the core DRL (Chapter 4 on page 8) syntax.

```

1 SemanticRepository repo      = locateSemanticRepository();
2 SemanticModule[]   modules  = repo.getSemanticModules();
3 SemanticModule     someModule = repo.lookupSemanticModule( someUri );

```

Figure 13.9: Usage of the SemanticRepository


```
1 SemanticRepository repo    = DefaultSemanticsRepository.getInstance();  
2 SemanticModule[]   modules = repo.getSemanticModules();
```

Figure 13.10: Retrieving and using the DefaultSemanticsRepository helper

13.8 The DefaultSemanticsRepository helper

The `DefaultSemanticsRepository` helper class is useful in that it contains all conforming semantic modules available on the classpath. Each module that has a module descriptor (Section 13.6 on the page before) located within the `META-INF` directory. Each `drools-semantic.properties` will be automatically discovered by the `DefaultSemanticsRepository` upon first use.

Being a help class that is initialized once, it follows the singleton pattern. To use the `DefaultSemanticsRepository`, the `getInstance()` method will retrieve the singleton instance (Figure 13.10).

Chapter 14

Rule Assembly API

14.1 Introduction

The DRL (Chapter 4 on page 8) constructs rules using the *Rule Assembly API*. This API may also be used directly by applications that wish to programatically construct `Rule` instances without using the DRL.

The `org.drools.rule` package contains the classes and interfaces that form the Rule Assembly API. The two concrete classes of note are `Rule` and `RuleSet`. Other concrete classes used in the assembly of rules are pulled from various semantic module implementations.

14.2 Requirements

Each `Rule` is required to have at least one `Parameter` and a `Consequence`. Zero or more `Condition` objects may be attached to each `Rule`¹.

¹Though, a `Rule` without a `Condition` is not much of a rule.

Part V

Examples

Chapter 15

Sisters and Pets

Chapter 16

Fibonacci Calculation

Chapter 17

Trouble-Ticket Escalation

Chapter 18

Manners Benchmark

Chapter 19

Fish Monger

19.1 Java Code

19.1.1 Creating a Rule Base

19.1.2 Loading the Rules—XML file

19.1.3 Adding the Rules to the Rules base

19.1.4 Feature Request—Two step Rules base

I hope all the above can become something like

```
final URL url = FishMonger.class.getResource("fishmonger.drl");  
final RuleSet rules = RulesSetLoader.getRuleSet(url);
```


19.2 Working Memory 101

19.2.1 Telling the Engine about your Objects

19.2.2 Telling the Engine to Run!

19.3 Drools Rule Logic (.drl) File

19.3.1 Explode Cart

19.3.2 Free Fish Food Special

19.3.3 Suggest a Tank

19.3.4 Apply a Discount

19.3.5 Source code of Fish Monger.drl

19.4 Output

```
R U N N I N G  fishmonger.FishMonger
```

```
loading: jar:file[this is specific to your computer!!!]
```

```
PRE
```

```
[ShoppingCart:
    gross total=90.92999999999999
    discounted total=90.92999999999999
    [CartItem: name='tropical fish'; cost=12.99]
    [CartItem: name='tropical fish'; cost=12.99]
    [CartItem: name='tropical fish'; cost=12.99]
    [CartItem: name='tropical fish'; cost=12.99]
    [CartItem: name='tropical fish'; cost=12.99]
    [CartItem: name='tropical fish'; cost=12.99]
    [CartItem: name='tropical fish'; cost=12.99]
]
```

```
Examining each item in the shopping cart.
Adding free tropical fish food sample to cart
Adding free tropical fish food sample to cart
Adding free tropical fish food sample to cart
Adding free tropical fish food sample to cart
Adding free tropical fish food sample to cart
Adding free tropical fish food sample to cart
Adding free tropical fish food sample to cart
*** SUGGESTION: Would you like to buy a tank for your7 fish?
Applying 15% discount to cart
```

```
POST
```

```
[ShoppingCart:
    gross total=90.92999999999999
    discounted total=77.2905
    [CartItem: name='tropical fish'; cost=12.99]
    [CartItem: name='tropical fish'; cost=12.99]
    [CartItem: name='tropical fish'; cost=12.99]
    [CartItem: name='tropical fish'; cost=12.99]
    [CartItem: name='tropical fish'; cost=12.99]
    [CartItem: name='tropical fish'; cost=12.99]
    [CartItem: name='tropical fish'; cost=12.99]
```

```
[CartItem: name='tropical fish food sample'; cost=0.0]
[CartItem: name='tropical fish food sample'; cost=0.0]
[CartItem: name='tropical fish food sample'; cost=0.0]
[CartItem: name='tropical fish food sample'; cost=0.0]
[CartItem: name='tropical fish food sample'; cost=0.0]
[CartItem: name='tropical fish food sample'; cost=0.0]
[CartItem: name='tropical fish food sample'; cost=0.0]
]
```

19.5 Sample Code

19.5.1 Cart Item

19.5.2 FishMonger

19.5.3 Shopping Cart

Chapter 20

Family Tree of Baggins




Figure 20.1: Baggins Family Tree

See <http://www.tolkienonline.de/etep/Tolkien/FamilyTrees/Bilbo.html>
and <http://www.tuckborough.net/baggins.html> for more info
Please note that *Belladonna* is female. This will be important later!

20.1 Java Code

20.2 Sample Code

20.2.1 Person

20.2.2 Main

20.2.3 Source code of Family Tree.drl

Part VI

Appendices

Appendix A

Algorithms

A.1 Efficient Matching

While it may be simple to create a rules engine that allows specification of business logic in a format that is comfortable to business analysts, the matching of the rules may still be problematic without a good algorithm.

The rules engine must be made aware of its environment, typically through a process called *fact assertion*. Fact assertion consists of the program asserting facts into a rules session, or *working memory*.

Whenever a fact is asserted, retracted or modified within the working memory, many rules may become candidates for firing, or may have become invalidated. A simplistic approach is to reevaluate all rules against the entirety of the working memory. This method is guaranteed to be correct but will also certainly be sub-optimal. Any individual fact modification only affects a small number of conditions in a small number of rules.

Variations of the Rete algorithm allow the rules engine to maintain a memory of the results of partial rule matches across time. Reevaluation of each condition is no longer necessary, as the engine knows which conditions might possibly change for each fact, and only those must be reevaluated.

A.2 Rete

Charles Forgy created the original Rete algorithm [?] around 1982 as part of his DARPA-funded research. Compared to many previous production-matching algorithms, Rete was very advanced. Even today, there have been few improvements to it in the general case¹. Variations on Rete, such as TREAT [?], may have different performance characteristics depending on the environment. Some perform better with large rule sets but small numbers of objects, while other

¹Both ILOG and Haley claim to have optimized Rete algorithms, but details are not currently public.

perform well for steady-state environments, but react poorly to numerous successive changes in the data.

A *Rete network* is a graph through which data flows. Originally, data was specified using Cambridge-prefix tuples since Lisp-like languages were in style for logic programming.² The tuples were used to express attributes about objects. For example, tuples may be used to express a person's name and her pets. The tuples are dropped into the Rete network, and those that reach the far end cause the firing of a rule. The original production-matching was based upon matches against tuple patterns.

The Rete network is comprised of two types of nodes:

- **1-input/1-output nodes**

The *1/1* nodes are constrictive nodes that only allow matching tuples to flow through. Any tuples that do not match are discarded by the node.

- **2-input/1-output nodes**

The *2/1* nodes simply connect the output arcs from two other nodes (either *1/1* nodes or *2/1* nodes) merging tuples from both the left and right incoming arcs into a single tuple on the outgoing arc. Maintains a memory of tuples for matching against future facts.

A forest of *1/1* nodes acts as the entry-point into the entire Rete network for any incoming tuple. The network-entry nodes filter tuples purely by their type. Tuples about dogs and tuples about cats may each have a different type and may be differentiated from each other by the *1/1* network-entry nodes.

Each condition of a rule is merely a pattern for a particular tuple type. The condition describes the attributes that a tuple must have and acts as a filter. Each condition is transformed into a *1/1* node that only allows tuples matching the specified attributes to pass. An attribute value may be specified as a variable and implies that the variable must hold the same value in all occurrences. The *1/1* filter nodes are attached to the network downstream from the *1/1* entry-node that differentiates their tuple type.

Consider a condition such as “For any person who has a dog that has the same name as that person's sister's cat, then...” This could be expressed with the condition patterns of:

```
(1) ( person  name=person?  sister=sister?  )
(2) ( person  name=person?  dog=petName?   )
(3) ( person  name=sister?   cat=petName?   )
```

Condition #1 models the sister relationship so that the rule only applies to two people who are sisters. The **person?** and **sister?** tokens are variables that must be consistent across any set of tuples that match this rule.

Conditions #2 and #3 serve two roles. The **dog** and **cat** attributes share the same **petName?** variable and serve to identify two people who have a cat

²As it is for many artificial intelligence projects.

and a dog with the same name. They each contain a **name** attribute with either the variable **person?** or **sister?** which ties the last two conditions back to the first two.




Figure A.1: Rete network

<i>type</i>	person	sister	cat	dog
<i>tuple set # 1</i>				
person	rebecca	jeannie	zoomie	<i>null</i>
person	jeannie	rebecca	<i>null</i>	zoomie
<i>tuple set # 2</i>				
person	rebecca	jeannie	zoomie	<i>null</i>
person	jeannie	rebecca	<i>null</i>	toby

Figure A.2: Example tuple sets

If two sets of tuples (see Figure A.2) were asserted against the rule, *tuple set #1* would cause a firing of the rule, where *tuple set #2* would not. In both cases, the two tuples would pass node *condition(1)*, as the nodes simply associate the *person?* and *sister?* variables with the appropriate values from each tuple.

The *join(1)* node would allow both tuples to merge and propagate past it in both the first and second case. Additionally, for both cases, the *rebecca* tuple would pass node *condition(2)* and the *jeannie* tuple would pass node *condition(3)*.

The *join(2)* node is where the two cases differ. In the first case, nodes *condition(2)* and *condition(3)* have each associated the value of “zoomie” to the *petName?* variable. In the second case, the two nodes has assigned different values to the variable. The *join(2)* node only allows those tuples that have consistent associations with all variables to pass.

A.3 Rete-OO

The Rete algorithm works wonderfully in language systems such as Lisp where pertinent attributes about objects are directly asserted to the rules engine. In an object-oriented language, such as C++ or Java, and entire graph of objects can be reachable from a single named root object. Expressing highly complex relationships between entities using Cambridge-prefix notation may require many separate assertions. In an OO language, the single root object is all that should be asserted, since attributes and relationships can be *extracted* using normal language constructs.

Bob McWhirter of The Werken Company adapted Forgy’s original Rete algorithm to object-oriented constructs, creating the Rete-OO algorithm. As with Rete, there are *1/1* nodes and *2/1* nodes. Unlike Rete, there are nodes that exist simply to extract reachable attributes and add columns to passing tuples. Rete always constructs the condition *1/1* nodes toward the root of the tree leaving the bottom portion to be comprised of purely aggregating *2/1 join* nodes. Rete-OO must interleave both *1/1* and *2/1* nodes.

The same example as in Section A.2, the conditions could be expressed in terms of object-oriented language boolean and assignment expressions. The choice of Java as the expression language is purely arbitrary.

```
(0)  Person personOne, personTwo
(1)  personOne.hasSister(personTwo)
(3)  petName = personOne.getCat().getName()
(3)  petName = personTwo.getDog().getName()
```

Rete-OO adds the concept of *root object declaration*, where the root objects of the condition are declared with a name and type. The object’s type maps directly to the tuple type in Rete. The root object name has no direct mapping in Rete and causes the addition of a *parameter node* in Rete-OO. Boolean expressions in Rete-OO conditions are equivalent to Rete’s condition patterns




Figure A.3: Rete-OO network

against attributes. The assignment expressions map to place-holder variables in Forgy's algorithm.

The types of nodes used in Rete-OO graph construction are listed here. Those that are new or different from Rete are denoted with a '*'.

- **Object type**

Object type nodes differentiate objects by filtering on their defined type.

- **Parameter***

Parameter nodes create a tuple with a single entry binding the object to the name.

- **Condition**

Condition nodes simply tests a tuple against an a boolean expression.

- **Extraction***

Extraction nodes extract new attributes, create new columns on tuples, and store the results.

- **Join**

Join nodes connect the output arcs from two other nodes and allows consistent tuples to be merged and passed through.

- **Terminal**

Terminal nodes fire to indicate a successful match for the rule.

The resulting Rete-OO graph is constructed in a different manner than the equivalent Rete graph, due to the addition and rearrangement of some nodes.

Appendix B

Frequently Asked Questions

B.1 Why are “or” conditions not allowed?

B.2 Can I nest rules?

B.3 What happened to version 1.0?

Appendix C

Project Information

C.1 Web Site

All development resources related to Drools are hosted by **The Codehaus**, the open-source arm of The Werken Company. Drools maintains a website at:

<http://drools.org/>

C.2 Mailing Lists

The drools project maintains two mailing lists. The first, known as **drools-interest** is for general discussion by users and developers of drools. The second list is **drools-cvs** which simply tracks changes made to the source-code through the CVS repository. For information about subscribing to each list or access to the list archives:

<http://lists.codehaus.org/listinfo/drools-interest>
<http://lists.codehaus.org/listinfo/drools-cvs>

C.3 Source Repository

The drools project maintains a revision control repository using CVS. To check-out the latest sources, you must issue two CVS commands. The first is used to login. When presented with a prompt for a password, simply press *ENTER*.

```
cvs -d:pserver:anonymous@cvs.codehaus.org:/scm/cvspublic login
cvs -d:pserver:anonymous@cvs.codehaus.org:/scm/cvspublic co drools
```

C.4 Internet Relay Chat

There is a dedicated channel on The Werken Company's IRC server for drools:

```
address  irc.codehaus.org
port     6667
channel  #drools
url      irc://irc.codehaus.org:6667/drools
```

C.5 Bug, Issue & Feature Tracking

For bug, issue and feature tracking, the Drools project uses the Jira project management system provided by The Codehaus.

```
http://jira.codehaus.org/
```

C.6 Project Team

C.6.1 Bob McWhirter

Bob McWhirter originally founded the Drools project in 2000 and developed the Rete-OO algorithm used by the engine. Bob is also the founder of The Werken Company and the chief architect behind the commercial **Fluxtapose** suite of tools which build upon Drools to provide a complete solution for implementing business rules.

C.6.2 Thomas Diesler

Thomas Diesler researched and supplied the JSR-94 Rule-Engine API bindings for Drools.

Editor's Note

*Thomas, please send more details to
bob@werken.com.*

C.6.3 Roger F. Gay

Roger F. Gay devised the XML Schemas for the core DRL syntax and each semantic module.

Editor's Note

*Likewise, Roger, please send more details to
bob@werken.com.*

C.6.4 Contributors

Others have contributed ideas, patches and testing assistance over the years:

- Dave Cramer (*eBox*)
- Martin Hald
- Matt Ho
- Pete Kazmier (*iBasis*)
- Christiaan ten Klooster
- James Roome
- Bart Selders (*iBanx*)
- James Strachan (*CoreDevelopers Network*)
- Tom Vasak

Editor's Note

Hello to contributors: send your current affiliation information to bob@werken.com if you wish it to be included.

Appendix D

Licensing

D.1 Drools license

D.1.1 The license

Copyright 2001-2003 (C) The Werken Company. All Rights Reserved.

Redistribution and use of this software and associated documentation ("Software"), with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain copyright statements and notices. Redistributions must also contain a copy of this document.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name "Drools" must not be used to endorse or promote products derived from this Software without prior written permission of The Werken Company. For written permission, please contact bob@werken.com.
4. Products derived from this Software may not be called "Drools" nor may "Drools" appear in their names without prior written permission of The Werken Company. "Drools" is a trademark of The Werken Company.
5. Due credit should be given to The Werken Company. (<http://werken.com/>)

THIS SOFTWARE IS PROVIDED BY THE WERKEN COMPANY AND CONTRIBUTORS **AS IS** AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE WERKEN COMPANY OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

D.1.2 Summary

Drools is provided under a license similar to that used by The Apache Software Foundation. It is a commercial-friendly license in that it allows you to modify and distribute Drools in either source or binary form. While you are *encouraged* to contribute changes back to the project, you are by no means *required* to do so. The Drools license is not viral or infectious. It does not alter how you license your own product. If you have any questions regarding the licensing of Drools, please contact info@werken.com.

Appendix E

JSR-94

JavaTM Rule Engine API

E.1 Introduction

The JavaTM Rule-Engine API, known as *JSR-94* is a specification developed by the Java Community Process (JCP) which attempts to standardize the way that applications interact with rule-engines.

E.2 Scope

JSR-94 only attempts to standardize the actual Java API for working with rule-engines and does not address the authoring of rules at all. Each rule-engine still has its own syntax for writing rules. JSR-94 *does* proscribe a method of loading a rule definition document using a DOM **Document** but does not speak as to the content of the document being loaded.

The specification also does not address any details regarding the underlying implementation or algorithms. It simply provides a consistent manner in which to obtain a rule session, manipulate knowledge within it, and to fire rules.

E.3 JSR-94 and Drools

The JSR-94 binding for Drools is an optional module that provides a thin wrapper around the client (Chapter 5 on page 12) and administrative (Chapter 6 on page 16) APIs.

Index

- activation queue, 20
- agenda, 14, 20
- assert, 6
- assertion, 13

- business rules, 7

- <class> tag, 24, 25
- Codehaus, The, 65
- ComplexityConflictResolutionStrategy class, 17, 21
- Condition class, 29, 32, 34, 35, 38
 - getRequiredTupleMembers, 34
 - getRequiredTupleMembers(...), 34
 - isAllowed(...), 34
- <condition> tag, 24, 25
- ConfigurableCondition class, 35
- ConfigurableConsequence class, 35
- ConfigurableExtractor class, 35
- ConfigurableObjectType class, 34
 - configure(...), 34
- Configuration class, 35, 36
- conflict resolution, 20
- Consequence class, 29, 32, 35, 38
- <consequence> tag, 24–27
- Cramer, Dave, 67

- Declaration class, 32–34
- <declaration> tag, 10
- declarative, 5
- DefaultSemanticsRepository class, 37
 - getInstance(), 37
- Diesler, Thomas, 67
- Document class, 71
- <document> tag, 26
- Drools rule language, 8
- <duration> tag, 11

- <expr> tag, 26
- Extractor class, 35
- <extractor> tag, 24–26

- FactHandle class, 13, 14
- FAQ, *see* Frequently Asked Questions64
- Fluxtapose, 7
- Forgy, Charles, 59

- Frequently Asked Questions, 64

- Gay, Roger F., 67

- Hald, Martin, 67
- Ho, Matt, 67

- Java Community Process, 71
- JAXP, 16
- JCP, *see* Jva Community Process71
- JNDI, 12, 17
- JSR-94, 71

- Kazmier, Pete, 67
- Klooster, Christiaan ten, 67
- knowledge, 6

- L^AT_EX, v

- McWhirter, Bob, 62, 67
- modification, 13
- modify, 6

- ObjectType class, 33, 34

- Parameter class, 38
- <parameter> tag, 10
- predicate object, 29
- procedural, 5

- Rete, 7, 34, 59, 62
 - network, 60
- Rete-OO, 62
- retract, 6
- retraction, 13
- Roome, James, 67
- Rule, 32
- Rule class, 20, 29, 30, 33, 38
- rule
 - activation, 14
 - firing, 14
- <rule> tag, 9, 27
- rule base, 6
- rule set, 6
- rule-engine, 6
- <rule-set> tag, 9

- RuleBase class, 12, 13, 16–20, 29, 30
 - newWorkingMemory(), 13
- RuleBaseBuilder class, 18
- RuleSet class, 16–18, 29, 30, 38
- RuleSetReader class, 16, 36
- RuleType class, 33

- SaliencyConflictResolutionStrategy class,
 - 17, 20, 21
 - getInstance(), 20
- SAXReader class, 16
- Selders, Bart, 67
- semantic module framework, 32
- SemanticModule class, 36
- SemanticsRepository class, 16, 36
- Serializable class, 18
- SimplicityConflictResolutionStrategy class,
 - 17, 21
- Strachan, James, 67

- Tuple class, 32–34

- Vasak, Tom, 67

- working memory, 6
- WorkingMemory class, 13–15, 30
 - assertObject(...), 13, 14
 - fireAllRules(...), 15
 - modifyObject(...), 13
 - retractObject(...), 13

- XEmacs, v