

Knowledge Builder for Capturing, Maintaining, Deploying Business Rules in eBusiness Systems

A White Paper by Attar Software

What are Business Rules

"Business Rules" is the name used to describe the knowledge underpinning an organisation, allowing it to perform its functions and potentially giving it its competitive advantage. It covers a wide spectrum of knowledge, skills, know-how and expertise which define an organisation. "Business Rules" represents the most important corporate asset and covers the following organisational knowledge:

- Documented procedures, methods & policies
- Compliance with external rules, regulations and legislation (e.g Tax and environmental regulations)
- Employee's know-how and expertise relating to customers, products, services, resources, processes, operations and risks.

The capture and automation of business knowledge can deliver significant competitive advantages to an organisation which include:

- Preserving the expertise of key employees
- Applying business knowledge widely, accurately, consistently & rapidly
- A key enabling technology for ecommerce; selling products and services and supporting clients and partners over the World Wide Web
- The ability to respond fast to business changes made possible by the ease of maintenance of captured knowledge.

Technologies for Capturing & Automating Business Rules

The technology for capturing and delivering business rules has, over the last 20 years, been given many "labels" including "Rule Based Systems", "Knowledge Based Systems" and "Expert Systems". These technologies share one basic design concept, that of separating the business rules from the rest of the application running these rules. The business rules are maintained in a separate file which is run by a generic software component called the "Inference Engine". The Inference Engine is designed to search for decisions based on facts (input data) it receives and by consulting the rules in the rules base file. The rules usually have the IF THEN format as in the example below:

IF Grade is Director
THEN decision is Pass expense claim

IF Grade is Senior Manager
AND Hotel class is A
THEN decision is Reject expense claim

IF Grade is Senior Manager
AND Hotel class is B
AND Department is Accounts
THEN decision is Pass expense claim

The main advantage of separating the rules from the inference engine is the ease of development and maintenance of the business rules. With user friendly rules maintenance tools, the business rules can be developed, reviewed and maintained by the experts and decision makers in a particular area of expertise (e.g risk underwriters or process maintenance engineers). This ease of rules maintenance enables an organisation to respond fast to business changes and is a key enabling technology for ecommerce.

Rules based technologies can suffer from two limitations; the lack of effective tools for capturing knowledge from domain experts and the limited deployment options for the Inference Engine (integration & platforms). The XpertRule Knowledge Builder technology has a number of unique features which overcome these limitations:

An enterprise knowledge development environment which is multi-user, project based and with a graphical knowledge object explorer

Extensive support for the capture & maintenance of rules featuring graphical knowledge representation / structuring, knowledge acquisition features and support for knowledge discovery from historic data

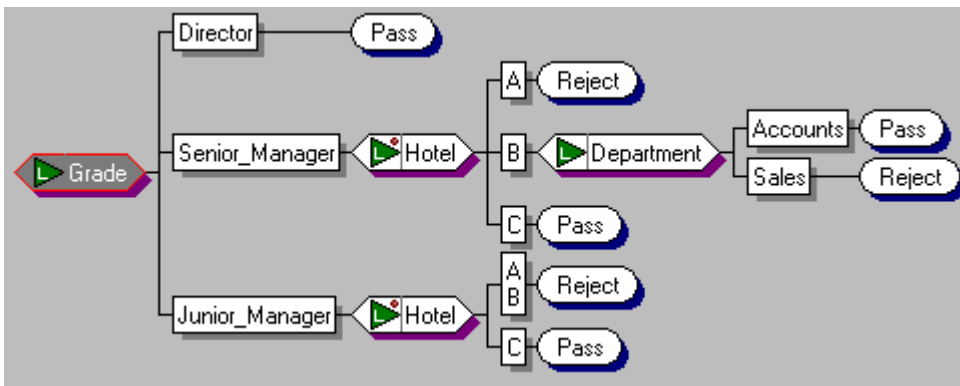
Flexible, Scalable & efficient deployment and integration featuring COM+ and Java rules server, XML data exchange and Thin web client architecture

Graphical Knowledge Representation in Knowledge Builder

In most business rules systems, the knowledge is represented in the following format:

- | | |
|---|---|
| <ol style="list-style-type: none"> 1. IF Grade is Director
THEN decision is Pass 2. IF Grade is Senior Manager
AND Hotel is A
THEN decision is Reject 3. IF Grade is Senior Manager
AND Hotel is B
AND Department is Accounts
THEN decision is Pass 4. IF Grade is Senior Manager
AND Hotel is B
AND Department is Sales
THEN decision is Reject 5. IF Grade is Senior Manager
AND Hotel is C
THEN decision is Pass 6. IF Grade is Junior Manager
AND Hotel is A or B
THEN decision is Reject 7. IF Grade is Junior Manager
AND Hotel is C
THEN decision is pass | <ol style="list-style-type: none"> 8. IF Cost < 50
THEN Hotel is C 9. IF Cost >= 100
THEN Hotel is A 10. IF Cost > 50
AND Cost < 100
AND InLondon is Yes
THEN Hotel is A 11. IF Cost > 50
AND Cost < 100
AND InLondon is No
THEN Hotel is B |
|---|---|

Once the number of such rules exceeds 20 or so, it becomes very difficult to get an overall picture of the knowledge. Furthermore, it is very difficult to assess if a list of such rules is incomplete or if it contains conflicting rules. XpertRule Knowledge Builder in contrast, represents knowledge as set of decision trees as shown below:



Note that the above two decision trees contain the same knowledge as the previous list of rules. However, decision trees are more understandable, give an overall picture and ensures no knowledge gaps or conflicting rules. Decision makers also find it easier to use trees than rules to express / maintain their know-how.


Furthermore, Knowledge Builder assists the developer in maintaining the knowledge structure by displaying the overall hierarchy of decision trees and attributes in a "knowledge map" as illustrated in the following map:



In addition to decision trees, Knowledge Builder also allows experts to express their know-how in tabular rules format called decision cases as illustrated in the example below:

	Grade	Department	Hotel	Expenses
1	Director	*	*	Pass
2	Junior_Manager	Accounts	A	Reject
3	Junior_Manager	Accounts	B	Reject
4	Junior_Manager	Accounts	C	Pass
5	Junior_Manager	Sales	A	Reject
6	Junior_Manager	Sales	B	Reject
7	Junior_Manager	Sales	C	Pass
8	Senior_Manager	Accounts	A	Reject
9	Senior_Manager	Accounts	B	Pass
10	Senior_Manager	Accounts	C	Pass
11	Senior_Manager	Sales	A	Reject
12	Senior_Manager	Sales	B	Reject
13	Senior_Manager	Sales	C	Pass

Finally, Decision Trees in XpertRule Knowledge Builder can also contain "procedures" which are scripting blocks (in VB script or XpertRule format) for calculations, string manipulations and for calling external programs. Below is an example of a tree with an embedded procedure:

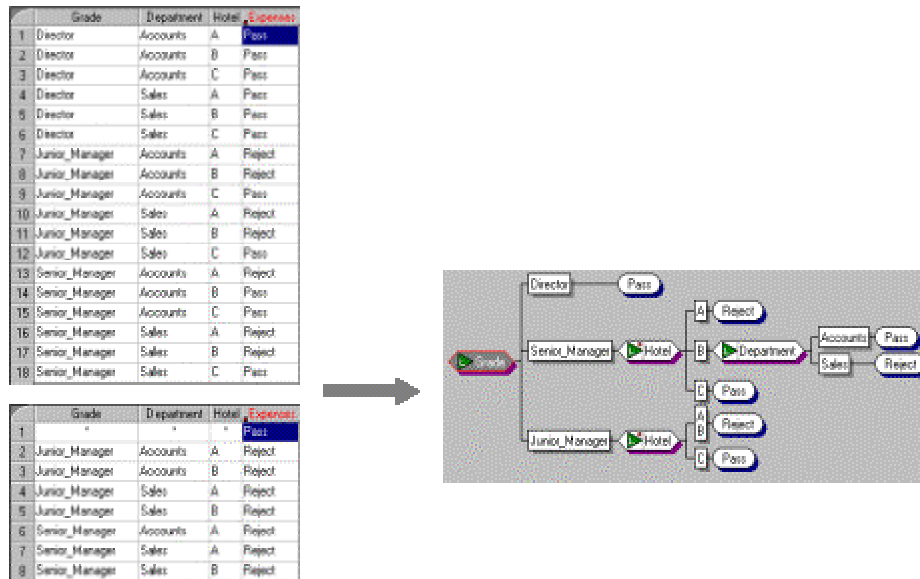


Knowledge Acquisition from Experts in Knowledge Builder

People with expertise gain such expertise through years of experience in applying their skills. Such people can find it difficult to articulate their expertise as a set of rules or a decision tree. Knowledge builder allows experts to express their know-how in one of two easy to express knowledge representations:

- A table of Decision cases representing examples of how they make decisions
- A table of Exception cases from a common decision

Knowledge Builder features a unique tree induction algorithm which automatically derives decision trees from decision or exception cases as shown in the example below:



Knowledge Discovery from data

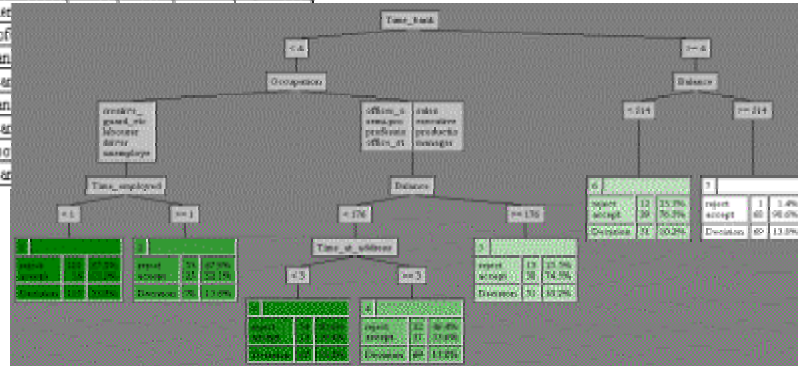
Historic data captured and archived by an organisation can be a source of new knowledge. Two types of knowledge can be discovered from such data;

- Data relating to decision making by experts such as risk underwriters or trouble shooters can be used to discover the rules these experts use in their skilled decision making process.
- Data relating to the performance of various business processes can be used to discover business performance rules which can be used to improve these processes. For example discovering the profiles of bad credit applicants or suspicious insurance claims.

The XpertRule Miner technology uses tree induction from data to generate decision trees with outcomes that have probability figures attached. Such discovered trees can be used as knowledge modules within Knowledge Builder.

The example below shows a decision tree generated from historic data relating to Accept/Reject decisions made by loans underwriters. The data table shows the application data and the decision made. The discovered decision tree reveals 7 leaves (rules or profiles) each with a probability of being rejected:

S	A	Time	ResStat	occup	Time	Time	House	Decision
e	g	Addr			Emp	Bank	Exp	
x	e							
M	50	0.5	owner	unemploye	0	0	00145	reject
M	19	10	rent	labourer	0.8	0	00140	reject
F	52	15	owner	creative	5.5	14	00000	accept
M	22	2.5	rent	creative	2.6	0	00000	accept
M	29	13	owner	driver	0.5	0	00228	reject
F	16	0.3	owner	unemp				
M	23	11	owner	prof				
F	27	3	owner	man				
F	19	5.4	owner	guar				
F	27	0.3	owner	man				
M	34	4	rent	guar				
M	20	1.3	rent	labo				
M	34	1.3	owner	guar				



Deploying Business Rules Applications using Knowledge Builder

In addition to its advanced knowledge representation and acquisition features, Knowledge Builder supports the deployment of knowledge components and applications on Wintel and Linux / Unix platforms with full integration with other applications.

For Windows 98 / 2000 / NT operating system, Knowledge Builder applications are run using a COM plus inference engine (Rules Server) with full ODBC, DLL & COM connectivity. This Rules Server can be used in two modes:

- Interactive Q&A Inference mode with a PC display Client or a thin web client (browser)
- Batch Inference mode whereby the rules server receives data and returns decisions using XML. In this mode the Rules Server provides rules processing for other applications or for a transactions server.

For deployment on Linux / Unix and other Java platforms, Knowledge Builder can generate the business rules and the inference engine as a Java Knowledge Class (source code) which can be compiled to produce high performance Java Rules Server. The Generated Java Knowledge Class can be wrapped as a Servlet or an EJB. As with the Windows deployment, the Java Rules Server can be used in two modes:

- Interactive Q&A Inference mode with a thin web client (browser)
- Batch Inference mode whereby the rules server receives data and returns decisions using XML. In this mode the Rules Server provides rules processing for other applications or for a transactions server

Applications of Knowledge Builder in eCommerce

Customer Need Analysis

In order to recommend the best products and services to potential customers it is important to understand their needs. This gives the customer the confidence that we understand his requirements and will offer a suitable product / service. This task needs knowledge of the customer requirements and how they relate to the Company's products and services. Customer need analysis involves an interactive session of guided Questions and Answers whereby the next question to the customer is driven by his/her profile and previous answers. In certain areas, such as financial services, the line of questioning has to ensure compliance with regulation regarding what information has to be given to customers.

Customer Need Analysis application can either be run over the internet or run by mobile sales force on a laptop computer. A Customer Need Analysis module would normally format the captured data in XML format for passing to other business modules such as "Product Selections" or "Risk underwriting".

Product Selection & Recommendation

Having captured customer needs through a simple form or a knowledge driven sequence of forms, the next step is to recommend to the customer the most suitable products and services matching his requirements. This can involve recommending a single product/service or ranking a number of products/services by suitability (matching the needs). A product selection Rules Server is normally passed an XML string with customer requirements and products features data and it populate this XML string with a list of products ranked by suitability. The Rules Server may also select any optional product features and values (for example which type of a Card and the credit limit on a Current Bank Account).

Credit Risk Underwriting

The automation of the processing of applications for credit, loans or mortgages requires the automation of credit risk assessment. The knowledge for risk assessment can be captured from underwriters and / or from historic data on applications and their subsequent credit performance. Once captured, this knowledge can be deployed as a Rules Server receiving XML applications data and returning risk decisions. Note that credit risk decisions can be Accept, Reject or Refer to manual underwriting. Risk decisions can also include a credit limit figure.

Insurance Underwriting

This covers include Life, Motor and other insurance underwriting application. Application (proposal) data in captured using XML forms and is passed to the Rules Server. The initial decision made by Rules Server is whether the application can be underwritten by the server or if it need referral for manual underwriting. The next decisions is to derive the premium (multiplier) and any special conditions.

Insurance Claims Processing

These applications decide if an insurance claim can be settled or if it should be referred for manual processing. The knowledge involves capturing common problem patterns and norms for various types of claims.

Customer Support and Help desk

Many organisations start their ecommerce initiative with developing front end systems for "e-selling". However, these organisations soon realise that a very import part of ecommerce is e-support, which is the ability to provide systems for helping customers with solving their problems relating to products and services. Such e-support systems need to capture the knowledge required to provide help, trouble shooting, problem solving and support. The knowledge captured in e-support system falls into two categories; diagnostic knowledge to identify the problem and advisory knowledge to recommend a solution. The decision cases, decision trees and tree induction features of Knowledge Builder are ideal for representing diagnostic and advisory knowledge.

Trouble-shooting and advisory systems can be used directly by the customers or by call centre agents to increase their problem solving skills.