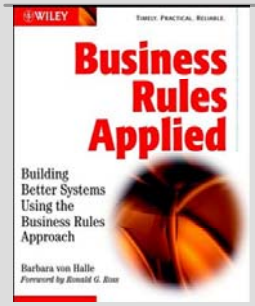

Business Rules Applied: Building Better Systems Using the Business Rules Approach

Barbara von Halle
John Wiley, 2001
$58.50 546 pages

Reviewed by Paul Harmon

This book was published in 2001. It is probably the best business rules methodology book – it describes a very systematic approach to developing business rule systems. On the other hand, the author, Barbara von Halle has just co-authored a new book, *The Decision Model*, with Larry Goldberg, and it suggests that she has shifted her approach to rules development. When I get the new book, I will review it, but I am reviewing this book because it represents a significant step in the development of the business rules approach.

There is a lot of interest in business rules among those engaged in Business Process Management (BPM) at the moment, and there is quite a bit of confusion about how business rules and business processes should be used together. A little history can clarify things a bit.

In the Eighties, organizations began to shift to relational databases. The good news is that relational databases provided organizations with much a much more flexible way of storing their data. Previously databases had been linked to the applications that used them. Thus, different applications maintained different databases and organizations were not in a position to examine all their data in a single database. Relational databases freed data from it's linkage with specific applications and made generic database searches (4th Generation Languages) possible. One of the problems uncovered during the transition to relational databases was where one might store procedures. Procedures, or rules if you would prefer, didn't fit the relational model. Various relational vendors came up with different ways of handling procedure storage – all deviations from the pure relational model.

One group that was particularly active in the struggle to solve the "procedure storage" problem was IBM GUIDE's Business Rules Project. In the middle of the Nineties they issued a report recommending how business rules should be handled. That report led to the establishment of the Business Rules Group, a non-profit group that continued to extend the work begun by the GUIDE project. Most of the individuals involved today in the business rules movement, including Ron Ross and Barbara von Halle, were active in the Business Rules Group.

To summarize, those initially involved in developing relational databases -- in 1994 von Halle co-authored the *Handbook of Data Management* – began to develop the business rules approach as a way of dealing with the procedural rules that were omitted from relational databases. In other words, their interest was in developing an alternative kind of software application that could be used in conjunction with relational databases. A series of software tools were created to manage the storage and use of business rules.

Now let's consider an alternative history. In 1981 Barr and Feigenbaum published the first volume of the *Handbook of AI*. The Handbook described recent work that AI researchers had

undertaken to create expert knowledge-based systems – software systems that could solve problems and make decisions that were otherwise solved or made by human experts. Initially these systems relied on knowledge rules. In other words, an expert's knowledge was acquired by means of interviews and encoded in rules. Expert systems often had thousands of rules and were manipulated, at runtime, by an inference engine, an interpreter that could sort through the rules and assemble those that applied into a decision tree.

It quickly emerged that it was inefficient to store all the knowledge in an expert system as rules, and most expert systems software tools shifted to an architecture that divided the knowledge between a semantic net – an early version of what today we would call an object-oriented database that supports a set of objects each with attributes and values and relationships. (Thus an object like animal, has the child object, bird, which, in turn, has the child object, robin.) By storing most knowledge in an object network and using rules to handle deductions, the expert systems could be made much more efficient.

In the end, however, expert systems proved too hard to maintain. Imagine a system with 10,000 rules that offered a medical diagnosis. A human doctor attends medical conferences, reads journals and constantly improves what he or she knows about a specific type of diagnosis. Any expert system that was to remain really effective would need to be constantly updated. By the mid-Nineties, most companies decided that rule-based knowledge systems were too hard to maintain and shifted to different technologies, including case-based reasoning and neural networks – which we won't pursue here.

Thus, in the mid-Nineties, there were a large number of software vendors who had expert system building software tools that were no longer in demand. At the same time, there was a growing community of business rules enthusiasts who wanted to build rule systems that were, by the standards of those in the expert systems world, very modest applications – systems with a few hundred rules. By the late Nineties the business rules practitioners, coming from the relational database tradition and begun to use the inference-based knowledge system building tools.

In spite of this marriage of convenience, there remained a stark difference in the approach of the two traditions. This difference is well defined by von Halle on page 45 of *Business Rules Applied* and I have represented it with the continuum figure below. (See Figure 1) In essence, the software applications conceptualized by the relational-based rules tradition imagines that the organization has defined clearly stated business policies. From those policies one defines rules, and then sub-rules. Because the rules begin as explicit injunctions, like "Gold Class customers get a 10% discount" the elucidation of rules can proceed in a very logical manner.

Contrast this with expert systems that were concerned with knowledge possessed by human experts. The expert may or may not have thought in terms of explicit statements or rules. In effect, knowledge acquisition involved asking the human expert to describe dozens of specific cases he or she had solved and involved the gradual formalization of the heuristics or rules of thumb that were involved in solving each case. One typically began a knowledge acquisition session by defining the procedure the expert went through to solve a case. In any case, the rules were not "under the control of the company," they were knowledge that had been accumulated in the mind of the human expert.

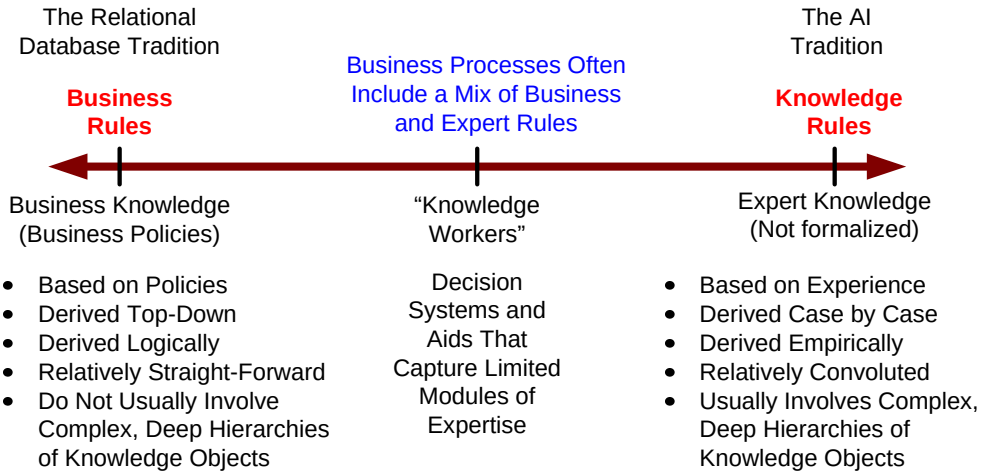

Figure 1. A Continuum from Business Rules to Knowledge Rules

The early books in the business rules/relational database tradition, include Ross's early books on rules and von Halle's *Business Rules Applied*, are aimed at defining the rules used in software applications. The main driver was to separate rules from data and other applications to make them explicit and easy to maintain. Thus, in essence, the rules methodology that von Halle proposed is a software development methodology.

Figure 2 illustrates the overall approach von Halle proposes. In essence, von Halle proposed a parallel development of different systems. Process is conceptualized as workflow, and rule development is defined as its own development process. Von Halle's book is some 546 pages in length and most of it is concerned with defining precise steps and guidelines for each phase.

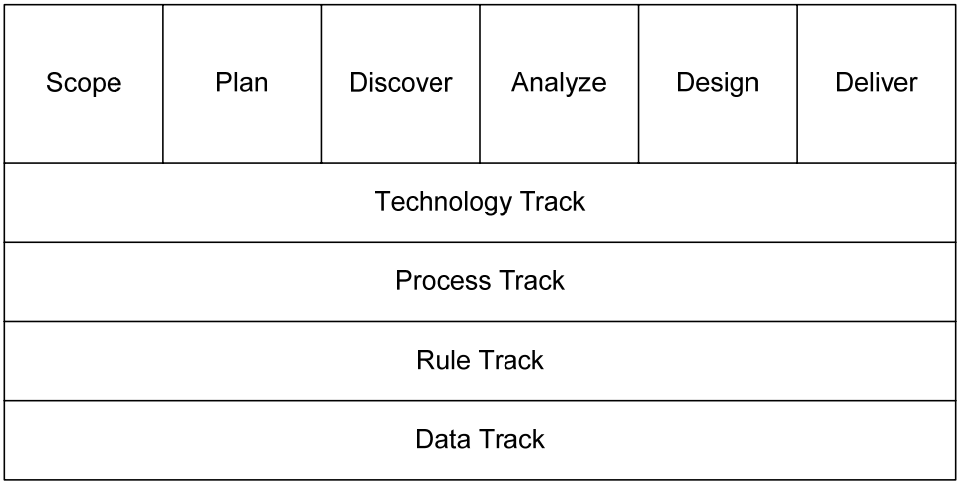

Figure 2. Business rules system methodology phases

To be explicit about von Halle's approach consider the outline for the Discovery Phase. (I have omitted out some detail about case studies which is also in the outline.)

Chapter 6 Discovering Initial Requirements
- What is the discovery of initial requirements?
- How is discovery of initial requirements different for a business rules approach?
- What is the purpose of discovering initial requirements?
- What are the deliverables of discovering initial requirements?
- What are the steps in discovering initial requirements?
  - Step 6.1: Describe the event response process details by creating use-case descriptions
    - Guideline 6.1.1 Simply use whatever technique assists you in soliciting from a business person a reasonable sequence in which the business event may be serviced
  - Step 6.2: Add Concrete Scenarios
  - Step 6.3: Identify Decisions
  - Step 6.4: Complete the conceptual model
    - Guideline 6.4.1 At a minimum, the conceptual models should consist of three deliverables: a conceptual data model, a conceptual process model, and a CRUM matrix showing information (entity), usage (create, update, delete, read, or usage) of the lowest levels of the conceptual process model.

This small portion of a table of contents that runs over 20 pages provides the systematic detail that von Halle provides. At the same time, a glance at the portion above shows that the entire approach is being conceived as a software development effort. Moreover, it is assumed that the process model has somehow been driven from a high level process model down to a very concrete procedural description that will allow, during Discovery, access to the "lowest levels of the conceptual process model."

As I suggested earlier, today there is considerable interest in how Business Process Management relates to Business Rules efforts. The business rules people are talking more and more about how rules can be used in process redesign efforts. I have been reading books like *Business Rules Applied* to see how the two approaches can be combined in an effective manner.

I am concerned with three major BPM tasks: (1) Defining an enterprise business process architecture and ongoing process management approach, 2) Undertaking major process redesign projects, and 3) Supporting continuous process improvement efforts. I would be happy for a business rules methodology that could be used in any of those three contexts.

Before I would consider any rules methodology, I would require that the methodology support the diversity that I include in process. I do not use the term "process" to refer to a workflow system, or to any kind of software system. I am increasingly focused on high-level business processes – processes like cross-organizational supply chain systems, or customer service systems that interface with customers throughout their stay at a hotel. I examine existing processes to see how they can be improved. If I identify problems, I consider how I might change the existing process – and am as happy to achieve my goals by changing how managers or employees perform their tasks as I am to consider changing or developing software systems to improve the processes performance.

At the highest level I do not focus on decisions very much, but as I drill down to Level 3 or Level 4 processes I find that lots of subprocesses or activities consist of analysis and decision tasks. In those cases, it is often most convenient to define the tasks by means of rules. If the task involves human performers, then the rules will probably end up embedded in job aids, guidebooks, or

training classes. If the task can be supported by a software application, then I am happy to consider developing a rule-based system. Any rule methodology that only focused on developing a rule-based software system is of limited use to me, because it has made a decision about the nature of the solution without the benefit of studying the actual process to be improved. Similarly any rule management system that is, in fact, a management system designed to manage rule-based software systems is only concerned with managing a portion of the rules used by the organization, and is similarly limited. Finally, there is no consideration of the fact that when I do a redesign project, I do not necessarily develop a comprehensive rule system, but only discover rules that are used in decisions in the subprocesses I am focused on. In essence, a rule approach that worked well with process redesign would have to provide insights into how to work backward, from specific rules to a broader rule framework.

Business rules constitute an element within the broader business process picture. Some companies are completely focused on rules (mostly financial organizations and government agencies) and on capturing and automating those rules for those organizations, the current Business Rules approach, and von Halle's methodology will probably work reasonably well. For those of use involved in the broader movement toward Process Based Management, however, a rule automation approach that is implicit in von Halle's methodology is too narrow and too constraining. We really need a rule methodology that is integrated within a business process redesign methodology that we can access when we need it – when we discover decision points – and that we can then apply to either human performance problem or automation problems.

As I suggested earlier, von Halle has already moved on and has just written a new book that I am eager to study.

Meanwhile, I don't know of any other business rule methodology book that is as complete or as well thought out as *Business Rules Applied.* This is a very process-oriented approach to how one creates and installs a business rule system.

I certainly recommend it to anyone who is trying to sort through the issues of how one integrates rules and business processes redesign. The methodology does not provide the flexibility that today's BPM practitioners require, but there are no others available that do a better job, and it contains much detail and wisdom that will apply in whatever approach eventually evolves to support process work.