

## Informavores Spark Studio

An InDetail paper by Bloor Research  
Author : Philip Howard  
Publish date : April 2007

... we have no hesitation in recommending the product and we urge any companies with relevant applications to seriously consider the Fire products

**Philip Howard**

Informavores offers three products: Firebox, Firefly and Firestorm, each of which is based on the same technology but which are targeted at different application environments. That technology is not easy to describe because we have not seen any other product with comparable capabilities.

On the one hand, Informavores provides rules-engine capability. However, it does not work like any other rules engine we have seen: with Informavores you build a tree or flow diagram that describes the process you are interested in and then the software generates all the relevant rules (the application if you will) in the background. In effect, this is a model-driven environment except that, in this case, it is the tree diagram (rather than the model) that is the application.

On the other hand, Informavores is also a knowledge management application that can be used for knowledge capture and transfer, in that you can capture scripts and other details within your tree (flow) diagram to lead call centre operatives (say) through a given procedure.

In so far as Informavores' individual products are concerned, Firebox is aimed at supporting on-line mortgage or insurance quotations and similar applications; Firefly is targeted at call centres, help desks, support, telesales and so forth; and Firestorm is used when you want to do mass tests against a rule or rules, such as in VAT audits or in ETL (extract, transform and load) environments.

## Key findings

In the opinion of Bloor Research the following represent the key facts of which prospective users should be aware:

- All of Informavores' products are based on the same underlying technology, which uses a model-driven approach to application development, except that, in this case, it is not models that drive the development but decision trees or flow charts. When using any of the Fire products you build your decision tree (or flow chart) and that is, to all intents and purposes, the application, which is automatically generated from the decision tree.
- This development process is the easiest we have seen from any company within the rules engine space or, indeed, many other environments.
- While model-driven development is much faster than conventional development it is still a technology for developers. Informavores' approach, on the other hand, is aimed directly at business analysts and users and not at developers: it is thus both faster (typically by an order of magnitude) and easier to use.
- Informavores' products are suitable for building a wide range of applications and we are sure that there are use cases where it would be appropriate that neither we nor Informavores has yet thought of. That said, the Fire products are not general purpose tools: in particular they will not be appropriate for data-centric applications.
- A substantial part of Informavores' target market is within the financial services space. The product includes audit capabilities that meet the requirements of the Financial Services Authority.
- Automated document generation is very extensive and is much more complete than in many other products.
- There are substantial capabilities for capturing expert knowledge to build walkthroughs and similar facilities.
- Informavores supports the development of offline as well as online applications.

## The bottom line

We have not seen any other products quite like those of Informavores: we have seen model-driven development tools but they are not really suitable for business users, we have seen rules engines but they are similarly not amenable to development by end users, and we have seen knowledge capture tools but these typically exist in a vacuum. Informavores, on the other hand, has developed a technology that is easy to use, highly productive, clearly focused and is especially good at handling complex business rules: we are very, very impressed.

## Background information

Informavores' technology derived from work that was previously done at CERN but the company itself was founded in 2003 and it brought its first product to market two years later. Since then it has gained clients such as BT, NHS, GE, AMD, National Westminster Bank, Norwich Union and the US Department of Defence (DoD), which is impressive for such a young company.

The company is privately owned and funded, with 15 employees and sales partners, and it is based in Wales, though it also has an office in the United States and has partners in other parts of the English-speaking world. That said, the product is not limited to working in English as a language and the company has customers using it, for example, in Hebrew. The company has technical partnerships with Sun, Microsoft and Salesforce.com (the company is part of the AppExchange programme) amongst others. Typical implementations range from £15,000 to £50,000.

Web address: [www.informavores.com](http://www.informavores.com)

## Product availability

Informavores uses Microsoft-style product numbering so Firefly, Firebox and Firestorm are all currently in their 2007 versions. They run on standard platforms such as Linux Red Hat 9, Fedora 9, Solaris, AIX, Windows 2000, Windows XP, Windows 2003, Vista, Mac OSX 10.3 and up.

The decision trees that you build (known as 'Sparks') require the use of a database to store them in (they are stored in XML format) as does the Audit database, which is required for certain environments. In principle any JDBC compliant database will be suitable though Informavores' experience with Oracle means that the company does not recommend it.

If version control is required then the Informavores products will integrate with CVS or Visual SourceSafe.

## Architecture

The overall architecture of the various Informavores products work is illustrated in Figure 1. Firebox provides the base functionality shown in the middle column, Firefly adds the interactive guide and firestorm adds the ability to run rules against data. We will discuss the various elements shown, in turn.

## The Decision Tree

The first thing that you do with Informavores (regardless of the product) is to build a decision tree or Spark, using conventional graphical, drag-and-drop techniques, an example of which is illustrated in Figure 2, using the Spark Studio Designer.

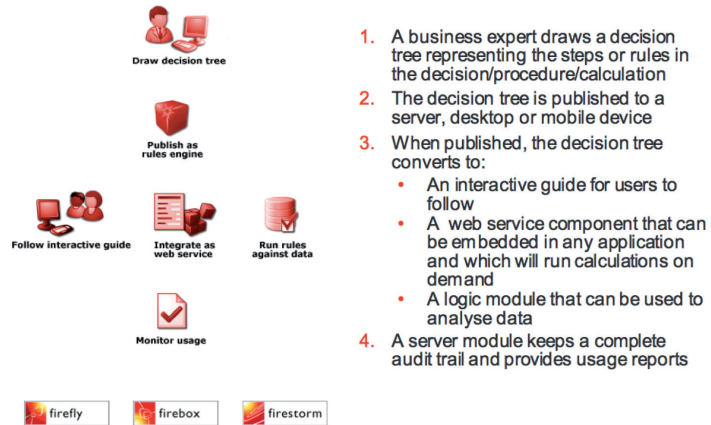
As can be seen the result is a decision tree. Note that each node in the tree may be expanded into a dialogue box (see later) where required, as opposed to using icons.

The important thing to note here is that these Sparks represent complex sets of business logic and, indeed, the software is particularly effective when used with complex sets of rules.

This Spark is your application. That is to say, once it is defined you have completed all the development work you need to do, other than testing, because, at this point, you simply generate the required application from the Spark. In technical terms, the product offers bi-directional synchronisation, when means that any change to the Spark will be reflected in the application and vice versa.

In so far as testing is concerned there is an automated testing tool built into the product, as well as a step-through debugger and there are data integration capabilities provided for accessing external data sources either for retrieval or update. The software itself ensures consistency: that is, it will not allow you to build a Spark in which there are inconsistent decisions.

Finally, application documentation is also generated automatically from the Spark and this includes both the procedural and functional specification. This is unusual (typically you get one or the other but not both) and indicates the power of the product. Documentation can be both printed and published to a web site. As an example, in the case of the company's integration with Salesforce.com [and also Siebel [Oracle], Microsoft CRM, SAP and others] the documentation generated might include automated sales documentation, product quotations, service diagnostics and, when used in the Firefly product (see later), sales scripts.



1. A business expert draws a decision tree representing the steps or rules in the decision/procedure/calculation
2. The decision tree is published to a server, desktop or mobile device
3. When published, the decision tree converts to:
  - An interactive guide for users to follow
  - A web service component that can be embedded in any application and which will run calculations on demand
  - A logic module that can be used to analyse data
4. A server module keeps a complete audit trail and provides usage reports

Figure 1: Overview of how Informavores works

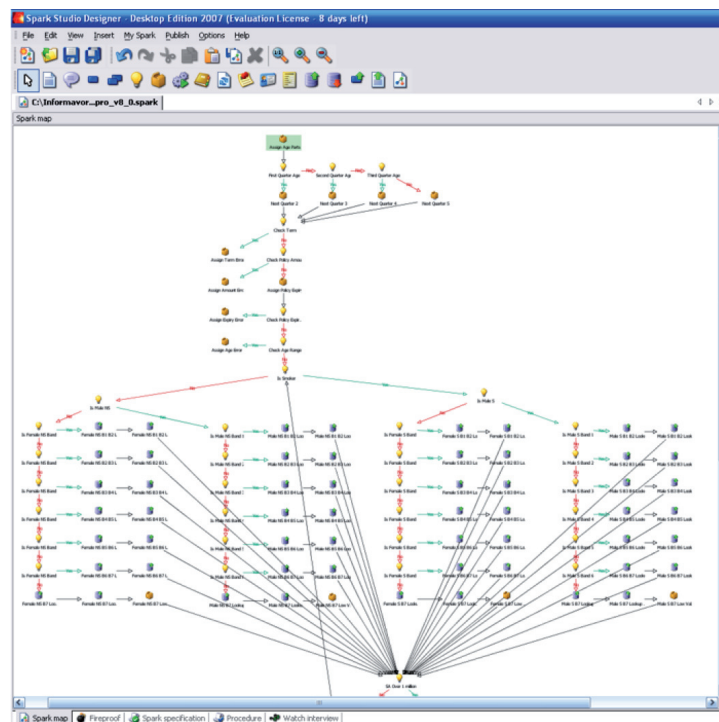


Figure 2: A decision tree in Spark Studio Designer

## Firebox

Firebox is the company's product aimed at on-line applications for, say, insurance quotes. In this case, what typically happens is that a web form is automatically generated by Firebox from the defined Spark (using web services). AJAX (asynchronous Java and XML), which will allow dynamic updating of web forms, is not yet supported but the company is working on such support for a future release.

However, while not all such applications are necessarily within the financial services sector, most of them are and, as such, these are regulated by (in the UK) the Financial Services Authority (FSA). This body mandates that a detailed and time stamped audit trail of all customer entered information is maintained, which is where Firebox's Audit database comes into its own, as it complies with the FSA's rules.

However, the Audit database is not limited to supporting an audit trail; it is also used for the analysis of information about user behaviour so that, for example, you can see the most commonly used path through the decision tree, or where people spent the most time within the decision tree, or where most people dropped out of the process. There is also an option to export data to Microsoft Excel for further analysis.

## Firefly

Firefly is aimed at help desks, contact centres, telesales and so forth, where the user needs to be guided through a script that will vary depending on the responses provided. In this case, the Spark will typically be built using a flow diagram using a graphical interface that looks similar to Microsoft Visio, as in Figure 3 (where it is being used for land planning decisions for aerials), rather than a decision tree.

Note that in this sort of environment the actual text that you capture is a form of knowledge management application to the extent that you are capturing expert information and then allowing its use by non-experts. In appropriate environments, this potential for knowledge transfer could be valuable as an application in its own right.

## Firestorm

Informavores' third product is Firestorm. Here the target market is for environments where a set of rules needs to be run against a (large) database. For example, one of the problems with traditional ETL (extract, transform and load) tools is that business users (as opposed to developers) cannot easily understand how the said tools extract information from the source databases. Firestorm can resolve this

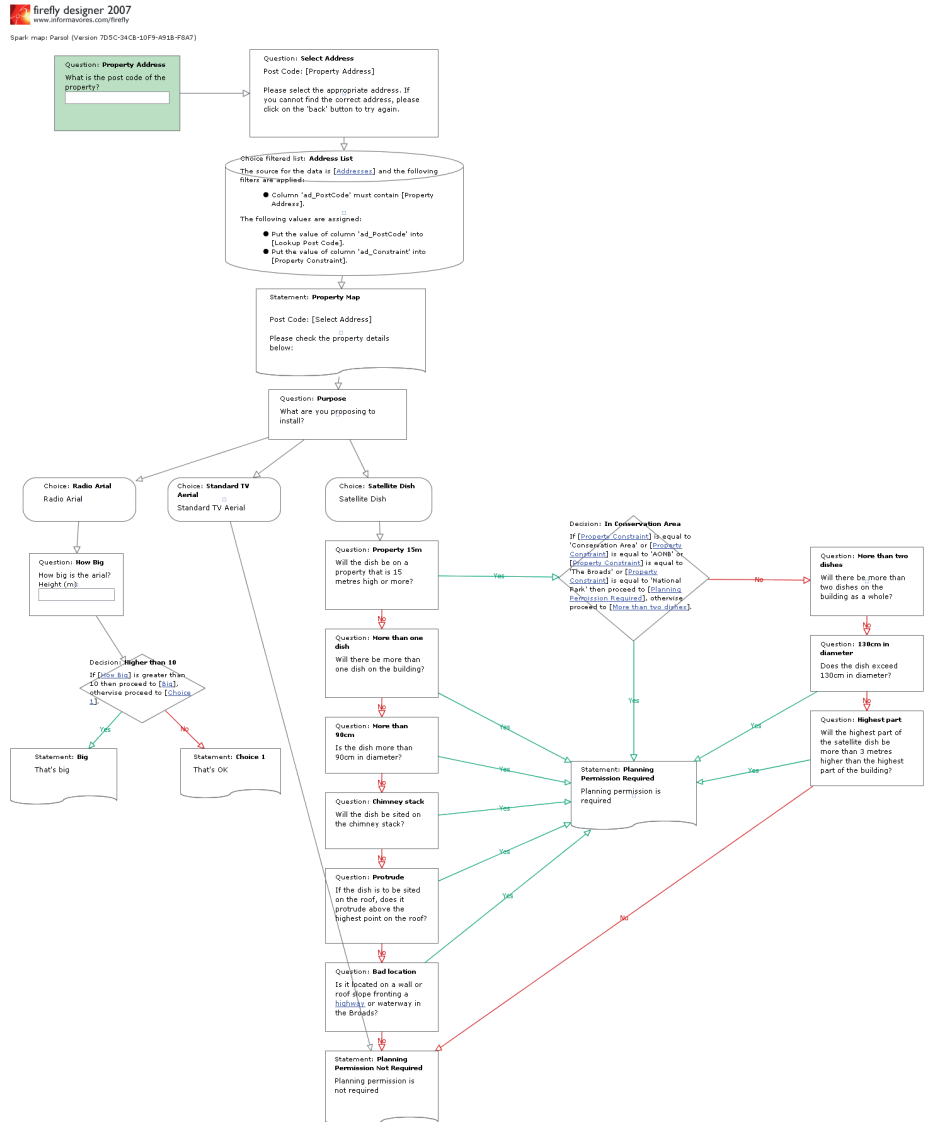


Figure 3: Building a Spark in a similar way to the Microsoft Visio GUI

problem by allowing a business analyst to define the rules (using a decision tree Spark) that source data must meet for it to be extracted.

Another example, of the use of Firestorm is in VAT audits. For example, PriceWaterhouseCoopers has built just such an application, using Firestorm, which audits SAP-based data against international VAT rules. Bear in mind that such rules change on a regular basis and that a manual approach to such auditing would be both time consuming and error prone. Once the comparison of the data with rules is complete, exception reports are generated by the software.

## Summary

It should be clear that we believe that Informavores has got a lot going for it. This is borne out by the blue chip client list that the company has acquired during its short time in business. As a result, we have no hesitation in recommending the product and we urge any companies with relevant applications to seriously consider the Fire products.

Bloor Research has spent the last decade developing what is recognised as Europe's leading independent IT research organisation. With its core research activities underpinning a range of services, from research and consulting to events and publishing, Bloor Research is committed to turning knowledge into client value across all of its products and engagements. Our objectives are:

- Save clients' time by providing comparison and analysis that is clear and succinct.
- Update clients' expertise, enabling them to have a clear understanding of IT issues and facts and validate existing technology strategies.
- Bring an independent perspective, minimising the inherent risks of product selection and decision-making.
- Communicate our visionary perspective of the future of IT.

Founded in 1989, Bloor Research is one of the world's leading IT research, analysis and consultancy organisations—distributing research and analysis to IT user and vendor organisations throughout the world via online subscriptions, tailored research services and consultancy projects.



**Philip Howard**  
Research Director - Data

Philip started in the computer industry way back in 1973 and has variously worked as a systems analyst, programmer and salesperson, as well as in marketing and product management, for a variety of companies including GEC Marconi, GPT, Philips Data Systems, Raytheon and NCR.

After a quarter of a century of not being his own boss Philip set up what is now P3ST (Wordsmiths) Ltd in 1992 and his first client was Bloor Research (then ButlerBloor), with Philip working for the company as an associate analyst. His relationship with Bloor Research has continued since that time and he is now Research Director. His practice area encompasses anything to do with data and content and he has five further analysts working with him in this area. While maintaining an overview of the whole space Philip himself specialises in databases, data management, data integration, data quality, data federation, master data management, data governance and data warehousing. He also has an interest in event stream/complex event processing.

In addition to the numerous reports Philip has written on behalf of Bloor Research, Philip also contributes regularly to [www.IT-Director.com](http://www.IT-Director.com) and [www.IT-Analysis.com](http://www.IT-Analysis.com) and was previously the editor of both "Application Development News" and "Operating System News" on behalf of Cambridge Market Intelligence (CMI). He has also contributed to various magazines and published a number of reports published by companies such as CMI and The Financial Times.

Away from work, Philip's primary leisure activities are canal boats, skiing, playing Bridge (at which he is a Life Master) and walking the dog.

## Copyright & disclaimer

This document is subject to copyright. No part of this publication may be reproduced by any method whatsoever without the prior consent of Bloor Research.

Due to the nature of this material, numerous hardware and software products have been mentioned by name. In the majority, if not all, of the cases, these product names are claimed as trademarks by the companies that manufacture the products. It is not Bloor Research's intent to claim these names or trademarks as our own.

Whilst every care has been taken in the preparation of this document to ensure that the information is correct, the publishers cannot accept responsibility for any errors or omissions.



Suite 4, Town Hall,  
86 Watling Street East  
TOWCESTER,  
Northamptonshire,  
NN12 6BS, United Kingdom

Tel: +44 (0)870 345 9911  
Fax: +44 (0)870 345 9922  
Web: [www.bloor-research.com](http://www.bloor-research.com)  
email: [info@bloor-research.com](mailto:info@bloor-research.com)