DECISION MANAGEMENT SOLUTIONS

Decision Management Systems Platform Technologies Report

Version 2, Update 1, August 9, 2012

James Taylor
CEO

Organizations are adopting a new class of operational systems called Decision Management Systems to meet the demands of consumers, regulators and markets because traditional systems are too inflexible, fail to learn and adapt and crucially cannot apply analytics to take advantage of "Big Data."

Decision Management Systems are agile, analytic and adaptive. They are agile so they can be rapidly changed to cope with new regulations or business conditions. They are analytic, putting an organization's data to work improving the quality and effectiveness of decisions. They are adaptive, learning from what works and what does not work to continuously improve over time.

Decision Management Systems are built by focusing on the repeatable, operational decisions that impact individual transactions or customers. Once these decisions are discovered and modeled, decision services are built that embody the organization's preferred decision-making approach in operational software components. The performance of these components, and the impact of this performance on overall organizational performance, is tracked, analyzed and fed back into improving the effectiveness of decision-making.

Decision Management Systems offer high ROI because they improve the management of risk and the matching of price to risk; because they reduce or eliminate fraud and waste; because they increase revenue by making the most of every opportunity; and because they improve the utilization of constrained resources across the organization.

Decision Management Systems are different from traditional enterprise applications and from business process or event-based systems. Established

CONTENTS

approaches and technologies play a role in the development of Decision Management Systems. Used alone, however, these technologies and approaches tend to deliver systems that are inflexible, static and opaque. To fulfill the promise of agile and adaptive systems that fully leverage "big data", organizations will need to expand their enterprise architecture to include capabilities from the proven technologies described in this report. Tested and established in many industries, technologies suitable for developing Decision Management Systems include Business Rules Management Systems, data mining or Predictive Analytic Workbenches and Optimization suites as well as new in-database analytic infrastructure and more. Organizations will need to select those that have the capabilities they need, that demonstrate Decision Management best practices and that fit the organization's architecture and use cases.

This report describes these product categories and identifies the key capabilities of these technologies. Best practices in their use and key use cases are identified and discussed. A complete list of vendors in the market is provided and an appendix provides more detail on Decision Management Systems.

Navigating this report

If you are new to Decision Management Systems, the *Appendix - Decision Management Systems* provides an overview of this class of systems. Also, the section on *Use Cases* will illustrate some potential uses of Decision Management Systems and the author's recent book, Decision Management Systems: A Practical Guide to Using Business Rules and Predictive Analytics (IBM Press, 2012) has a more complete list. Those familiar with Decision Management Systems and considering new technology choices as part of building their own should begin with *Decision Management Systems Platform Capabilities* and then move to *Vendors* to find candidate vendors to consider. The section on *Key Characteristics* as well as the section *Selecting Vendors* should also be helpful. Whether familiar or not with Decision Management Systems, the section *Best Practices in Decision Management Systems* is worth reading before any new project.

Those looking for an overview should read *Decision Management Systems Platform Capabilities* and then skip the detailed capability chapters, moving directly to *Key Characteristics*, *Best Practices in Decision Management Systems* and *Use Cases*.

Scope

This report is prompted by the growing interest by organizations in Decision Management Systems. It is always challenging to draw the boundaries around such an exciting and growing area, but for practical purposes it must be so. For this report, we are focused on platform technologies used to build custom Decision Management Systems and our goal is to be comprehensive within this scope.

Many vendors have developed powerful pre-configured Decision Management Systems focused on solving specific decision problems such as loan underwriting, claims handling or cross-channel marketing. For many organizations these solutions are ideal but they are not the focus of this report. Similarly there are vendors that build custom Decision Management Systems for their customers and that have developed powerful platforms for doing so. If such a platform is not for sale to those building their own solutions then it is out of scope for this report.

In both these scenarios the report's discussions of what kinds of functionality is useful, best practices and characteristics for suitable products may well be useful in the selection of vendors but some interpretation will be necessary. For instance, when evaluating pre-configured Decision Management Systems, a discussion of business rule management capabilities may well be relevant while one about connecting to data sources may not be.

If you know of other products you believe should be included or have other feedback, please let us know by sending us an email info@decisionmanagementsolutions.com.

Current Version

This is the second version of this report. As such several sections have been enhanced and extended since the first version and others will be enhanced and extended in future versions over the coming months.

In this version the sections on the key characteristics of the technologies and best practices have been expanded to help you select and successfully adopt technologies for building Decision Management Systems. Like this paragraph, these new sections show an edit mark on the right hand side to help you identify them. Some additional guidance on navigating the various product categories available in the market has also been added.

In addition this version lists the products offered by each vendor that are relevant to building Decision Management Systems along with a link to the most recent "First Look" description of the product. First Looks are also posted to www.JTonEDM.com as they are completed.

Future versions will add more material on use cases as well as expanding the section on criteria for selecting suitable vendors. Vendors and products within the report scope will be added on an ongoing basis.

Each new version of the report will be made available at www.decisionmanagementsolutions.com/decision-management-technology.

Sharing this report

Decision Management Systems Platform Capabilities

Overview

Four aspects of building a Decision Management System drive organizations to adopt new, Decision Management System specific, technologies:

- Managing decision logic for transparency and agility
- Embedding predictive analytics for analytic decision-making
- Optimizing results given real-world trade-offs and simulating results
- Monitoring and improving decision-making over time

In this section we will introduce these four capabilities and put them in a broader context. Subsequent sections will describe the capabilities in more detail.

Managing Decision Logic

Like all information systems, Decision Management Systems require the definition of the logic to be applied during operations. In Decision Management Systems this logic is primarily that of decision-making—how a particular decision should be made given the systems understanding of the current situation. Decision Management Systems must be more agile than traditional information systems, however, so this logic cannot be managed as code. The use of code to define decision-making logic makes that logic opaque to those on the business side that understand how the decision should be made. It also makes it hard to record exactly how a decision was made as recording exactly what code was executed is often problematic. To manage logic in this way most organizations will adopt a Business Rules Management System or a product that contains equivalent functionality.

> Decision Management Systems require that decision logic is managed in a way that delivers design transparency, so it is clear how the decision will be made, and execution transparency, so it is clear how each specific decision was made.

Embedding Predictive Analytics

The management of decision logic is a foundational capability for Decision Management Systems. Most Decision Management Systems should also take advantage of the information available to an organization to improve the accuracy and effectiveness of each decision. Unlike human decision-makers, Decision Management Systems cannot use visualization and reporting technologies to understand the available information. In addition, while people have a great ability to extrapolate from information about the past to see what might happen in the future, systems treat data very literally.

To maximize the value of available information in terms of improved decision-making, Decision Management Systems must therefore embed predictive analytic models derived from historical data using mathematical techniques. Such models make assessments of the likelihood that something will be true in the future and make this assessment available to the decision logic in a Decision Management System, allowing decisions to be made in this context.

This shift from presenting data to humans, so that they can derive insight from it, to explicitly embedding analytic insight in systems using predictive analytic techniques means that organizations will need to adopt additional technologies to analyze their data.

Specifically they will need to adopt a Predictive Analytic Workbench or equivalent functionality. They may also choose to adopt additional analytic infrastructure.

Optimization and Simulation

Many decisions rely on resources that are not unlimited. Whether these resources are staff, product inventory or service capacity, decisions must often be made in the context of a constrained set of resources. Organizations will generally want to optimize their results given these constraints and this means that trade-offs must be made. Organizations will need to adopt optimization and simulation technologies to manage trade-offs and to ensure that decisions are made in a way that produces the best possible results given the constraints on decision-making. These technologies allow modeling of the constraints and trade-offs and then use mathematical techniques to pick the set of outcomes that will maximize the benefit to an organization. These models can also be used to drive simulations of various scenarios to see which will produce the best outcome for the organization.

An organization that is established in developing Decision Management Systems will ultimately adopt technologies for all these capabilities—decision logic management, predictive analytic insight and optimization and simulation. Some will find it useful to have more than one product with the same kind of capability, some will standardize on a single product. The products do not need to be adopted all at once and some Decision Management Systems require only some of the capabilities.

Monitoring Decisions

The nature of decision-making is that it is often not possible to tell how good a decision will turn out to be for some time. As a result the ongoing monitoring of decisions made and their outcomes is really important. Such monitoring allows for decision-making to be systematically improved over time both by tracking decision performance and making changes when this performance is inadequate and by conducting experiments and analyzing the results of these experiments. Most organizations will find that they will use their existing Performance Management and data infrastructure to conduct much of this analysis. However, the use of the decision logic and predictive analytic capabilities discussed above will also be necessary. This will allow for the explicit logging of decision-making approaches and outcomes as well as allowing for easy management of experiments in decision-making. In general this ongoing decision analysis requires design decisions and integration with existing infrastructure rather than additional technologies.

Overall Architecture

These capabilities come together in an overall platform for building Decision Management Systems as shown in Figure 1: The capabilities of a Platform for Decision Management Systems below.

Decision Logic capabilities allow for the editing of the business rules that represent the decision logic. These business rules are deployed to a Decision Service for execution.

Predictive Analytic capabilities allow for data to be analyzed and turned into either additional business rules (representing what has worked in the past and is likely to work in the future) or predictive analytic models that can be deployed either to a Decision Service or to the operational datastore being used by the Decision Service.

Figure 1: The capabilities of a Platform for Decision Management Systems

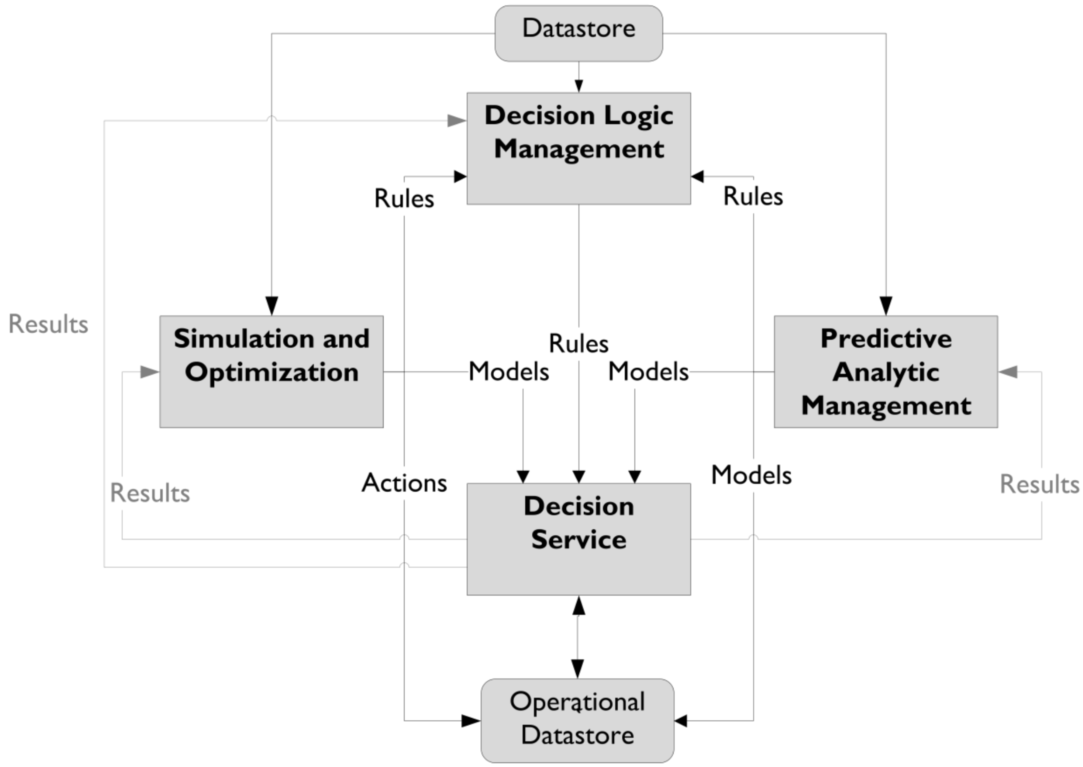

Simulation and optimization capabilities are used to manage tradeoffs and constraints and can result in business rules that have been optimized, optimization models that can be solved in a Decision Service, or an explicit set of actions to be taken that can be pushed into an operational datastore to drive behavior.

All three sets of capabilities rely on data infrastructure to deliver test and historical data while the predictive analytic capabilities can take advantage also of in-database modeling and scoring. The Decision Service itself can execute business rules, score records using predictive analytic models, solve constraint optimization problems and potentially tune predictive analytic models to improve their predictive power while in use. All the capabilities can consume actuals information generated by a Decision Service's ability to log its decision making.

Vendors with all these capabilities could produce a product that provides all four capabilities in a single, integrated environment. However, because the capabilities discussed above can be used for more than building Decision Management Systems, it is likely that some vendors will continue to package up each capability as a separate product while integrating them ever more tightly to make it easier to use them as a set. Other vendors will remain focused on a specific area of capability and will work with partners and standards organizations to ensure that other capabilities can be integrated with theirs. What organizations will need to build Decision Management Systems is an ability to manage decision logic, create and embed predictive analytic insight and simulate/optimize outcomes. How best to assemble this functionality will be different for different organizations.

Decision Services

All the various deployment capabilities described above result in code or packages of definition being deployed to a Decision Service execution environment. This is typically a conceptual environment being in practice made up of elements of multiple products. This environment needs to be able to execute various elements, typically when invoked using a standard API. The decision itself is made by executing generated code on the underlying platform, business rules on a deployed business rule engine, optimization models on a solver and predictive analytic models on a model execution engine.

Decision Services also need to be able to log what happened each time a decision was made—which rules fired, what model scores were calculated and which outcomes were selected by the optimization model. Again this logging often involves elements of multiple products but conceptually a single log can be generated.

Finally model tuning may be available in the decision service, with a piece of analytic modeling code deployed to monitor the performance of deployed models and conduct experiments to see how the predictive performance of those models may be improved.

Decision Services in Context

These technologies are used to develop and manage the decision logic, predictive analytic insight and optimization models required by a Decision Management System and to deploy this to a Decision Service. Decision Services operate in a broader enterprise IT context, however, as shown in Figure 2: Decision Services and Decision Analysis in a broader architectural context below.

Decision Services are invoked for decision-making in an application context. This application context is increasingly a process being managed by a Business Process Management System. Decision Services can also be invoked by enterprise applications both packaged and legacy. While this is less common than invocation from a business process it is by no means an uncommon pattern. A further growing pattern is the use of a Decision Service to support an event processing context, making a decision in response to a pattern of business events and then kicking off a business process or other service as a result. In each scenario, the application context fulfils an overall business need and the invoked Decision Service improves effectiveness and efficiency.

Decision Services are not stand-alone systems that run on specialist hardware or unique platforms. Instead they run on the standard enterprise platforms in use today. Different technologies support the various application servers, service-oriented platforms and programming metaphors that are common and Decision Services can be developed that run on any such platform.

Decision Services rely also on a modern data infrastructure. This data infrastructure supplies operational data to the Decision Services and can also provide in-database analytic scoring. Business intelligence capabilities typically use this same data infrastructure provide insight to human decision makers. While this is not needed to develop Decision Services, business intelligence often complements these systems by supporting the handling of exceptions.

Figure 2: Decision Services and Decision Analysis in a broader architectural context

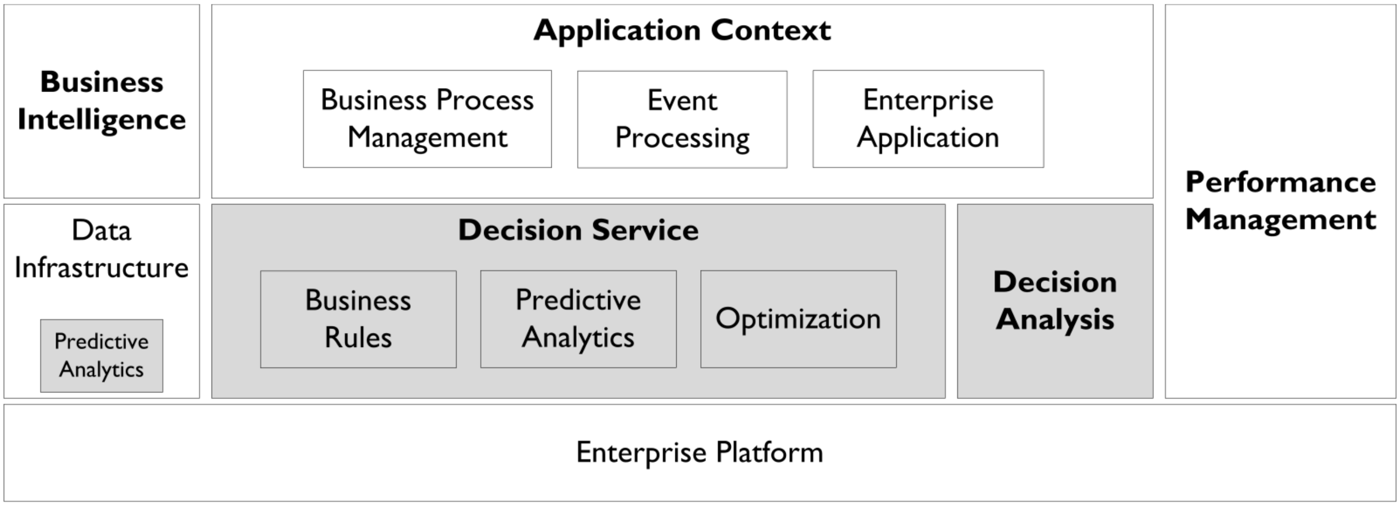

Decision Services may also use business intelligence infrastructure to provide additional context when returning a number of options.

Finally the performance of Decision Services must be monitored to support continuous improvement. This ongoing decision analysis requires both capabilities within the Decision Service, such as support for both a champion and a challenger approach to a single decision, and integration with more typical business or corporate performance management capabilities such as dashboards and alerting systems.

Navigating Product Categories

Multiple product categories exist in the market for decision logic management, predictive analytic modeling, optimization and simulation. As Figure 3: Overlapping product categories below shows, these product categories often overlap.

Figure 3: Overlapping product categories

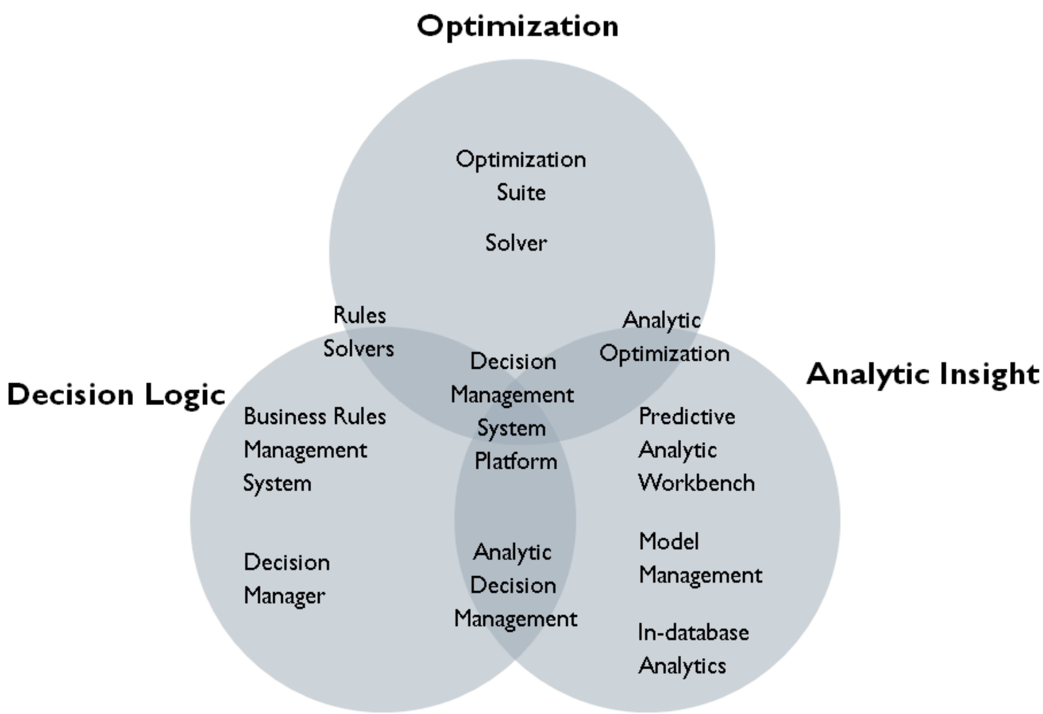

While there are many Business Rules Management Systems that just manage decision logic, there are also products that combine the management of decision logic with optimization or with building predictive analytic models. There are also products that are called Decision Managers or Business Decision Management Systems that manage decision logic and other Decision Management products that manage decision logic and build predictive analytic models. Some Predictive Analytic

Workbenches include in-database scoring capabilities, some package this separately while model monitoring and tuning is similarly sometimes packaged separately.

Navigating products for managing decision logic

When considering products for managing decision logic, there are two main areas of potential confusion—is decision logic the primary focus and is the product rules-centric or decision-centric?

- Decision Logic as a primary focus.
  The first is the degree to which the product is explicitly focused on managing all the logic for a potentially complex decision rather than managing some decision logic so that it can be combined with some analytic insight. For instance a number of products focused primarily on building and deploying analytic models also allow you to manage some business rules. These are typically either focused on eligibility or cut-offs. Eligibility rules might select a subset of all possible records in a set before applying analytic models to them or determine that only certain outcomes are allowed for a given record regardless of what the model may predict. Cut off rules generally turn predictive scores into simple actions based on clearly defined values.
  Such capabilities are much to be desired in analytic products but they will not allow an organization to manage the number of business rules involved, for instance, in complex eligibility decisions. Because they assume that analytics are at the core of a decision they are also not they likely to be effective if used to manage decision logic for decisions driven entirely by policy, regulation and best practice that therefore have no analytic component. For these decisions a product that is either primarily focused on decision logic or that regards explicit logic and analytic insight as peers in decision-making will be more appropriate. Such products are more likely to be referred to as a Business Rules Management System or a Decision Manager. Some Decision Management platforms treat the two as peers while products focused on Analytical Decision Management are more likely to be focused first and foremost on analytics.
- Rules-centric or decision-centric.
  The second is a more nuanced consideration. Some products for managing decision logic began as expert systems or focused on the management of business rules. While users of these systems always used them to automate decisions, this was often *implicit* in the deployment of the rules rather than *explicit* in the design. Other products have begun with a focus on business process and added business rules capability and similarly evolved towards a more decision-centric focus. Such products used to refer to themselves as business rules engines then as Business Rules Management Systems and now, increasingly, as Decision Management products.
  In contrast other tools have begun with an explicit focus on decisions. These typically allow a decision to be documented as such, including its inputs and

outputs, before the logic of the decision is specified to fill the gap between inputs and outputs. These have always used "decision" in their names and are often called Decision Managers or use Decision Management in their names. Some, from companies with a strong analytic focus, might refer to Decision Analytics also.

This focus can be just a naming thing, with equivalent products having different names, but it can also reflect a subtle yet important emphasis on decisions over business rules as the primary organizing principle of a product.

Navigating products for developing analytic insight

In many ways the products that offer support for the development of analytic insight are more straightforward. Most such products are described either as Data Mining Workbenches or Predictive Analytic Workbenches. These are often easy to compare with each other and offer broadly comparable capabilities. Some such products are narrowly focused, offering a small number of analytic algorithms or support for a particular kind of data. More common are workbenches that support a wide range of such techniques and data sources.

The only areas of potential confusion come in the support of model management and in-database analytics and the potential support of some general purpose decision management platforms for limited development of analytic insight.

- Model management.
  Some analytic products include capabilities for model management in the same product that is used to build analytic models, some package it up as a separate capability. There is also a small group of products designed explicitly for model management. There is no particular advantage or disadvantage to the approaches though having a more web-based and less analytic professional-oriented environment for model management can be appealing. Some of these model management capabilities support models built using a variety of analytic tools. This support for monitoring and managing models that were built using multiple tools is a significant differentiator, regardless of whether it is packaged with an ability to build models or not.
- In-database analytics.
  Similarly some analytic workbenches package up support for in-database analytics (either development or deployment or both) with the workbench while others sell it as a separate capability. When considering a workbench from a functional perspective, what matters is the tool's support for the databases in use at your organization and the depth of integration. Packaging may affect pricing but it generally does not affect capability.
- Decision-centric analytics.
  Some products that are primarily focused on managing decision logic provide capabilities for developing analytic insight. Some offer what can be described as "data mining for business rules", allowing data mining algorithms that produce

decision trees or association rules to be used within the tool to find suitable rules from historical data. Some offer data mining algorithms integrated with a decision tree editor for "data-driven strategy design". Both such capabilities are highly desirable and the use of data mining to find business rules is a clear best practice (discussed in the section on analytic and it cooperation). Nevertheless these tools do not offer the same range of analytic insight capabilities as a specialist tool.

Some of these platforms go a little further and offer automated analytic model building capabilities also. These start to compete more directly with pure-play analytic workbenches, especially for organizations focused on decisions outside of regulated credit industries that are comfortable with automated modeling approaches. Most users of these tools will still find at least occasional reasons to use a pure play analytic workbench however.

Navigating products for optimization

The big differences between optimization products are a solution focus versus a tool focus and the degree of tooling available.

- Solution or tool focus.

  Because optimization can be complex to configure and use, many organizations adopt optimization technology as part of a solution. In this approach the optimization model is pre-configured with integrated reporting and simulation interfaces focused on the solution. These might address a scheduling problem, supply chain issues or product configuration. In contrast a tool focus means delivering an optimization product that can be used to solve any problem but that must be configured before it does anything.

  Because many of the pre-configured solutions are provided built on a specific tool, organizations can often begin with a pre-configured solution and then expand usage by also acquiring the underlying tool. For some organizations, however, there is only one problem that seems to justify optimization and they are likely to be happy with a single solution-focused offering
- Solver or workbench.

  Some optimization products are really just a set of solvers with well defined APIs while others offer a complete workbench with debugging tools and graphical interfaces. The solver-only approach allows the tool developer to focus on performance and scalability while supporting practitioners who want to use a particular problem definition language or editing environment. A more complete workbench tends to be more supportive of less technical users and to involve less work to set up at the expense of being somewhat more limited in terms of how a practitioner can approach defining the problem.

Managing Decision Logic

The first requirement is for a complete set of software components for the creation, testing, management, deployment and ongoing maintenance of the logic of a decision—the business rules—in a production operational environment. The most common product category name for this capability is a Business Rules Management System.

> For the purposes of this report we are concerned only with executable logic, executable business rules—that is with business rules defined at a level that allows them to be executed in a computer system. Business rules may be defined and managed as a requirements approach and to ensure consistency and accuracy in manual decision-making, but this is not the focus of this report.

Generally, an executable business rule is simply a statement of what action should be taken if a given set of conditions are true. Each rule has a conditional element that can be assessed at a moment in time to see if it is true or false as well as one or more actions to take if it is true. These actions could be as diverse as sending emails or invoking functions but generally involve setting data values. In most Business Rules Management Systems each rule can also have an owner, notes, version history and other metadata that describes it.

Managed, executable business rules offer many advantages over traditional code, especially when automating and managing decisions:

- Business rules are easier for non-technical business experts to read, improving business/IT collaboration and improving the accuracy of business rules relative to code. This is especially true because business rules can also be represented in a variety of graphical and tabular metaphors.
- Business rules are declarative, allowing each to be managed independently and so simplifying the management and reuse of decision-making logic while also allowing more precise and granular assessment of consistency, completeness and quality.
- Business rules either "fire" (the conditional element evaluates to true) for a particular transaction or they do not. This can (and should) be recorded each time a decision is made and represents a precise description of how a decision was made. This supports subsequent analysis and improvement of decision-making.

A Business Rules Management System or equivalent functionality gives business users and analysts the ability to make routine changes and updates to critical business systems while freeing IT resources to concentrate on higher value-add projects and initiatives. Even when used by an IT organization in a more traditional way, a

Business Rules Management System allows for more rapid change by making it easier to find, make and test changes to decision-making logic.

Managing decision logic requires software that supports a range of activities:

- Integration with other applications and services and linking business rules to data sources so that business rules can be developed that will use the data available in existing systems and processes.
- The development and testing of business rules by both technical and non-technical users so that all those involved in defining a decision can participate in writing the business rules.
- Identification of rule conflicts, consistency problems, quality issues and more for both technical and non-technical users so that full advantage is taken of the declarative nature of business rules.
- Assessment of the business impact of changes to the business rules through simulation and reporting to ensure the right changes are being made and to understand the business consequences of changes that must be made.
- Deployment of a defined package of business rules to Decision Services in different computing environments.
- Measuring and reporting of decision and business rule effectiveness based on the results of executing business rules in decision services.

Such a system requires the following capabilities. In a future release of the report a set of specific items to look for in each category will be identified.

Rule Management

A business rule management environment suitable at least for technical users is essential. This environment typically also includes design tools to integrate the deployed business rules with the rest of the enterprise computing environment. Generally this is provided as part of an Integrated Development Environment or IDE, often one based on Eclipse or Visual Studio.

Technical users are generally not the only ones who will need to edit business rules. Interfaces to allow business analysts and business users to manage business rules directly and in-context, or tools to allow such interfaces to be built and maintained, are critical elements of a robust approach to managing decision logic. These interfaces could be part of an IDE, though this is less common, and a thin-client interface is more likely. Some products provide editing environments for non-technical users based on the Microsoft Office products, specifically Microsoft Word and Microsoft Excel.

A variety of metaphors are often used to author business rules. A rule flow or decision flow is used to lay out multiple steps within a decision. Business rules can be specified for each of the steps or tasks in such a flow as a decision tree, decision table, rule sheet, decision graph, decision model, rule family or simply as a list of independent rules. The differences between these metaphors and the value of each will be discussed in a future version of the report.

Verification and Validation

Verification and validation tools that check business rules for completeness, consistency and logical errors help ensure that valid business rules are being written. Such tools should be suitable for both technical and business users to use and should be integrated with the various editing interfaces provided. These tools should ensure that the business rules being authored are at least *potentially* valid. They cannot tell if the business rules are the right ones for the business or if they handle every business scenario but they can tell if they are structurally and logically complete and that they handle known variations in data such as lists of values.

Testing and Debugging

Testing and test management tools that support unit, system and acceptance testing are a necessity. While there are circumstances in which business rules change so rapidly that formal testing is not part of the release cycle, most organizations will still have a set of tests they wish to run before allowing a new set of business rules to be deployed. Managing these tests should be straightforward. Business rules must sometimes be tested with other new components in the context of a broader application deployment and being able to test the business rules in this context is useful. Many products support integration with open test management standards such as xUnit.

Technical users, and ideally less technical ones, should also be able to debug business rules. They should be able to walk through the business rules executing in a decision to see what happens in specific cases. This may be supported only for a local test environment or for both local and production environments.

Impact Analysis

Impact analysis and business simulation tools to allow non-technical users to see the impact of a set of rule changes on their business outcomes are an increasingly important part of managing decision logic. Business analysts and business users will not generally be willing to make changes to business rules unless they can see what impact a change will have. Similarly when a change must be made to the business rules, due to a regulatory or policy change for instance, business users will want to see the likely impact of this change. The results must be presented in business terms to be useful—an increase in profitability, a reduction in fraud, etc.

These facilities may be provided as a batch tool for running historical or sample data through a set of business rules or as a more interactive tool allowing a business user to select the data they care about and running new or changed rules against that data. The best practice is clearly moving this closer to the editing of the rules themselves with the potential business impact of a change being shown automatically as a change is made in the editing tools.

Data Management

Decision logic must be integrated with the data that will be available when the business rules are deployed. It needs to provide tools that at least allow technical users to integrate the business rules with the organization's data. In addition it is useful for a product to be able to bring in large amounts of historical data as well as large test datasets to support effective testing and impact analysis.

Deployment

A set of deployment tools that support the deployment of a set of business rules either as executable code or as a package that can be executed by a high performance Business Rules Engine, ideally on multiple enterprise platforms, is required.

> One point of confusion is the difference between a Business Rule Engine and a Business Rules Management System. A Business Rule Engine can be part of a complete system for handling all the things involved in working with business rules. It is clearly an important part, but it deals only with execution. It determines which business rules need to be executed in what order. A Business Rules Management System is concerned with a lot more.

Business rules can be executed in a number of different ways once deployed. Some Business Rules Management Systems support inferencing execution. Based on various algorithms, many derived from the original Rete algorithm[1] these determine the correct execution sequence based on the structure of the business rules and the data available when they must be evaluated. As business rules fire and change data the engine reassesses which business rules might need to be fired next. While there are some scenarios that are very difficult or even impossible to handle without inferencing support, they are not common. The key advantages of inferencing in normal use are that it allows the business rules to be written in any order and that it

1 The original algorithm is described in "On the efficient implementation of production systems" Charles Forgy's Ph.D. Thesis, Carnegie-Mellon University, 1979. Extensive development and refinement has resulted in multiple, distinct versions of the algorithm.

ensures business rules are re-evaluated when the data used in their conditions changes.

Business rules can also be executed in a sequential way, using the order specified for the business rules at design time. In many scenarios, especially when most business rules in a set will be executed for most transactions, this approach is faster. It also allows business rules to generate code, which can result in smaller and more portable deployments.

Finally a number of products offer designed execution where the rules are executed sequentially but the order is determined by automated analysis of the business rules at deployment time. This simplifies execution but allows business rules to be written and edited in any order without any unexpected impacts on their behavior as the deployment time analysis will sequence the new and changed business rules appropriately.

For most business scenarios all these approaches work well. Each approach has its own set of best practices in business rule writing.

Repository

Last, but by no means least, products should offer an enterprise-class repository for storing and managing business rules. This repository may be a complete decision management repository that also stores predictive analytic models and optimization models. It is more likely to be one that only manages business rules. It should provide access control and security, audit trails for changes made to the business rules and versioning at a number of levels. An extensible repository that allows additional information to be added as well as an API for repository access can improve the integration of the product with other enterprise components. Some products provide integration with source code control systems, allowing business rules to be stored and managed alongside code used in the rest of the application.

Embedding Predictive Analytics

Embedding predictive analytics requires a software component for the creation, validation, management, deployment and ongoing re-building of predictive analytic models. Such a Predictive Analytics Workbench allows a data miner, data scientist, analytics professional or business analyst to explore historical data and use various mathematical techniques to identify and model potentially useful patterns in that data.

> For the purposes of this report we are not concerned with the use of data mining or predictive analytic workbenches for one-off research projects to answer a specific question or with the construction of statistical models per se. Only models that can be applied to a specific transaction or item to classify it or make a prediction about it are included. Other forms of data mining and predictive analytics can have tremendous value to an organization but they are not relevant to this discussion of Decision Management Systems.

The predictive analytic models created can predict a binary outcome (yes or no), provide a number (often representing a probability or ranking of likelihood) or a selection from a list (of products for instance). They might also cluster or group based on likelihoods and may identify what item is associated with what other items.

> Data mining and predictive analytics allow organizations to turn historical data into useful, actionable analytic insight.

Data mining and predictive analytic models are often grouped with business intelligence, reporting and visualization under the general term "analytics". Data mining and predictive analytics differ from business intelligence capabilities in a number of ways:

- They are focused on extracting meaning about the likely future rather than summarizing or understanding the past—they use historical data to make predictions about what is likely in the future.
- They are probabilistic rather than definitive in that they rarely if ever make a prediction that something concrete is definitely going to happen. Generally they say how likely something is, make a prediction with a certain degree of confidence, or rank order a set of possible outcomes from most to least likely.
- Rather than relying on the visual processing power of humans to see patterns in data, they rely on mathematical algorithms to explicitly extract these patterns from the data.

This last point has an important consequence for predictive analytic workbench products being used to develop Decision Management Systems. These products must do more than simply define the right mathematical models. Presenting the results of a predictive analytic project as mathematics or even as visualizations and reports is not sufficient. It must be possible to use the product to both produce an effective predictive analytic model and embed such a model into an operational system. Unless the predictive analytic models produced can be effectively embedded they will not be useful for Decision Management Systems.

A predictive analytic workbench needs to support a range of activities that are generally performed in a highly iterative way:

- Integration with a wide range of data sources so that data can be brought into a modeling environment for analysis. These data sources might be systems that are internal to the organization or external data. Increasingly these sources go beyond traditional relational data sources to unstructured and semi-structured data.
- Cleaning, integration, summarization and exploration of this data including sampling, identifying outliers, providing distribution statistics and more.
- The creation of an analytical dataset suitable for analysis including identifying and creating potentially useful derived variables, and managing very large datasets with thousands of attributes (both original and derived).
- Automated or mostly automated analysis of very large numbers of records using a variety of algorithms such as classification, decision trees, linear and logistic regression, clustering, neural networks, nearest neighbor and more. Increasingly the use of ensemble methods, where multiple techniques are applied in combination, must also be supported.
- Creation of analytic representations, models, based on this analysis such as predictive scorecards, functions or business rules.
- Validation of these models to prove they will be predictive with data not used to build them as well as assessment of their effectiveness in making predictions.
- Deployment of these models into an execution environment or as code that can be independently executed.
- The definition and management of repeatable processes or workflows to handle all these steps so that they can be repeated with new data, as part of assessing multiple possible approaches or with minor edits as the user evolves their approach.

One of the most important facets of these kinds of workbenches is their support for an industrial scale process for building predictive analytic models. Predictive analytic model building used to be something of a cottage industry, with each modeler making their own choices for scripting language and a largely manual process. This approach relies heavily on the skills of the modeler and is hard to scale. With organizations increasingly needing dozens or hundreds of models, a more industrial process is called for. This does not eliminate the skill of a modeler, but it does require more repeatability, automation and scalability in the way predictive analytic models are built and managed. This is where a predictive analytic workbench is essential.

A predictive analytics workbench gives data miners and possibly business analysts the ability to derive useful probabilities about the future from potentially large amounts of data about the past. These probabilities may group or segment customers or other records, identify the propensity of someone to do something (buy, churn, respond, visit), determine the strength of an association between two records or identify what is likely to be the best combination among many possible ones.

Embedding predictive analytics requires the following capabilities. In a future release of the report a set of specific items to look for in each category will be identified.

Data Management

Predictive analytic models are typically built from a large amount of data, often pulled from multiple data sources. A predictive analytic workbench must be able to connect to and retrieve information from a variety of structured and unstructured data sources as well as flat files of various kinds.

Data Preparation

The data available is often not immediately suitable for the construction of predictive analytic models. A predictive analytic workbench provides a variety of tools to allow the clean up and integration of data prior to modeling. These tools include renaming and re-categorizing data fields, imputing missing values, filtering outliers, extracting samples and transforming data to make it more suitable for modeling. The end result of this data preparation work is what is often called an analytical dataset—a large set of data attributes (some original, some derived) with any hierarchical structure "flattened" into a single list of attributes.

Data Visualization and Analysis

Modeling efforts typically begin with exploration of the data available to develop some understanding of the data and of the patterns in that data. A rich set of visualization and graphical tools as well as statistical analysis routines help find the hidden patterns and relationships that might drive an effective model. These tools are often used in conjunction with the data preparation tools so that problems found in graphing the data, for instance, can be corrected in a data preparation routine. The same visualization and analysis tools will also be used to assess model outcomes once models have been developed.

Modeling

At the core of a predictive analytics workbench is a model creation environment suitable at least for data miners and other analytic users. The modeling environment might also allow business analysts to create and manage the modeling process—typically through a combination of automation and simplified interfaces.

Some predictive analytic workbenches are designed for expert users. Some are primarily aimed at these experts but provide simplified interfaces that aim at a broader audience. Some are designed with a single environment that works for both expert and less expert users. While the style of interface and its expectations can vary, all these workbenches create predictive analytic models and related resources in some form of shared repository.

The modeling environment typically involves laying out a series of steps that will result in the construction of a model or models that can be evaluated for performance. Steps will include data preparation and analysis as well as the execution of one or more algorithms from an extensive set. Algorithms supported include clustering, association, linear and logistic regression, decision trees, support vector machines, Bayesian modeling and nearest neighbor techniques to name a few. It is increasingly common to find ensemble models where several techniques are applied, or one technique is applied with different parameters, and the results aggregated in some fashion to create a single, overall ensemble model[2].

Some predictive analytic workbenches can take advantage of in-database modeling engines that can handle some of the data preparation tasks as well as execute the modeling algorithms themselves on the database server that contains the data being analyzed. This improves performance by eliminating the need to move data from the database to a separate analytic server and takes advantage of the increasingly powerful servers supporting data infrastructure.

Model Validation

Regardless of which technique or set of techniques is used, model performance assessment and comparison tools are used to see how well a model performs. Different models can be compared and tools such as lift curves (comparing selection using a model to a random distribution) used to see how effective the model would be in production. These tools typically use new data (data that was not used to build the model) to see how predictive the model would be once deployed.

2 Ensemble models are explored extensively in “Handbook of Statistical Analysis and Data Mining Applications” (Elsevier, 2009) by Nisbet, Elder and Miner.

Deployment and Scoring

Once the final model or models have been identified they must be deployed. A predictive analytic workbench may allow multiple approaches to deployment:

- Models can be used to score data in a batch mode, applying the results back to the database that contained the data from which the model is built.
- Some predictive analytic workbenches can act as a real-time scoring server using their own scoring engine and providing a web services or other API to allow it be called during decision-making.
- Scoring code can also be generated (as C or Java, as SQL or as business rules) so that it can be deployed to a Decision Service for real-time scoring.
- In-database scoring is also available, with the definition of the model being pushed to the analytic infrastructure where the scoring engine is running.
- A number of predictive analytic workbenches also allow models to be generated using the Predictive Model Markup Language[3] (PMML), allowing the model to be executed by any business rules or scoring engine that supports this standard.

Model Monitoring

Models are built from a snapshot of data. As such they "age"—as time passes the data being fed into the deployed model may look less and less like the data from which it was built. A predictive analytic workbench needs tools to monitor deployed models to see how their performance is varying over time and to identify variations in performance or in data distributions. Many new models are initially deployed to challenge an existing model and the performance of both the original "champion" model and the new "challenger" model need to be compared to see if the challenger is good enough to replace the champion. Model monitoring tools need to identify opportunities to refresh and retrain models and to provide tools to make it easy for users to rebuild models to take advantage of new data.

Model Tuning

Some predictive analytic workbenches provide components for automated model tuning and updating. These machine learning techniques monitor the performance of a model as it is used in deployment and automatically adjust its underlying equation based on that performance. Some of these environments can start with no model and gradually build a predictive model based on the results of random experiments while others are designed to be used with pre-defined models. Model Tuning can be left to run forever or it can tune the model within defined boundaries and flag a model for re-building if its performance starts to drift outside those boundaries.

---

3 PMML is an open standard managed by the Data Mining Group that provides a standard XML representation of predictive analytic models so that they can be exchanged between multiple products.

Model Tuning capabilities are often deployed in a Decision Service if that is where the model is being executed.

Repository

A predictive analytics workbench should offer an enterprise-class repository for storing and managing predictive analytic models. This repository may be a complete decision management repository that also stores business rules and optimization models. It should provide access control and security, audit trails for changes made to models and versioning.

> There is a growing category of software products that allow business rules to be specified and managed alongside predictive analytic models built in the same product. The degree to which large numbers of business rules can be managed and the range of predictive analytic models that can be built varies and such a combined product may not therefore support the complexity required for a specific Decision Management System. These products typically allow models built in other predictive analytic workbenches to be integrated also.

In-database Analytic Infrastructure

Two main kinds of analytic infrastructure are relevant to Decision Management Systems—in-database modeling and in-database scoring. In-database modeling infrastructure allows predictive analytic models to be developed using algorithms embedded in databases and data warehouses to improve the performance of predictive analytic model creation. This infrastructure is often integrated with predictive analytic workbenches so that it can be called as part of a modeling exercise.

Some in-database modeling capabilities support the same wide range of modeling techniques as a predictive analytic workbench. Some support only a subset and are used by a predictive analytic workbench to offload some of the work of building predictive analytic models to the database server. This allows, for instance, cleaned and transformed data to be created in the database and passed to the predictive analytic workbench rather than requiring all the raw data to be moved and then processed by the workbench.

In-database scoring infrastructure takes models developed using some combination of in-database modeling infrastructure and a predictive analytic workbench and executes them in an operational datastore so they are available to operational systems accessing that datastore. This generally involves turning models into functions that can be called using SQL and that take database fields as input. These functions can be used in standard SQL statements and embedded into database views as though it was a database column, making the scored data available widely.

Optimization and Simulation

An optimization suite is an environment for defining and solving mathematical models and for simulating the differences between multiple similar mathematical models. An optimization suite allows a modeler or business analyst to define a business objective and a set of constraints and then "solve" this problem to see how best to run the business. Optimization suites support what is sometimes called Operations Research or Management Science.

> There are really three uses of optimization in the context of Decision Management Systems:
>
> - When a decision has a potentially complex answer that involves multiple elements it may be effective to optimize the selection of these elements.
> - When a decision answer is a single element then it may be useful to optimize across many decisions to allocate the available answers to each specific decision most effectively.
> - When reviewing possible decision-making strategies as part of decision analysis it may be possible to use optimization to tune or select between these strategies.

Optimization allows organizations to either find a feasible solution to a heavily constrained problem or to maximize the value gained from a constrained set of resources by finding the most profitable, quickest or cheapest combination of resources that are allowed. Optimization differs from both business rules and predictive analytics in a number of ways:

- Business rules are absolute where optimization need not be. For instance business rules allow an offer to be made to someone only if certain conditions are true where an optimization model might allocate offers based on where they will be most effective.
- Optimization can be effective when business rules are numerous and potentially contradictory as it allows for trade-offs between values where business rules require defined sets of conditions.
- An analytic model is created through analysis of historical data while an optimization model is built explicitly from business know-how and historical data may be used to see how the model would have worked in the past (though this is not necessary).
- Because predictive analytic models are built and executed separately they are often very quick to execute. Optimization models in contrast must be solved each time they are used and this can require significant time and resources.

An optimization suite needs to support a range of activities:

- Defining a constrained optimization problem as a mathematical model using variables, an objective function and constraints both hard and soft.
- Solving this problem, often multiple times as elements of the problem are changed and re-assessed.
- Integration with a wide range of data sources so that data can be brought in and run through a defined optimization model. These data sources might be systems that are internal to the organization or external data.
- Simulation and comparison of different scenarios by a non-technical user to see what the best choice is likely to be going forward.

An optimization suite gives modelers and possibly business analysts the ability to manage tradeoffs and constraints to find the optimal action to take. An optimization suite requires the following elements:

Optimization Model Development

At the core of defining an optimization model is a modeling language of languages. Some optimization suites have their own such language but a number of popular ones exist and some solvers (see below) can support several languages. Most optimization suites will provide an optimization model development environment suitable for modelers to specify models in one of more of these languages. This environment may be based on a commercial available IDE such as Eclipse or Visual Studio.

Optimization Model Debugging

Debugging and profiling tools allow modelers to review and change the model to correct for identified problems – find conflicts, relax constraints or profile performance. Models can be complex and even unsolvable so profiling and debugging tools are essential to allow a viable model to be defined.

Solvers

Most optimization suites include multiple engines or solvers that apply mathematical techniques to the developed models to "solve" the problems defined in those models. These solvers can be specific to different kinds of problems such as linear programming problems, mixed integer problems, quadratic-problems and combinations such as mixed integer quadratic problems. These solvers may be used to run scenarios, to find optimal actions that can be loaded into a production system as a batch or can execute in a Decision Service to solve an optimization problem as part of a single decision. In addition, many standalone solvers are available.

Data Management

Optimization models are coded or constructed by hand but scenarios typically involve a large amount of data, often pulled from multiple data sources. An optimization suite must be able to connect to and retrieve information from a variety of structured and unstructured data sources as well as flat files of various kinds and present this data for scenario analysis.

Scenario Analysis

Many optimization problems require an interface that allows a business analyst or business user to run and compare scenarios based on these models and associated data. Such scenario analysis involves rich visualization and the ability to bring real world historical data into the system to run through the model. Optimization suites include either scenario analysis interfaces or the ability to rapidly generate such interfaces for a given model.

Deployment

The results of optimization can be deployed in a number of different ways. Deployment tools in an optimization suite may support the deployment of a model as results or recommendations, the packaging of a model to run against a solver running in another environment at run time or the conversion of optimal actions into rules that mimic the assignment of an optimal action.

Repository

An optimization suite should offer an enterprise-class repository for storing and managing optimization models and associated scenarios. This repository may be a complete decision management repository that also stores business rules and predictive analytic models. It should provide access control and security, audit trails for changes made to models and versioning.

Monitoring Decisions

The final area of capability is that of monitoring and improving decisions over time. These capabilities are essential for Decision Management Systems both because decisions are high change components and because the time it takes a decision to come to fruition can be extensive, making it hard to tell good ones from bad ones.

There are many drivers of change in decision making. Regulations change so organizations must change how they make eligibility decisions to remain compliant with those regulations. Policies change so organizations must, for instance, change their validation of suppliers to track new data requirements. Competitors change so organizations that wish to remain competitive must change their discounts or pricing. Markets, such as the financial or credit markets, change so organizations must constantly change the way they assess risk. Consumer behavior changes regularly and continually so organizations working with consumers must constantly address these changes in their decision-making. Finally, of course, fraudsters adapt and seek new loopholes to exploit so organizations must change how they detect and process fraud to focus on new fraud as it develops.

In addition to outside changes that explicitly drive changes to decision-making, organizations want to continuously improve their decision-making. The challenge for some decisions is the time it takes for decisions to play out—it may be weeks or months before an organization knows if the decision was a profitable one for instance. To continuously improve in these circumstances it is essential to be able to conduct experiments and compare their results. Such an experiment makes the same decision in two or more different ways, applying the different approaches to different transactions and comparing the results. Sometimes called adaptive control, champion-challenger or A/B testing, these approaches drive continuous improvement in decision making.

As shown in Figure 4: Continuous improvement in decision making below, this approach requires that the results of a decision be evaluated, predictive analytic models and business rules updated and refined and new "challengers" or alternatives developed. These are fed back into the decision-making loop and used to make future decisions. The results of these decisions are evaluated in turn with successful experiments being adopted, unsuccessful ones dropped and new ones developed in a continuing cycle.

Figure 4: Continuous improvement in decision making

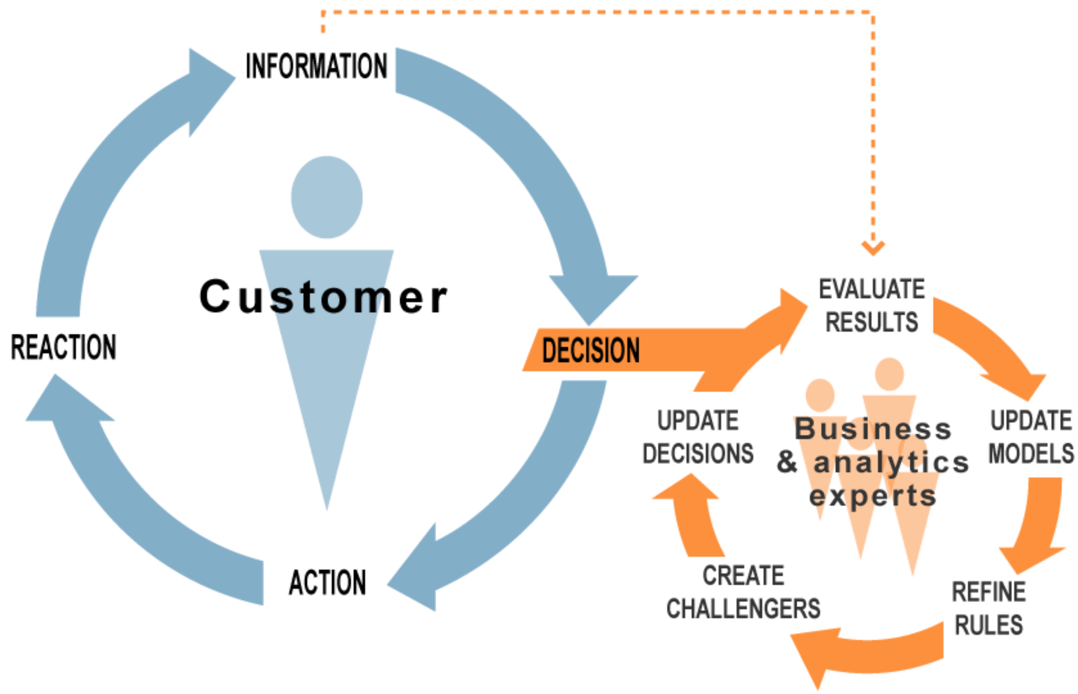

The capabilities to support monitoring and improving of decisions are not typically found in a single software product. Instead these capabilities drive the requirements for products used for both decision logic management and embedding predictive analytics.

Logging

The primary capability required for decision monitoring is that of logging decision execution. When a decision is made by the Decision Management System it must be possible to log how that decision was made, what business rules fired. This log should include any predictive analytic model scores calculated during the decision as well as the specific action recommended by the Decision Management System.

In addition, these decision-making logs should be stored in a way that allows them to be integrated with information about the response of customers and others to the decision—did the customer accept the offer, did the salesperson override the price with an additional discount, was the deal closed and so on. The long time results such as orders placed or customers retained that can be attributed to these responses are logged by other systems. It should be possible also to tie the decision-making specifics to these results.

While logging is essential for ongoing improvement of decision making, logging also supports compliance and audit needs by providing complete execution transparency.

When an audit or compliance review is conducted it will be possible to tell exactly how a decision was made and whether or not that decision followed the correct guidelines.

Experimentation

To ensure continuous improvement of decisions it will often be necessary to conduct experiments. These experiments typically involve multiple approaches to either the decision logic of the decision, the predictive analytic models used in the decision or both[4]. Additional decision logic must be managed to determine which of the approaches should be applied to a specific customer or transaction and it must be possible to record this as part of the decision itself.

All products suitable for managing decision logic can manage experiments in this way. Some products for decision logic management have additional capabilities built in to make it easy to manage, review and compare the various approaches being used within a decision.

Decision Performance Management

The performance of a decision can and should be managed and monitored in the same way any other aspect of business performance is managed and monitored. Generally it is straightforward to apply the standard performance management capabilities of an organization to decision logs to see trends, hotspots, etc.

Business User Decision Logic Management

While changes to decisions and decision logic are sometimes extensive, requiring all the capabilities described above, sometimes more localized and focused changes are required. These should generally be made by business users so that a full IT cycle can be avoided for what could be regular, minor updates. To make this work it must be possible to use the decision logic management capabilities for non technical users to present a business person with their own business rules, in context. Ideally this environment will only allow them to make changes that make sense and will present no unnecessary information. Most products for managing decision logic either include suitable interfaces or allow suitable interfaces to be developed.

Impact Analysis

As noted above it is important to provide impact analysis tools to allow non-technical people to rapidly see the business impact of any changes they make. This should cover both design impact and execution impact and involves more business-centric functionality than is required for testing. Impact analysis tools may need to

[4] Designing suitable experiments can be complex. The classic reference for experimental design remains Fisher's book "The Design of Experiments" first published in 1971.

consider changes to decision logic, to predictive analytic models or to both. See above under *Overall Architecture*.

Alternatives Assessment

When multiple decision making approaches are being used in parallel it will be essential that the effectiveness of these alternatives can be assessed. Capabilities such as swapset analysis (showing which customers, for instance, would get offer B rather than offer A) as well as more general comparison of business performance metrics are critical. In addition, simulation and what-if analysis tools that can use each alternative approach and compare the outcomes of multiple simulations based on the approaches will be required.

Model Monitoring

See above under *Embedding Predictive Analytics*.

Model Tuning

See above under *Embedding Predictive Analytics*.

Key Characteristics

Experience in working with organizations that are developing Decision Management Systems shows that while there are many ways to develop them effectively, certain key characteristics come up repeatedly as critically important. These characteristics fall into a number of areas including the completeness of the platform, engagement of business users, architectural flexibility, organizational scale and decision monitoring.

This set of characteristics is neither a definition of a complete set of features and functions required to build a Decision Management System nor a complete list of characteristics for any of the product categories. It is intended as a set of characteristics you can look for in products you are purchasing or using that will support a focus on Decision Management Systems.

Platform Completeness

A small number of vendors offer a complete platform for building Decision Management Systems. These platforms handle decision logic or business rules, support data mining and predictive analytic modeling, include constraint-based optimization and provide monitoring and integration capabilities for deployed systems. While it is not necessary to buy a complete platform from a single vendor, it is valuable for products to see themselves as part of a broader ecosystem. For instance Business Rules Management Systems that are "aware" of predictive analytics and offer integration with such systems and predictive analytic workbenches that offer business rules-friendly deployment options are more suitable for Decision Management Systems than more narrowly focused products.

Complete Platform

A complete platform is an integrated set of offerings that allow for the management of decision logic, the building and deployment of predictive analytic models and the mathematical optimization of decisions. These offerings are either a single product or a product set with a common user interface, shared repository and common tooling that operates across the products. Support for decision monitoring and analysis is either provided or the data is made available to standard reporting and dashboard components.

Complete Ecosystem

A company may not offer a complete platform for Decision Management Systems themselves while still supporting a complete platform through their ecosystem. By supporting open standards such as PMML and by partnering with other vendors that offer more pieces of the puzzle, vendors can offer a Complete Platform Ecosystem.

Openness

Some companies are not focused on Decision Management Systems but on providing a specific component. They may be focused only on managing decision logic, building predictive analytic models or constraint based optimization. They may not even think of themselves as participants in the development of Decision Management Systems. These companies are not likely to have a complete platform nor are they likely to actively partner to develop a complete platform ecosystem. Their products can still be easy to integrate and use alongside other products and can support standards such as PMML for predictive analytic models or JSR-331[5] for constraint-based optimization. For standalone products focused on a specific technology market this kind of openness is critical in being part of a complete platform for Decision Management Systems.

Business User Engagement

The agility and adaptability of Decision Management Systems crucially relies on the engagement of business users. The extent to which products being used to build these systems can bring business users into the development team is therefore critical. Products that focus on allowing business users to read and write decision logic, participate actively in building or reviewing analytic models and allow non-technical users to run through scenarios are more likely to be successful that those focused only on technical developers.

Business User Analytic Modeling

Analytic tools can engage business users by providing an environment designed for non-technical users to create and use analytic models. This might be a complete environment that uses automation and machine learning algorithms to build predictive analytic models with minimal data mining expertise required. Such use of machine learning algorithms and automation should be complemented by user interfaces, reporting and automated checks designed to support a less knowledgeable user. These additional capabilities can ensure that problems such as overfitting are avoided and that test and validation data is automatically set aside, for instance.

Alternatively a product might be a business user friendly environment layered on top of a more typical data mining/predictive analytic workbench. This would use wizards and other simplifying features to make it possible for non technical users to do data mining and create predictive analytic models. Generally these environments integrate with the more traditional modeling environment and store models in the same repository such that modeling specialists can refine or enhance the models built using these less technical interfaces.

---

[5] More information on JSR 331 can be found at http://jcp.org/aboutJava/communityprocess/final/jsr331

Business User Rule Management

The management of decision logic by non-technical users, non programmers, is a key element in delivering the agility required of a Decision Management System. While any product focused on managing business rules or decision logic might be said to allow some business user rule management, true business user rule management requires a number of elements.

First the rules themselves must be approachable. Using declarative statements in place of procedural code so that each rule can be considered and edited independently as well as the use of a business user vocabulary not technical data element names in rules will ensure readability and clarity. Supporting a verbose, readable syntax in near natural language rather than terse programmer-centric constructs helps as does emphasizing graphical editing of rules in decision tables, rule sheets, decision trees, decision graphs and decision flows so that as little as possible has to be written out "longhand."

Because business users are not programmers, testing tools need to be accessible to them and excellent completeness and correctness checks are essential. Ideally these tests are performed inline, as business rules are edited, to ensure that obvious mistakes and omissions are avoided.

Business users do not like technical environments designed for programmers so support for the editing of business rules outside these environments through a web interface or a point and click, business friendly editing environment makes rules more accessible. Such editing environments might include support for editing using standard Microsoft Office products also. Ideally these editors will be embeddable or mashable so that rule management can be embedded in environments focused on a specific business task rather than on rule editing.

Support for a learning curve

Most organizations will not be able to jump straight to either business user analytic modeling or business user rule management. They will need to bring business analysts on board first, exposing some of the tasks previously performed by IT or analytic teams to these semi-technical users. Over time the role of business analysts can be expanded and true business users brought in to work on certain elements. Products that provide way stations and gradual increases in complexity through multiple editing environments will be easier to adopt than those with a more limited set of options.

For example, a product that allows users to bring analytics to bear incrementally rather than all at once will improve analytic adoption. If simple interfaces allow access to candidate association rules and proposed splits in decision trees then users will become gradually accustomed to the power of analytics to improve their decision logic. Over time they can be exposed to completely data-driven decision trees, unsupervised clustering and ultimately more complex predictive analytic models. Similarly a variety of rule editors that allow a user only to change numeric values in a locked-down rule can get users used to the idea that they can change decision logic. Over time editors that allow new rules to be built based on templates and perhaps new rules to be built using a point and click editor can bring users up the learning curve gradually.

Impact Analysis

One of the biggest barriers to business users taking ownership of their decision logic, in fact probably the biggest single one, is an inability to see what the impact of a change will be. Products that provide strong impact analysis tools, especially tools that allow a non technical user to see how the change they are considering will impact their business results, will be more able to drive successful business user engagement.

Impact analysis can be done using functions designed for regression or performance testing. The ability to do impact analysis and simulation using real data in a business friendly way is invaluable, however. This involves simple to use interfaces for loading data, an ability to assign business value to different outcomes and generally the ability to get results into Excel for further analysis. Ideally the environment should continuously perform impact analysis as changes are being made so that all edits are made in the context of their business impact.

Impact on application context

Simulating the impact of a change on the application or process context in which the decision is being made is sometimes critical. For instance, the impact of a change to fraud detection decisions may be best considered in terms of the workload create for fraud investigators and the average handle time for fraud cases. These measures are not measures of the decision performance alone but include process / application design issues. An ability to determine the broader business impact of specific changes made to decision-making is a potentially very powerful capability.

It may seem that this requires the application or process context and decision-making technologies to share a vendor. Certainly this capability is easier to provide in those circumstances but most organizations will ultimately find themselves in a more heterogeneous environment as noted above, so more flexible capabilities that allow decision simulation to be integrated with simulation capabilities from other tools should be valued as well as more packaged capabilities.

Architectural Flexibility

Decision Management Systems automate decisions that must often be used in multiple channels where these channels may be supported by different applications and architectural approaches. Decisions may need to be made as part of business processes, in response to events detected or in support of legacy environments. A degree of architectural flexibility is therefore very useful in products used to develop Decision Management Systems. Support for multiple platforms and deployment styles as well as a wide range of integration options helps a lot.

Big Data support

There is tremendous interest in "Big Data" at the moment as the rapid growth of social media, weblog data, sensor data and other less traditional datasources creates new challenges for managing this information. While much of this interest has been around supporting queries and reporting, organizations are beginning to use these new data sources in their Decision Management Systems. Products that can support both traditional and newer Big Data sources therefore offer increased scope for organizations going forward.

Support for Big Data involves being able to bring potentially very large amounts of data stored in NoSQL systems such as Hadoop into analytical modeling as well as into the operational environment. As many of these sources are less structured it also involves supporting text analytics and operations that use text operators. Flexibility in data definition, so that the variety and velocity of these data sources do not disrupt operations will also make a big difference.

Cloud ready

The recent growth of the cloud as a platform for enterprise applications means that more organizations are increasingly relying on cloud-based solutions for CRM, HR and other applications. Because Decision Management Systems must integrate with these systems it is becoming increasingly important for products used to build Decision Management Systems to be cloud-ready. This means being able to connect to cloud-based systems to access data and being deployable to the cloud so that decision services can be easily integrated into cloud-based systems.

The cloud also has a lot to offer for the development of Decision Management Systems. Many tasks, such as building predictive analytic models and running simulations or impact analysis, require a great deal of computing power. Being able to push analytic modeling tasks, simulation runs and impact analysis execution onto cloud-based resources means these can be run in background while the user works on something else using their personal computer and can greatly increase the scope of what is possible in these tasks. For products being evaluated for Decision Management System construction the ability to integrate cloud resources for these high compute power tasks offers great productivity increases.

Heterogeneous environment

Most organizations of any size have a heterogeneous environment with multiple operating systems, multiple databases, different communication protocols etc. Different channels have different systems, mobile devices and in-store or kiosk/ATM machinery is unique and organizations often have layers of computing equipment of different ages. No organization ever has a single, coherent architecture across all its systems, at least not for long. Because decision-making components must often support multiple channels and be consistent across multiple systems, products for Decision Management Systems should have multiple deployment options and be easy to deploy and integrate with these different operational environments.

Organizations are often heterogeneous in another way. Some organizations use multiple business rules management systems, many use multiple predictive analytic workbenches as analytic modelers choose their own or use a tool to get access to a specific algorithm. Tools that recognize they must operate in this environment will generally be preferred therefore, especially in the ongoing evolution and management of decision management systems e.g. in analytic model management.

Embeddable management and control components

Decision Management Systems do not stand alone. In particular the management and control of Decision Management Systems should be easy to integrate into other management and control interfaces. For instance it should be easy to integrate analytic model management reporting with more general business performance reporting and business rule management components should be embeddable in other interfaces. A key criteria then for products used to build Decision Management Systems should be how easy is it to embed management components into portals and dashboards built using other tools, feed analytic model management data or rule performance data into a regular performance management environment and so on.

Decision Monitoring

Most products used for developing systems are unconcerned with the operation of those systems once they are deployed. Products used to develop Decision Management Systems, in contrast, offer much more value if they are able to support the ongoing monitoring and improvement of the decision-making embedded in those systems. Products that provide analysis and other tools that integrate with deployed systems are particularly useful in this regard.

Decision Performance

Measuring overall decision performance by tying decision outcomes and decision-making approaches to business results is an important aspect of Decision Management Systems. In practical terms this means be able to easily log the decisions made, including those made as part of A/B or Champion/Challenger tests, so that they can be integrated with overall business performance data in a reporting environment. Products that allow this kind of recording to be done automatically or with flags and settings rather than code are preferred as they create a lower maintenance overhead and are more likely to stay up to date as time passes and the decision making in the system evolves.

Model Performance

Predictive analytic models are generally built at a point in time and so their performance, in terms of how predictive they are, degrades over time. Predictive analytic workbenches that provide automated facilities for monitoring model performance, for identifying models whose performance is degrading, are to be preferred over those that require a development team to hand code this kind of model performance monitoring. In addition it is often helpful if model performance monitoring tools can support models built in multiple environments as this is a common situation.

Rule Execution

One of the most important ways in which decisions can be monitored is through logging the rule execution involved in the decision. While this kind of logging can be hand coded into almost any system, a tool that allows this to be turned on and off for different parts of the decision, that handles this automatically as a background task and that supports it without a significant performance impact is highly desirable. Logs that can be easily stored in database tables and used for reporting and logs that can be easily converted into or viewed in their more verbose format (using actual rule names for instance rather than ids) are also more useful.

Performance and Scalability

As with any technology, performance and scalability should be considered as part of a product selection. Most of the products listed in the appendix are scalable and perform well enough for most if not all scenarios. Organizations with specific and very challenging performance and scalability requirements should be sure to consider these explicitly. For most organizations it is enough to look for solid scalability and for performance adequate to support real-time decisioning.

Scales up and out

In general it is more important to consider if a tool scales well than to assess its particular performance on a given piece of hardware. If a product scales up and out well then more or more powerful hardware can be bought and used effectively as demands increase. If a product does not scale then, even if its initial performance is superior, an organization runs the risk that future demands cannot be met.

Products that support multi-core processors, in-memory processing and distributed processing will scale better than those that do not. This is especially important in high compute power functions such as analytic modeling, optimization, simulation and impact analysis.

Real-time

There is a general move from batch to real-time decision making in organizations of all sizes and types. Many initial Decision Management Systems, however, are batch oriented or have demands that are not truly real-time with several seconds allowed for responses. Over time most organizations should expect to see more demand for real-time decision making as well increasing needs to support streaming/event-based systems. Products that have the kind of low-latency, time based capabilities these solutions need and therefore to be preferred over those that do not.

Organizational scale

Organizations adopting Decision Management Systems generally start with only a single project or two. Over time they become aware of the ROI of Decision Management Systems and the potential for them to change how their organization, its systems and business processes operate. At this point they begin to scale up their plans for Decision Management Systems. Most organizations do not wish to replace the tools with which they are familiar with new tools during this expansion. As a result products with characteristics that support organizational scale will be usable longer. In particular products that support industrialized analytics and enterprise rule management will scale to organization-wide use.

Industrialized analytics

When organizations first adopt predictive analytic models they generally only build one or two models. These models are often hand crafted by an analytic practitioner and then deployed by hand into an operational environment through batch updates or manual re-coding of the model. As the use of predictive analytics expands, however, hundreds or even thousands of analytic models may be required by the organization. These models must also be monitored and regularly updated if they are to maintain their level of predictability. Given that most organizations cannot simply recruit many more analytic professionals, a more scalable process is required.

An industrialized analytic process emphasizes the use of automation in model construction, both to prepare and analyze data sources and to perform some or all of the modeling itself. It focuses on rapid deployment of models to real-time operational environments and monitors these models automatically to identify when they need to be re-built. Analytic professionals are engaged to handle difficult problems, to check on models that show problems or otherwise to supervise and manage a largely automated production line for analytic models. Supporting this environment requires analytic tools that emphasize scale and automation not just model precision.

Enterprise rule management

For decision logic the problem is slightly different. Reviewing business rules, comparing them to new regulations or policies and making appropriate changes are still manual activities, even when scaling business rules to the whole organization. The challenges come in being able to find the business rules that matter, ensuring the business rules that should be reused are reused, and in handling governance and security policies. When there are many rules that are owned by different groups and when reuse means that no one organization handles all the business rules in a decision, enterprise-scale management capabilities become essential.

A product that allows federated storage of business rules in multiple repositories, that provides robust integration options with other repositories such as those for services and business processes, and that supports a variety of repository structures will be better able to scale. Similarly support for approval workflows, integrated security and good user management capabilities will be important.

Best Practices in Decision Management Systems

There are four key principles of Decision Management Systems:

- Begin with the decision in mind
- Be transparent and agile
- Be predictive not reactive
- Test, learn and continually improve

Within each of these principles it is possible to identify a number of specific best practices in analysis and design, in development, in deployment and in operation.

Begin with the decision in mind.

Decision Management Systems are built around a central and ongoing focus on automating decisions, particularly operational and "micro" decisions. Developing Decision Management Systems with a focus only on business processes, only on events or only on data is not effective. Understanding the business process or event context for a decision is helpful but the development of Decision Management Systems requires a focus on decisions as a central component of enterprise architecture. Focusing on operational or transactional decisions—those that affect a single customer or single transaction—is a significant shift for most organizations and requires a conscious effort. In particular where the operational decision in question is what is known as a "micro decision", one that focuses on a how to treat a single customer uniquely rather than as part of a large group, organizations must learn to focus on decision-making at a more granular level than previously.

It is also worth nothing that this focus on decisions must come first, before a focus on business rules or predictive analytic models. When it comes to developing Decision Management Systems, the right business rules and most effective predictive analytic models can only be developed if there is a clear decision focus. While the most basic best practice is encapsulated in this principle—begin with the decision in mind—there are some more specific best practices that should be followed.

Decisions as peers for Process

One of the most important aspects of building Decision Management Systems is to ensure that decisions start being treated as peers to business processes. Many organizations that are being successful with SOA and that are successfully adopting new and more advanced development technologies and approaches have done so using a business process focus. A focus on the end to end business process, not on organizational or system silos, and the tying of these business processes to real business outcomes represents a significant improvement in how information technology is applied to running an organization.

To move forward with Decision Management Systems, however, it is necessary to do more than regard decisions as just part of a business process. Our work with clients as well as the evaluation of results from multiple companies shows that organizations that can manage decisions as peers to business processes do better. While it is true that decisions must be made to complete most business processes, simply encapsulating the decisions within the business process is not enough.

Decisions are true peers for processes. Decisions are often re-used between processes and how a decision is made has a material difference on how the process executes. Failing to identify decisions explicitly can result in decision-making logic being left in business processes making them more complex and harder to change. Identifying high level decisions at the same time as you identify high level processes allows your understanding of both to evolve in parallel, keeping each focused and simpler.

Link Decisions to business outcomes and results

Your business can be thought of as a sequence of decisions over time. Organizations make strategic decisions, tactical decisions and operational decisions but each decision, each choice, affects the trajectory of the business. In fact, given that each choice you make about products, suppliers, customers, facilities, employees and more is a decision it is clear that decisions are the primary way in which you have a impact on the success or failure of your business. If there is no decision to make then there is no way for the organization to affect its destiny.

One of the first steps, then, in understanding your decisions so that they can drive the development of effective Decision Management Systems, is linking them to business outcomes and results. For each decision you identify it is important to understand what key performance indicators, objectives, or business performance targets are impacted by the decision. Understanding that a particular decision has an impact on a particular measure and understanding the set of decisions that impact a measure has two important consequences. First it enables you to tell the difference between good decisions and bad decisions. A good decision will tend to move the indicators to which it is linked in a positive direction, a bad one will not. Second it enables you to see how you can correct when a measure gets outside acceptable bounds or moves in a poor direction. Understanding which decisions could be made differently gives you an immediate context for solving performance problems.

Building links between decisions you identify and your performance management framework is important as you identify and design your decisions. It is also important to use this information to present options and alternatives to those who are tracking the objectives in a performance management context.

Understand decision structure before beginning

Identifying decisions early, considering them as peers to processes and mapping them to your business performance management environment are all great ways to begin with the decision in mind. Before you start developing a Decision Management System, however, you should understand the structure of your decisions.

The most effective way we have found to do this while working with clients is to decompose decisions to show their dependencies. Decisions are generally dependent on information, on know-how or analytic insight, and on other (typically more fine grained) decisions. Having identified the immediate dependencies of a decision you can the evaluate each of the decisions you identified and determine their dependencies in an iterative fashion.

The dependency hierarchy you develop will actually become a network as decisions are reused when multiple decisions have a dependency on a common sub-decision. This network reveals opportunity for reuse, shows what information is used where and identifies all the potential sources of know-how for your decision making whether regulations, policies, analytic insight or best practices.

For more on this approach please see the author's book *Decision Management Systems* and Alan Fish's *Knowledge Automation* in the list of works cited.

Be transparent and agile.

The way Decision Management Systems make each decision is both explicable to non-technical professionals and easy to change. Decision-making in most organizations is opaque—either embedded in legacy applications as code or existing only in the heads of employees. Decisions cannot be managed unless this decision-making approach is made transparent and easy to change or agile. As noted in **Managing Decision Logic** above this need for both design and execution transparency is the primary driver for the use of a Business Rules Management System to manage decision-making logic.

Three main best practices are relevant in this area—design transparency, execution transparency, business ownership and explicable analytics.

Design Transparency for business and IT

The first best practice in transparency is that of ensuring *design* transparency for both business and IT practitioners. Most code that is written is completely opaque as far as non-technical business users are concerned. Much of it is even opaque as far as programmers other than the one that wrote it are concerned. This lack of transparency is unacceptable in Decision Management Systems.

Design transparency means writing decision logic such that business practitioners, business analysts and IT professionals that were not involved in the original development can all read and understand it. This allows the design of the decision-making to be transparent as everyone involved can see how the next decision is going to be made. This supports both compliance, by allowing those verifying compliance to see how decisions will be made, and improves accuracy by ensuring that everyone who knows how the decision *should* be made can understand how the system plans to make it.

From a practical perspective this means writing all business rules so they can be read by business people (even those that will be edited by IT going forward) by avoiding technical constructs such as ++ and terse programmer centric variable names for instance. It means ensuring that a business friendly vocabulary underpins the rules—the use of IT-centric names for objects and properties is one of the biggest reasons business people cannot understand business rules. It also means using graphical decision logic representations such as decision tables and decision trees whenever possible and following rule writing best practices like avoiding ORs and writing large numbers of simple rules instead of a small number of large complicated ones.

Design transparency is the fundamental building block for all other kinds of transparency and for agility.

Execution transparency and decision logic logging

It is essential to understand how the next decision will be made. Once decisions have been made, however, it will also be necessary to understand how they were made. The approach to the next decision will change constantly as business situations change or new regulations are enforced. The way the next decision will be made therefore diverges steadily from the way a decision was made in the past.

Execution transparency means being able to go back and look at any specific decision to determine exactly how it was made. The decision logic and predictive analytic models used to make the decision must be recorded, logged, so that the decision-making sequence is clear. Ideally this should be "left on" all the time so that every decision is recorded rather than being something that is only used for testing and debugging. When every decision can be analyzed, ongoing improvement becomes much easier. In an environment where any decision can be challenged, by regulators for example, then such ongoing logging may be required.

Most products support logging to a fairly technical format designed for high performance and minimal storage requirements. This will need to be expanded to be readable by non-technical users and integrated with other kinds of data (such as customer information or overall performance metrics) to deliver true execution transparency.

Explicable analytics

While the use of well formed business rules to specific decision logic makes the biggest single contribution to transparency, explicable analytics have a role also. When decisions are made based on specific predictive analytic scores it will be important to be able to understand how that score was calculated and what the primary drivers of the score were. Just like decision logic, the way a score is calculated is likely to evolve over time so it is important that the way a score was calculated at a particular point in time can be recreated.

Some predictive analytic models are more explicable than others. The use of predictive analytic scorecards based on regression models, for instance, allows the contributions to a predictive score to be made very explicit and supports the definition of explanations, reason codes, that can be returned with the score. Thus a customer may have a retention score of 0.62 with two reason codes "Never renewed" and "Single product" that explain where that low score comes from. Decision trees, association rules and several other model types are also easily explicable. In contrast models such as neural networks and other machine learning algorithms as well as compound or ensemble methods involving multiple techniques are often much less explicable.

The value of explicable analytic techniques varies with the kind of decision involved with regulated consumer decisions putting a premium on explicability while fraud detection, for instance, does not.

Business ownership of change

The final best practice is to focus ownership of change in the business. This means empowering the business to make the changes they need to the system when they need those changes made or when they see an opportunity in making a change.

Business ownership of change is not essential for a successful Decision Management System. Many, most, of such systems still use IT resources to make changes when necessary. Often these are less technical resources, business analysts rather than programmers, but it is still IT that makes and tests and changes.

Over time most organizations will find that business ownership will improve the results they get from their Decision Management Systems. By empowering business owners to make their own changes (using capabilities like business user rule management and impact analysis) organizations will increase their agility and responsiveness, eliminating the impedance of the business/IT interface. Empowering the business to own their changes is not a trivial exercise, however, and cannot be simply asserted ("here you go, here's your new business rules interface now please stop calling us"). An investment in suitable user interfaces and tools will be required along with time and energy invested in change management.

Be predictive, not reactive.

Decision Management Systems use the data an organization has collected or can access to improve the way decisions are being made by predicting the likely outcome of a decision and of doing nothing. Decisions are always about the future because they can only impact the future. All the data an organization has it about the past. When information is presented to human decision-makers it is often satisfactory to summarize and visualize it and to rely on a human's ability to extract meaning and spot patterns. Humans essentially make subconscious or conscious predictions from the historical data they are shown and then make their decisions in that context.

When building Decision Management Systems, however, this approach will not work. Computer systems and Business Rules Management Systems are literal, doing exactly what they are told. They lack the kind of intuitive pattern recognition that humans have. To give a Decision Management System a view of the future to act as a context for its decision-making we must create an explicit prediction, a probability about the future. Technology for this is described in **Embedding Predictive Analytics** and **In-database Analytic Infrastructure** above.

Three best practices relate to this focus on turning data into insight. The use of data mining and other analytic techniques to improve rules and analytic/IT cooperation are best practices in development approaches. A focus on real-time scoring will make for more powerful Decision Management Systems.

Using data mining with business rules

Many organizations building Decision Management Systems keep their rules-based development of decision logic and their use of analytics completely separate. At best they only bring the two disciplines together when they reference a predictive score in a business rule. This is a pity and a clear best practice is to do more to drive collaboration in this area, specifically by engaging data miners and data mining approaches in the development of business rules.

To get started with this best practice the first step is to use analytical techniques to confirm and check business rules. Many business rules are based on judgment, best practices, rules of thumb and past experience. The experts involved in defining these rules can often say what the intent behind them is—that a rule is to help determine the best customers or to flag potentially delayed shipments for example. Historical data can be used to see how likely these rules are to do what is intended. For instance the number of customers who meet the conditions in a "best customer" rule or the correlation between the elements tested in the delayed shipment rule and actual delays in shipments. Using data in this way both improves the quality of business rules and helps establish the power of data to improve decision-making. While reporting and simple analysis tools can help in this area, the use of data mining is particularly powerful for these kinds of checks.

More sophisticated organizations can also use data mining to actually find candidate business rules. Many data mining techniques produce outputs that can easily be represented as business rules such as decision trees and association rules. Using these techniques to analyze data and come up with candidate rules for review by those managing the decision-logic can be very effective. Because the output is a set of business rules it is visible and easy to review, breaking down the kind of reluctance that more opaque forms of analytics can provoke.

At the end of the day the best practice is simple to define—organizations should regard their historical data as a source of business rules just like their policies, best practices, expertise and regulations.

Analytic and IT cooperation

The power of predictive analytics is sometimes described as the power to turn vertical stacks of data (data over time) into horizontal information (additional properties or facts). Analytics professionals almost always look at data this way, seeking patterns in historical data that can be turned into probabilities or other characteristics, using analytics to simplify large amounts of data while amplifying its meaning.

The challenge is that IT people do not think of data in the same way. IT departments tend to think of historical data as something to be summarized for reporting and as something to moved off to backup storage to reduce costs or improve performance. They are very familiar with the design of a horizontal slice of the data—its structure—but not with how it ebbs and flows historically. They will often change data structures to improve operations without considering how it might affect historical comparisons, clean data to remove outliers and to include defaults, or overwrite values as time passes and data changes. Many of these kinds of standard IT tasks are very damaging from the perspective of an analytics team.

A clear best practice then is to improve analytic and IT cooperation around data governance, data storage and management, data structure design and more. In this context the analytic team cannot just be the Business Intelligence, dashboard and reporting team but must include those doing data mining and predictive analytics. While the former are often part of the IT department and well integrated with the rest of the IT function, the latter are often spread out in business units or focused in a risk or marketing function. Building cooperation over time between analytic specialists and IT will reduce costs, improve the value and availability of data for more advanced analytics and make integrating analytics into Decision Management Systems easier.

Real-time scoring not batch

A clear majority of organizations applying predictive analytic models today do so in batch. Having developed a predictive analytic model they run daily or weekly updates of their database, adding a score calculated from a model to a customer or other record in the database. When a Decision Management System needs access to the prediction it simply retrieves the column that is used to store the score. Integration is easy because the Decision Management System accesses the score like it does any other data item.

The problem with this is that batch scores can get out of date when data is changing more rapidly than the batch is being run. For instance, a customer propensity to churn score that does not include the problem the customer had this morning or the inquiry they made about cancellation penalties is not going to be accurate. In addition this arms-length integration may be technically simply but it also keeps the IT and analytic teams from needing to work together and is therefore potentially damaging in the long term.

For long term success with Decision Management Systems, and in particular to develop the kinds of Decision Management Systems that will allow you effective response to events and new more mobile channels, organizations need to develop systems that use real-time scoring. A real-time score is calculated exactly when it is needed using all the available data at that moment. This might include recent emails, SoMoLo (Social Mobile Local) data, the opinion of a call center representation on the mood of the customer and much more. Ultimately being able to decide in real time using up to the second scores, or even score data as it streams into a system so that predictions are available continuously, will be a source of competitive advantage.

Test, learn and continually improve.

The decision-making in Decision Management Systems is dynamic and change is to be expected. The way a decision is made must be continually challenged and re-assessed so that it can learn what works and adapt to work better. Supporting this kind of ongoing decision analysis requires both design choices in the construction of Decision Management Systems and integration with an organization's performance management environment. Both Business Rules Management Systems and Predictive Analytic Workbenches have functionality to make this easier while Optimization Suites can be used to develop models to manage the potentially complex trade-offs that improving decision-making will require.

This kind of continuous improvement relies on many of the features noted earlier such as being able to link decisions to business outcomes and results, having execution transparency and decision logic logging and support for real-time scoring not batch. In addition the development of integrated environments for ongoing decision improvement, broad use of experimentation and moving to automating tuning, adaptive analytics and optimization are all best practices worth considering.

Integrated decision improvement environment

To provide an integrated decision improvement environment, organizations should bring together the logs they have on how decisions have been made in the past, information about the business results they achieved using these decisions and the decision logic/analytic management environment itself. Each piece of this environment typically involves a different piece of technology to develop with everything from a business rules management system to an analytic model management tool to traditional dashboard and business intelligence capabilities being used. Providing an integrated, coherent environment where all this is brought together around a particular decision offers real benefits to an organization. When business results can be compared to the decision making that caused them and when the business owner can navigate directly from this analysis to editors allowing them to change future decision-making behavior, organizations will see more rapid and more accurate responses to changing conditions.

Broad use of experimentation

Relatively few organizations are comfortable with experimentation. For most, experiments are confined to the marketing department or to low volume experiments where customers and prospects are quizzed on preferences or likely responses. Some organizations use experimentation to determine price sensitivity and a growing number of web teams use experimentation for website design.

Yet without experimentation it is very hard to see if what you are doing is the best possible approach or to truly see if a new approach would work better. Unless the behavior of real customers or prospects (or suppliers or partners) is evaluated for multiple options, those options cannot really be compared. Asking people what they would do if they got a different option rarely results in data that matches what they actually do when they get that different option.

Organizations that wish to succeed in the long term with analytics and with Decision Management Systems will invest in the organizational fortitude and expertise required to conduct continuous and numerous experiments.

Moving to automated tuning and adaptive analytics

The logical extension of a focus on real-time is to focus on automated tuning and adaptive analytics. Today most Decision Management Systems and the analytics within them are adapted manually, with experts considering the effectiveness of the decision and the making changes to improve it. As systems become more real-time, however, this becomes increasingly impractical and suboptimal. Especially in very high volume, quick response situations such as ad serving, the system is continually gathering data that shows what works and what does not. Waiting until a person considers this data before changing the behavior of the system means allowing the system to make poor responses long after the data exists to realize this is going on.

The best practice is to consider the use of machine learning and adaptive analytic engines in these circumstances. Building trust in the organization that analytics work will increasingly allow analytic systems to be left to make more of the decision themselves. Allowing analytic engines to collect performance data and respond to it, perhaps within defined limits, will improve the performance of real-time decision making while reducing the length of time it takes to respond to a change.

Not all decisions are suitable for these kinds of engines. For instance those decisions that have a strong regulatory framework or where the time to get a response to a decision is long will not work well. Where a decision is suitable, however, a clear best practice is to integrate these kinds of more adaptive engines into Decision Management Systems.

Optimization

One final best practice in this area is to increase the use of optimization over time. A powerful approach, optimization is often siloed into specific parts of the business and regarded as a little bit of a side bar to “core” analytic efforts. In part this is because the mathematics can be very complex and because the solutions can take a long time to develop. A lack of business user friendly interfaces for reviewing results and a need to integrate optimization with simulation tools also limit the use of optimization in many organizations.

This is beginning to change, however, as more business friendly interfaces are developed and as optimization tools become more integrated into the overall stack for developing Decision Management Systems. Faster and more stable optimization routines, standard templates and integration with both predictive analytics and business rules are also helping. Organizations should regard the use of optimization as part of their decision design and improvement processes as a best practice and should seek therefore to bring it out of its silos and into the mainstream.

Use Cases

There are many compelling use cases for Decision Management Systems. In future releases of this report we will expand on these and include real customer stories to illustrate the power of Decision Management Systems. For now here are some of the most effective use cases for Decision Management Systems.

Fraud Detection

Many organizations suffer from losses caused by fraud and abuse. These range from fraudulent claims for services that were never performed to applications for credit for people that don't exist to orders that include bribes and illegal payments. In every case an organization must decide whether to accept the transaction as valid, reject it or investigate it for fraud. These decisions are high volume as they must be made for each transaction and are ideal for automation using a Decision Management System. Fraud detection systems typically involve business rules for compliance with policies and regulations as well as predictive analytics to match the current transaction to patterns known to be fraudulent or identify that the current transaction looks very different from legitimate ones.

Customer Next-Best-Action

Organizations focused on becoming customer-centric are increasingly turning to an approach known as next best action or next best activity (some more grammatically precise organizations talk about best next action). Such an approach involves considering every action that the organization could take towards a customer—making a cross-sell offer, collecting new information about customer preferences, reminding them to use a product they already own—and ensuring that each opportunity for interaction uses the "best" one for long term customer value. This focus on actions—not just offers—and a desire to centralize and systematically improve the selection of the best action drives a need for a Decision Management System focused on this decision. These systems involve business rules about who to contact and when as well as definitions of product or action eligibility while predictions of propensity to accept and of likely future profitability are at the heart of effective choices.

Targeted Marketing

Organizations are trying to ensure that their marketing is more relevant and targeted. They are personalizing with everything they know about a prospective customer. Moving beyond just using names and locations, these systems are making a micro decision about each prospect, generating messages and contact strategies specific to that prospect. Combing business rules and predictive analytics to effectively target every prospect, this approach to targeted marketing relies on a Decision Management System at its core.

Underwriting

One of the most established areas for Decision Management Systems is that of underwriting. Whether underwriting loans, mortgages, credit or insurance, underwriting decisions offer a strong use case for Decision Management Systems. Often regulated and constrained by policy, underwriting decisions can be effectively managed using business rules. An assessment of risk is often critical to deciding what price or terms to offer. Higher risk customers must provide more documentation or pay a higher interest rate. The use of predictive analytic models to predict risk in these circumstances is well established. Combining the business rules and predictive analytic models into a Decision Management System is a very effective tool for automating the underwriting decision.

Eligibility and Compliance

The use of a business rules-based system to determine eligibility or to ensure that a transaction is being handled in a compliant way is increasingly common. Decisions such as "is this person eligible for this product/service" or "what is the required next step" are policy and regulation-heavy and the use of a Business Rules Management System to handle all the business rules is very effective. While eligibility and compliance decisions can seem fairly static, changes are often outside of the control of an organization and can be imposed at short notice.

Suitable Operational Decisions

Besides these specific use cases, there are some key characteristics of decisions that make them suitable for automation using Decision Management Systems. These are discussed more fully in the book but suitable decisions have four characteristics:

- Repeatable.
  If a decision is not made in a repeatable way and made regularly it will not be possible to automate it nor to show a return on doing so.
- Non-trivial.
  A decision must have a degree of complexity to make it worth the investment in additional capabilities discussed above. There must be policies or regulations that drive and control the decision, a degree of expertise involved in making it well or some analysis of information required.
- Measureable business impact.
  It must be possible to tell what the business impact of improving the decision will be, and even what a good decision is relative to a bad one. If the value of improvement cannot be described or worse yet the value of a decision cannot be measured at all then it will not be possible to show the value of a Decision Management System.
- Suitable for automation.
  Every organization has a different attitude to automation. Unless the organization is willing to consider a system to make the decision there is no point in building a Decision Management System. A decision that must be taken by a person might involve dependent decisions that are suitable candidates for automation but building a Decision Management System to automate a decision that an organization believes should be taken by a person will result in a system that does not get used.

Suitable decisions often break down into rules-centric decisions such as eligibility, validation and calculation and analytic-centric decisions such as those relating to risk, fraud and opportunity.

Selecting Vendors

Each Decision Management System requires different subsets of the capabilities described above. The right set of vendors and products is going to vary, depending on the requirements and needs of both the project and the organization as a whole. There are many vendors large and small to choose from, with more being added every day. The products they offer have great breadth and depth of functionality in every area. Every one of the products listed continues to evolve and grow, adding new functionality and enhancing existing capabilities. Some vendors are merging to create complete product sets or suites under one umbrella while others are collaborating to allow their products to be used together more effectively. PMML already provides some standards support for this collaboration and new standards are on the horizon that will extend this.

There are a wide range of vendors available for each of the product categories you need to develop Decision Management Systems. Many organizations will have existing relationships with vendors and will use other software products they provide. Experience with clients shows that familiarity and comfort with a vendor, confidence that you can work well with them, is a strong predictor of success. Some organizations will work with Systems Integrators or other service providers who have strong vendor relationships that will likewise contribute powerfully to a successful project. The "fit" of the vendor(s) you select with your organization is often much more important than the specifics of their functionality.

All the vendors in this report have paying customers who are successfully using their technology to build some kind of Decision Management System. There is no magic or "best" set of vendors or products. There is a rich set of vendors and products and most, if not all organizations could pick from multiple vendors or vendor combinations and be successful.

In future releases of the report we will drill into more specifics on how to pick vendors for Decision Management System projects. Some things to consider:

- When both decision logic and analytic insight must be combined in a Decision Management System, those Business Rules Management Systems and related products that are "model-aware" and can consume and integrate with predictive analytic models are more likely to be successful than those that are not.
- A Predictive Analytic Workbench that supports a range of deployment options for predictive analytic models into production and can monitor and manage these models will generally require less integration work than those that do not.
- Predictive Analytic Workbenches that support in-database modeling and in-database scoring (directly or through partnerships with others) are increasingly valuable.

- An optimization environment that supports the generation of business rules (often through integration with data mining capabilities) as well as solving to produce a set of actions will provide more deployment options.
- Those components that support standard platforms and provide a rich set of APIs and thin client interfaces are generally preferred.
- Depending on the system involved, a focus on real-time or on batch (or on a mixture of the two) will be essential as will an understanding of the need for support for.Net, Java or legacy platforms.
- The view taken by the organization of open source products may constrain or focus your selection process.

Vendors

As noted, this report is prompted by the growing interest by organizations in building decision management systems and/or adopting decision management domain-specific applications. It is always challenging to draw product boundaries, but for practical purposes it must be so. For this report, we are focused on platform technologies used to build custom decision management systems. This leads to some specific criteria for inclusion:

1. While there are many innovative companies developing complete solutions – pre-configured Decision Management Systems – we will be focusing in this report on tools for building custom systems. To be included it must be possible to build a custom Decision Management System with the product.
2. The technologies must be productized, released, for sale, and must have at least one production customer who can be contacted. Technologies used exclusively by the vendor themselves to develop solutions are therefore not eligible, even if completely custom solutions are being developed for specific customers.
3. The product should be sold generally, not only to companies that have also bought a pre-configured application – it must have a market presence as a genuine platform technology.
4. The products must have decision management, business rules, analytic or optimization functionality that has not been licensed from a vendor that is included separately in the report. They can include licensed technology as part of the product but they must have genuinely distinct functionality in addition.

If you know of other products you believe should be included or have other feedback, please let us know: info@decisionmanagementsolutions.com.

This lists vendors and basic company contact information alphabetically. Although every attempt has been made to validate this information, Decision Management Solutions accepts no liability for the content of this report, or for the consequences of any actions taken on the basis of the information provided.

Advertisements are the responsibility of the vendor concerned and have not been reviewed by Decision Management Solutions.

The vendor section has two parts. First there is a list of vendors with their contact information. Many vendors listed above have multiple products some of which meet the criteria for inclusion in the report and some of which do not. A single eligible product is enough to get the vendor listed.

The table of products that follows includes listings of individual products as well as links to First Look reviews, if any, on JTonEDM.com. Only the most recent product First Look is linked, though previous First Looks are generally referenced when a new one is written. Some First Looks cover multiple products and in these circumstances both products are listed and linked to the same First Look.

11Ants Analytics

www.11antsanalytics.com

Hamilton, New Zealand

Contact

Angoss

www.angoss.com

Toronto, Canada

Contact

Be Informed

www.beinformed.com

Apeldoorn, the Netherlands

Contact

BigML

bigml.com

Corvallis, OR

Contact

Bosch Software Innovations

www.bosch-si.com

Chicago, IL

Contact

Clario Analytics

www.clarioanalytics.com

Eden Prairie, MN

Contact

Code Effects Software

codeeffects.com

Alpharetta, GA

Contact

Dymatrix Consulting Group GmbH

www.dymatrixconsultinggroup.com

Suttgart, Germany

Contact

EigenDog

www.eigendog.com

Mountain View, CA

Contact

Experian Decision Analytics

www.experian.com/powercurve

San Diego, CA

Contact

FICO

www.fico.com

Minneapolis, MN

Contact

Fuzzy Logix

www.fuzzyl.com

Charlotte, NC

Contact

GDS Link

www.gdslink.com

Dallas, TX

Contact

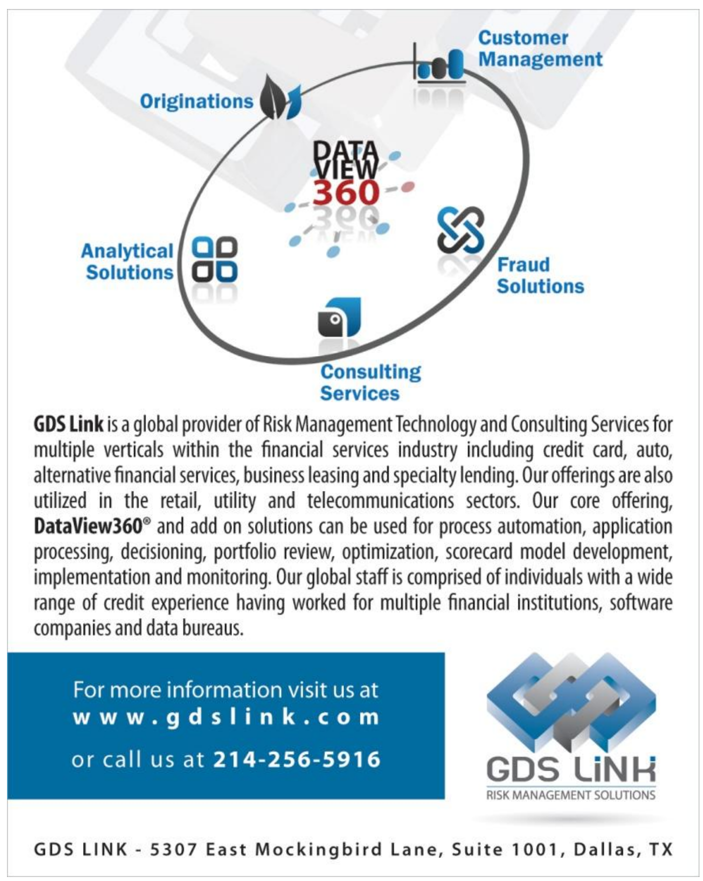

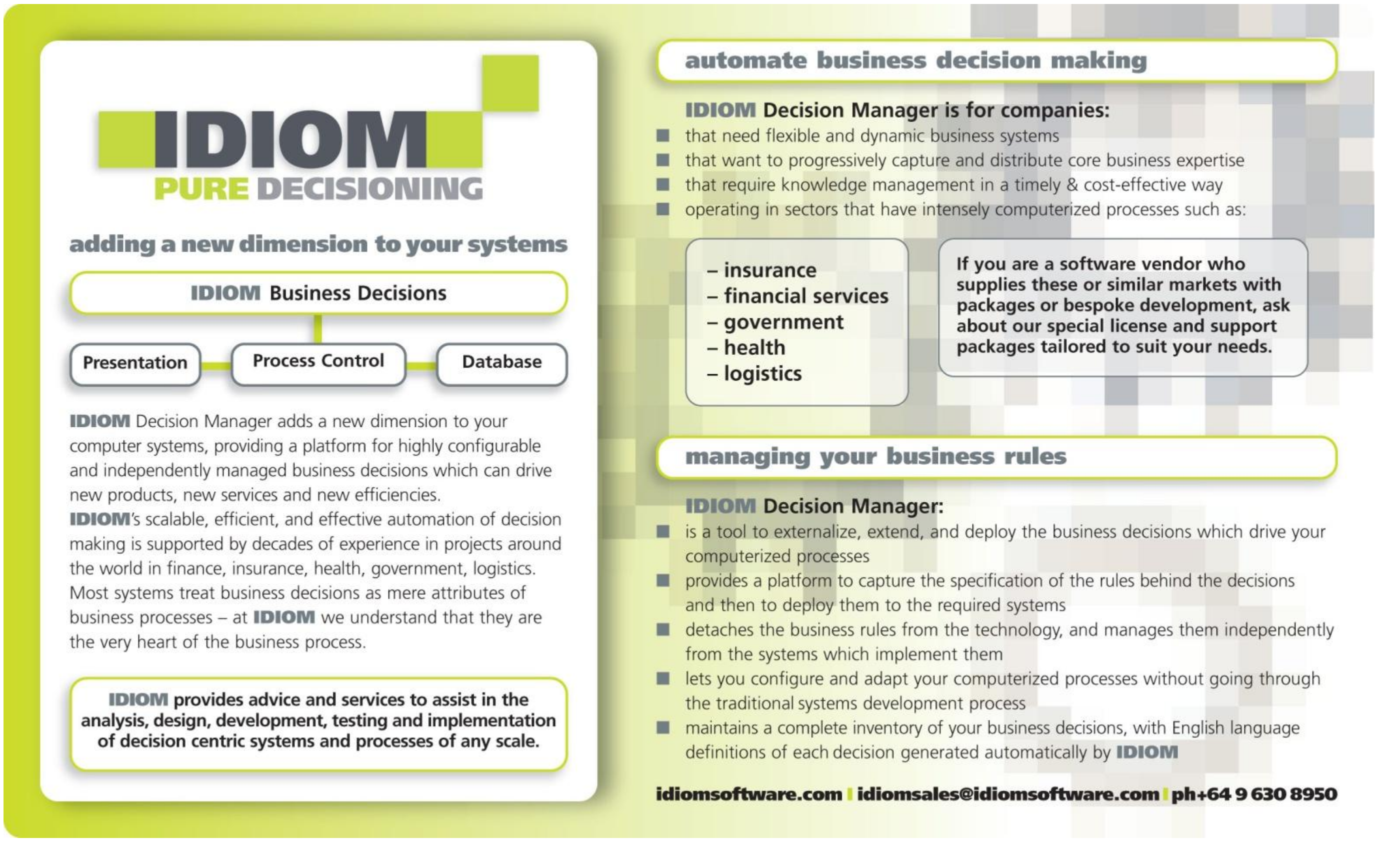

Gurobi Optimization

www.gurobi.com

Houston, TX

Contact

IBM

www.ibm.com/decision-management

Armonk, NY

Contact

IDIOM Software

www.idiomsoftware.com

Auckland, New Zealand

Contact

Inkiru

www.inkiru.com

Palo Alto, CA

Contact

InfoCentricity

www.infocentricity.com

Novato, CA

Contact

InRule Technology®

www.inrule.com

Chicago, IL

Contact

KNIME

www.knime.org

Zurich, Switzerland

Contact

KXEN

www.kxen.com

San Francisco, CA

Contact

OpenRules

www.openrules.com

Edison, NJ

Contact

Oracle

www.oracle.com

Redwood Shores, CA

Contact

Pegasystems

www.pega.com

Cambridge, MA

Contact

Pervasive Software

www.pervasive.com

Austin, TX

Contact

Predixion Software

www.predixionsoftware.com

San Juan Capistrano, CA

Contact

Progress Software

www.progress.com

Bedford, MA

Contact

Quiterian

www.quiterian.com

Barcelona, Spain

Contact

Rapid-I

www.rapid-i.com

Dortmund, Germany

Contact

Rapid Insight Inc.

www.rapidinsightinc.com

Conway, NH

Contact

Red Hat

www.redhat.com

Raleigh, NC

Contact

Revolution Analytics

www.revolutionanalytics.com

Palo Alto, CA

Contact

Salford Systems

www.salford-systems.com

San Diego, CA

Contact

SAP

www.sap.com

Walldorf, Germany

Contact

Sapiens

www.sapiensdecision.com/

Cary, NC

Contact

SAS

www.sas.com

Cary, NC

Contact

Skytree

www.skytreecorp.com

San Jose, CA

Contact

Sparkling Logic

www.sparklinglogic.com

Sunnyvale, CA

Contact

StatSoft

www.statsoft.com

Tulsa, OK

Contact

Starview

www.starviewinc.com

San Jose, CA

Contact

Teradata

www.teradata.com

Dayton, OH

Contact

Tibco

www.tibco.com

Palo Alto, CA

Contact

USoft

www.usoft.com

Baarn, the Netherlands

Contact

Yottamine Analytics

www.yottamine.com

Bellevue, WA

Contact

Zementis

www.zementis.com

San Diego, CA

Contact

Zoot Enterprises

www.zootweb.com

Bozeman, MT

Contact

Table 1: Product Reviews

Vendor	Product	Review
11Ants Analytics	11Ants Model Builder	First Look
	11Ants Predictor	First Look
Angoss	KnowledgeSEEKER	First Look
	StrategyBUILDER	First Look
	KnowledgeSTUDIO	First Look
	In-database analytics for Teradata	First Look
Be Informed	Be Informed 4 Business Process Platform	First Look
BigML	BigML	First Look
Bosch Software Innovations	Visual Rules Suite for BRM	First Look
Clario Analytics	Clario Core	First Look
Code Effects Software	WebRule	First Look
Dymatrix	DynaMine	First Look
EigenDog	EigenDog	First Look
Experian	PowerCurve™	First Look
FICO	FICO® Blaze Advisor® business rules management system	First Look
	FICO® Model Builder	First Look
	FICO® Xpress Optimization Suite	First Look
	FICO® Model Central™ Solution	First Look
	FICO® Decision Optimizer	Pending
Fuzzy Logix	DB Lytixs™	First Look
GDS Link	DataView 360	First Look
	Modellica Decision Engine	First Look
	DecisionIntelligence	Pending
Gurobi	Optimizer	First Look
IBM	Analytical Decision Management	First Look
	Operational Decision Management	First Look
	In-Database Analytics	First Look
	SPSS Modeler	First Look
	CPLEX	Pending
IDIOM	Decision Manager	First Look
Infocentricity	Xeno	First Look
Inkiru	Predictive Intelligence Platform	First Look
InRule Technology®	InRule	First Look
KNIME	KNIME	First Look
KXEN	InfiniteInsight® Express	First Look
	InfiniteInsight®	First Look

Vendor	Product	Review
OpenRules	OpenRules BDMS	First Look
	OpenRules Rules Solver	First Look
	OpenRules Rules Learner	First Look
Oracle	Advanced Analytics	First Look
	Oracle Data Mining	First Look
	Real-Time Decisions	First Look
	Business Rules	First Look
Pegasystems	Decision Management	First Look
Pervasive	RushAnalyzer	First Look
Predixion Software	Insight	First Look
Quiterian	Analytics	First Look
Progress	Corticon	First Look
Rapid-I	RapidMiner	First Look
Rapid Insight	Analytics	First Look
Red Hat	JBoss Enterprise BRMS	First Look
Revolution Analytics	R Enterprise	First Look
Salford Systems	Salford Predictive Modeler	First Look
SAP	BRFplus	First Look
	SAP NetWeaver Business Rule Management	First Look
	SAP NetWeaver Decision Service Management?	Pending
	Predictive Analytics	Pending
Sapiens	DECISION	First Look
SAS	High-Performance Analytics	First Look
	Enterprise Miner	First Look
	Model Manager	First Look
	SAS/OR	First Look
	Rapid Predictive Modeler	First Look
Skytree	Server	First Look
Sparkling Logic	SMARTS	First Look
Starview	Starview Active Analytics Platform	First Look
StatSoft	STATISTICA	First Look
Teradata	Asterdata	First Look
Tibco	Spotfire	First Look
USoft	URequire Studio	First Look
Yottamine	Predictive Platform	First Look
Zementis	ADAPA	First Look
	Universal PMML Plug-in	First Look
Zoot	zDecision	First Look

Appendix - Decision Management Systems

Decision Management Systems are a new class of system. Decision Management Systems bring together two kinds of systems—operational systems that manage the transactions of the business and analytic systems that help you understand how to run the business better— to deliver systems that actively work to help you run your business or organization. Decision Management Systems are agile, analytic and adaptive and are built using a three step process of decision discovery, decision services and decision analysis. Decision Management Systems deliver high ROI by reducing fraud, managing risk, boosting revenue and maximizing the value of scarce resources.

There is more information on Decision Management Systems, their characteristics and the process of building them in the author's book "Decision Management Systems: A Practical Guide to Business rules and Predictive Analytics" (IBM Press, 2012). This book is organized into three parts.

Part I: The Case for Decision Management Systems

The first three chapters make the case for Decision Management Systems—why they are different and how they can transform a 21st century organization.

- **Chapter 1, "Decision Management Systems are different"**: This chapter uses real examples of Decision Management Systems to show how they are agile, adaptive and analytic.
- **Chapter 2, "Your business is your systems"**: This chapter tackles the limits of manual decision-making, showing how modern organizations cannot be better than their systems.
- **Chapter 3, "Decision Management Systems transform businesses"**: This chapter shows that Decision Management Systems are not just different from traditional systems—they represent opportunities for true business transformation.
- **Chapter 4, "Principles of Decision Management Systems"**: This chapter outlines the key guiding principles for building Decision Management Systems.

Part II: Building Decision Management Systems

Chapters 5 through 7 are the meat of the book, outling how to develop and sustain Decision Management Systems in your organization.

- **Chapter 5, "Discover and model decisions"**: This chapter shows how to find, describe, understand and model the critical repeatable decisions that will be at the heart of the Decision Management Systems you need.

- **Chapter 6, "Design and implement Decision Services"**: This chapter focuses on using the core technologies of business rules, predictive analytics and optimization to build service-oriented decision-making components.
- **Chapter 7, "Monitor and improve decisions"**: This chapter wraps up the how-to chapters, focusing on how to ensure that your Decision Management Systems learn and continuously improve.

Part III: Enablers for Decision Management Systems

The final part documents people, process and technology enablers that can help you be successful.

- **Chapter 8, "People Enablers"**: This chapter outlines some key people enablers for building Decision Management Systems.
- **Chapter 9, "Process Enablers"**: This chapter continues with process-centric enablers, ways to change your approach that will help you succeed.
- **Chapter 10, "Technology Enablers"**: This chapter wraps up the enablers with descriptions of the core technologies you need to build Decision Management Systems.

The Characteristics of a Decision Management System

Decision Management Systems have three critical characteristics. These characteristics strongly differentiate them from current, mainstream business applications. Most such business applications are difficult, expensive and time consuming to change. Decision Management Systems are agile and transparent. The business applications that support most organizations are entirely separate from the analytic environment of those organizations. Decision Management Systems are both operational *and* analytic. Finally most business applications are designed and built to meet a specific set of requirements that is known and not expected to change. Decision Management Systems are adaptive, learning and improving as they are used.

Decision Management Systems are Agile

Business Agility is an overused expression in corporate IT with all manner of approaches and technologies being promoted as delivering business agility in some way. Decision Management Systems are agile because the logic in them is easy to change and easy to adapt to changing circumstances. When new policies or regulations are issued the logic that implemented them can easily be found and safely be changed. These changes don't undermine compliance because Decision Management Systems are transparent—it is clear how they will work in the future and also clear how they acted in each specific historical situation. This agility allows more stable business processes, as changes are easy to make to the Decision Management Systems that support those processes, and ensures that rapidly changing know-how and experience can also be effectively embedded in systems without the danger that it will become stale and out of date.

Decision Management Systems are Analytic

Analytics is a hot topic in many organizations today. Yet most analytic systems are completely separate from the operational systems that run the business. These analytic systems rely on data extracted from the operational systems but are otherwise quite standalone. In contrast, Decision Management Systems deeply embed analytic insight to improve their operational behavior. Analytics are used to divide customers or transactions into like groups to allow actions to be effectively targeted. Analytics are also used to make predictions of the degree of risk involved in a transaction, the likelihood of fraud or the extent and type of opportunity available. These predictions are used to select from the available alternatives in a way that will manage risk according to the organization's guidelines, reduce fraud, maximize revenue and effectively allocate resources across competing initiatives.

Decision Management Systems are Adaptive

Business systems, like business people, need to constantly adapt and learn. They need to experiment and see if a new approach might work better than a long established one, challenging conventional wisdom. They must manage trade-offs in an ever changing business climate. They must allow their performance to be monitored in terms of how effective the decisions they make turn out to be. In this way Decision Management Systems are adaptive, built to respond to changing conditions and to support a process of continuous improvement through testing and experimentation.

Building a Decision Management System

Building Decision Management Systems involves many of the same techniques, tools and best practices that building any reliable, high-performance operational system involves. All the skills and experience an organization has in developing information systems apply. The new and changed activities required fall into three phases—decision discovery, decision services and decision analysis.

Decision Discovery

Decision Management Systems are focused on automating and improving decisions. Most organizations do not have a well defined approach for finding, modeling and managing the decisions they make. To effectively build Decision Management Systems, then, the first step is to find the repeatable, non-trivial decisions in the organization that have a measurable business impact and are therefore candidates for automation and improvement. Examples of suitable decisions include checking the eligibility of a person for a government benefit or commercial product, validating that an organization can become a supplier or meets some defined criteria, pricing a loan or other financial instrument based on an assessment of the risk involved, and making an offer to a consumer to maximize the value of an opportunity to interact with them.

There are a number of ways to find these decisions. They can often be found explicitly simply by interviewing and working with business experts. The tasks in business processes where choices are being made or where there is a pause for review are typically decision-making tasks. Many branches in processes are preceded by decision-making. Decisions can also be found by analyzing Key Performance Indicators and other metrics to see what choices make a difference to those metrics. It is unusual for something to be tracked as a metric if there are no choices made that cause it to go up or down. The top level decisions that these approaches find should be described, primarily by defining a question that must be answered to make the decision along with the allowed or possible answers. For instance, a claims review decision might answer the question "Is this claim fraudulent and what should we do about it?" with allowed answers including routing it to the fraud investigators, putting it through a regular claims review or fast tracking it for immediate payment.

Top level decisions can and should be decomposed into the subordinate decisions they are dependent on—the smaller decisions that must be made before the top level one can be made. This decomposition is recursive and provides necessary detail on how these decisions are actually made day to day.

Decision Services

The decision discovery step enhances traditional analysis and requirement gathering tasks. Once the decisions are identified and modeled, an iterative process of development can begin. The objective is to develop Decision Services—coherent, well defined components that make a decision for the other processes and system components in the solution. Beginning by defining simple interfaces that allow these services to be asked a question and give back one of the allowed answers, an iterative approach is used to flesh out the decision making. The decomposition of the decision shows the sequencing and structure of the decisions and the business rules, predictive analytic models and optimization models required can be developed and added to this structure.

A complete Decision Service will require some combination of business rules, predictive analytic models and optimization models. Most will not require the deployment of optimization models. Optimization is more likely to be applied to a large number of similar decisions with the optimal actions identified for each decision used to derive new business rules that are more likely to result in optimal decisions in the future. Some decision services will not require predictive analytic models, especially those primarily concerned with eligibility and compliance where business rules dominate. Even when analytic insight is important to a decision, sometimes that insight can be best represented with a set of business rules mined from historical data. When probabilities are needed, however, predictive analytic models will either need to be executed by the decision service, executed in the database the decision service is using or stored in the database that the decision service used if a batch update of the prediction is acceptable.

Decision Analysis

Decision services are developed and deployed as part of an overall systems development effort. Once deployed they must be monitored and analyzed to see if changes are required going forward. Decision services should be monitored for both proactive and reactive changes—changes that might help improve performance as well as necessary changes for compliance for instance. Performance management and other analytic tools can be used to assess the effectiveness of the decision making embedded in the system. As changes are identified and proposed it must be possible to effectively assess the impact of these changes before they are deployed. It may also be necessary or desirable to design new approaches and conduct new experiments to gather new data about what works and does not. Any changes made should be monitored to make sure they work as expected.

Making the case for a Decision Management System

There are many ways to make a case for a Decision Management System. The cost of development of a Decision Management System plus the additional software required to develop it must be offset by a satisfactory return if a case is to be made. The top-line benefits of a Decision Management System typically come in a number of areas:

- Reduced losses by eliminating fraud and waste.
  Predictive analytic models that establish the probability of a transaction being fraudulent or wasteful can be combined with business rules for known fraud schemes and waste prevention policies to determine which transactions to reject or to refer for investigation. Decision Management Systems have an excellent track record in dramatically reducing fraud.
- Reduced risk exposure and better matching of risk to price.
  It has been said that there is no such thing as a bad risk, only a bad price. Using predictive analytic models to predict the risk of a loan or a policy and then applying business rules to correctly price for the predicted risk is a well established use case for Decision Management Systems. Decision Management Systems can manage more fine-grained risk models, with dozens or hundreds of segments, and more closely match risk to price.
- Increased revenue from targeted marketing and correct opportunity identification.
  When an organization has an opportunity it can be unclear how to maximize its value. As opportunities migrate across channels and as windows of opportunity get narrower, this gets even harder. Decision Management Systems can apply

campaign and offer rules with predictive analytic models that estimate the propensity for specific offers to be accepted and to be profitable. The offer that is most suitable, most likely to be profitable can then be made even if only a short time is available. Optimization can be used when there are constraints on how many offers can be made or on capacity, ensuring maximum return on those limited assets. Deciding on the next best action across channels ensures a focus on long term customer value and increased revenue over time.

- Productivity and maximized use of business assets—physical and people.
  Decision Management Systems handle large numbers of routine decisions, freeing up people to focus on more complex, higher value tasks. They can also assign people to those tasks most likely to see a return on that investment, such as assigning collections activities based on the likely total return. These same approaches can maximize the use made of limited assets, deciding how best to handle them at each moment.
- Faster time to market and shorter time to respond.
  In many of these different areas the time it takes an organization to get to market with something or to respond to a change is critical. By making it easier and quicker to add new business rules, Decision Management Systems can improve time to market and so add additional value. This can be particularly effective as part of a legacy modernization effort where hard to change components are upgraded to Decision Management Systems for increased agility and lower management costs.

Finally, when calculating the costs of a Decision Management System, it should be noted that these technologies often result in cheaper development relative to the equivalent traditional coding approaches. Decision Management Systems are also often dramatically cheaper to maintain. This both reduces the total cost of ownership of the system when maintenance costs are included and makes it more likely that any given improvement in the system (to match a competitor's move or to take advantage of a fleeting opportunity) will actually be attempted.

Bibliography

Fish, Alan. *Knowledge Automation: How to Implement Decision Management in Business Processes.* New York, NY: John Wiley & Sons, Inc, 2012.

Fisher, Ronald A. *The Design of Experiments, 9th Edition.* Macmillan, 1971.

Forgy, Charles. *On the efficient implementation of production systems.* Ph.D. Thesis, Carnegie-Mellon University, 1979.

Nisbet, Robert, John Elder, and Gary Miner. *Handbook of Statistical Analysis and Data Mining Applications.* Burlington, MA: Elsevier, 2009.

Taylor, James. *Decision Management Systems: A Practical Guide to Using Business Rules and Predictive Analytics.* New York, NY: IBM Press, 2012.

Taylor, James, and Neil Raden. *Smart (Enough) Systems: How to deliver competitive advantage by automating hidden decisions.* New York, NY: Prentice Hall, 2007.

Contact Us

If you have any questions about Decision Management Solutions or would like to discuss engaging us we would love to hear from you. Email works best but feel free to use any of the methods below.

Email : info@decisionmanagementsolutions.com

Phone : +1 650 400-3029

Fax : +1 650 352-9247