Link corporate policy and data design with

Decisioning

by Mark W Norton, February 2009

Abstract: The concept of decisioning recognizes that decision-making is the proprietary know-how that drives value and differentiates one organization from another in the marketplace. Because of this fundamental importance, decision-making is usually apparent in corporate policy, from where it can be distilled into precise decision models that can be used to drive data analysis. Data is collected for a purpose, and that purpose has its roots in decisioning. This paper discusses the methods which can be used to discover and model decisions in a structured manner, analogous to data normalization. Data and decisions both have important and complementary roles in this decision-centric approach. Data models show the valid states of the system at rest; decision models describe the valid transitions between the states. However, it is the state transition described by the decision models that generates value for any business, giving the decision model a primacy that is not shared by data.

About the Author

Mark Norton has 30 years development experience, mostly on enterprise scale, mission critical systems in finance, insurance, government, and health administration. In 2001 Mark established Idiom Ltd. to develop and market tools and techniques in support of the decisioning concept. Through joint project participation with integration partners, these tools and techniques have been personally developed and tested on dozens of projects to deliver the pragmatic and conceptually sound decision centric development approach as described in this article.

Contact: +64 21 434669
mark.norton@idiomsoftware.com
www.idiomsoftware.com

Decisions and Data.

The discovery and normalization of data is a proven and accepted design approach that underpins systems design activity. This article proposes that we preface this approach with a new and fundamental analysis artifact the decision. Like both data and process, decisions can be the subject of original discovery and can drive further design activity.

Further, we will argue that decision modeling is a higher order analysis that provides context and guidance to both data and process analysis. Data is collected for a purpose, and that purpose has its roots in decisioning. Similarly, processes exist to service decisions by connecting them to events. A further purpose of this paper is to discuss the methods which can be used to discover and model decisions in a structured manner, analogous to data normalization. Data and decisions both have important roles in this decision-centric approach each showing a different and complementary view of the same system. Data models show the valid states of the system at rest; decision models describe the valid transitions between the states. However, it is the state transitions described by the decision model s that generate value for any business, giving the decision model a primacy that is not shared by either data or process.

In fact, we argue that it is the corporate strategy itself that defines when and how 'value creating' state changes occur, and that this is the ultimate foundation upon which decisioning approaches are built.

Decisioning

Defn: The discrete and systematic discovery, definition, assembly and execution of decisions.

We assert that decisions are first order citizens of the requirements domain. To provide a conceptual basis for this, let us go back to the foundations of any requirements analysis.

A business derives value from changing the state of some object or 'thing' that the business values (literally, usually on its balance sheet). This focuses us on the very core of the business what is traded, managed, leveraged, used, built, or sold in order to achieve the objectives of the business. Until something happens to the object, the purpose cannot be achieved and no value can be derived; therefore, in order to generate value, we must have a state change on the object. Note that we are interested primarily in the objects that are the basis of value and which are fundamental to the purpose of the business. We are not interested in data entities per se, including such traditionally important entities as customer (we do not generally derive value from changing the state of customers). Changing state implies that there is activity against an object, and this observation gives rise to the traditional process (activity) and data (object) approaches. But if we look closely at any process causing state change, we will see that the change is always the result of a decision within the process rather than the process itself. Many processes can execute against the object, but value is only derived when a state change occurs. Whenever state change occurs, the process can be shown to be a container for decisioning the state change is a result of the decisions rather than an inherent characteristic of the process. This confusion between decision and process is a systemic failure in today s methodologies and hides the importance of decisioning to the extent, for instance, that UML does not include business rules (the closest existing concept we have to decisioning) within the UML standard[1]. Processes are merely mechanisms by which data is supplied to a decision, and by which the system responds to the decisions made. If these pre and post decision activities are wrapped together with the decisioning logic, then we have a complete end-to-end process that takes the business from one valid state to another, and which generates value in doing so. But it is clear that the entire process is built around the decisioning kernel.

Decisions are the critical mechanisms by which we choose one action from the set of all possible actions. The purpose of decisions is to select the action that is most beneficial to the decision maker in a given situation.

A decision is made when we resolve the available facts into a single definitive outcome; a decision cannot be multi-valued, ambiguous or tentative. A decision results from applying discipline, knowledge and experience to the facts that describe the decision context. This application of discipline, knowledge and experience within decision-making is the characteristic that most closely defines the unique character of any business. Because of this fundamental truth, decision-making behavior is the only truly proprietary artifact in any system development. All other artifacts can be inferred from industry practice, and do not differentiate each business specifically.

Decision

Defn: A proprietary datum derived by interpreting relevant facts according to structured business knowledge

The power of automated decisions to regulate and control system responses in a zero touch processing environment is driving increasing interest in decisioning.

Businesses do not make decisions merely because they have available data; nor do they act without making a decision. Decisions are clearly identifiable at the heart of any automated systems response, giving purpose to the data and direction to the process. If decision centric analysis is used to rigorously identify and describe decision-making behavior prior to systems development, then it can be used to drive the discovery of data and its relevance to the business it is the need for the decision that is the primary driver of the need for the data. Decisioning is therefore a primary data analysis tool, a precursor to formal data modeling. When data

---

[1] The Object Management Group (OMG) is a non-profit computer industry specifications consortium. Its members define and maintain the Unified Modelling Language (UML). The UML is the OMG s most-used specification, applying to application structure, behaviour, and architecture, as well as business process and data structure. OMG s UMLv1.5, section 3.22.3 appears to exclude the entire class of business rules from the methodology: Additional compartments may be supplied as a tool extension to show other pre-defined or user defined model properties (for example to show business rules) ...but UML does not define them...

analysis is driven by decision modeling, it gives rise to concise and provably relevant data models. And because decisions also predicate process responses, decisioning also implicitly drives process definition.

Decisions are also drivers for the discovery of other decisions. For instance, the decision to accept insurance business depends on precursor decisions regarding the acceptability of the customer, the insured risk, pricing calculations, etc. Each additional decision can also depend on further decisions until a tree structured model of inter-dependent decisions is formed (the decision model).

Decision Model

Defn: An ordered assembly of decisions that creates new and proprietary information to definitively determine an optimal course of action for the business.

It is the Decision Model that is the highest order analysis artifact and the only one to stem directly from the strategy itself. The concept of a model that defines order and context for decisions is a critical differentiator between the disciplines of decisioning and those of business rules . A decision is atomic, and like its data counterpart, is only given meaning by its context. The model is the critical concept in decisioning that provides this context. Further, it is only a complete assembly of decisions within the context of the model that can drive the business responses   such that the decision model must operate as a single unit of work in order to transition the business objects from one valid state to the next. The decision model representation of this transition is the only artifact in the system design world that directly describes the act of value creation for the business. It is this definition of value creation that defines each business in a way that is different from every other business   the decisions by which it recognizes, values, and responds to business events. We can liken this knowledge to the business DNA, given that the complete business operation can be cloned if its core decision making IP can be replicated. From this perspective, the requirements of the decision are the primordial requirements that drive the data and process requirements respectively, notwithstanding that there may be other factors that need to be accounted for in any given solution.

When the decision model is used to fast-track the discovery and modeling of both data and processes, it gives rise to a decision oriented system   a system that has business decision making as its core architectural consideration as indicated in Figure 1
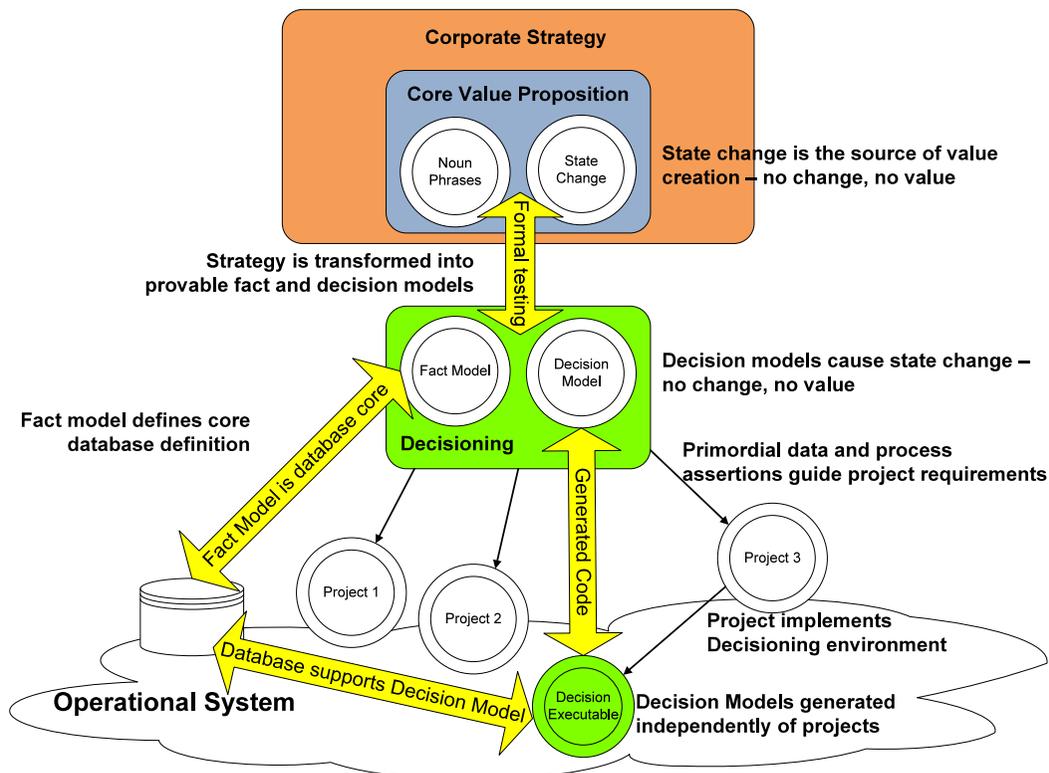

Fig 1

Decision Analysis

Decision analysis provides a unique and lightweight opportunity to analyze, define, and test the primordial requirements of many projects. This process itself provides valuable insights that reduce risk and increase value; however, the products of the process represent more than simply new found perspective. The process produces two critical models that can be mutually reconciled   the decision model and the fact model (or schema).
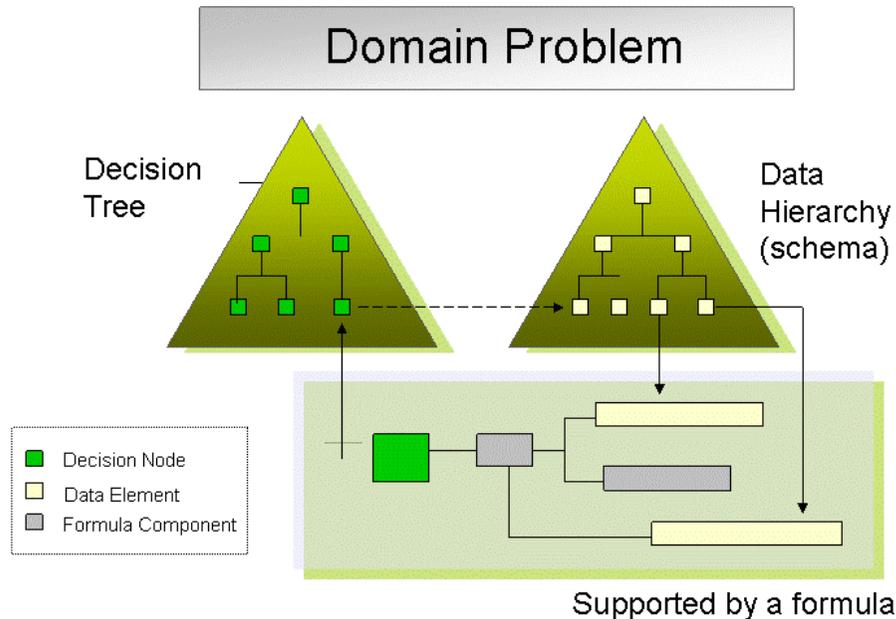
Decision-centric Development

Fig 2

These models validate each other and form a mutually consistent whole that can, with the correct tools, also be represented in the form of executable assets and actually tested to provide an assurance that the core value transitions that are proposed for the system can be implemented. The critical input to decision modeling is a clear understanding of the strategy that governs what the business will and will not do to create value, so this exercise can occur as soon as the corporate strategy clearly describes the business value propositions.

It is entirely practical to undertake such a strategic decision analysis as part of strategy definition rather than as part of a systems development project. Our experience includes a major logistics operator who spent considerable effort defining their proposed rating system in policy terms before contemplating a major re-development of their charging systems. By using decision analysis they were able to prove the viability of the proposal (to the extent of modeling revenue from actual operations) and to understand the project development requirement in terms of both data and process before funding the development project itself.

It is our experience that simply identifying the true value targets within the proposed system generates insight within the business. The extent of de-risking prior to project commitment is difficult to measure, but is intuitively significant. In order to de-risk a project, the decisioning analysis should occur not later than during inception.

How do we Start?

The business derives value or achieves its purpose by changing the state of primary business objects (or things), whether they be insurance policies, loan accounts, patients, trains, or any other object of value (including decision models!). These primary business objects are usually self-evident, and a cursory review of strategy documentation will highlight and clarify any ambiguities. But if the primary objects cannot be defined quickly (minutes, not days) and with certainty, the business should not be building a system   it has much bigger issues to worry about! A primary business object is by definition both tangible and discrete; therefore, it can be uniquely identified. Also by definition, it will be a source of value for the business and so the business will usually assign its own unique identifier to it   its first data attribute. In fact, a useful way to find these objects is to look for proprietary and unique identifiers assigned by the business (even in manual systems). Because it exists and generates value, external parties may also be involved and cause us to add specific additional non-discretionary attributes for interface or compliance purposes (e.g. a registration number), and we might also add other identifying attributes to ensure that the primary object can be uniquely found using real world data. So there is a small set of non-discretionary data that can be found and modeled simply because the primary object exists; this set of data is generic and will be more or less common across all like businesses. We can think of this as plumbing   it cannot be a source of business differentiation. So what other data is there that will allow us to differentiate our business? The amount of data

that we could collect on any given topic is unbounded   how do we determine what additional data is relevant? Decisioning! How we make decisions is the key differentiator of any business   how we decide to accept customers, approve sales, price sales, agree terms and conditions, etc. The important additional data will now be found because it is needed for decision making. We are now ready to do decision analysis; that is, after the initial strategic scoping, and prior to either detailed data or process analysis.

Decision Discovery

Decision analysis is both simple and intuitive because this is the primary activity of the business   this is what the business knows best. A decision is a single atomic value   we cannot have half a decision anymore than we can have half an attribute. So we are looking for explicit single valued outcomes, each of which will be captured as a new data attribute. Lets start with our business object (say, insurance policy), and with a state change that will cause a change in the value of that object (say, approve insurance policy). If we have a set of governance policies, we will interrogate them looking for noun phrases, and related assertions and conditions. Noun phrases can be interpreted directly into a fact model. The assertions and conditions will be reduced to a set of operations and declared as decision logic. We also need to access the domain expert within the business. The domain expert is the person charged with responsibility for strategy directed decision making, and who can readily answer the question: What decisions do you make in order to   approve this insurance policy? There is a pattern to most commercial decision making that helps structure the decision discovery process (see Figure 3). The green circles are the most important in the cycle. We can start analysis with the first of the primary (green) decision components (Authorization or Acceptance in Figure 3)   these are the will I or won t I decisions. For our insurance policy, will I or won t I accept this risk? (for other problem domains simply replace the business object and state as required; e.g. for a hospital, will I or won t I admit this patient? , etc.). Determining this decision may identify many precursor decisions that need to be considered first. For instance, for our underwriting question we might need to ask:

- What is the inherent risk of the object?
- What is the customer risk?
- What is the geographic risk?

These decisions in turn may give rise to further decisions so that we develop a tree of decisions   this is the decision model for authorization. Now we can move on to the next in the primary class of decisions: at what price? (or cost if a cost centre) (Pricing or Costing in Figure 3). In this case, the question is what decisions do you make to determine the price... of this risk? Again, this may result in a tree of decisions (for instance, pricing based on various optional covers, pricing for the different elements of risk, channel pricing, campaigns and packages, etc.). Following at what price? we can repeat the process for the Terms and Conditions and then the other pre and post decisions:

- Pre-Processing Check: Do I have sufficient information to start decision making?
- Context Based Validation: Is the supplied data valid?
- Enrichment: What further data can I derive to assist with the primary decision making?
- Product or Process Selection: Do I need to determine one decision path from other possible paths for the primary decision making?

And after the primary decision making...

- Workflow: Now that I have made my primary decisions, what do I need to do next?
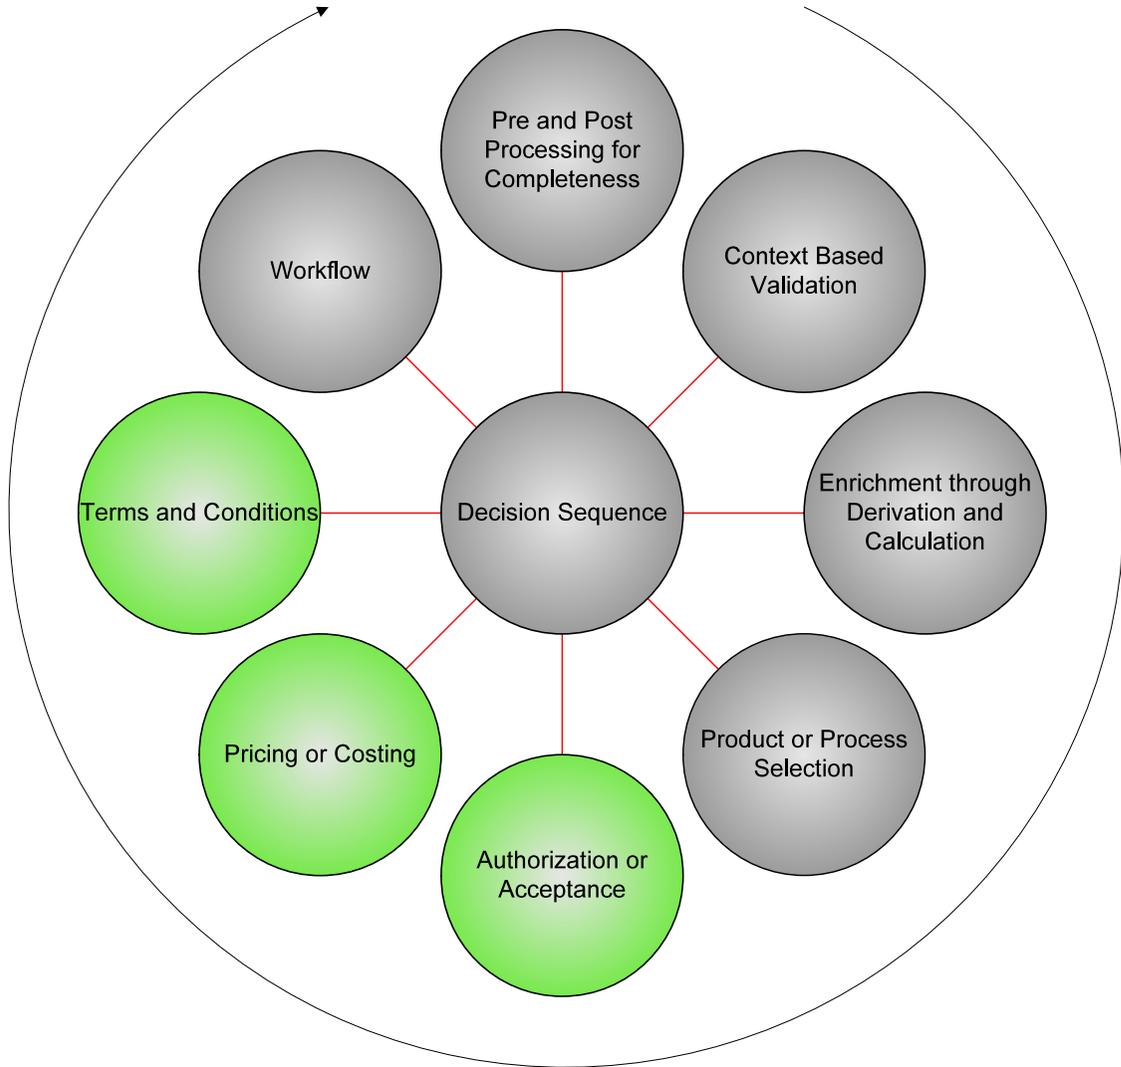- Post-Processing: Am I finished and complete without errors? Are my decisions within acceptable boundaries?

Fig 3

Normalization

Data normalization is a semi-rigorous method for modeling the relationships between atomic data elements. It is semi-rigorous because normalization is a rigorous process that is dependent on several very non-rigorous inputs, including:

- Scope: Determines what is relevant — we don t normalize what is not in scope.
- Context: Determines how we view (and therefore how we define) each datum.
- Definition of each datum: This is highly subjective yet drives the normalization process.

Data normalization is based on making a number of subjective assertions about the meaning and relevance of data, and then using the normal forms to organize those assertions into a single coherent model. Normalization with regards to decisions is similar. Each decision derives exactly one datum, and is atomic in the same way that data is. Similarly, each decision has relationships to the other decisions in the model. The decisions are related by both sequence and context. In this regard, context plays a similar role to the primary key in data normalization. Some of the inter-decision relationships within a model include:

- All decisions have a data context that is derived from the normalized placement of its output datum.
- Each decision definition may precede and/or follow exactly one other decision definition.
- Each decision may belong to a group of decisions that share the same context.
- In a direct parallel of 4th normal form, unlike decisions that share context should be separately grouped.
- A group of decisions itself has sequence and may also be grouped with other decisions and groups according to shared context.

Decision Driven Data Design

We have suggested that decision and fact models both evolve from strategy directed decision analysis. Following the initial discovery and elaboration of the primary business objects, we worked through the decision discovery process by analyzing the

strategy defined policies, which in turn identified new data attributes (the decision outputs). By locating these decisions around the data elements in the Fact Model, we built an integrated decision/data model (see Figure 2).

This combined model is self-validating and improves the overall rigor of the analysis. Following the discovery of the decisions, we can then elaborate them with formulas. Formulas provide additional detail regarding the consumption of data by decisions, thereby driving further demand for data. If the system cannot provide this data, then by definition the business cannot make the decision and the objective of the process cannot be achieved. In this way, decisioning can be shown to define the scope and context of the data, which then compounds with the decision usage to complete the data definition, which in turn drives the normalization and modeling processes for both. Note that the fact/data models used in decision modeling are subsets of the domain data model   they need to contain only the data required by the decisioning that is currently in focus. They are, in effect, normalized models of the key transactions rather than a model of the entire business. The XML schema standard is a useful standard for describing normalized transactional data. The set of fact models can then be merged to synthesize the primordial data model that will underpin all further project related data analysis. The resultant model will only contain the data actually required by the business   there is no second guessing of the data requirements as often occurs in a traditional approach, with significant and positive implications for development cost.

Decision Driven Process Design

Decisioning also drives process. For a decision to execute, the decision model must be supplied with its fact model by a process. This data, and therefore the process, cannot be defined until the decision requirements are known. Then, for the decision model to have any effect, a process must implement some response. While it is possible to define and build processes in anticipation of the decisioning that will drive them, it is sensible to analyze the decisioning in order to determine the range of inputs and outcomes, and then to normalize the process responses to them. In practice, bespoke discretionary processes do not occur in a vacuum, and should only occur when required to support value creation for the business   as described by Decisioning (regulation and industry practice may require additional processes, but these are non-discretionary by definition and may not add value). Processes that are found to exist that do not have this fundamental decision driven requirement are not good subjects for analysis   certainly their mere existence does not make the process requirement a necessity, and they are candidates for removal or re-engineering. There may be many process options for supplying data and responding to decisions made. In particular we should look for opportunities for direct integration with external parties to create integrated industry solutions. This 're-engineering' opportunity offers significant value, but is likely to be missed if analysis of existing use cases is conducted as the primary analysis. For this reason we do not consider processes or even use-cases to be attractive first order development.

Handover and the SDLC

We can achieve a verified and tested decision model, and it s integrated and co-dependent data model, with relatively modest effort   often only a fraction of the cost of traditional approaches. No system design or coding is required to this point. Most architectural and design options remain open   at this stage even platform is irrelevant. We have, in fact, defined and constructed a requirements model of the core functionality of the system from the business strategy perspective without constraining the technology options. Further, this requirements model is testable and can be proven to be complete, consistent and correct. Even better, we can retain the separate identity of this critical requirement over the long term, even across multiple implementations   there need never be another legacy system. Using a biological analogy, this is the DNA that defines how the organization is operated irrespective of how it is implemented in terms of platform or system. Implementation is now of little interest to those business users who are responsible for decisioning   we have a clear and well defined handover point to begin the traditional development cycle. It is feasible, even desirable, to simply hand over the decision model and fact model to systems designers as their starting point. This highlights a critical point: Decision models live outside of any project or process   decision analysis belongs in the inception phase of traditional SDLC s, or precedes it entirely. By definition the subsequent system design must be able to supply and receive data that complies with the decisioning schemas. And the system must provide appropriate process responses to the decisions made. While these processes remain undefined, this is secondary analysis and is tightly bounded by the decision design that precedes it; therefore, it is of comparatively low risk, and at the same time, it offers rich opportunities for business process re-engineering. The traditional systems development cycle is then used to design, build and/or reuse various software components (the plumbing) as appropriate to support the decisioning requirements.

Conclusion

The Decision-Centric development approach offers a significant advance on traditional development methodologies by focusing on a missing link between business strategy and operational systems   the decision model. The Decision Model is a new and important requirements artifact that is accessible and understood by both business and development practitioners, giving rise to a previously unobtainable level of shared understanding that can help bridge the gap between business strategies and the systems and databases that support them.