A Decision Architecture Whitepaper
Part 1/2

What Is Decision Architecture And Why Is It Important to Making Agile, Acquisition, Gap Resolution and Other EA Decisions.

by
Dr. David G. Ullman
Decision Architect

August 2016

I Purpose of the Whitepapers

This is the first of two whitepapers on Decision Architecture (DA). This first one focuses on defining DA and supporting why it is important for many types of enterprise decisions. The second white paper explores DA in government and open standards. Standards like TOGAF, FEA and DODAF give structure and process to enterprise design, change and evolution. The second paper also explores how DA affects an agile environment and the success of an RFP effort (Request for Proposals).

II Goals of this Whitepaper

The goal of this paper is to develop the basic structure of Decision Architecture and tie it to existing EA Standards, associated architectures and decision-making support methods. This is accomplished in detail to serve as a foundation for proposing Decision Architecture extensions to TOGAF, DODAF, FEA, et al and assisting Agile and RFP practitioners.

Decision Architecture is important because:

- A decision is a commitment of resources. The quality of what is done during the decision-making process and the buy-in for the results directly affect the resources used to make the decision and all downstream efforts.
- Decision Architecture structures the information that is key to the vetting and reuse of critical decisions. In the ideal, decisions would be captured in a manner such that the rationale[1] behind them is easily studied and understood.
- Decision Architecture can incorporate risk into decision-making. Decisions are always based on projections into the unknowable and risky future and ambiguous desires by stakeholders. A good architecture can manage uncertainty and ambiguity while making decision risks transparent.
- Decision Architecture supports codifying those decisions that can be reduced to decision models and rules (i.e. structured), and supporting the resolution of those needing human intervention and characterized by uncertain information (unstructured).
- For structured situations Decision Architecture:
    - Helps structure the logic of the business model.
    - Supports implementing automated decision-making.
- For unstructured decisions Decision Architecture:
    - Enfranchises stakeholders by making them part of the decision-making process. A decision is only as good as the buy-in by the stakeholders.
    - Provides a systematic architecture that can make decisions rational rather than emotional or political.
    - Helps manage information that is uncertain, incomplete, evolving and conflicting so that decision-makers can arrive at the most robust decision possible.
- Decision Architecture supports choosing a decision-making process that well matches the problem, leading to significant ROI.

To set up this paper, fill in Table 1, a ten measure form. Each measure is a major topic in this white paper. Answers are on an always-to-never scale. When judging these, think of the last few acquisition, gap resolving, agile or design decisions. The ten measures in the Table apply whether making a large system acquisition decision that takes months or years, or an agile design decision that takes minutes or hours. Regardless of the granularity, these ten are important for a robust Decision Architecture. Each of the measures are amplified and discussed in following sections. These ten are also used in the follow-on white paper to aid in evaluating EA standards and methods. Finally, a more detailed version of this form is given in Appendix A which can also serve as an outline for the topics in this paper.

Project or Organization:		5=Always, 0=Never
	Measure	
1	There is an attitude that decision-making is an important part of all processes.	
2	For each decision to be made, the stakeholders and ownership is clear.	
3	The objective of decision-making activities is clearly known.	
4	Multiple alternatives are generated for each decision to be made.	
5	Information and analysis used to evaluate alternatives clearly supports the decision-making process.	
6	An appropriate decision-making method is used for each decision to be made.	
7	Risk consideration is a core part of the decision-making process and based on information uncertainty and ambiguity.	
8	It is clear when a decision has been made.	
9	Decisions are recorded, reviewed and reused.	
10.	There is decision buy-in.	
	Total in Column	

If your total score was less than 25 then there is a good potential for Decision Architecture to enhance your processes. If some of the measures were unclear to you, then read on as, for each of the ten measures, there are two to four sub-measures developed in the following sections, Before detailing these, the background and definition of Decision Architecture is refined.

III Decision Architecture Background and Definition

Enterprise Architecture (EA) is the practice of planning, analyzing, designing and implementing business strategies. It requires the analysis of complex structures and processes under potentially uncertain operating conditions with the goal of helping organizations make the best possible business, application and technology decisions. Over the past twenty years the EA community has built many methodologies aimed at understanding and managing system complexity and business alignment.

During each phase of architecture development and application, EA success is a direct function of the alternatives considered and the decisions made. While EA methods are strong on developing models, they are generally weak from a decision architecture viewpoint as is shown in the follow-on white paper.

Decision Architecture sees modeling, analysis and information management as needed activities that support the decision-making process. This process, whether automated or human based, requires the comparison of multiple alternatives relative to identified measures with the goal of either choosing the best alternative, ranking the potential alternatives, or deciding what else to do next to improve the likelihood that the best choice can be made. This applies equally to large acquisition decisions and small agile decisions.

There are two broad classes of decisions made in EA: structured and unstructured. Structured decisions are made in situations which are fully understood and can rely on deductive reasoning: If-then-else. Many of the decisions made during the EA process can be reduced to business rules and structured.

Unstructured decisions, by contrast, rely on abductive reasoning, the testing of hypotheses to discover what is true. Unstructured decisions require human intelligence to create and manage the uncertain and ambiguous information, and associated risks.

Where structured decisions are generally made for routine, operational tasks, unstructured decisions are made for one-off tactical or strategic situations. In EA, structured decisions are generally reduced to business rules. Much effort has put into the development and management of business rules (see Section 3.6 of the second white paper). However, the critical decisions in EA are generally unstructured.

Regardless of whether structured or unstructured, the decision-making process can be diagramed as shown in Figure 1. This figure shows the basic vocabulary and relationships needed to describe all types of decision-making processes. Here the solid lines connect activities that are the essential parts of all decisions and the dashed lines are the result of the meta-decision – “decide what to do next”. Although not shown in the figure, each activity is affected by uncertainties and ambiguities that create decision risk, the probability that a poor decision is being made. Risks come from information variance, stakeholder inconsistencies, external influences and differences in stakeholder values. Also, not shown is that each decision may be the work of many people and thus reflect the values of many stakeholders. In order for “Choose

Figure 1: Decision-making Basics for Unstructured Decisions

alternative and take action" (the lower-right activity) so that the resulting decision is robust, i.e. still looks good a year later, requires effort on all the activities.

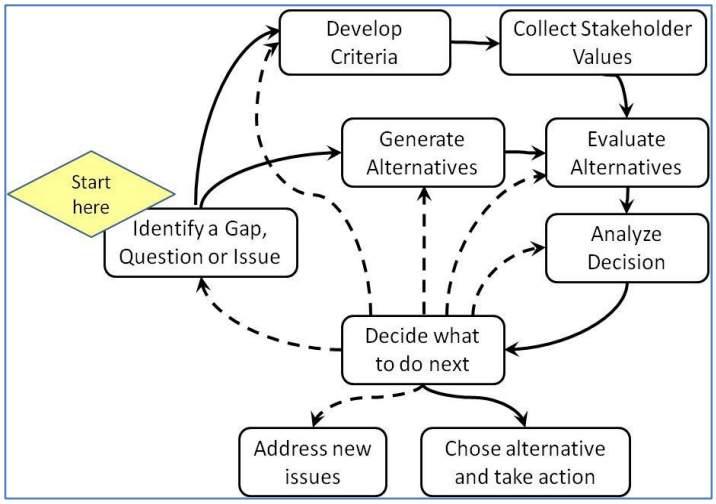

Each decision begins with a need. In EA, this is usually a **gap**, the desire to get from an "as is" business, application or technical situation to a "target" state. There are generally many paths or systems that could bridge the gap and the goal is to choose one of them. There are other **questions or issues** that arise during EA such as portfolio management decisions, RFP decisions, code development decision, and analysis of alternative (AoA) situations which are really other forms of gaps. Regardless, all need a set of **criteria** (i.e. measures) that define an ideal solution and all have multiple **alternatives** that might potentially meet the ideal for resolution.

A characteristic of unstructured decisions is that, since there are generally multiple **stakeholders**, there is need for them to reflect their **values** in terms of indicating what is most important to them. Collecting and managing these values is a part of Decision Architecture.

A major part of decision-making is the need to **evaluate the alternatives** relative to the criteria. For both structured and unstructured decisions this means collecting input data, modeling the situation, and conferring with subject matter experts, analytical results or rules to find how well the alternatives meet the criteria.

For any Decision Architecture there needs to be a methodology to combine the evaluation results and **analyze the decision** to supply reduced results so the meta-decision can "**decide what to do next**". There are many decision analysis methods ranging from the use of rules, to formal optimization, multi-attribute decision theory, and Bayes Nets. The "what to do next" meta-decision plays the role of controlling the iteration (the dashed lines in Figure 1) needed to get a good decision in a timely manner.

It is important to realize that once the decision is made, it is a commitment to use resources. As soon as you decide how to resolve a gap, how code a story, or which system to purchase, people's time and effort is focused on the choice made at the exclusion of all others. If the decision was a poor one and the issue needs to be revisited, or if stakeholders do not buy into the decision and do activities inconsistent with it, then resources are wasted.

Left out of the discussion above is "naturalistic" decision-making[2]. A naturalistic decision is one made with limited time and is often called intuition or gut feel. A late 1980s study of individual engineers designing products included recording their activities and dissecting them on an utterance and line-by-line basis[3]. One finding of this study was that, at their finest, detailed design decisions occur on average once per minute, some faster and some slower. This is keeping with human's cognitive limitations and the generally accepted understanding that, in order to solve problems and make decisions we must divide them into very small pieces, address each one and then fuse the results together. This applies to a single designer sitting at a computer or a management team meeting to make a system level decision. This study also found that the basic building blocks shown in Figure 1 occur, whether conscious or not, even at the finest levels of granularity. True, some situations such as the house is on fire, or the airplane is diving toward the ground require faster decisions where the alternatives and criteria are built in by nature or trained in, but these situations are not common in an EA environment.

Of course, Figure 1 shows an ideal. To be useful, any Decision Architecture must combine the ideal with the practical, a balance between what should be done and what is done. So, the goal of following sections that detail the ten measures in Table 1 is to strike this balance using the notation in Figure 1 as a guide.

1 There is an attitude that decision-making is an important part of all processes.

Decisions are the punctuation points in any process and the commitment of resources for what is to follow. They tell us when to continue on, to pause to gather more information or stop because we are done. They tell us when to iterate because we are not ready to proceed.

Decisions range from those made by a team that affects the very architecture of the entire system, to those made by an individual on the very smallest detail. Where a major system level decision may require a management team many months to make, every high level decision is composed of literally thousands of finer grained decisions.

There are three clear indicators of decision-thinking; 1) The decision points in a project are known, 2) It is clear when decisions are needed and 3) The decision-making process is aligned with the organization's operating model.

1.1. Decision points or gates in the project are clearly identified.

In most projects, choices are made at gates in a stage/gate process or at each iteration in an agile process. These decision points are actually the sign-off on much effort leading up to them. In EA practice, the decision points occur in the all layers: Business, Application and Technology.

There are two key elements in identifying the decision points, what is needed and when it is needed. The "what" or the objective of each point is important enough that it is

addressed as a separate measure, Measure 3. "When" it is needed is the next topic in this measure.

1.2. For a specific issue, it is clear when a choice, a commitment, is needed.

Often what drives the decision-making process is an awareness of when a decision is needed. If the house is on fire, a decision is needed now. If the issue is just exploratory, the current situation works, but could be better, there is no end date and the process can drag on for months or years, or never be resolved.

One agile philosophy is to make decisions at "the last possible moment". This does not mean to procrastinate until this moment arrives and then make a choice. Rather, knowledge of this moment, to a Decision Architect, means being able to plan how to make the best possible decision in the time available. This requires understanding the stakeholders, the information available and the risks as early as possible and then using the resources on hand so the choice and risks in it are evident as time runs out. In fact, doing the proper pre-work may enable the decision sooner, at the "first responsible moment". This is what Decision Architecture strives for.

1.3. The Decision Architecture is aligned with the organization's operating model.

In these white papers a push is made to structure the decision-making process. But, if this is not aligned with the organization, this effort is like that of a "seagull consultant" dropping in, ruining everything and then leaving. In a paper published in *Insight*, The ENCOSE Journal, titled "Decisions of the 3rd Kind" [4], four different levels of decision-making maturity were defined. These levels rank organizations from the use of informal, ad-hoc methods to very structured, probabilistic methods. It is key that any process proposed by a Decision Architect fits the organization while encouraging it to a higher level of decision-making maturity without getting in the way.

Meeting this implies two assessments: what is the current decision-making environment, and what is an ideal environment for the organization. This ideal should include an ROI, an evaluation of how any reduction in the gap between "what is" and "what could be" will benefit the organization. This is not always easy to estimate, but is essential.

2 For each decision to be made, the stakeholders and ownership is clear.

2.1. It is known who is has the final authority to make each decision (who has signature authority) and who is governing its implementation.

Every decision has one person or a small group of people who actually "make" the decision; they give OK to accept the choice made and commit resources to its implementation. There is a wide range of awareness of who is this final authority. Those following a formal process such as the decisions in a large procurement effort, will have a fully spelled out signature path. For Agile code development or other design efforts it's usually the engineer or his/her supervisor who signs off on decisions made.

Often, however there is no formal authority for decisions made. Sometimes this is best and sometimes it is not. When choosing which algorithm to use to solve an issue when writing code, there is no needed authority beyond the analyst. This is as it should be for most design decisions. However, if the choice reaches well beyond the individual and the issue at hand, then there should be an authority beyond the individual to approve it.

Besides concern for who is making the decision, who is governing its implementation? In "Why Decisions Fail"[5] the author, Paul Nutt, studied 400 decisions made by senior managers in medium-to-large organizations. He found that fully half the decisions had failed—either action was not taken; or if taken, it did not stick. Of the decisions that failed, the only effect observed two years later was used resources—time, money, personnel, equipment, etc.—all expended without achieving success in attaining the original goals. Besides following poor decision-making practices, these failures had little or no governance, or follow through by someone or some group to ensure the choices made were implemented and resulted in ROI for the organization.

2.2. For each decision to be made, the stakeholders are known (those making the decision, those affected by the decision and those with knowledge affecting the decision).

There are three classes of people important to the decision-making process; those making the decision (or advising those who do), those affected by the decision and those with knowledge affecting the decision. Decision Architects consider them all as stakeholders. It is important that these people be identified as early in the decision-making process as possible.

Where stakeholders contribute to all the decision-making elements shown in Figure 1, the one that is generally underrepresented and managed is stakeholder values. Values are how each stakeholder or group of stakeholders hold some criteria as more important than others. Allowing each stakeholder to indicate how important each criterion is enfranchises them as they know their voices are being heard. It also, reduces disagreements because all party's values are being captured and utilized. This often conflicting (i.e. ambiguous) view about what is important generates risk in that what is a good alternative for one stakeholder may be wrong for another. A goal is to satisfy all the stakeholders as best as is possible to build buy-in for any decision. This is an important part of a Decision Architect's task.

3 The objective of decision-making activities is clearly known.

As mentioned previously, EA methods like TOGAF support finding gaps which are clear statements of the need for a decision or multiple decisions. Without this guidance it is sometimes fairly obvious what the question or issue might be, but it is not always so. Spending time on the wrong issue leads to waste of time and money. The risk of

focusing resources before the team agrees on the issue can result in people wasting time and resources.

3.1. For each decision point, the issue, problem or question is written down and clearly understood by all stakeholders

It is amazing how often team members think they're working together when in fact they're not even discussing the same problem. Even if you're working alone, you will generate ideas and evaluate them while developing a clear statement of the issue you are addressing. The only way to ensure that your effort is focused is to explicitly state the issue, question or problem being addressed.

A Decision Architect will attempt to get the team working to agree to a written statement that includes a call for action directed at some feature of an object or process.[6] For example; "Choose the best method to close the gap on the keyless capability", or "Select the best proposal to fund from the RFP responses".

3.2. The objective or criteria that help discriminate amongst alternatives have been developed.

Criteria are the measures used to discriminate amongst the alternatives. Criteria are not the same as requirements and constraints as generally developed and as seen in TOGAF, but may include some of them. For example: In specifying a system to resolve a gap, one of the requirements is: "must process 3k orders/hour" and a constraint is "it must be compatible with the existing system." If these and other requirements and constraints are listed in an RFP the resulting proposals will probably meet all of them. These are the filter requirements. If a proposal doesn't pass the filters then it is rejected during alternative evaluation. Those that pass the filter requirements and constraints must then have criteria to help the decision team discriminate one alternative from the others.

For example, some years ago Motorola released an RFP with over 60 specifications for an electro-mechanical device[7]. These related to cost, mechanical and electrical performance, reliably and other engineering and business measures. Twenty proposals were submitted. The reviewer quickly filtered these into two piles, those that met the requirements and those who did not. Then the hard part began - how to differentiate the five that did meet the requirements. The sixty requirements were all "filters", but Motorola was missing discriminating criteria, criteria that allow the close front-runners to be differentiated. The filtering requirements were only part of the picture needed to find the best proposal. A Decision Architect can help facilitate developing both requirements that filter and criteria that discriminate.

Secondly, the more time spent developing criteria, the better the decision[8]. This statement is supported by experimental evidence. This isn't to say that all criteria need to be developed at the beginning of a project, as some coevolve with other information, but that a specific effort needs to be made to generate and refine criteria during the decision-making process.

3.3. For each criterion, its goal or target; the ambiguity of it; and the sensitivity to being off-target is known.

Most lists of requirements include a goal or target (underlined): “The project should take less than 3 weeks”, or “The maintenance costs should be less than $5,000 per month”. Or informally, when no set list is developed, there is at least a target in mind. Still, more is needed to really understand each goal.

Consider a simple example. You and your wife go into an auto dealership to buy a new car. You are sure you want to spend less than $30,000 USD. She thinks you can go as high as $35,000. Thus the target for our cost criterion is ambiguous. Compounding this ambiguity, as you start to look at cars you see one at $29,000 but it doesn’t meet your desire to have high efficiency (another criterion). You see another at $32,000 that has good efficiency and you are willing to consider it. A third, at $38,000 is just too expensive and you walk away from it. These three, the targets at $30,000 and $35,000, the willingness to consider $32,000 and the lack of consideration for $38,000 paint a picture of the sensitivity to being off-target.

Managing the targets, and their ambiguity and sensitivity is an important part of making choices considered by Decision Architecture.

3.4. The objective includes compliance with principles, policies, laws and regulations.

Finding a solution to an issue that violates organizational principles, policies, laws and regulations is pretty much a waste of time. However, many projects are undertaken with little knowledge of these and only wrestled into compliance later. Not a very good use of resources.

4 Multiple alternatives are generated for each decision to be made.

Alternatives are the potential resolutions for the gap, issue, question or choice situation. There is always more than one alternative because if there is only one, then there is no decision to be made, only justification (see next section). In many situations there is a single alternative to consider with “do nothing” as a second alternative. There are costs that must be considered when evaluating the doing nothing alternative.

4.1. An effort is made to discover multiple solutions to the problem

The importance of developing a good set of alternatives is evidenced in the results of a GAO study of large system acquisitions[9]. Of the thirty two projects studied, ten did not do any development of alternatives. Of these, seven were updates of earlier projects or were supported by other analyses. For the other three, high cost and schedule growth occurred. Of remaining twenty projects (shown in Figure 2), thirteen considered a narrow scope of alternatives. Nine of these had moderate to high cost or schedule overruns while the nine that included a broad scope of alternatives only had two with moderate to high overruns. These results clearly show the risk when multiple alternatives are not generated and evaluated.

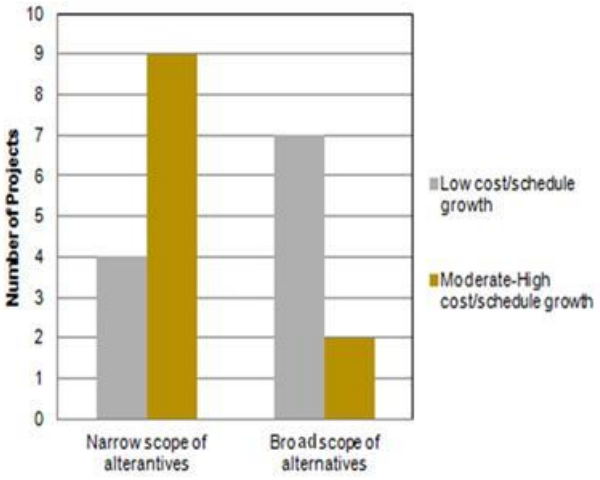

Figure 2: Partial results of GAO AoA study.

In another study, this one on single engineers solving design problems, it was found that increasing the number of alternatives was linearly correlated to increased technical quality of the results[10]. Dramatically, doubling the alternatives considered nearly doubled the quality of the result.

Decision Architecture provides methods to encourage exploring the “alternative space”. It helps the stakeholders develop multiple potential solutions regardless of the granularity of the problem. Even in Agile projects, having more than one option to consider can greatly improve the quality of the results[11].

4.2. “Justification situations", issues with a single pre-chosen result, are not mistaken for "decision situations".

It is amazing how often people have a preconceived answer to a question that has many potential answers. They then go about justifying their “pet” solution. Sometimes this “answer” comes from above (either down the chain of command with the boss’s pet solution or by some divine message). Sometimes the single solution comes from below. One manager who insisted on multiple alternatives received three proposals from his engineers, their favorite and two other cobbled together variations that they knew were inferior. Regardless, a good Decision Architecture can help flag when there is justification rather than decision-making.

5 Information and analysis used to evaluate alternatives clearly supports the decision-making process.

A key part of any decision is evaluating how well the alternatives meet the objectives. This activity includes estimating costs, time, performance, and other measures by gathering information and analyzing it. For example, assume one criterion is "Initial cost less than $32,000". To estimate the cost of Alternative A will require estimating the time needed to develop the various parts of it, the cost of equipment needed, the costs of other resources and some analysis of how all this information should be combined to give a cost estimate.

Let's say you estimate Alternative A's initial cost at $30,000 then you can compare it to the target of < $32,000. But before you do, some details you might want to think about:

1. The assumptions that went into the estimate (detailed in Section 5.1).
2. The dependences on other issues (detailed in Section 5.2).
3. The match between the issue and the evaluation fidelity (detailed in Section 5.3).
4. The ability to tradeoff meeting one goal for another (detailed in Section 5.4).
5. The information and analysis uncertainty (detailed in Section 7).

5.1. The "known knowns" and "known unknowns" are itemized and "unknown unknowns" searched for.

In February 2002, Donald Rumsfeld, the then US Secretary of State for Defense, stated at a briefing: "*There are known knowns. There are things we know that we know. There are known unknowns. That is to say, there are things that we now know we don't know. But there are also unknown unknowns. There are things we do not know we don't know.*" He was initially lampooned as uttering some form of government double-speak, but in reality he was stating the basis for all assumptions about information.

There are known knowns, facts that can be stated and are not going to change. Question the assumptions behind all facts because sometimes they are really not as fixed as assumed and are really known unknowns. Known unknowns refers to things we know about but are not sure of the details (see more on this in Section 7 on Risk). Finally the unknown unknowns are the black swans, the unexpected.

Decision Architects help decision-makers understand their level of knowledge about known factors and lead exercises to search for the unknowns, lowering the risk of poor decisions.

5.2. It is clear how the information is dependent on other issues and other issues are dependent on the current decision.

One method that helps in understanding the stakeholders is to map the issue interactions to find out who the decision results affect, and what other activities affect the information on which the decision at hand is dependent. Understanding information dependencies is important not only in finding the stakeholders, but in developing decisions that lead to useful activity. Many projects have a high level of interactions and methods in TOGAF, for example, which provides tools to identify and track them (e.g. Business Interaction Matrix And Dependency Priority Matrix).

Strong decision architectures spend time developing the interactions prior to finalizing a choice so the decision-makers are sure that no important dependency has been overlooked.

5.3. Information evaluation is consistent with the issue resolution need.

It is foolish to over-analyze an issue, yet the term "paralysis by analysis" has entered the vernacular. This refers to situations where team members keep gathering data, generating code and exploring deeper. The result; a decision not made in a timely manner. It is often easier to do more work that it is to reach a decision. At the other extreme is the use of highly subjective gut-level evaluations when more analysis is needed. The balance between the needed evaluation and what is actually done is to match the fidelity of the analysis to the need. Decision Architects look for this balance.

5.4. Tradeoffs are managed.

If one alternative meets the cost goal but not the time goal then maybe some money can be traded off to buy more help and reduce the time, helping the alternative meet all the needs. This is the essence of tradeoffs. Managing them can range from very informal to detailed mathematical methods[12].

What isn't always understood about tradeoffs is that the better the issue is understood, the better the tradeoffs can be managed. Paying attention to the other items in this white paper leads to better understanding and better tradeoff management, one goal of Decision Architecture.

6 An appropriate decision-making method is used for each decision to be made.

6.1. There are many different decision support methods known to the organization that can be used in a situation.

Most organizations do not think much about their decision-making methods, they don't think of activities as a series of decisions punctuating processes. There are four levels of decision thinking maturity[4] (Introduced in Section 1.3) and most companies are in the lowest two levels with scant decision-thinking knowledge. A Decision Architect can help improve the maturity and suggest the best methods for different situations.

6.2. Issues that are often repeated and can be codified to a set of if-then rules, have been considered for codifying using a rule based method.

The concept of business rules has found good traction in recent years. Business rules can be developed when the decision points are repeated and the structure of the stakeholders and information in them is consistent and can be reduced to a set of if-then-else rules[13]. Further, there are standards for modeling rule based systems such as OMG's Decision Model and Notation (DMN) as discussed in Part 2 of this white paper series. A Decision Architect can help identify situations that can be reduced to business rules and help in developing a system to support them.

6.3. There is a go-to person, a Decision Architect, in your organization.

Few companies have a Decision Architect, a person to help match the best decision-making methods to the needs. Some rely on consultants or trainers for methods. The scant use of Decision Architects is a reflection of the decision-thinking maturity discussed in Section 6.1.

7 Risk consideration is a core part of the decision-making process and based on information uncertainty and ambiguity.

Risk management is integral with decision-making. While generally not recognized as such, they are, and thus this section is the longest in this white paper. Dennis Stevens[14] of Leading Agile, in a power point presentation titled "Agile and the Nature of Decision-making" notes that "Traditional" risk management is limited by:

- Creating bureaucratic overhead
- Managing point solutions mean that the risk impact may not be closely connected to objectives
- Significant gaps in ability to handle ambiguity and emergence
- Ineffective integration of risk-management

The primary goal of this section is to show that taking the decision-centric view espoused here resolves these and other limitations. Secondly, terms associated with "risk" are used very inconsistently in EA, business, and other literatures. Thus, a second goal here is to clarify these terms. Two types of risk are discussed, object/event risk and decision risk.

Object/event risk is the tradition use of the term "risk". It amounts to answering three questions:

1. *What can go wrong?* Some event or change in an object (e.g. the code (the object) doesn't work (the event) or the customers (the object) can't place orders for sweaters (the event).

2. *How likely is it to happen?* Probability that the event or change happens based on past statistics, an analytical model, or best guess based on experience. The terms "probability" and "likelihood" are used interchangeably in the literature.
3. *What are the consequences?* The implications or outcomes of the event or change usually in terms of money, time, quality, damage or possibly even lives wasted. The terms "consequence", "implications" and "outcomes" are used interchangeably in the literature.

In common usage, people often shortcut these three with statements like; "There's a risk that it might rain tomorrow." The event "rain tomorrow" is clear, its likelihood is "might" which is not very clear, and the consequences are only inferred, e.g. you will get wet or the picnic will be cancelled. Weather forecasters try to improve on this with statements like "chance of rain tomorrow is 40%", refining in the likelihood.

Formally, risk is an expected value, the probability of the consequence. In the engineering and business worlds, risk assessment is a methodology that uses numerical representations of the likelihoods and consequences to calculate the severity of the risk. For example, if I want to know the risk of the light over my desk burning out (a change in light bulb) I can find the probability of failure per unit time for light bulbs (the likelihood) either by experiment or by looking it up in a table (based on past experiments or collected data). This likelihood already contains a probability and the consequence and is, in itself, the risk. This type of risk assessment has been formalized in many areas and is generally referred to as Probabilistic Risk Assessment (PRA) well documented in the NASA PRA handbook[15].

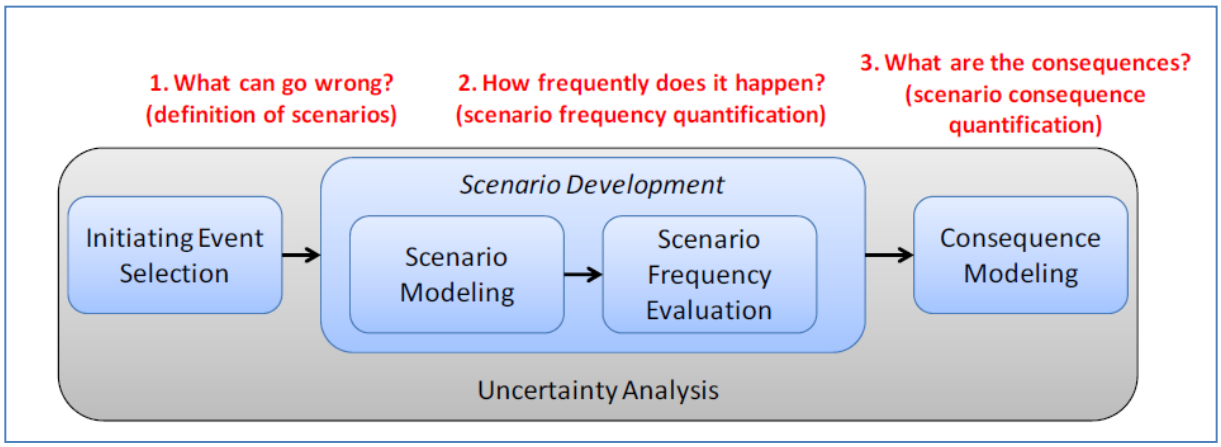

Figure 3: How PRA addresses risk and uncertainty

Figure 3 shows the steps through the PRA as they address the three questions that define risk. Also shown is that uncertainty affects all the steps. The likelihood of the initiating event is, for most, uncertain; the scenarios unclear. Modeling the scenario can range from gut feel to detailed numerical analysis. Regardless, the model will not be perfect with its level of fidelity to reality only as good as modeling and time will allow. Any mismatch between the model and reality is uncertain. Finally, the consequence may also be uncertain, e.g. if it does rain, will the picnic really be ruined.

Even though methods like PRA are used extensively in high-technology areas by NASA and others, often in EA situations there is no data on which to base detailed analysis. In these cases a simpler method can be used that is based on a risk matrix. This method is widely used in engineering, the military[16], and in some businesses. Figure 4 shows the probability of occurrences and consequence. They are both estimated and then a risk matrix used to determine the level of risk and follow-on action.

Probability of Occurrence

Level	Level	Definition
5 (Very High)	>70%,	almost certain
4 (High)	>40%,	More likely than not
3 (Moderate)	>30%,	Significant likelihood
2 (Low)	>1%,	Unlikely
1 (Very Low)	<1%,	Very unlikely

Consequences or Impact

Level	Mission	Implementation
5	Mission failure	Overrun reserves
4	Significant reduction	Consume all reserves
3	Moderate reduction	Significant reduction in reserves
2	Small reduction	Small reduction in reserves
1	Minimal (or no) impact	Minimum reduction in reserves

Risk = Probability * Consequences

>=6 - OK, monitor for changes

6-12 - Aggressively manage, consider alternatives

>12 - The risk is high, implement other alternatives

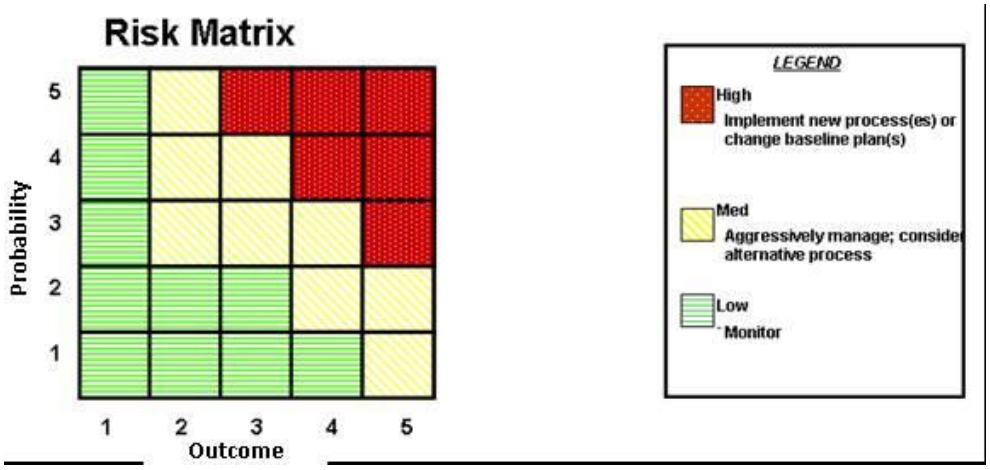

Figure 4: Simple risk matrix

Using this method to estimate the risk of a project not meeting its cost budget for example; first, based on what is the currently known estimate, the probability that the project will go over budget. Say, for example, that it is moderate, there is significant likelihood, and the probability is at Level 3. Then estimate the consequences. If there is a small reduction in reserves due to being over budget, then Level 2. Finally, the risk is the product of the levels of occurrence and consequences, or, using the matrix, the level of risk can be found. Here the risk is medium as seen on the matrix. This situation should be "aggressively managed" and other alternatives considered.

The object/event risk evaluation logic can be extended to the decision-making risk. During decision-making, risks are inherent in uncertain knowledge, information and models; ambiguity across stakeholders (i.e. poor agreement); and ignorance about what might happen. All of these create the risk that a poor decision will be made. This doesn't say that the alternative chosen will fail, that is object/event risk which was discussed above. Drawing analogy to object/event risk, decision-making risk focuses on:

1. What can go wrong? – A poor choice is made within the resources available.
2. How likely is it? – Here the probability is dependent on uncertain knowledge of outcomes and probabilities, the fusion of the team's interpretation of information and models, and the management of ambiguity and ignorance
3. What are the consequences? – Money, time and possibly lives are wasted

In assessing what can go wrong there is always a balance between the quality of the decision and resources expended. More time and money leads to better evaluation of the alternatives, and potentially a more likely successful decision. The problem with decisions, is it is usually impossible to know that the best possible decision has been made.

The assessment of likelihood is confounded by the uncertainty, ambiguity and ignorance. The better these can be taken into account during the decision-making

process, the more unlikely that there will be a waste of money, time and other resources.

For example, say a decision must be made about the acquisition of a new system to fill a gap identified by TOGAF. Clearly the risk here is a decision risk - you want to ensure that you don't make a poor system choice. This is not asking the question, "Will System A fail?" It is asking, "If we choose System A, have we made the best choice within the resources we have available." System A failing (object/event risk) is clearly part of the decision, but so is going over budget or over cost, or having poor buy-in from the team.

The criteria to assess the alternatives can be developed, but since these systems have never been built before, some of the measures of effectiveness and the costs will be uncertain. Further, the stakeholders have inconsistent values, the decision team may not agree on some of the evaluations and so there is ambiguity. Finally, there is some ignorance; there may be some factors that are unknown unknowns.

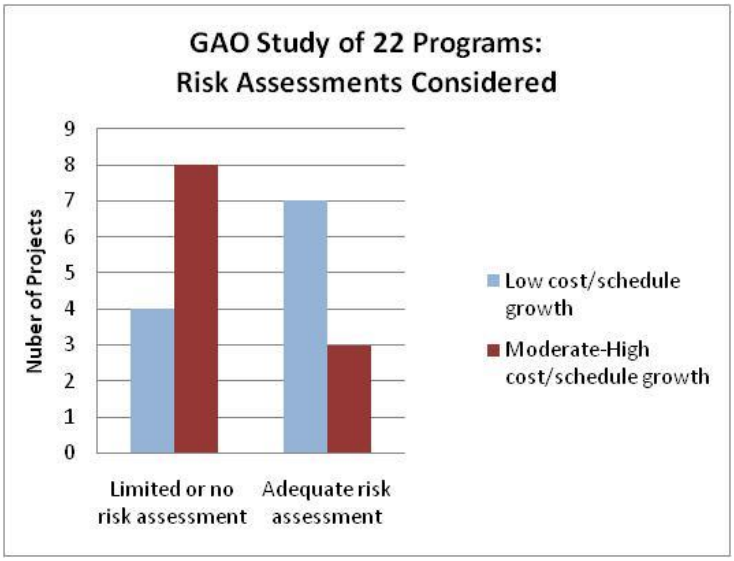

Figure 5: The effect of risk assessment on project success

Although there are no studies about the effectiveness of analyzing decision-making risks, there is about object/event risk, a sub-set. In the GAO study introduced in Section 4.1, twelve of the twenty two projects that used Alternatives of Analysis (a structured Decision Architecture) conducted limited or no risk assessment for each alternative (Figure 5). Of these twelve, eight experienced moderate to high cost and schedule overruns. The other ten projects were judged to have adequate assessment of risks. Of these, three had moderate to high cost or schedule overrun. The other seven had low cost/schedule overrun. Acquisition projects that do not examine risk, present overly optimistic assessments of the alternatives leading to cost and schedule growth.

There are three important areas to consider when assessing how well an organization manages object/event risk and decision risk. They are detailed in the following sections.

7.1. The risk attitude of the organization is known.

Some organizations are very risk averse and others are risk takers. Knowledge about where the organization falls on the continuum between the extremes is important in managing both object/event risks and decision risks. TOGAF methods like Business

Transformation Readiness Assessment[17] and various Technology Readiness Assessment methods[18] give some insight into the readiness for change. If an organization is not keen about change then it is surely risk averse and the risk assessment and mitigation methods need to be matched to this aversion.

7.2. There is effective integration of risk management in projects both large and small.

By forcing risk assessment into the decision-making process it becomes part of the culture and not a bureaucratic overhead lumped onto projects when things are not going well. It integrates risk-management into the process both on the object/event level and the decision risk level. Once the risk is understood how (and if) it will be mitigated can be managed. Mitigation can range from acceptance of the risk (knowledgeable acceptance) to elimination or reduction of it. By tying risk to the uncertainties in the decision-making process it is clear where to use scarce resources for mitigation.

The concepts here apply regardless of project granularity. Even agile projects need to have integrated risk management[19]. As evidence consider that the Standish Group[20] estimates that only 39% of IT projects were successful in 2012. This is based on self-reported data and has not improved with the adoption of agile methods. Further, they estimate that 1.5 decisions are made for every $1,000 in labor costs. Clearly, project failure - the ability to see and manage the risks - is not totally caused by poor decision-making practices, but they contribute. The role of the Decision Architect is to minimize the risks and improve the odds of project success.

7.3. There is an awareness of information and model uncertainty, ambiguity and ignorance.

All decisions are fraught with uncertainty (in data, information, and models), ambiguity (i.e. poor stakeholder agreement) and ignorance (the unknowns). These are the root causes of risks. The book “Making Robust Decisions” which has been referenced many times in this paper is focused on the inclusion of uncertainty in the decision-making process. In this work the term “robust” is defined as:

> *“A robust decision is the best possible choice, one found by eliminating all the uncertainty possible within available resources, and then choosing with known and acceptable levels of satisfaction and risk.”*

This definition states that uncertainty (hence risk) is eliminated within resources and then a choice made that is as insensitive as possible to that which remains. This is an important concept in Decision Architecture which strives to give you a window on the uncertainties and the ability to manage them.

8 It is clear when a decision has been made.

Amazingly it is often not clear when a decision is made. Some stakeholders assume there is closure when there has been none, and others go about their own agenda even though the team has agreed to a common path. Remember, a "decision" is a commitment of resources so either of these paths is a waste of time and money.

8.1. "Final decisions" are really final and only revisited in planned situations.

If there is poor stakeholder buy-in, then some will ignore decisions and keep working on other alternatives, waiting for the decision to fail. Stakeholder buy-in is refined in measure 10. All of the measures in this paper focus on building the right information at the right time so buy-in can be achieved and final decisions are cleanly made.

8.2. Decisions are not made too soon, ignoring evidence that should have been considered.

Making the "final" decision too soon without fully developing the criteria and alternatives is another cause of people going their own way and wasting resources. Again, the ten measures in this paper are an effort to ensure the right information is developed so that decisions are made cleanly.

9 Decisions are recorded, reviewed and reused.

Recording decisions so they can be reviewed and reused can be very important to an organization. In the early 1990s this author performed NSF sponsored research on the capture and use of rationale for mechanical design projects[21]. Here "rationale" implies the reason for the choice made which includes all the decision elements in Figure 1. Efforts were made to capture and index information about: 1) the objects being designed, 2) the process followed and 3) the decisions made. It was found that capturing information at the decision points was key to indexing and reuse.

More recently, a European group applied the same logic to Enterprise Architecture. They found that capturing the rationale was critical as not doing so resulted in a lack of transparency that can cause design integrity issues when architects want to maintain or change the current design[22]. Further, they interviewed practitioners and found that while 70% claimed to capture rationale only 27% captured rejected alternatives and only 42% captured decision impact. The 70% who claimed they captured the rationale is suspect based on the other data. Exactly what information was captured as part of the rationale?

The capture and reuse of information is important for EA projects where the work may be distributed, where there may be high employee turnover and consultants are often used to develop key pieces of code. There are four sub sections that refine the concept.

9.1. Decisions are recorded or captured in a way they can be revisited later on.

Few companies actually record any detail about their decision-making. Many record the results, the choice made, but few capture the alternatives, the criteria, the evaluations and the arguments for and against the options. This is not easy to do and not easy to search and reuse (see next section). While it is intellectually ideal to capture all the detail, with the current tools it is not yet practical. However, as the ability to capture and manage information continues to evolve, it gets easier to automatically glean this information from communications and work-day activities. Decision Architects can help identify what information is worth capturing.

9.2. The organization keeps a searchable record of past decisions for reuse in similar situations

If you capture decision information, then you should be able to reuse it in similar situations. However, few organizations do this. What usually happens is based on the memories of stakeholders with statements like “we tried that when…” and “my buddy did this and…”. Reuse is very difficult as it is often hard to know when two situations are similar enough that the decision-making process of one can help with the other.

9.3. Past decisions are revisited to evaluate their success or failure and the lessons learned integrated into future efforts.

Hindsight is notoriously inaccurate. We tend to inflate the good and downplay the bad. This is well known human psychology yet we rely primarily on hindsight when evaluating the quality of past decisions.

Decisions are funny things, you do not know if they were good or bad until much later and then you have little to compare them to as only one path was followed. Few organizations make the effort to revisit decisions to learn what went well and what did not. A Decision Architect can help an organization learn from their past failures and successes.

9.4. Stakeholders and governance know of the decision and rationale for it.

Often decisions are not well communicated. The goal here is that they should be to both the stakeholders and those responsible for ensuring that they are carried out. Partially this is to gain buy-in (see next section) and partially this is a check that a governance process has been set up.

10 There is decision buy-in.

The term buy-in has been used throughout this white paper. Buy-in ensures accountability. People feel accountable for a decision if they play an active role in the process leading up to it. The choice that was made may not have been their first choice,

but they contributed to its selection. Accountability is born from collaboration. You know you have collaboration when[23]:

- Everyone can paraphrase the issue to show that he or she understands it.
- Everyone has a chance to contribute to solving the problem. This can be accomplished by participating in refining the issue, developing alternative solutions, building criteria, or contributing evaluation information.
- Everyone has a chance to describe what is important to him or her.

Those who do not agree with the final decision will still be likely to support the team because they have been included in the decision-making process and appreciate the work required to reach a decision.

The three sub-sections below reveal symptoms of poor buy-in. A good Decision Architecture will help in developing buy-in for any decision.

10.1. Stakeholders buy into and comply with decisions with little do-your-own-thing activity after the "decision" was made?

One symptom of a lack of buy-in is that stakeholders do tasks that are not consistent with the activities needed to support the decision. As said before, decision is a commitment of resources. If people are doing work inconsistently with a decision that has been made, then resources (i.e. time and money) are being wasted.

10.2. On hearing the choice made, nobody says: "Who decided that?" "Why wasn't I involved?" or "Why didn't they ask me, that decision affects my work (or unit)?"

A second symptom of a lack of decision buy-in is to opt out stakeholders. The examples in the headers are signals that certain stakeholders have been left out of the process. In order to have buy-in, the right stakeholders need to be involved. If they are not, then how can you expect them to feel responsible for the result?

10.3. On hearing the choice made, nobody says "Why didn't they ask me, I have information about that?"

Missed information increases the odds of a poor decision and is another symptom of poor buy-in. Having the right stakeholders involved from the beginning is important to decision success.

Summary and Lead In to Other White Papers

The goal of this paper was to develop the basic structure of Decision Architecture and tie it to existing EA Standards, associated architectures and decision-making support methods. The paper was built around ten measures, each a major topic in this white paper. These measures can be applied to acquisition, gap resolving, agile or design decisions. They apply whether making large systems acquisitions decision that takes months or years, or an agile design decision that take minutes or hours. Regardless of the granularity, these ten are important for a robust Decision Architecture.

In the second paper, the ten measures are applied to EA standards like TOGAF, DODAF, FEA, et al and agile methods.

Appendix A

1. There is an attitude that decision-making is an important part of all processes.
   1.1. Decision points or gates in the project are clearly identified.
   1.2. For a specific issue, it is clear when a choice, a commitment, is needed.
   1.3. The Decision Architecture is aligned with the organization's operating model
2. For each decision to be made, the stakeholders and ownership is clear.
   2.1. It is known who is has the final authority to make each decision (who has signature authority) and who is governing its implementation.
   2.2. For each decision to be made, the stakeholders are known (those making the decision, those affected by the decision and those with knowledge affecting the decision)
3. The objective of decision-making activity is clearly known.
   3.1. For each decision point, the issue, problem or question is written down and clearly understood by all stakeholders
   3.2. The objective or criteria that help discriminate amongst alternatives have been developed.
   3.3. For each criterion, its goal or target; the ambiguity of it; and the sensitivity to being off-target is known.
   3.4. The objective includes compliance with principles, policies, laws and regulations.
4. Multiple alternatives are generated for each decision to be made.
   4.1. An effort is made to discover multiple solutions to the problem
   4.2. “Justification situations", issues with a single pre-chosen result, are not mistaken for "decision situations".
5. Information and analysis used to evaluate alternatives clearly supports the decision-making process.
   5.1. The “known knowns” and “known unknowns” are itemized and “unknown unknowns” searched for.
   5.2. It is clear how the information is dependent on other issues and other issues dependent on the current decision.
   5.3. Information evaluation is consistent with the issue resolution need.
   5.4. Tradeoffs are managed.
6. An appropriate decision-making method is used for each decision situation.
   6.1. There are many different decision support methods known to the organization that can be used in a situation.
   6.2. Issues that are often repeated and can be codified to a set of if-then rules, have been considered for codifying using a rule based method
   6.3. There is a go-to person, a Decision Architect, in your organization.
7. Risk consideration is a core part of the decision-making process and based on information uncertainty and ambiguity.
   7.1. The risk attitude of the organization is known.
   7.2. There is effective integration of risk management in projects both large and small.
   7.3. There is an awareness of information and model uncertainty, ambiguity and ignorance.

8. It is clear when a decision has been made.
   8.1. "Final decisions" are really final and only revisited in planned situations
   8.2. Decisions are not made too soon, ignoring evidence that should have been considered.
9. Decisions are recorded, reviewed and reused.
   9.1. Decisions are recorded or captured in a way they can be revisited later on.
   9.2. The organization keeps a searchable record of past decisions for reuse in similar situations
   9.3. Past decisions are revisited to evaluate their success or failure and the lessons learned integrated into future efforts.
   9.4. Stakeholders and governance know of the decision and rationale for it.
10. There is decision buy-in
   10.1. Stakeholders buy into and comply with decisions with little do-your-own-thing activity after the "decision" was made?
   10.2. On hearing the choice made, nobody says "Who decided that?" "Why wasn't I involved?" or "Why didn't they ask me, that decision affects my work (or unit)?"
   10.3. On hearing the choice made, nobody says "Why didn't they ask me, I have information about that?"

Project or Organization:		5=Always, 0=Never
	Measure	
1	There is an attitude that decision-making is an important part of all processes.	
2	For each decision to be made, the stakeholders and ownership is clear.	
3	The objective of decision-making activities is clearly known.	
4	Multiple alternatives are generated for each decision to be made.	
5	Information and analysis used to evaluate alternatives clearly supports the decision-making process.	
6	An appropriate decision-making method is used for each decision to be made.	
7	Risk consideration is a core part of the decision-making process and based on information uncertainty and ambiguity.	
8	It is clear when a decision has been made.	
9	Decisions are recorded, reviewed and reused.	
10.	There is decision buy-in.	
	Total in Column	

References

[1] Ullman, D. G. and R. Paasch, *Issues Critical to the Development of Design History, Design Rationale and Design Intent Systems*, ASME Design Theory and Methodology Conference, DTM94, Minneapolis, Sept 1994, pp 249-248.

[2] Klein M; "Capturing Design Rationale in Concurrent Engineering Teams", IEEE Comp. V26 No.1 93-47,Jan1993.

[3] Stauffer, L.A., D.G. Ullman, "Fundamental Processes of Mechanical Designers Based on Empirical Data," Journal of Engineering Design, Vol. 2, No. 2, 1991, pp. 113-126.

[4] David Ullman, Decisions of the 3rd Kind, Insight: Special Issue; Systems of the 3rd Kind, July 2012, Volume 15, Issue 2, pp 21.

[5] Paul C. Nutt, *Why decisions Fail*, Berrett-Koehler Publishers, 2002

[6] David Ullman, *Making Robust Decisions*, Trafford Publishing, 2006, Chapter 5.

[7] David Ullman, *Making Robust Decisions*, Trafford Publishing, 2006, page 153.

[8] David Ullman, *Making Robust Decisions*, Trafford Publishing, 2006, Page 143.

[9] GAO: "DEFENSE ACQUISITIONS: Many Analyses of Alternatives Have Not Provided a Robust Assessment of Weapon System Options", Report to the Chairman, Subcommittee on National Security and Foreign Affairs, Committee on Oversight and Government Reform, House of Representatives, Sept 2009. http://www.gao.gov/products/GAO-09-665

[10] David Ullman, *Making Robust Decisions*, Trafford Publishing, 2006, pages 125 - 126

[11] Strategies for generating creative and innovative ideas in Scrum teams, Fransisco Attanasio, Agile alliance http://agileatlas.org/articles/item/strategies-for-generating-creative-and-innovative-ideas-in-scrum-teams

[12] David Ullman, Trade Studies with Uncertain Information, Sixteenth Annual International Symposium of the International Council On Systems Engineering (INCOSE) 2006, http://www.davidullman.com/images/DecisionPDFs/TradeStudies.pdf

[13] *Decision Management Systems*, James Taylor, IBM Press 2012.

[14] Agile and the Nature of Decision-making, Dennis Stevens, LeadingAgile*http://www.agilealliance.org/files/session_pdfs/Agile%20and%20the%20Nature%20of%20Decision%20Making.pdf*.

[15] "Probabilistic Risk Assessment: procedures guide for NASA managers and practitioners" NASA/SP-2011-3421, 2nd edition, December 2011, http://www.hq.nasa.gov/office/codeq/doctree/SP20113421.pdf

[16] MIL-STD-882E, Department Of Defense Standard Practice: System Safety (11-May-2012), http://everyspec.com/MIL-STD/MIL-STD-0800-0899/MIL-STD-882E_41682/

[17] TOGAF, Chapter 30., Business Transformation Readiness Assessment

[18] David Ullman, The Mechanical Design Process, 5th edition, McGraw Hill, pages 236-239.

[19] Five Simple Steps to Agile Risk Management, Michale Lant, June 2010, http://michaellant.com/2010/06/04/five-simple-steps-to-agile-risk-management/

[20] Chaos Report, The Standish Group, http://www.versionone.com/assets/img/files/CHAOSManifesto2013.pdf

[21] Ullman, D. G. and R. Paasch, "Issues Critical to the Development of Design History, Design Rationale and Design Intent Systems", ASME Design Theory and Methodology Conference, DTM94, Minneapolis, Sept 1994, pp 249-258.

[22] Plataniotis, G. et al, An Empirical Evaluation of Design Decision Concepts in Enterprise Architecture, IFIP, PoEM 2013, LNBIP, pp34-38

[23] David Ullman, *Making Robust Decisions*, Trafford Publishing, 2006, Page 96