



RISK MANAGEMENT: ISO31000 Better, but still not there yet!

By
Jeff Popova-Clark

Principal Partner
Data Analytics Management Consulting
72 Valley Drive
TALLEBUDGERA QLD 4228
Phone: +61 (0)7 5534 8770
Mobile: +61 (0)421 960 048
Email: JeffP@dataanalytics.com
www: www.dataanalytics.com

© Data Analytics Management Consulting, 2011

This document or any part of this document may be freely quoted or distributed either electronically or in hard copy format, provided the identity and contact details of the author (ie Jeff Popova-Clark) and this sentence are included and that none of the contents of the document are altered.

Further articles are freely available at the Data Analytics web site: www.dataanalytics.com

RISK MANAGEMENT: ISO31000 Better, but still not there yet!

Written by: Jeff Popova-Clark

Australia has long been blessed with an excellent Risk Management Framework (currently AS/NZS ISO 31000:2009). In fact the International Standard Organisation's own Enterprise Risk Management Framework Standard borrows heavily from it. However there are still some basic conceptual errors and difficulties that the process and framework neglect. In fairness, a couple of the key issues I have been raising for over a decade appears to have been finally addressed in the 2009 version: the first is "what is being affected by the risk". Essentially damage to a house doesn't matter if the house continues to perform its function of shelter. However the concept of severity of damage (i.e. the cost to reinstate the house to its former state) was always conflated with the impact on the entity's ability to fulfil its function. The entity should not be the house, it should be the house's objectives that are put at risk: e.g. provide shelter and/or earn capital return. However ISO 31000 has incorporated this concept and although I'm not convinced many risk practitioners have understood and implemented this concept appropriately in their risk assessment processes, it is nonetheless captured in the new standard. The second was integration of risk management into an overall governance framework, which I'm pleased to see is a core component of the updated framework. So I've pulled these issues out of my criticisms of the current standards (and how they are frequently implemented) and left the remainder as is. Essentially the remaining issues are:

1. Understanding that frequency and severity of events are not point estimates of risk but co-varying distributions
2. Risk events have a chain of causation. One event begets another event which begets another, but the third in this chain may also be caused by other drivers. Few practitioners tease this issue out clearly during risk assessment processes.
3. Understanding whether you are asking your risk estimators to estimate intrinsic risk (i.e. sans controls), residual risk (i.e. remaining risk if all controls are functioning as designed), current risk (i.e. risk as it currently stands given the current performance and design of controls) and finally target risk (i.e. level of risk acceptable to the business).
4. We allow risk assessors to lump different risks into a single catch-all or too broad a category. For instance, the risk of "understating revenues" is too broad. There are so many ways that revenues can be understated that it is meaningless to assess the likelihood or severity of "understating revenues". In addition when we go to develop mitigating controls such a broad category is unlikely to generate effective responses.
5. Too many practitioners conflate uncertainty of an actual current state with

probability of future risk realisation over time

6. We often ask unqualified managers to estimate risks and consequences when they are simply not knowledgeable enough to provide a reasonable estimate. Different risks require different expertise to assess the potential cost to the business and possibly different expertise again to assess the likelihood of occurrence and then different expertise again to identify the most cost effective mitigation strategies.
7. Sometimes our controls create further vulnerabilities that also need compensating controls. Brakes on a car may decrease the likelihood of a crash and therefore operate as an effective mitigating control, but they can also fail if they are not maintained. Control generated risks are often not explicitly captured.
8. As a result of the above glaring failures of traditional risk assessment processes, managers go through the process and do not value the outputs of the risk assessment processes beyond meeting compliance requirements. The risk assessment process should be a valuable assessment of business opportunities, but it is often left unactioned. Actions arising from a risk management effort should result in actual reallocation of resources, changes in management KPIs, modification of stated aims and missions and changes to reporting and monitoring processes.
9. We have the opportunity to take the risk assessment process one step further and calculate an annual (or some other period) cost of the risk, which helps us assess the cost effectiveness of our controls and the long run cost-benefit of our business

Traditionally risk assessment practices involved identifying a context/entity/objective, identifying their risks, estimating the frequency of each risk, estimating the severity if the risk is realised and then treating the risks. However this process does not really reflect reality and therefore produces contrived results more often than not. Contrived results might be acceptable to demonstrate compliance with some regulation, but they do nothing to actually cost effectively improve the business. Properly performed risk management should be a key driver of strategy and operations and should definitely be visible in the budget and financial reporting of your organisation.

As a first example of the paucity of the standard approach let us first consider “What is a risk?”. Often a risk is seen as an event like an earthquake or misappropriation or a blackout. Sometimes its seen as the consequences of the event like “incorrect capitalisation of an expense” or “inability to power computer network”. Sometimes its seen as the vulnerability “senior executive don’t communicate with staff” or “system allows same invoice to be processed twice”. Few practitioners take the time to lay the ground rules for risk assessment

workshops to ensure that participants are producing assessable risks at the beginning of the process. So what should we do to identify consistent assessable risks?

Firstly, participants need to understand that risks are essentially the threats to the mission or purpose of the entity being assessed. The very first question to ask is “What is it that the entity does that can be put at risk?”. For many companies the key *raison d’être* is to generate a return on investment for shareholders. Public sector agencies may well exist for other reasons. However, making money is too generic for a risk assessment process. The key question is how the entity makes money...what is its unique purpose. BHP might be “Identify subterranean mineral resources of value and then to reliably and efficiently extract them and deliver them to market”. Once the group feels it has identified a unique purpose for the entity, it is then time to move on to identifying the threats to the entity achieving its purpose. Threats include both risks and issues. Risks are potential events that, if they occurred, could impact on the entity’s ability to achieve its objectives, whereas issues are problems that may be in existence now that may be impacting the ability for the entity to achieve its purpose. An issue may be “insufficient qualified personnel in the labour market to meet project needs” whereas a risk is a potential future event like “price of iron ore falls below cost of extraction”. Many practitioners allow participants to include issues (sometimes identified as “opportunities”) within their risk assessment (although they then let them assign a probability to these things!). I agree with this approach with a note that the framework’s name needs to reflect the fact that it deals with both risks and issues.

One risk or many?

So with that out of the way we can now turn to the first criticism of the standard. We still hold on to this quaint simplification that we can estimate a single frequency and a single severity of a risk. This is clearly impossible and yet it has been central to risk management standards since the early 90s. Engineers understand that risks can have differing consequences and differing frequencies of occurrence. For instance they often talk of a 1 in 20 year flood as compared to the much more severe 1 in 100 year flood. The 1 in 20 year flood threatens infrastructure and dwelling very close to standard waterways, but a 1 in 100 year flood can inundate entire neighbourhoods. Note that if we are to talk about the frequency and severity estimate of a flood risk, we need to determine which flood are we talking about, the relatively frequent 1 in 5 year flood or the rare but catastrophic 1 in 100 year flood. Enterprise and Corporate risk management standards do not handle this concept during the risk assessment processes.

For example if we consider misappropriation of commercial assets, risk workshop leaders will ask participants: “How often does misappropriation occur in your organisation? Once a year, Once every 5?” Well it depends on how big a misappropriation you’re referring to. If you include pilfering office stationery, I’d

say the risk is triggered daily. However if you mean more than \$1M secreted out of the company into an employee's bank account, I'd say this is a rare event indeed. We could split the different levels and types of misappropriation into the various sizes of loss and then estimate frequency across each size. However, if we do this for every risk we are going to result in hundreds and even thousands of risks to assess (i.e. 1. Misappropriation < \$100, 2. Misappropriation Between \$10 and \$100, 3. Misappropriation between \$100 and \$1000, 4. Misappropriation between \$1000 and \$10,000, 5. Misappropriation Between \$10,000 and \$100,000 and finally 6. Misappropriation > \$100,000). This is clearly unworkable. The answer is Risk Scenarios. We need to identify the likelihood and severity of Ghost employees/contractors, stationery and other non-cash asset pilfering, KPI manipulation, incomplete cash reconciliation, skimming, commission fraud, self-dealing, and any other identifiable method of misappropriating assets. Obviously these will have the same frequency/severity problem as misappropriation generally, but it will be much more effective to identify the size and frequency of most import to your organisation when dealing with an actual scenario.

Which Risk?

The second criticism is the issue of choosing the right risk event. For instance a major storm may cause a lightning strike, which brings down a power line, which cuts power to the business premises, which stops the lights, air conditioning and computers from working, which makes the staff unable to perform their tasks. Which is the risk event and which is the consequence? Should a risk assessment assess the likelihood of a major storm? But other things may bring down the power line: a mistake by the electricity distributor, a major traffic accident, a flood, vandalism etc etc. In addition a storm may produce other consequences: flooding of premises, inability of staff to get to work, overloading of electrical equipment, hail damage of vehicles, wind damage of external infrastructure etc etc. Which point on the risk causality chain do you choose to use as the risk event?

The answer comes down to three considerations: 1. Which point on the risk chain allows the identification of the most cost-effective mitigating controls, 2. Which point on the risk causation chain is most able to be assessed in terms of severity and frequency and 3. Which risk has the most unique consequences (i.e. loss caused that is not caused by other risks). Other than my own risk assessment sessions I am yet to attend any other workshop that considers these fundamental questions. In my view this renders most risk assessment sessions down to a do-the-process session rather than actually achieving the purpose of risk assessment.

With or Without Controls?

A common problem I see at risk assessment events is the confusion between both participants and facilitators of what type of risk we are asking attendees to estimate: intrinsic, residual or current risk. Noting that this is different again from target risk (the point at which the business accepts any remaining risk). Intrinsic

risk is the risk that the risk event would represent if we had no mitigating controls in place: a car without brakes for example. Residual risk is the risk that would remain if all of the currently planned controls were fully implemented and working perfectly. Current risk is the actual risk that the risk represents to the business at this point in time, knowing that not all controls may be fully implemented nor are they all always going to be working perfectly. Too often have I seen risk assessment sessions identify what they believe is the residual risk and then nominate the current controls as the risk mitigation. When doing the risk assessment they are presuming the controls are in place and then when planning the mitigation they are presuming they are not there. Its a kind of risk management double dipping.

I recommend that risk assessments should explicitly have an assessment of all four of these kinds of risk. This way there is no confusion. Don't presume the car has no breaks or rear vision mirrors when measuring intrinsic risk, this is the way they are delivered by default. Intrinsic risk is the risk that the organisation would face if management had done nothing. The intrinsic risk is the risk faced without any controls in place. However few organisations leave risks totally unmitigated even if they have no formal risk management process in place. It is important to understand the intrinsic risk, because a little understood outcome from a risk assessment process is that some current controls may be too expensive to justify the risk reduction achieved. I've seen some organisations decide not to increase controls even though the risk was assessed as higher than they initially thought, but (other than ones facilitated by me) I am yet to see any risk assessment processes where one of the outcomes was a decision to cease an existing control to save the operational cost. This is difficult to do if you haven't assessed the intrinsic risk.

The difference between Residual Risk and Current Risk is a little more subtle. In most cases, management will have made some investment in mitigating controls for risks. Some may be awaiting the implementation of some IT system, or some may be currently out of service, but most will be in place and operational. Whatever the case the key is that the investment has been there to mitigate the risk. In essence, residual risk is the risk level that has been planned to be achieved using the to date investment in mitigating controls. By contrast, Current Risk is the risk level that is currently being borne by the entity. This may be higher than Residual Risk, because controls are not functioning as designed, not yet implemented or have unmitigated risks of their own (think of the unmaintained brakes). In general Intrinsic Risk is always larger than or equal to Current Risk which is always larger than or equal to Residual Risk (i.e. $IR \geq CR \geq RR$).

Target Risk, however, is the risk level that is tolerable to the organisation. Often this is left out of assessments. It could be that the organisation is willing to tolerate the risk of a catastrophic but extremely unlikely event. Do we take mitigating action against space alien invasion? A very unlikely event (it certainly hasn't ever happened before), but undoubtedly very severe consequences were it to occur.

However, by doing nothing to prepare for or mitigate against space alien invasion, our society is demonstrating that its target risk level for this risk is above the perceived current risk level and therefore there is no justification of any costs of controlling for or mitigating against it. Similarly a very likely but low consequence risk (e.g. stationery pilfering) might be of such low consequence that, though it occurs frequently, the entity is willing to leave the risk unmitigated. It is important to assess Target Risk for a few reasons:

1. It is rarely the best idea to mitigate a risk to zero. Theoretically it is impossible to render a risk to zero and it would cost infinite resources anyway. It is best to render a risk mitigated to a point where the cost benefit to the entity over a certain time horizon is optimised. Sometimes it is hard to convince legal people, for instance, that an entity breaks the law sometimes despite its best efforts. But to guarantee that the law will never be broken, even accidentally, is prohibitively expensive. So effectively all businesses choose to break the law a certain amount in order to stay in business. In other words setting a target risk is a good way to educate management that a mitigated risk is not a prevented risk, just a mitigated one.
2. Setting a target risk may result in your organisation realising that it is over-investing in expensive mitigating controls. According to wide ranging analyses of human decisions, a human life in the developed world is estimated to be worth roughly US\$6-8M and an extra year of quality human life is worth roughly US\$50-129K. If the organisation is spending \$20M per annum to reduce the risk of life lost from 1 in 100 per year to 1 in 1000, it may be over investing. Bad example? Ok if its investing \$20M to reduce the number of product defects from 1 in 1000 to 1 in 10000, when the cost of fully replacing the extra 9 in 10000 defects is less than \$20K, it may be time to review if that expensive control is worthwhile. Similarly, is it worth paying flood insurance for that mountain top transmitting station?
3. Setting a target risk also allows you to estimate the cost of that risk to the organisation on a per annum basis. This allows a comparison to the cost of the mitigating controls. For a simple example, if an intrinsic risk is “on average once in 10 years a \$1M loss is expected to occur” and the target is that “on average once in 20 years a \$200,000 loss occurs”, then the business should be willing to expend up to \$90K per annum to achieve this (i.e. $\$1M/10 - \$200K/20$) outcome.

Uncertainty vs Probability

A third criticism with modern risk management processes is that risk practitioners conflate uncertainty with probability. Especially when risk process facilitators allow issues to become included in the list of “risks”, the likelihood becomes the chance that something is true, rather than probability that something may occur in

the future or that something may affect the achievement of objectives. Of course when making any estimates of the probability of a future event there are two parts to the estimate: 1. Is the uncertainty inherent within the event itself and 2. the level of certainty management has about the accuracy of its estimate. For instance if the occurrence of a >7 Richter scale earthquake is estimated to be 1 in 100 years, what is the confidence that our 1 in 100 year estimate is accurate. If we have data going back several thousand years and can verify the 1 in 100 has held up fairly nicely over that time, we can feel fairly confident that our 1 in 100 is a good estimate, because we have good knowledge of the risk. However, if we have no historical data for the region, our 1 in 100 year estimate is little more than an educated guess and we should up the risk a little to reflect this uncertainty (say to 1 in 50). Another area of uncertainty is the level of consequences that the risk event will cause. How much will the >7 Richter earthquake affect our operations? We may estimate that it will cost \$1M in lost operational capacity and a further \$1M in repair costs, but it might cause more or less impact than that. The level of uncertainty we have over these cost estimates should also be reflected in our estimates of the overall risks.

However as mentioned previously most risk management facilitators allow management to nominate issues during the risk identification process. This happens because there is no equivalent “issues management” process in most organisations where issues that participants feel are not getting the attention required can be raised. Therefore they are raised during risk management workshops. Issues might include “insufficient training of operational staff” or “increased reliance on extended staff workhours”, or “Silo mentality between divisions”. Some risk facilitators will try to get workshop participants to turn the issue into risks like “staff resignations” or “more workplace accidents” or “customer complaint about poor service”. But in reality these issues may cause many potential negative consequences and possibly ones that are too numerous and small when cast as risks but significant and manageable when cast as an issue. So, many risk facilitators allow issues to be included in their list of risks even though they are technically not risks in the true sense of the word. But when we have issues (instead of risks) and we are asked to assess how much risk (i.e. risk being used here in its other meaning) these issues represent, we are not estimating the probability that an event may or may not occur, we are instead assessing the uncertainty that the issues really exists and the uncertainty we have over the size and impact of the consequences.

Should the ignorant reign?

Although senior management risk identification and assessment workshops have a place, they are relied upon far too much in modern risk management. Agreed that it is important to get your senior managers to overtly identify and own risks...get them thinking about it. This means they are more likely to deal with the risks appropriately. However, if risk management is meant to be more than just a

method of getting your senior managers "thinking about risk"; if you intend to get a reasonable handle on the portfolio of risks facing your organisation then you need to approach risk identification and analysis as a project. To illustrate the point, I was running a risk workshop for the senior management of a local government entity. One of the attending senior managers was Director of a Division which included a team of qualified hydrologists (although he himself was a civil engineer) whose primary task was to model flood risks for the local area. The risk workshop was regarding risks to the local council and inevitably the subject of flood risk was raised. I was the facilitator, so tried to keep my personal judgment of the risks to myself (as much as possible), but I did find it astonishing that, despite my goading, the relevant director stayed non-committal on estimating the frequency and severity of flood risks to council property. In the end, although the result was generally agreed by the entire group, the most influential executive on the likelihood of a disruptive flood event was an accountant from the Office of the CEO, someone who had next to no experience or expertise in hydrology. Should we be asking an accountant about flood risk? Should we even be asking a civil engineer who has qualified hydrologists reporting to him. Shouldn't we just go straight to the hydrologists themselves and get a scientifically valid estimate? And then with regard to the potential impact of the flood, why not use the information available in the asset management system about the value and locality of assets? This kind of information is not going to be available in a quarterly senior management risk workshop.

The key is that a risk assessment project should include a fact gathering exercise that may include a risk workshop but by no means should be confined to the outcomes of such a workshop. A risk facilitator should use a range of sources to identify the threats to an organisation. These include one on one interviews with a cross section of staff, staff/supplier/partner/customer surveys, transactional records (e.g. legal settlement payments), industry literature (i.e. what's happened to other similar organisations), organisational history/incident logs (what's happened to this organisation in the past), review of major procedures (what would happen if this bit doesn't work), scenario enactments, risk portfolio prompters, etc. A comprehensive risk assessment process would build up a risk register over time and have this validated by senior management. Additionally, the assessment of the quantum of risk may, in some cases, be left to those with relevant expertise or to an analysis of the available data, rather than polling the opinion of senior executives.

Jeff Popova-Clark is the Principal Partner of Data Analytics in Gold Coast, Australia. He consults in areas as diverse as governance, strategy, risk management, market analysis, marketing strategy, datawarehousing, performance management, and human resource management. Mr. Popova-Clark has also published works in the areas of criminology, psycholinguistics, management strategy, and recruitment and selection. His current focus is on unleashing the creative potential already within organisations in order to develop strategy for competitive advantage. Mr. Popova-Clark has numerous post graduate qualifications including a Master of Business Administration and a Masters of Taxation and Financial Planning and has qualified for Mensa.