GRC - One Out of Three Ain't Good

Why your GRC investment may be letting you down, and what you can do about it

For years now we've been hearing about the three promises of Governance, Risk, and Compliance (GRC) solutions:

1. That they would help us to more cost-effectively govern our risk landscape,
2. Make better-informed risk decisions, and
3. Maintain compliance with whatever standards or regulations we might be subject to.

In this post I'll share five reasons why many organizations are, at best, realizing only one of these objectives. I'll also discuss approaches that can help an organization realize all three objectives and significantly improve the return on its GRC investments.

GRC Solutions in a Nutshell

Generally speaking, GRC solutions provide the following features and functions:

- A place to document the "risks" an organization wants to measure and track (commonly referred to as a "risk register")
- A place to identify policy and other compliance requirements an organization is subject to
- A place to document audit findings and other reported control deficiencies, and a way to associate those deficiencies to the organization's "risks"
- A place to document and track compliance and remediation activities, as well as the stakeholders associated with those efforts
- Reporting capabilities for the four bullets above, and
- Workflow capabilities to maximize efficiency in everything listed above

This doesn't seem like rocket science, so how can organizations be failing to achieve the three GRC objectives? To understand this, we first need to look a little more closely at each of the objectives so that the problems described in this post can be connected to how they affect the value of GRC.

Governance

Ultimately, leadership is expected to cost-effectively govern the organization's risk landscape. Accomplishing this requires setting and communicating expectations (policies, etc.), overseeing and

facilitating the achievement and maintenance of those expectations, and managing conditions that don't align with expectations. GRC solutions are supposed to assist with this by recording and reporting where these expectations are and are not being met, within a meaningful business context -- i.e., in terms of risk.

The "R" in GRC

This objective is all about making better-informed risk decisions, which boils down to three things: 1) identifying "risks", 2) effectively rating and prioritizing "risks", and 3) making decisions about how to mitigate "risks" that are significant enough to warrant mitigation. (You'll understand shortly why I've been putting quotation marks around the word "risks.")

Compliance

Of the three objectives, compliance management is the simplest -- at least on the surface. On the surface, compliance is simply a matter of identifying the relevant expectations (e.g., requirements defined by Basel, PCI, SOX, etc.), documenting and reporting on how the organization is (or is not) complying with those expectations, and tracking and reporting on activities to close any gaps. Most GRC products we've seen seem to do a pretty good job with these functions, which by itself is often of significant value. Chalk one up for the "C" in GRC -- maybe. More on this shortly...

The problems and some solutions

1) When "risks" aren't

Let's cut to the chase. Every GRC implementation we've been exposed to is populated, sometimes to a large degree, with "risks" that aren't risks. What do I mean? Let me give you an example. Which of the following would be a legitimate entry in a risk register?

- Failure to change smoke detector batteries
- Smoke detector batteries fail
- Building catches fire

If you chose the last one, congratulations. Why? Well, the last entry is the only one that represents a loss event scenario and therefore is the only one you can put a meaningful likelihood and impact estimate to. Unfortunately, much of what we see in GRC risk registers more closely resembles the first two entries -- potential control deficiencies. Don't get me wrong -- it's important to document control deficiencies but these should be recorded in a separate part of the GRC product, sometimes referred to "Findings".

(NOTE: if your GRC solution doesn't provide a way to differentiate between risk entries and control deficiencies it's broken in a very fundamental way.)

The bottom line is that being well informed about risk and making cost-effective risk mitigation decisions (the "R" in GRC), and effectively governing an organization's risk landscape ("G") cannot be achieved if the foundation is seriously flawed. For example, we have seen more than one GRC risk register populated with scores of control deficiencies rather than loss event scenarios -- each with a likelihood and impact rating. When we asked the risk professionals how they came up with these ratings, they typically shrugged and said those were required fields so they entered values that felt right. When we probed to understand the assumptions underlying some of these entries, it became clear very quickly that many of the ratings are indefensible because the entries represented control deficiencies rather than loss events. As a result, the risk information being used to make decisions and inform executive management was inaccurate and misleading.

The solution to this problem can simply be a matter of reviewing the entries in the organization's GRC risk register and throwing out or refining any entry that doesn't represent a loss event scenario. This can be challenging in some organizations however, because the notion that control deficiencies are "risks" can be pretty firmly ensconced and difficult to weed out. Once this weeding out is accomplished, what's left is to accurately set likelihood and impact ratings, which we'll discuss below.

2) How likely is likely?

If you asked someone to estimate the likelihood that our sun will go supernova, by most accounts the only reasonable answer is "very high" (if not certain). Likewise, if you asked someone to estimate the likelihood of a malware infection occurring in an organization the answer would almost have to be "very high" (if not certain). Simply stated, a likelihood estimate by itself is rarely very useful. In order to be useful it has to be cast within the context of a timeframe -- e.g., likelihood this year, in our lifetime, or whatever timeframe is relevant. Unfortunately, many of the likelihood scales we see in GRC implementations do not include a timeframe reference. As a result, there can be a lot of ambiguity and inconsistency in the likelihood values for risk register entries, which severely affects the ability to effectively prioritize and report on risk. And without proper prioritization and reporting, effective governance and risk decision-making is a long shot, at best.

This is often the easiest of the problems to fix, as it should simply be a matter of redefining your likelihood scales to include a timeframe reference. Of course, once this has been done all of the risk entries in the risk register will need to be reviewed to ensure that likelihood estimates are still appropriate.

3) Likelihood/impact incongruity

So, let's assume that your organization has either fixed or wasn't subject to the two problems described above. Before you congratulation yourself too exuberantly you might want to have

someone review the likelihood and impact estimates to ensure they're logically aligned with one another. An example might make this clearer.

Let's say that the event in question is personal injury in the workplace. Very often, a risk entry on this topic will have a likelihood estimate of "Very High", reflecting the fact that people suffer paper cuts, carpal tunnel and similar injuries frequently in the workplace. No problem with that. The problem occurs when the impact estimate for a risk entry reflects a worst-case outcome -- an event very different from what the likelihood estimate is based on. In this instance the worst-case outcome for personal injury is "death", which logically equates to a "Very High" or "Critical" impact rating. Taking these "Very High" likelihood and "Very High" impact ratings together, the logical inference is that people are frequently dying in the workplace. If this were true, it would have to be one of the most critical concerns in the organization. Odds are, however, it's not even remotely accurate.

This occurs more often than you might imagine and it can obviously pose a huge problem in terms of effective prioritization and reporting. Sometimes risk professionals will "hedge" these situations by overriding a GRC product's programmatically derived risk rating and downgrading something like our example to "Medium." While this may be an improvement in terms of accuracy, the rationale is often not well documented. This kind of adjustment also shouldn't be necessary if sufficient thought went into the GRC implementation criteria in the first place.

The bottom line is that an organization needs to decide whether their likelihood and impact estimates are going to reflect worst-case or most-likely outcomes. This need to choose between worst-case or most-likely outcomes reflects a fundamental weakness associated with the simple qualitative scales used in almost all GRC implementations. One alternative is to have separate entries and ratings for the common scenarios and worst-case scenarios. An even more robust and useful alternative is discussed in the "Taking GRC to the next level" section at the end of this post.

4) What does "High" mean?

The fourth GRC challenge has to do with rating scales. It isn't unusual to encounter impact scales where each qualitative rating (High, Medium, Low, 1, 2, 3, etc.) is described purely in other qualitative terms. For example:

- High Impact: Any outcome where significant damage occurs to the organization's finances or reputation

In this example we defined the qualitative term "High" with the qualitative term "Significant". Generally, these kinds of descriptions are too open to interpretation, which results in inconsistent ratings and an inability to reliably prioritize risk. Descriptions that include quantitative ranges of loss and specific characteristics of how it materializes (e.g., "Results in a significant reduction in stock price") are much more useful and drive more critical thinking into the ratings that are chosen.

5) Multiplying red times yellow

Last but not least, most people would laugh at the idea of multiplying red times yellow yet that's exactly equivalent to multiplying 3 times 2 when the numbers are based on an ordinal scale. We don't need to get into the mathematical details of why this is a bad idea, but it's an incredibly common mistake in risk rating systems that often contributes to inaccurate assessments and poorly informed decisions.

There are better options, one of which I'll discuss in the following section.

Obviously, the problems described above can severely impact the quality of risk decision-making and governance in an organization. Perhaps not surprisingly, they can also affect compliance because the black-and-white veneer of compliance sits atop an underlying (and sometimes overlooked) need to prioritize compliance gaps and optimize gap mitigation choices. Consequently, if the R in an organization's GRC implementation is bad enough the organization may be checking compliance boxes but it may not be addressing the most important gaps first or optimizing its gap mitigation choices. When this is the case, an argument can be made that the organization isn't fully realizing even the C in GRC.

It's worth pointing out that the problems described here are sometimes less a function of poorly designed GRC products and more a matter of organizations slapping GRC into place without giving enough thought to these kinds of issues. While it is true that many (if not most) GRC products would benefit from refinements that better facilitate and guide a more mature implementation, their clientele have to set the bar higher before that's likely to happen. It's also crucial to recognize that GRC is a set of processes and not a technology. The technology simply facilitates the processes, and both the processes and technology must support the organization's risk management objectives, else it's all wasted time, money, and energy.

Taking GRC to the next level

Organizations that have managed to avoid or fix the problems described above will almost undoubtedly be in a better place, risk management-wise, than many others. And for some organizations, at least for now, that may be enough. For those organizations wanting to defensibly claim risk management as a strength, or where external stakeholders (e.g., regulators) have set a higher bar, there are additional steps that can make a significant difference.

Adopt a risk model

There are models, and there are "models." Many of the risk-related "models" we encounter (particularly in the information security and operational risk arenas) are actually checklists of required controls, or controls that are considered common practice. Examples include COBIT, ISO,

PCI, and the like. And although these lists can be very useful they aren't models in terms of measuring risk in business terms.

The type of models I'm suggesting are analytic models that not only describe risk elements but also the contextual relationships between those elements. (The science geeks in the crowd sometimes refer to these as ontologies.) Note that I am NOT talking about those highly complex mathematical formulas often used in financial risk analysis. A good example of a simpler, more practical analytic model[1] is shown below:

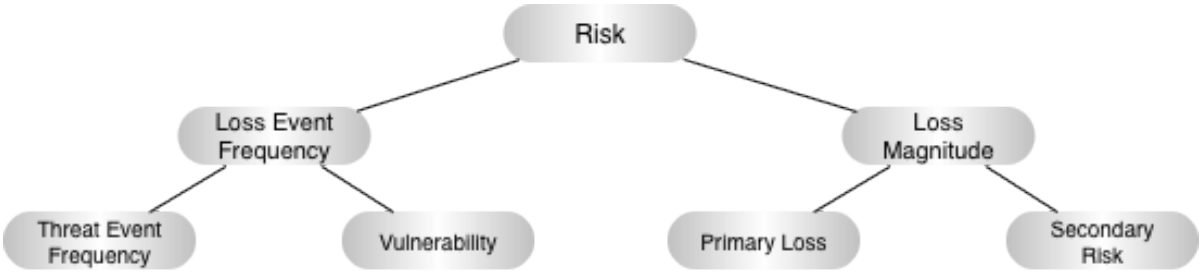

This sort of model is, first and foremost, a framework for critical thinking. In other words, when faced with a risk-related question, the model helps a person think through the problem logically and consistently. An organization using a decent risk model like the one above would be unlikely to fall victim to most of the problems described earlier.

Another significant benefit to adopting such models is that they can normalize risk terminology within an organization. This is important because if you ask ten people in an organization to define risk (or the elements that drive risk), you will undoubtedly get several different -- sometimes very different -- definitions. It isn't a stretch to say that it is impossible to effectively manage what you haven't clearly defined.

Last but not least, a model like this enables practical measurement of risk in business terms.

It's hard to be discrete

If qualitative estimates are problematic, then what are the options? The obvious option is to use quantitative likelihood and impact values like "40% probability (in the next year) with an impact of $1.2M if it occurs." Most people shy away from discrete values like those however, because there

[1] Factor Analysis of Information Risk (FAIR) ontology - An international standard first established by CXOWARE and adopted as an international standard by The Open Group

aren't usually enough data to support discrete estimates and it implies a level of precision that doesn't really exist. A better option is to use ranges or distributions. With ranges, instead of saying there's a "40% probability" we might say the probability is between 30% and 50%. This allows us to reflect the fact that we're working with imperfect data.

An even more powerful option is to use PERT [2] distributions based on calibrated[3] subject matter expert estimates. Not only does this allow us to more accurately and faithfully represent data-related uncertainty, it also allows us to deal effectively with the best-case/worst-case conundrum because we can shape distributions to reflect both ends of the spectrum.

Built for uncertainty

Let's face it -- any time you're estimating the likelihood or impact of a loss event, you're dealing with uncertainty. This is true whether you're using qualitative or quantitative estimates. This being the case, the ability to account for that uncertainty can spell the difference between a well-informed decision and one that, well, isn't. Monte Carlo methods were developed specifically to account for such uncertainty and have been used for decades. Especially with today's readily available computing power, Monte Carlo and other stochastic functions, combined with good models, available data, and calibrated subject matter expert estimates, can be very pragmatically applied to improve the accuracy and utility of risk analyses.

Of course, the entire reason for more accurate risk analysis is to provide an improved level of risk management decision support. In practical terms this means an organization can finally:

- Understand the significance of control deficiencies in business terms,
- Prioritize risk based on loss exposure, and
- Optimize risk management expenditures based on cost/benefit statements that will stand up under close scrutiny

At the end of the day, a thoughtfully designed and implemented GRC solution truly can live up to its promises. To achieve this though, we first have to avoid common blunders like those described above. That gets us part way there. Getting all the way there requires that we leverage more evolved models and methods like FAIR to provide a much higher level of decision support.

[2] PERT - Program Evaluation and Review Technique
[3] "How to Measure Anything" by Douglas Hubbard (ISBN-13 978-047053939)

About Us: CXOWARE is a quantitative risk software company providing cybersecurity risk analysis solutions. CXOWARE's applications provide C-level executives, risk managers, auditors, and cybersecurity professionals with a decision-analysis solution that dramatically improves their ability to measure the monetary loss associated with cybersecurity risk. Decision-makers can then efficiently tailor risk mitigation budgets to the highest loss exposure areas, thus saving money and reducing liabilities to the business.

866-936-0191
www.cxoware.com
http://www.linkedin.com/company/cxoware-inc.
https://twitter.com/cxoware