

Elasticity: Will your organization bend or break?

By Benjamin Tomhave – ISSA member, Phoenix, USA Chapter

Bad data, poor communication, excessive reliance on technology, and bias all impact how elastic the enterprise will be when faced with pressure from increasing risks.

Abstract

Enterprises often jump into risk assessment and mitigation (treatment) with both feet, but to what end? Just because an enterprise assesses and mitigates “risk” does not mean that a risk-tolerant program is in place. Bad data, poor communication, excessive reliance on technology, and bias all impact how elastic the enterprise will be when faced with pressure from increasing risks.

There seems to be a common misconception in risk management these days. In particular, there seems to be a lot of focus on the assessment and treatment of risk, but on what basis? That is, how can organizations effectively manage their risk exposure without first defining what risk means in their context and understanding their tolerance for risk.

This concept of tolerance has been gaining some headway in the last couple years. You might have heard about it, but possibly by a different name, such as survivability or resiliency. Whatever you call it, the same fundamental principles apply:

- How do you define risk in your context?
- What levels of risk are palatable in your context?
- What prioritization approach is optimal in your context?

Risk management itself can be problematic for organizations, even just as a concept. Looking into the space, most of the focus has historically been on financial and business risk management (notice how well that worked out for Wall Street). It seems that the basic concepts are similar and sound, but there are some very complex definitional challenges that can make or break your own practices (not the least of which being gathering and using quality data).

In the context of this article, we are talking about information risk management. As a sub-field within risk management, information risk management is relatively young and under-developed. As luck would have it, there are a few efforts in play that can be leveraged to help address some of these foundational concerns, such as the NIST Risk Management Framework¹ the

1 “Risk Management Framework (RMF) Overview,” National Institute of Standards and Technology – <http://csrc.nist.gov/groups/SMA/fisma/framework.html> (accessed 15 August, 2009).

COSO Enterprise Risk Management Framework,² and the EDUCAUSE/Internet2 Framework.³ ISO 27005 also lays claim to “risk management” guidance, but we will exclude that here for the sake of clarity.

What most frameworks share in common is an approach where you first model the risk management program, then perform an assessment, remediate, and finally analyze the results to evaluate the effectiveness of your controls (essentially gap analysis between expected and actual results from the implemented controls).

A quality risk management program provides margins of error that allow bad things to happen without endangering the whole of the business.

Of particular interest here, however, is the tendency for most risk management programs to skip the first and last steps in that process.⁴ If you perform an Internet search, you will see many examples of risk assessment methodologies, but the search results pertaining to formal, complete risk management programs are sparse. Look at a variety of professional services firms and the services they offer. What you will find is a strong tendency toward assessment and remediation without first putting risk into a properly customized context. Case-in-point, consider a penetration testing report received from a consultant that assigns “risk” ratings (minimally High, Medium, and Low) to each finding. How does this consultant know what is or is not a “high” risk in your context? Did he define it in a way that was specific to your environment? More often than not, the risk level is based on generalizations and has no basis in your organization.

Situations like this lead to a couple problems. First, the use or misuse of terms without proper contextual definition will lead to term confusion. What is a high risk to me may not be a high risk to you, depending on the requirements of each respective context. Second, failure to put risk into a proper context leads to faulty, bad, and oftentimes biased risk decisions. Decisions made with bad data generally succumb to the “garbage in/garbage out” dilemma.

Defining success

Context definition is vital within any risk management program, but it is not exclusive to risk management. Even within the context of this article, proceeding without clear definitions would be detrimental to clarity. As such, let’s spend a

couple minutes putting some scaffolding around the structure.

Risk tolerance itself is not a new concept, but in fact has been around the block a time or two. Previous incarnations may have leveraged terms other than *tolerance* such as *resiliency* or *survivability*. Regardless of the word used, there are common traits. Risk must be well-defined, leveled (or baselined), and prioritized. It represents a mentality or cultural artifact as much as an action or definition.

In essence, risk tolerance introduces a degree of elasticity to the organization. That is, rather than being brittle and easily broken, a quality risk management program provides margins of error that allow bad things to happen without endangering the whole of the business. It is the ability to “bounce back” (or recover) from bad things happening that makes risk tolerance about elasticity and flexibility.

In application, risk tolerance pulls in many attributes from an overall assurance management program, including risk management, operational security controls and practices, policy development and enforcement, business intelligence, incident response management, and security testing (to name a few). The objective is to build an organization that will flex or break well when faced with major threats and vulnerabilities, and that will recover efficiently and effectively. After all, it is not a question of if but when your organization will have a security incident.

To achieve a risk-tolerant organization that can demonstrate a reasonable degree of elasticity, it is then important to first get a comprehensive risk management program in place. This program must establish a working model for managing risk within the organization, including defining what risk levels will be used (i.e., not just High, Medium, Low, but definitions in business terms for each of those levels), how risk will be assessed or measured, what prioritization approach will be leveraged in treating risk, and what sort of key metrics will be tracked to evaluate the overall effectiveness of the program. Putting all these pieces together will lead to an organization that understands risk and can work toward maintaining acceptable levels of risk within their given context. From risk management flows improved practices flows quality data results in risk tolerance.

Self-Healing Hulls⁵

One of best modern examples of engineering in risk tolerance is the application of smart materials to the hulls of racing yachts. On the open seas, a breach of the hull can pose a major threat. In racing, it minimally means introducing drag, and thus slowing down the ship. To help deal with that issue, engineers have developed a special carbon-fiber composite that automatically heals hull breaches. While the immediate application is for competitive vehicles, there are potentially

² “Enterprise Risk Management Frame,” The Committee of Sponsoring Organizations of the Treadway Commission – <http://www.erm.coso.org/> (accessed 15 August, 2009).

³ “Risk Management Framework,” EDUCAUSE – <http://www.educause.edu/security/riskframework> (accessed 15 August, 2009).

⁴ Assertion based on independent research for a forth-coming white paper.

⁵ Sally Adee, “Self Healing Hulls,” *IEEE Spectrum* – <http://www.spectrum.ieee.org/consumer-electronics/gadgets/self-healing-hulls> (accessed 15 August, 2009).

broader applications for safety and security, such as with aircraft.

This engineering marvel provides a great analogy for the desirable state of risk tolerance and management within the enterprise. At its most ideal, the enterprise should have a degree of elasticity that allows it to bend, but not break, when stressed by risk. Having elasticity frees the enterprise to absorb an increase in risk long enough to adjust controls to bring risk back into a preferred operating range. This notion may seem abstract and academic, but in practice it is important to understand and accept. There are too many flash issues that have the capability to crush the enterprise (issues ranging from so-called 0-day attacks to flash popularity to natural disasters) if proper planning around risk tolerance is not performed.

Put into practical terms, flash popularity⁶ provides a ready example of where elasticity can be designed into the enterprise. In the late 90s it was not uncommon for a site to have a link posted to a site like slashdot.org and subsequently crash or become unavailable. Today, there are ways to better handle flash popularity, such as through geo-routing and virtualization. Assuming that adequate network bandwidth is allocated, it is now possible to have idle server space available to dynamically absorb a sudden increase in traffic in order to better distribute the load. Thus, it is possible to create a self-healing design that addresses a certain condition.

DiD vs the Möbius Defense⁷

Unfortunately, for now we do not have equivalent special materials in our networks and systems that can automatically detect and respond intelligently to a breach (though they are coming – see Joel Weise’s white paper “Designing an Adaptive Security Architecture”⁸). As such, we are left with the old practice of defense in depth (DiD), which seems all good and fine until you meet Pete Herzog of ISECOM. In June 2009, Pete introduced a new idea called the “Möbius Defense” to counter the focus on DiD techniques. In his presentation, he points out that the concept of DiD is perhaps faulty, failing to provide the elasticity that we need in our environments. In essence, rather than bending, our environments are likely to break given the right pressures in the right places. In fact, it seems likely that there is not any flexibility in the average environment.

While not everyone agrees with the conclusions of the presentation, it highlights a fundamental concern that impacts your organization’s ability to tolerate risk. Thinking about

6 Flash popularity is an incident of greatly increased site traffic, sometimes associated with the so-called “slash-dot effect” named after the resulting spike in traffic that can occur when a link is posted to slashdot.org. Flash popularity can result in a denial of service from good, legitimate traffic, if the site is not designed to handle certain levels of traffic.

7 Pete Herzog, “Anti-Guerrilla Tactics,” ISECOM – Herzog, Pete. “Anti-Guerrilla Tactics,” ISECOM, http://www.isecom.org/events/The_Mobius_Defense.pdf (accessed 15 August, 2009).

8 Joel Weise, “Designing an Adaptive Security Architecture,” Sun Microsystems – <http://wikis.sun.com/display/BluePrints/Designing+an+Adaptive+Security+Architecture> (accessed 15 August, 2009).

Modeling is...often overlooked in lieu of gut-based risk management decisions driven by non-contextual risk assessments that leverage generalizations and vagueness to drive risk treatments that may or may not truly benefit the organization.

the risk management life cycle (Model=>Assess=>Treat=>Analyze), it seems that there is a failure in analyzing the effectiveness of controls. Moreover, while we often talk about *controls*, those controls are focused on defending threats instead of interactions. When you consider your risk tolerance, make sure you are measuring the right things, or else you may end up surprised when everything falls apart.

In practical terms, the core problem is this: the second you poke a hole in your firewall, you effectively obviate having the firewall.⁹ The second you expose a service to the outside world is the second that service can become compromised. Moreover, if your service has access to other services behind it that are accessed in response to external and programmable stimuli, then one must assume that those back-end services, which may be thought of as protected, are perhaps not as protected as we would like. So goes the argument that defense in depth is a misnomer.

Model and define first

Toward that end, risk tolerance has real-life applicability in modeling and defining your risk management framework. Through this process you define key metrics and sensitivities that are ideally based on key business requirements. One key area of definition is risk ratings – deciding what the threshold is for crossing into each escalating tier. The NSA’s INFOSEC Assessment Methodology (IAM)¹⁰ provides a nice approach for tackling this challenge, which in turn leads to better risk decisions. In IAM, one of the first steps is defining your risk rating levels in terms of business impact (lost revenue, down time, etc.). Everything else follows from these initial definitions.

So it is with risk tolerance. Until you determine your thresholds of pain, you will be challenged to set a good approach for managing risk. How can you be tolerant of risks that are not well-known or well-defined? Moreover, how can you even understand what it means to be risk tolerant without good visibility and understanding into what risk is? And, lastly, you must beware the fruits of the poisonous tree. That is, if

9 The author concedes that this is an arguable point.

10 INFOSEC Assurance Training and Rating Program, “INFOSEC Assessment Methodology (IAM),” NSA – <http://www.iatrp.com/iam.php> (accessed 15 August, 2009).

you do not have good data, then how do you know that you are making good decisions?

A comprehensive risk management program, then, will start by modeling risk for the organization. That model will set business-based definitions for each risk level, determine the preferred approach to assessment, establish a prioritization strategy for remediation, and track key metrics that will measure the effectiveness of the overall risk management program. Modeling is not, however, an easy task, nor is it immediately tangible. As such, it is often overlooked in lieu of gut-based risk management decisions driven by non-contextual risk assessments that leverage generalizations and vagueness to drive risk treatments (remediation) that may or may not truly benefit the organization.

Analysis of key metrics is also commonly lacking in most organizations, not the least of which being because the metrics are not set or collected. While this attribute may seem like bean-counting, there are in fact important benefits to evaluating the effectiveness of the risk management program. Without identifying and tracking key metrics, such as the impact of given controls on the overall elasticity of the enterprise, there is no quality data upon which to refine the program and make well-informed decisions. Perhaps at first glance this is not deemed a heady problem, but in actuality it highlights the voodoo nature of assurance management. Without the clear definition of terms and metrics, and without the collection of metrics data, it is impossible to warrant quality risk management decisions.

Beware Biases¹¹

One of the problems that comes from not having adequate quality data is that it creates a void to be filled by other information, whether it is relevant, factual, or not. Enter the realm of cognitive bias. There are many types of bias, but all of them have one characteristic in common: they influence your decisions without necessarily being true or relevant.

From a risk management perspective, biases can be dangerous as they can create tremendous blind spots in your program. Visibility is key to acquiring quality data, and if you are blinded to certain important areas by bias, then you are in effect increasing the overall risk profile your organization is facing. Now factor in risk tolerance. Biases that lead to bad assumptions can cause an expectation for elasticity that in practice will not be true. When put under pressure, then, your organization may break rather than bend, resulting in much higher costs for the incident and the related recovery.

The “garbage in/garbage out” problem plagues many organizations in all industries. One need only look at key contributing factors for the current economic recession to understand just how important quality data is to making good risk management decisions. The situation also highlights how risk

Data gathering leads to transparency and honesty, addressing a major weakness inherent in most audit and compliance approaches today.

tolerance can be an illusion – a house of cards – if it is not oriented around good, true, honest data.

This problem of quality data is one that has plagued the security industry for quite some time. Even though Bayesian statistics gives us a path through the murky waters, allowing us to build working risk models with minimal data, we still suffer from not having as extensive data as the insurance industry. Enterprises can, however, begin to address these deficiencies, at least for themselves, by implementing comprehensive risk management programs that set a goal of building in elasticity.

There is also good news in a broader sense. Verizon Business has now released two annual reports on data breach incidents.¹² White Hat Security releases quarterly reports on website vulnerabilities seen through their assessments.¹³ The Center for Internet Security has released a set of core security metrics.¹⁴ And, a community has been created for the discussion and dissemination of security metrics approaches, data, and techniques, as well as for hosting the MetriCon conferences.¹⁵

But what do you do with the data once you have it? To what end are these statistics and metrics useful?

Acceptable level of compromise¹⁶

Risk tolerance, or even this notion of elasticity, is a somewhat abstract idea. How does it play IRL (in real life)? Consider the wisdom imparted recently by Jack Daniel (self-described security curmudgeon). In July 2009, Jack described how organizations appear to think about risk tolerance; or, more correctly, how they oftentimes explain away risks in order to avoid remediation. In the end, it comes down to the pain an enterprise feels. That is, if the enterprise does not feel the pain, then it will be less likely to do anything about the potential for pain in the future.

As such, *Acceptable Level of Compromise* (ALC) is defined as “the level of system compromise people and enterprises are willing to live with.” While perhaps a sarcastic and disparaging notion, the point is very clear: enterprises will choose to exist in an insecure state, relying on an artificial level of risk

12 See <http://www.verizonbusiness.com/products/security/risk/databreach>.

13 See <http://www.whitehatsec.com/home/resource/stats.html>.

14 See <http://www.cisecurity.org/securitymetrics.html>.

15 See <http://www.securitymetrics.org>.

16 Uncommon Sense Security blog, “Not that we need another acronym” – <http://blog.uncommonsensesecurity.com/2009/07/not-that-we-need-another-acronym.html> (accessed 15 August, 2009).

11 George Spafford, “Incident Decision Making and Cognitive Bias,” ITSMWatch.com – <http://www.itsmwatch.com/itil/article.php/3690326> (accessed 15 August, 2009).

tolerance, based on bad assumptions that stem from bad data and bias.

Introducing good data, then, frees the enterprise from biases and works against poor decisions. When the emperor finds out his new clothes are imaginary, the right and proper reaction is to address the problem forthwith. So it is for organizations that have suffered from an absence of quality data (if they were using any data at all). Actual metrics tracking actual performance can provide a path to enlightenment that can synchronize the ALC with the desired state of risk tolerance.

Temporal tolerance

If there is one lesson that history should teach us all, it is that what can be built up over the course of years or decades can be leveled in mere hours or days (e.g., Rome, London, Chicago, Hiroshima, Nagasaki). So it is that risk tolerance is about a veritable lifetime of planning and preparation in anticipation of single events that can make the difference between an incident and a major, enterprise-shattering breach.

All that has been described to this point is good and useful, but it cannot be implemented overnight. In fact, it could be argued that to reach a properly elastic state could take months, if not years. What, then, is the practical use of this information? Is there value in pursuing a risk tolerance state when the journey may be beyond the threat horizon?

The answer is a mitigated *yes* that acknowledges the steep road ahead, but also points out that the lessons learned along the way will help inform the enterprise in ways invaluable to its own survival. Consider, for example, the mere virtue of having good data with which to evaluate the effectiveness of controls and the overall assurance management program. Metrics can be identified and tracked without the presence of a formal risk management program. In fact, metrics should be identified and tracked immediately. In the future, perhaps other metrics will be deemed more useful, but some good data is better than none. Moreover, data gathering leads to transparency and honesty, addressing a major weakness inherent in most audit and compliance approaches today.¹⁷

¹⁷ Help Net Security, "Survey: 20% of IT security professionals cheat on audits," Help Net Security, <http://www.net-security.org/secworld.php?id=7659> (accessed on 15 August 2009).

Summary

A degree of elasticity is desirable within enterprise risk management. However, developing a risk management model that includes risk tolerance can be a daunting and difficult task to achieve. Organizations are continually plagued with breakdowns in communication, inadequate expectations for being risk tolerant, reliance on increasingly antiquated concepts for protection, incomplete risk model definition, numerous biases, and bad decisions stemming from all of these other fault states.

The path through these concerns is clear on the surface, but perhaps not straight-forward to address. Risk management programs, if they are formally defined, need to be evaluated to ensure that the quality of data generated is acceptable. Risk decisions themselves need to be measured and analyzed to ensure effectiveness and to identify and resolve the incursion of bias and data quality weaknesses. Information risk management programs should start with business requirements, integrating with existing risk management practices, and mapping key risk definitions to business concepts.

The first steps toward a future state of risk tolerance will identify and collect key metrics that create a culture of transparency and honesty upon which a firm foundation can be built. The more honest an enterprise is with itself, then the better visibility it will have into its own machinations, and the more flexible it will be in absorbing and recovering from incidents.

About the Author

Benjamin Tomhave, CISSP, is a security director in Phoenix, AZ. He holds a MS in Information Security Management from George Washington University and is a member of committees within the American Bar Association and OASIS. He may be reached at tomhave@secureconsulting.net.

