All.Net Analyst Report and Newsletter

Welcome to our Analyst Report and Newsletter

The accumulation of consequences over time

Recent regulatory interventions have finally caused consequences to top management levels of enterprises. Like Sarbanes Oxley, this will ultimately, if continued, result in better governance, which is a good thing.

However, one of these requirements is to report material events (for discussions we will say 5% loss of enterprise value, but it is not that well defined) within a limited time frame (4 days) of the occurrence of the event. The problem here is that it is impossible in many cases to know within 4 days of an event whether the consequences will accumulate to a material level.

Furthermore, up until today, there is little if any analytical framework for making such a determination or to operate a cybersecurity (or any other) program that seeks to anticipate and constrain consequences on such a basis.

This is largely because no systematic approach has been applied to measure the accumulation of consequences over time.

Risk management

Cybersecurity still lacks a good commonly applied definition of risk. The R word, the 4-letter word that ends in K, is defined differently by different people, many with a lot of expertise in cybersecurity. There is an ISO definition and a dictionary definition, and other definitions, and they all, or at least the good ones, go essentially like this:

Risk := Uncertainty about the future

To be clear, uncertainty, in financial terms, does not mean loss. Higher risk does not correspond to higher losses or expected losses. It does not mean higher threats (whatever they are), higher vulnerabilities, or higher consequences. It means less certainty.

Higher risk := less certainty

Lower risk := more certainty

When adding terms to ‘risk’, context is added. High risk of failure is a common phrase, which sort of represents a misunderstanding of the term risk. Higher chance of failure might be the better term. Of all the uncertainty associated with risk, more identifiable outcomes result in failure compared to some norm. In the startup world, the normal chance of failure is about 80% over the first 2 years. Higher chance of failure in that contest means substantially more than 80%, which actually reduces risk, in that failure is a more certain outcome.

Modeling

This brings us to modeling, and more specifically, model-based situation anticipation. In essence, in order to get a handle on risk, we need to have models of the current and anticipated situation, and those models need to cover a range of possible futures. The range of futures can then be translated into financial or other metrics to produce ranges of future outcomes over time. These ranges can be broad or narrow, and the narrower they are, the less risk there is. Similarly, the larger the variation in futures, the higher the risk.

To continue along the line of clarity, if a company is generating $1M +/- 10% per day in profit (or whatever financial metric you may choose), the financial risk on a daily basis is $200,000. If an event causes this number to remain within that range for any given day, it is within the envelope of expected profit. But if the result on a given day is $1.5M in profit, then the risk exceeded the envelope of expectations. Not a bad thing, but a thing nonetheless.

Time frames

When you look at any given business, variations of 50% in daily profits may be commonplace. Or they may not be. That is an inherent risk associated with the business and how it operates. There are variances at time scales. The same company might have a quarterly (90 day) expectation of $90M +/- $5M. Note that the daily expectation of variation of $200K over 90 days would come to a range of $18M, far more than the quarterly variation. That's because accumulation of consequences do not all go in one direction. If they did, there would be no uncertainty. So models tend to be based on statistics that match up the short term to the long term. Also note that we get more accurate longer-time frame predictions than shorter time frame predictions, because the timing and rate at which things happen gets away from the individual events toward the bigger picture.

In most companies of substantial size, the time frame for expectations being met are shorter the lower you are in the hierarchy. So a salaried worker doing sales will have daily, weekly, monthly quotas and be fired if they miss a minimum quota enough times and bonused if they exceed the bonus threshold(s). These decisions will be made on a weekly, monthly or sometimes quarterly basis. But at increasing levels of management hierarchy, the time scales go up both in terms of what they measure and the consequences to the worker.

What about security events?

In cybersecurity, we often hear things like expected loss. That is calculated, as an example, by multiplying the average frequency of an attack by the average loss suffered. 1 million events per day at a firewall with a cost of $0.001 each for the added overhead of dealing with them, produces $1,000/day in expected loss. Of course this number is commonly ignored as being too trivial to care about, but then again, it is rarely calculated. And even more to the point, without the firewall, the expected loss from these million attempts might be far higher. But we don't usually calculate that either. In fact, we don't usually calculate any of it. And we suffer the death of 1,000 cuts, except it's millions of cuts per day, if we don't do better.

Of course I didn't mention here that unit economics are not the only thing going. There is a fixed cost of having a firewall that gets added to by the events it handles. If a firewall costs $10,000 more and reduces the cost per attempted bypass by $0.0001 per event, it saves $100/day which means the $10,000 is paid off in 100 days. Anyone who looks at this will notice that lifecycle costs can and often do exceed acquisition costs. But almost nobody tries to do this calculation because it's too trivial to bother with… that hundredth of a cent... or is it?

Bandwidth also costs, and by moving the firewall to the other side of the connection, we can save or lose more… and on and on. We will not know what we are missing until we look for it.

And did I mention that workers tend to intentionally move things around a bit? If I get a big bonus for exceeding $1M in sales each year and it is near the end of the year and I just hit $1M in sales, I likely delay closing other deals till after the end of the year, so my bonus next year is easier to make. If the better firewall reduces response requirements, am I still needed?

Back to the board room

So here I am on the board of a corporation or two (or 10), and new regulations come down telling me the CEO has to tell me about material breaches. And I ask about it, and the CEO cannot tell me what would be a material breech or if one is underway right now or anything else of use in that context. The CEO asks the CFO, who asks the CISO who asks the … and nobody can actually answer the question.

Of course the CFO should know this answer right away and have it on their dashboard and in their head all the time. But I haven't found one yet that has the right numbers there, or any numbers at all around the issues that can be material to the enterprise. But perhaps an example would help here?

Where's the model?

One of the companies we are starting / have just started, is called Trust Architects. The idea here is to architect the trust-related decisions of the enterprise. To help out a bit, the word trust is commonly defined as:

Trust := The willingness to be harmed by another

We take (accept) the risk (uncertainty) of the harm we can suffer at the hands of a network engineer as part of the deal of hiring them to technically operate our network. We trust them because of the potential negatives that can result from their harmful actions (whether intentional or malicious or not). But how much do we trust them? What is the metric here?

The CFO almost certainly has a model they use for who can spend how much and what levels of authority and approvals are required for such expenditures. As network engineer you might have a company credit card for relatively small expenses, perhaps with a purchasing threshold of $1,000/month, and a purchase order process requiring an approval by a manager for up to $10,000 at which point it has to go further up the chain of command. So we trust you according to the financial model up to the consequence of $1,000/month.

But that model is not apparently applied to the engineering actions you are allowed to take. Because a typical network engineer, say in an automotive plant, can shut down all of the production lines at a factory with a few clicks on a keyboard. The loss might be $1M/hour, and if they do this maliciously, it could take days to fix. So we actually trust the network engineer to the level of perhaps $50M in their role as network engineer but only $1,000 in terms of purchasing power. This makes no sense.

Let's apply the financial approach

So what is the approval process required for a $50M consequence at the company? I bet it goes pretty close to the CEO, even for a major enterprise. And the discussions surrounding the $50M acquisition take place over months, involving internal committees, multiple levels of approvals, and so forth. The same should be true of a network change that could have the same level of consequence… shouldn't it?

The inconsistency can be resolved in many ways, and perhaps the one I would most appreciate as a network engineer would be to have purchasing power of $50M and the salary, office, and perks to go with it. But somehow I suspect that the disregard for technologists among executives would not let this happen. So… the executives better find another way.

But time…

Here's another perspective on that. Suppose we are willing to suffer $1M in consequence, or in other words, trust the network engineer to the level of $1M. We perhaps want the network engineer, who perhaps costs $200,000/y to employ (all in) to be able to make changes without requiring an additional network engineer to approve it every time (another $200,000/y), but not trust them up to $50M.

A different approach is to architect the trust so that the network engineer can shut down the manufacturing lines, but that the lines can be restarted within 1 hour regardless of what that network engineer does, whether by accident or intent. Suppose that change in architecture costs $1,000/day in efficiency, or $360K/year. $49M reduction in potential loss in exchange for $360K/y in costs is an interesting insurance policy. We can price it out against insurance and see if the insurer will give us the same deal or better.

The technical approach to achieving this tradeoff is feasible but not trivial. It involves change control systems that can be reverted, and thoughtful failsafe designs, and a lot of engineering based on different than typical assumptions. The knock-on effects of a network change might be substantial, and in the case of things like software updates and computer viruses, both of which may have larger-scale effects than a network change, the ability to revert to a previous state also tends to lose the intervening information, which can be problematic, etc. But if we try we can do it.

Anything is not everything

Another important thing to understand in this context is that anything is not the same as everything. Part of the problem we have is that the security-related approaches to identity management and resulting controls over access are not architected to address the difference between the two.

- **Anything:** The network engineer can change any setting or configuration of any router or switch or other network technology component in any way it can be changed.
- **Everything:** The network engineer can do anything to all network things.

Technology today typically allows access or not, and as a result, if you can do anything you can do everything. But that doesn't have to be the case. An alternative is a strong workflow mechanism that tracks what is done by whom against trust limits and does not facilitate (and thus prevents) access after thresholds are reached, and better yet prevents access that can lead to thresholds being exceeded. The workflow then requires approval processes commensurate with the activity undertaken, and as a result, additional network engineers, managers, and eventually, the CFO and CEO will be required to make a network change if the consequences require that level of approval.

Of course now we need to teach the CEO what that critical setting is and means so they can make an actual judgment based on understanding the true situation. And when it goes bad, they will be responsible instead of the network engineer.

Conclusions

Of course as you walk down this path, you start to see a whole new line of issues to be addressed. And that is the point of the new regulatory scheme. It more or less forces you to start looking at things this way. And that's a good thing.

