WHITE PAPER

Managing Cybersecurity Surprises:

The Executive's Perspective

By Jack Jones
Chairman
FAIR Institute

Executives hate surprises…

…at least the painful kind, which is the type most often associated with cybersecurity. This being the case, it's important for us to examine where cybersecurity surprises come from and what we can do to limit them.

In my experience, painful surprises occur from one or a combination of the following factors:

- Poor visibility into some part of the risk landscape — e.g., the threat landscape, control conditions, and/or unidentified or misidentified crown jewels.
- Mis-prioritization, which can take two forms: 1) something is characterized as lower risk than it actually is, or 2) characterizing things as being higher risk than they actually are. Both are harmful, but the second one is far more common in my experience, which prevents organizations from focusing on the things that matter most. Critical issues simply get lost in the noise.
- Failure to identify and treat root causes — i.e., putting some problem on the shelf thinking that it had been fixed, only to have it bite the organization in the backside later.
- Poor communication — i.e., the infosec team knew about the exposure and had "communicated" it to executives, but not in a manner that resulted in real understanding or informed decision-making.

My experience has been that most organizations suffer from all of these to some degree. Very often to a considerable degree. In fact, I strongly suspect that if you dug into contributing factors behind any of the major cyber breaches, you will find that one or more of these played a significant role.

Of course, sometimes a low probability but high impact event occurs even though it had been identified, properly measured, effectively communicated, and reasonably (but not perfectly) controlled. Although this is possible, it is by definition extremely unlikely. That said, this type of event shouldn't qualify as a surprise because nobody said it couldn't happen. "*Low probability*" is not the same thing as "*Won't happen*", although executives sometimes overlook that important distinction.

Let's examine each of these sources of unpleasant surprises, beginning with the last one in the list…

Poor communication

Surprises result when reality is misaligned with beliefs and expectations.

Executives and boards should intuitively and rationally understand that there is no such thing as perfect security. Furthermore, many within our profession have tried to convey the various challenges in cybersecurity:

- The ever-changing threat landscape
- The complex and dynamic range of business processes and technologies we're supposed to protect
- Users/employees
- Resource limitations

Nonetheless, how we communicate these challenges matters a lot in terms of a stakeholder's actual understanding and ability to appropriately compare cybersecurity concerns with every other imperative on their plate. A "red risk" in a heat map, a new threat vector, a NIST CSF score of 3.5, or thousands of "critical" and "high" vulnerabilities simply can't be compared effectively to the other things on an executive's radar, which are mostly expressed in economic terms. As a result, the decisions being made based on these non-economic security metrics are driven more by fear and hope than by understanding. Consequently, the chances are much greater that executives will feel blindsided if/when the organization experiences a major cybersecurity event.

There are those in the industry who believe the remedy is to add someone to the board who is a subject matter expert in cybersecurity. Although that might help to some degree, I believe to a certain extent it's just kicking the can down the road. Why insert a cybersecurity expert in between the board and the organization's existing cyber security expert (the CISO)? How does another layer help, especially when internal and external audit, and for some industries — the regulators — are already involved in overseeing the organization's state of cyber security? I believe the communication problem is more fundamental, which I discuss further in the section on mis-prioritization.

Poor visibility

Here’s the bad news — your organization will never have perfect visibility into the assets, threats, and controls that make up your risk landscape. That’s unavoidable. The key is to focus intently on identifying and tracking those assets that have the potential for truly painful outcomes, which are the surprises executives care most about. If you take the time to focus on the crown jewels, then it becomes much easier to have good visibility into the threats to those assets, and the controls protecting those assets. This doesn’t mean that you ignore the rest of your assets — it’s simply a matter of prioritizing where you spend the most time and effort.

By the way — when defining what constitutes your crown jewels do not make the mistake of thinking that every sensitive record or “key” system qualifies. The old saying, “*When everything is a priority, nothing is*” applies here too. For example, although nobody wants to have even a single customer record compromised, unless it’s your only customer or a customer that makes up the majority of your revenue, then the loss of a single customer record is not going to result in a material loss. I occasionally get pushback on this because some people are afraid that they’re being negligent by not having a “zero tolerance” for customer record compromise. What they fail to recognize are two things:

- Defining crown jewels as being those systems/technologies containing more than (for example) one million customer records is not the same thing as saying they’re okay with smaller volumes of records being compromised. An organization isn’t going to ignore smaller volumes of customer records. This is about focus. Once the organization has gotten a handle on strongly protecting systems/apps that handle larger volumes of records they can lower the threshold to 500k, or whatever. You have to start somewhere.
- “Zero tolerance” for customer record compromise is a pipe dream, and a sure way to lose focus and fail to adequately protect the assets that represent the potential for material harm to the organization and its stakeholders.

Of course, crown jewels aren't constrained to customer records or critical systems. Very often crown jewels will include intellectual property, sensitive corporate information, etc. Essentially, they should be anything with an exceptionally high value/liability proposition.

Toward the end of this document I’ll discuss the customer’s perspective in this matter.

Mis-prioritization

There’s an old saying in marketing that, “*Half of your marketing dollars are wasted — you just don’t know which half*.” I’d submit that a similar state of affairs applies to cybersecurity today because the vast majority of cybersecurity risk measurements are unreliable. This lack of reliability is due to several factors that are common to most cyber risk measurements today:

- Inconsistent understanding of fundamental risk-related nomenclature
- No explicit articulation or examination of key assumptions
- Ambiguous measurement scales
- Uncalibrated mental models of the persons performing the measurement
- Uncalibrated estimates
- Failure to capture or faithfully represent the uncertainty in measurements

Any one of these elements can result in unreliable risk measurements, yet in many organizations all six are deficient to some degree. Making matters even more challenging is the fact that many of the cybersecurity technologies in use today rely on risk models that inflate risk ratings. The resulting noise makes it that much harder to identify the things that truly deserve attention.

For the priorities an organization does choose to focus on, unless the organization is measuring risk in economic terms (i.e., dollars and cents) it is unable to reliably determine the value proposition of their cybersecurity investments.

That said, there are a number of commonly expressed concerns whenever the topic of quantitative cybersecurity risk measurement comes up, the two most common being:

- The threats are intelligent and can change with the drop of a hat, therefore we can’t predict the future
- There isn’t enough data to do quantitative analysis

Although the first concern is correct about bad actors sometimes changing how they operate and that we can't predict the future, it ignores the fact that, a) this is just as true for qualitative risk measurement, and b) organizations still have to prioritize based on the best information available at the time. There is no viable alternative.

Both the first and second concerns demonstrate a lack of awareness regarding quantitative methods. There are entire books written on this topic, my favorite being, “*How to Measure Anything*” by Douglas Hubbard. It also ignores the fact that qualitative measurements are invariably data driven — else how do they come up with “Medium”? At least with quantitative methods the analyst is able to faithfully represent the uncertainty in measurements, which is a crucial data point for decision-makers that’s entirely missing from qualitative measurements.

Failure to treat root causes

If your organization plays whack-a-mole, fighting the same issues over and over, you have this problem. Unfortunately, to-date I have yet to encounter an organization that doesn't struggle with this. Common examples include:

- An inability to keep up with patching
- Failure to stay on top of access privileges
- New shadow IT instances constantly popping up
- Weak passwords
- Frequently failing to meet audit resolution deadlines
- 3rd parties consistently failing to remedy the concerns your organization raises

If any of these sound familiar, then your organization isn't treating root causes.

What I commonly find is that organizations are great at applying Band-Aids, and that some organizations think they're doing root cause analysis when what they're actually doing is proximate cause analysis. Regardless, these organizations leave themselves open to surprises (and difficulty defending due diligence) by not taking the time to understand and treat the deeper causes underlying these challenges.

A Customer's Perspective

Some people may argue that the risk management approach described above only considers the executive point of view — but what about the individual customer? An executives's pain isn't a customer's concern. Customers simply want their information to be protected. I get that. After all, I'm a customer of many organizations. So, let's take a look at this…

Let's say that I'm one of 9 million customers of a retail organization. It is 100% certain that my personal information is contained within any system that contains all of that organization's customer records. Hopefully, there aren't very many of those "*all customer records exist here*" systems within the organization, but that's a topic for another discussion.

Let's also say that this organization has several other, smaller systems containing subsets of its customer records — some containing information on half of the customers (4.5 million customers), some containing information on 10% of the customers (900 thousand customers), and others containing information on 1% of the customers (90 thousand customers). The odds of my information being in any of these smaller locations are 50%, 10%, and 1%, respectively. That being the case, as a customer, wouldn't I want the organization to do as much as it can to protect the systems that I am 100% certain contain my information? Logically, as a next priority I'd want them to focus on the systems that have a 50% chance of containing my information, and so on. See where this is going?

This is NOT just about executive pain. It's simply a recognition that the size, complexity, and dynamic nature of the problem, combined with inherently limited resources, means that prioritization has to take place. It's in everybody's best interest.

A Final Note

There is a common belief within the cybersecurity field that it isn't "if" but "when" an organization will have a significant security event. This fatalistic perspective is true to some degree ("On a long enough timeline, the survival rate for everyone drops to zero" – from the movie The Fight Club), and it has driven some much-needed focus on detection and response capabilities. However, if we fail to deal with the sources of potentially major adverse events in the first place, then we expose our organizations to a greater likelihood of having to deal with such events.

The bottom line is that unless organizations get a firm grip on the issues I've discussed in this article, it doesn't matter how much money is thrown at cybersecurity or how much executive support is given to CISOs. Executives (and customers) will still be exposed to significantly unpleasant cybersecurity surprises.

Jack Jones is Chairman of the FAIR Institute (www.fairinstitute.org) and creator of Factor Analysis of Information Risk (FAIR), the leading model for quantitative analysis of cyber, technology and operations risk.

Read Jack's book: *Measuring and Managing Information Risk: A FAIR Approach* (www.fairinstitute.org/fair-book).