Fred Cohen & Associates - Analyst Report and Newsletter

Welcome to our Analyst Report and Newsletter

Risk management: There are no black swans

Karl Popper in his excellent work on the philosophy of science, used the black swans thought not to exist by British ornithologists until they went to Australia, to demonstrate that a universal statement about an infinite set (there are no black swans) cannot be proven by confirmations, but can be disproven by a single refutation. Today it is being used by risk managers to excuse the decision to accept risks, or the failure to do a thorough job of assessing them when hit by an incident. Who could have ever predicted 9/11? Lots of people did. It was not a black swan event any more than recent widely published security breaches.

There are indeed black swan events in the world, but they are very rare and even more rarely cause substantial negative consequences. The discovery of quantum entanglement was a black swan event when it was first discovered, but it is now known and is thus not a black swan event when it happens in the future. Once we knew that all swans were not white, the event passed and all of the risks that depended on all swans being white should have been reevaluated.

It's time to own up to your risk management process

There are four likely things and one very unlikely thing that happened when a serious negative consequence occurs and risk management did not mitigate it.

Likely: The risk was not identified in the risk management process

Likely: The risk was identified but not properly characterized (e.g., rated as low probability)

Likely: The risk was identified and accepted by management

Likely: The risk was identified but not mitigated adequately

Very unlikely: Nobody could have known - it was a black swan

If a realistic risk management process is to proceed, it must start with the people involved owning up to mistakes and working to correct them. The black swan excuse has been used again and again, and frankly, or perhaps we should call it Fredly, it is an embarrassment to see the information protection community laying off poor decisions and bad management practices on the innocent creatures of the Earth or the scientists who were surprised and then adapted their theories to compensate.

Who could have ever imagined?

With the Internet available to all and news stories on a minute by minute basis relating to information protection weaknesses, and scores of years of theory and experiments, the excuse of "who could have imagined a tape could disappear" is no longer viable, if it ever was. The fault does not lie in our data - it lies in ourselves. The failure of imagination is not the problem - the failure is in the approach to risk management.

By now, there is more than enough evidence to support the hypothesis that probabilistic risk analysis does not work in the information protection arena today for reasons that have been widely published by now. The simple charts we see for making risk-related decisions are fatally flawed, and to fix this flaw, there are a few relatively simple steps that have to be taken.

The chart at the right is an example of a typical layout used in risk management. It shows areas that rate risks using likelihood P(x) and consequence C(x) and associating them with regions for which the organization identified thresholds based on risk tolerance. But it has been modified in one simple way. The low probability, high consequence area has been eliminated. The thinking is this: The events that hurt have high consequences with risk accepted. Risk is accepted because, even though the event would be disastrous, the likelihood is so low that it can never happen - but of course it can. The 100 year flood - happens on average every 100 years - except that weather patterns change and as a result we cannot really get a good average that is stable over that period. Probability is skewed toward events that occur frequently - because in order to get good statistics events have to occur frequently enough to (1) have counts that provide defined error levels and (2) be properly detected and categorized. The reason for (1) is obvious to anyone who has taken a statistics course, and the reason for (2) is that errors in attribution of effects to causes happen and when the count is low, even a single attribution error causes wildly different outcomes.

High			
Med			
Low			NO!!!
P(x)/C(x)	***Low***	***Med***	***High***

Figure 1 - A modified risk rating chart

There is another major flaw in the chart, and if you haven't guessed it yet, you haven't read enough of my previous articles on risk management. I will harp again on the notion that the probability of events when malice is involved can not be accurately characterized in the same way as random stochastic processes that are the basis for the mathematics of probability. The probability is unknown in advance and 1 after the fact. A minimum standard of coverage is required because the choice to ignore things that can kill you will ultimately result in your death. No matter how many times you have closed your eyes and walked across the highway, it doesn't make it safe to do the next time.

Where to go from here

I have often rallied for people who work in information protection to tell management not to do foolish things, even at the risk of loss of job. And in the economy as it is today, loss of job is a real serious threat to many folks. I understand. But that does not excuse. There is a level of professional responsibility involved in fixing the risk management problem, and while this short paper is not going to solve that problem, perhaps it will at least help to eliminate the excuses that just don't hold water.

"Yes Virginia, there is a Santa Clause." - it might work for small children who have lost their faith in humanity, but it is no way to responsibly run a business that affects other people. There are no black swans in computer security - or at least I haven't seen one in the last 15 years or so.