Technology Risk Management: Scope and Scale

January 9, 2018

Richard J. Van Horn

rvanhorn@technologyatrisk.info

https://www.linkedin.com/in/rvanhorn

http://www.technologyatrisk.info

Table of Contents

Editors & Reviewers

Executive Summary

As explained in the first in this series of white papers, Technology Risk Management is a distinct and separate function from Cyber Security, IT Compliance and other technology control functions. Where Cyber Security typically reacts to malicious attacks targeting an organization and IT Compliance ensures the organization is meeting all its regulatory obligations, Technology Risk Management enables an organization to *proactively* manage potential technology events based on their *possible business impact*. The benefits of this approach go far beyond compliance with regulations. This approach enables proactive decisions to be made based on economic impact to the organization and allows business executives to directly manage technology from a risk perspective.

The proactive nature and ability to convey the business impact in dollars and cents distinguishes Technology Risk Management (TRM) from other technology control functions. That distinction has two implicit assumptions: all technology related assets, processes or services are within scope and that the risk equation - the inherent value of technology assets, the value of controls and resulting impact - can be quantified in economic terms.

These fundamental differences highlight the value that Technology Risk Management offers beyond the technology control functions performed today. Proactively managing ALL possible technology events AND relying on dollars and cents to convey business impact enables the executive office to make true risk based decisions involving the organization's technology control environment. This goes beyond the typical reactive posture of most Cyber Security functions.

This white paper further explains these assumptions and includes a compelling example - a poster child - of why the scope and scale of technology risk management is so important.

Scope of Oversight

As explained in the first white paper in this series, the primary function of Cyber Security is to defend against active threats and attacks to the organization's technology services, including preventing the theft of intellectual property or client data, ensuring services are available, etc. The perpetrators in these attacks can be malicious insiders, outside attackers including hackers, nation states and a variety of other players. Just like it's peer department in the real world - physical security - Cyber Security is typically focused on malicious activity that could

impact high value assets and the day to day operations of the firm. Of course, Cyber Security may also be proactive and try to prevent attacks. For example, some Cyber Security departments use behavioral analysis to look for anomalous activities before there is an actual attack. However, Cyber Security's primary focus is very clear: protect high value technology assets from actual or suspected malicious activity.

In comparison, Technology Risk Management is a risk function that looks to prevent events from happening in the first place rather than react to events as they are happening. Through a variety of assessments and approaches, TRM will assist in identifying high value assets of the firm and help determine if the controls in place to protect those assets are adequate and cost effective. In effect, TRM's purpose is to help allocate technology resources appropriately based on the potential negative economic impact of an event and controls to mitigate that event. Given this view, the functions performed by Cyber Security are in scope of oversight of TRM because of the need to protect the firm's most valuable assets in a cost effective manner. TRM helps identify the firm's most valuable assets, helps determine the cost of controls to protect those assets and as a result, may direct the activities of the Cyber Security department.

However, TRM goes beyond the scope of Cyber Security to include all technology related activities that may impact the organization. In contrast to Cyber Security that typically reacts to active attacks as they are occurring, TRM focuses on ALL technology assets and activities that *could* impact business services including passive threats and events that have never happened before - so called black swan events. Of course, Technology Risk Management's goal is to avoid those types of events or minimize their impact. However, it is important to note the difference in scope between Cyber Security and TRM. That difference is highlighted in this whitepaper with a poster child of Technology Risk Management. That example, where a single, black swan, non malicious event from a rudimentary process bankrupted a company, will highlight why the scope and scale of Technology Risk Management are so important.

Scope: Cyber Attack Versus 'Glitch'

To many, it is quite clear what constitutes a Cyber Attack. The examples are almost endless, including the well known Target Breach, the Sony attack, the attack on the Office of Personnel Management, etc. The list is long and continues to grow. An event associated with technology risk is not always so clear. However, the press has identified a term for the broad category of technology events that are not cyber attacks. They are called 'glitches'.

It is clear that not all technology issues and events are the result of a Cyber attack. Sometimes bad things happen for other reasons. The premier example of this occurred on July 8th, 2015. On that Wednesday, the media reported in quick succession that United Airlines was having a service outage, as was the New York Stock Exchange and finally, the Wall Street Journal. For different parts of the day, all 3 companies were out of service or 'dark' as they say. For many, the coincidence meant a coordinate attack, possibly by a nation state. Given the size of the organizations and impact - the NYSE could not process trades for the better part of the morning and United Airlines planes were grounded for most of the day - it was not an unreasonable assumption. That concern brought news trucks to the front of the Exchange - apparently to the surprise of their CEO.

However, by all public and off the record accounts, this trifecta of events were unrelated and not caused by malicious act or with malicious intent. As the press eloquently explained, they were 'glitches' where technology failed and services became unavailable. Of course it's safe to assume each firm did a root cause analysis to understand how the events occurred and to prevent repeat performances, but it's widely accepted they were not the result of a Cyber attack. Rather, they were the result of operational challenges or mistakes. Unfortunately, some of those lessons were not learned by peer companies. This past October 2016, Delta Airlines also had a service outage, where flights were grounded for a day and a half.

These examples show that the scope of Technology Risk Management is important: the services that went dark were important revenue sources for their organizations. The outages led directly to lost revenue. In addition, they were not the result of Cyber attacks: other factors outside the scope of Cyber Security lead to the outages. As summarized earlier, a Technology Risk Management program should include all technology services and processes that may have a negative economic impact on an organization. By including all activities that affect high value technology services, the TRM program can identify operational issues that might impact services and advocate for controls to mitigate those events. The scope of that view is beyond the current scope of Cyber Security.

Again, while a Cyber attack may have tremendous economic impact on an organization, an organization is equally likely to be the victim of a 'glitch'. Those glitches are also very costly with impact both on the bottom line and reputation of the firm. Looking beyond Cyber Security, a TRM program helps prioritize all technology related activities and events - including glitches - that may have a negative economic impact on an organization.

Scale of Impact

The second assumption associated with Technology Risk Management is that the impact of possible events can be conveyed in economic terms. That is, Technology Risk Management can differentiate itself from other technology control functions, which are typically qualitative in nature, and provide quantitative analysis showing the possible impact of an event and value of controls in dollars and cents. While those dollars and cents may only be indicative of the true economic impact rather than authoritative, being able to convey impact in dollars and cents is a significant change from best practices today. Quantifying the economic impact of possible events changes the landscape significantly and would put Technology Risk Management in the same category as other risk functions, such as market and credit risk. In those fields, risk and impact are conveyed in economic terms. TRM will do the same for technology risk where technology related issues can be conveyed in economic terms and managed directly by business executives with support from the technology division.

Quantitative Versus Qualitative

As a general practice, technology related issues are usually prioritized on a qualitative scale: high, medium or low. In some situations, the scale may be numeric but still finite: say 1 to 5.

Obviously prioritizing issues or control gaps is important, but qualitative measures such as high/medium/low or 1-5 cannot convey the true potential economic impact of a control gap. It also does not give enough context to determine the cost benefit of controls to address that gap. Given a high risk issue or control gap, does a 'high' rating justify a $1M dollar expense or a $10M dollar expense? By their very nature, qualitative scales are limiting in their breadth and scale.

A telling example of this limitation is with audit issues. Most large firms will have dozens if not hundreds of audit issues. If a firm has 500 audit issues, 100 may be ranked high, another 100 may be ranked low with 300 ranked as medium. In many firms, the 100 low issues will probably be accepted and not proactively remediated: over time they may be included in a larger initiative but they will not be addressed with a dedicated project. In most firms, the 100 high issues will be remediated quickly with an 'all hands on deck' mentality. However, for the remaining 300 that are ranked medium, the question is, are they all really 'medium'? Are they all really the same classification without any further differentiation? For many CIO's, a further way to prioritize those medium issues -

and possibly the high issues as well - would be valuable to help allocate resources and manage costs.

Determining the possible economic impact of an issue - the actual losses if the event occurred - is proposed here as the most effective way to properly prioritize control gaps, audit issues or other technology related issues. Being able to quantify the impact of possible technology events in dollars and cents enables the executive office to make informed decisions regarding technology risks - both the possible impact and cost to mitigate. That information allows the 'business' to make an informed decision whether to accept the risk or mitigate.

Those decisions, made by business executives and not technologists or security professionals, is a fundamental benefit of a Technology Risk Management program: moving the decision making on technology risks out of the technology organization and into the executive office. That goal has been a recent initiative in the financial services industry. If done with the proper tools, that migration allows business managers to proactively manage their technology environment based on risk and economic impact, alongside other management initiatives such as strategic direction, partnerships, etc. This migration would put technology risk alongside the other risk functions many firms have, including market and credit risk. As a peer function, technology risk management will enable the executive office to have a more holistic view of their organization's risks.

Measuring technology risk in economic terms is a significant leap forward from current best practices. It is such a significant leap that many say it cannot be done: there are too many variables, there is insufficient information to determine probability, etc. The only practical way to assess technology risk is via qualitative measures. Unfortunately, the need for better ways to manage technology events and outages has never been greater: obviously there are the data breaches - including one that affected the executive office at Target - with no sign of decline. There have also been high profile service outages at Delta, the New York Stock Exchange and United Airlines to name a few. These events and outages will only continue. The goal is a better way to prevent them or minimize their impact.

A Compelling Example

The purpose of this white paper is to highlight the scope of a Technology Risk Management program - what should fall under its purview - and that dollars and cents is the best scale to show the possible economic impact of technology events.

Unfortunately, there is an example - a poster child company - that shows the value of managing all technology processes that might impact high value services (scope) and that 'high' cannot convey true impact (scale).

Knight Capital Group

Knight Capital Group was an American global financial services firm engaging in market making, electronic execution, and institutional sales and trading. With its high-frequency trading algorithms Knight was the largest trader in U.S. equities, with a market share of 17.3% on NYSE and 16.9% on NASDAQ[1].

In a widely known event - called a glitch[2] by the press, which again is the simplistic adjective used to describe these types of technology issues - one of the firm's proprietary trading algorithms executed a series of orders over a matter of minutes rather than days, resulting in a loss of $460 million dollars and driving the firm into a merger that ended it's independence.

While the details have been discussed in many other publications[3] [4], the results of the event have made Knight Capital Group the poster child of technology risk management: a single technology related event - in this case an existential event for the firm - literally bankrupted the company in 30 minutes.

Wikipedia has the best summary of the technical details that led to the trading losses:

2012 stock trading disruption

> On August 1, 2012, Knight Capital deployed untested software to a production environment which contained an obsolete function. The incident happened due to a technician forgetting to copy the new Retail Liquidity Program (RLP) code to one of the eight SMARS computer servers, which was Knight's automated routing system for equity orders. RLP code repurposed a flag that was formerly used to activate the old function known as 'Power Peg'. Power Peg was designed to move stock prices higher and lower in order to verify the behavior of trading algorithms in a controlled environment. Therefore, orders sent with the repurposed

[1] https://en.wikipedia.org/wiki/Knight_Capital_Group

[2] http://dealbook.nytimes.com/2012/08/02/knight-capital-says-trading-mishap-cost-it-440-million/?_r=0

[3] http://money.cnn.com/2012/08/09/technology/knight-expensive-computer-bug/

[4] https://www.bloomberg.com/news/articles/2012-08-02/knight-shows-how-to-lose-440-million-in-30-minutes

> flag to the eighth server triggered the defective Power Peg code still present on that server. When released into production, Knight's trading activities caused a major disruption in the prices of 148 companies listed at the New York Stock Exchange, thus, for example, shares of Wizzard Software Corporation went from $3.50 to $14.76. For the 212 incoming parent orders that were processed by the defective Power Peg code, Knight Capital sent millions of child orders, resulting in 4 million executions in 154 stocks for more than 397 million shares in approximately 45 minutes.

This event is significant from a Technology Risk Management perspective for two reasons. For one, this was NOT an attack by a nation state or 'hackers' that normally makes headlines - other wise known as a Cyber attack and typically managed by the Cyber Security department. Rather, this was the result of a breakdown in well known, relatively mundane technology processes: quality assurance testing and change control. Those processes are well known components of the Software Development Life Cycle. Again, SDLC is not typically part of a Cyber Security program. However, in this example, given the potential impact on high risk services, it would be part of a Technology Risk Management program.

The second reason is that the cost of this event has been quantified in dollars and cents: $440 to $460 million dollars depending on the source. As discussed earlier, qualitative measures of risk cannot convey this magnitude of loss: 'high' does not begin to indicate the size of this economic impact. To repeat, this is not a hypothetical 'what if' scenario but an actual technology related breakdown that lead to significant dollar losses, the firm losing its independence and merging with another company only days later.

The conclusion is that if Knight Capital had a Technology Risk Management program that included SDLC activities, it might have ensured extensive testing of these software changes was performed, as the possible negative impact would be $460 million.

Conclusion

Over the past 25 years, information technology has become the lifeblood of almost every organization in the developed world. Information technology and services - managed directly by an organization or indirectly via a service provider - are indispensable when providing services to clients, partners and employees. Depending on the organization and degree of interconnectedness, the economic

impact of a technology outage or other issue can be significant - possibly existential - to the organization, its partners and possibly an entire industry.

For the majority of that time, information technology has been managed based on any number of criteria: return on investment, corporate strategy, pet projects of senior managers. Only recently have Compliance and, for lack of a better term, Cyber Security become drivers to manage and govern information technology.

However, as Cyber attacks and data breaches continue to occur, and the Knight Capital trading catastrophe shows, there is a need for a better way to manage technology and its related risks. Being fully compliant with industry or regulatory requirements is an important first step to being secure. However, technology events at firms that are compliant show compliance is just a baseline and additional controls were necessary.

This white paper showed that additional controls should be justified based on the possible economic impact of technology events. Nothing but dollars and cents can truly convey the scope of potential loss related to a data loss or service outage. While hard to determine, dollars and cents are the best way to convey that cost. In addition, this white paper showed that the scope of technology risk management and governance should include any technology related activity that could negatively impact - in a significant way - business services. While mundane and rudimentary, Knight Capital's existential trading error showed that change control and QA testing, when associated to a high risk business service, should be in scope for governance to ensure controls are appropriate. Otherwise, as Knight Capital showed, catastrophic events can take place.

This white paper has attempted to show that to minimize negative technology events, a Technology Risk Management program has to be all inclusive of all technology related activities. Even the rudimentary and mundane can have an impact on an organization. The other main point of this paper is that as best possible, the economic impact of an event - and the cost to mitigate that event - should be established to convey the potential cost to the firm, and allow for proper decision making. Only through dollars and cents analysis can an organization make proper technology risk management decisions.

Determining the inherent economic value of technology services will be the subject of the next white paper.

About the Author

Richard has been in the world of IT Governance, Risk & Control over 20 years, and is currently a Vice President at JP Morgan Chase. His career has evolved along with the field, from working as an IT Auditor at the Federal Reserve Bank of Boston, to implementing enterprise security solutions at Fidelity Investments, to managing IT Risk at Goldman Sachs, the CIT Group, DTCC and now JP Morgan Chase. He was certified as a Certified Information Systems Auditor (CISA) and is currently as a Certified Risk and Information Systems Control (CRISC) from the Information Systems Audit and Control Association (ISACA).