CyberCube

Three Degrees of Separation

Understanding Cyber Tail Risk with Counterfactual Analysis

CYBCUBE.COM

Three Degrees of Separation

Understanding Cyber Tail Risk with Counterfactual Analysis

Cyber analytics is a conversation between the world of risk modeling and the world of information security. The value of cyber insurance is unlocked in how these two worlds come together. For cybersecurity practitioners, risk modeling offers the potential to elevate considerations from the realm of "techspeak" into the bottom-line language of business insurance: dollars and cents, likelihoods and impacts, as well as frequency and severity.
Reciprocally, information security knowledge is what grounds cyber risk modeling, bringing it from the theoretical realm of probabilities and hypotheticals into concrete examples – not only the potential impacts to businesses but also the actors and methods by which such impacts can be realized. Each discipline informs and enriches the other. This exchange shapes how we work at CyberCube and allows us to holistically engage the cyber insurance market and move it forward.

One important area where cybersecurity informs our collective understanding is by bringing context to the tail of (re)insurers' loss distributions. While most people now grasp that cyber comes laced with catastrophic risk potential, it remains difficult for many to imagine what a large-scale industry loss of $30 billion might look like, for example.

For this purpose, one approach that can assist greatly is counterfactual analysis. Counterfactuals require looking at the past differently: not what ***did*** happen, but what ***could*** have happened, or what ***nearly*** happened. This approach provides a useful bridge between historical events and hypothetical events having similar characteristics but worse outcomes. By using counterfactuals in this way, we can understand cyber tail scenarios more tangibly – a 100-year event might only be a few degrees of separation from an event in recent memory.

The CyberCube team recently conducted a counterfactual analysis to understand cyber catastrophe events in this way. This research follows a similar path as our earlier counterfactual paper[1] but with more of a tail focus. This study is concentrated on widespread malware events, which certainly figure among the Realistic Disaster Scenarios or drivers of Probable Maximum Loss for every cyber (re)insurer.

In this study, we began with two past events that were certainly notable at the time but were ultimately inconsequential in terms of claims paid: ***SolarWinds*** and ***WannaCry***. These events allowed us to study two different ways that malware could propagate widely. For each of these events we considered the important factors shaping how it unfolded, and contemplated ways it could have played out differently. Finally, we quantified the results by leveraging CyberCube's Portfolio Manager v5 using our Industry Exposure Database to put numbers to each counterfactual scenario.

[1]Gallagher Re & CyberCube, *A History of Near Misses: Utilizing counterfactual analysis to understand cyber risk*, April 2024

Reconsidering SolarWinds

Event	*SolarWinds*
Occurrence	December 13th, 2020
Motivation	Espionage
Event Type	Software Supply Chain Injection Attack
Threat Actor	Cozy Bear aka APT29 or "Nobelium" (Russia)
Foot Print	US-focused event with ~ 30,000 user organizations and ~ 150 compromised
Description	A software supply chain injection attack aimed at espionage on US government systems connected to a Russian linked group known as Nobelium (widely believed to be a part of the Russian Intelligence Service SVR). The attack gave foreign operatives access to US government systems for pervasive espionage before the discovery and remediation of the breach by the company.

Background

The ***SolarWinds*** attack was part of a Russian cyber operation to gain privileged access to United States government systems. SolarWinds, a Texas-based company, makes software which allows businesses to observe and even automate parts of their IT and network infrastructure. This capability was used by thousands of organizations, including the U.S. government. Around 425 of the Fortune 500, the top ten U.S. telecom companies, all five branches of the military, and most U.S. government agencies and services relied upon this software (as reported by multiple sources before SolarWinds removed the listing on their webpage)[2]. In total, between 115 and 150 companies were reported to be compromised, with estimates ranging between 10,000 to 20,000 exposed to the malicious update of the total roughly 25,000 US and 30,000 global users.

The breach was the result of access to the development environment of SolarWinds' software engineering teams. The attackers gained access to SolarWinds systems and were able to inject a malicious backdoor into the code, nicknamed 'SUNBURST'. Instead of attempting to introduce changes that the development teams might detect, the attackers reportedly waited until the final code compilation before slipping their malware into the final product to avoid extra security checks. This level of planning and detail shows the sophistication and capability with which this threat actor operated. When the update was downloaded by end users, the backdoor would reach out to the threat actor's servers to establish a connection after a two-week waiting period. After the backdoor was established, the threat actors could use malware known as 'TEARDROP' to upload malicious software through the originally-named 'SOLORIGATE' attack channel.

[2]Newsweek and The Verge articles published on Dec 15, 2020

Discovery of the breach almost occurred a few times when anomalous activity was identified in networks and investigated. Even Palo Alto Networks, a major cybersecurity company, investigated such an incident relating to SolarWinds but eventually concluded it was an isolated event until a reinvestigation was triggered when the final news of the supply chain origination was widely reported. The malicious update was finally discovered when cybersecurity firm FireEye notified SolarWinds and the U.S. government to begin an investigation. CrowdStrike subsequently isolated a section of code identified as the maliciously-injected script. What followed was a coordinated response to isolate affected systems and begin remediation of the infections.

The discovery of this breach triggered a much larger software supply chain discussion across the government and cybersecurity sectors. The fact that this software and the supply chain attack method gave attackers elevated permissions with pervasive access across networks concerned many in the cybersecurity world. Many observers naturally asked the question: ***What if the threat actors decided to do something disruptive instead of just spying?***

Considerations

While the depth of access obtained by threat actors in this event was truly frightening to the information security community, it was equally frightening that such access could have been used for much more destructive ends. Several factors were considered:

- **Intent:** The SOLORIGATE attack was focused on espionage, as the Nobelium group was intent on access to the U.S. Government and prominent commercial company data and communications. However, we can easily imagine situations where the objectives were broadened to include or shifted toward financial and destructive ends. Had the focus been widespread and destructive, the level of access would have meant an event much larger than what occurred.

- **Targeting:** Much of the cybersecurity analysis of the attack has highlighted the targeted victim group. Most of the victims who reported pervasive access across their networks were U.S. Government services, branches and agencies – and related or connected public companies. This focus shows the threat actor's intentions were geopolitical: to establish intelligence and a future command-and-control foothold should a conflict arise. Had the intent been focused on destruction or financial gain, a payload would likely have targeted a wider audience outside of U.S. government-connected circles. Exploiting more of the access gained at Fortune 500 companies and corporations would have resulted in a larger financial and economic impact.

- **Larger SPoF:** Single Points of Failure (SPoFs) are technology hub points that provide services to many organizations as well as potential access to many organizations. SPoFs make widespread malware events far more feasible by creating economies of scale for threat actors. But SolarWinds, while a large vendor, is by no means as large as IT vendors get. Companies such as VMware have software that is ubiquitous across global networks for virtualization and network management. We can imagine the risk

if a much larger vendor with privileged network access exposure – similar to SolarWinds – were to be compromised. A similar attack deployed via a larger SPoF could have generated substantially larger losses.

Counterfactuals

With these factors in mind, our team considered several counterfactual variants for ***SolarWinds***. Note that each of the variants in this study compound, meaning that as each new change is described, it is assumed the prior changes have also occurred.

■ Damage/Financial Focus

A shift in intent for the threat actor could have meant a much greater impact and financial loss for companies affected by the malicious update. The use of wiper malware or ransomware would have meant a far greater impact and financial loss on the infected. The counterbalancing point is that such a disruptive payload would likely lead to faster detection and response to the attack, reducing the spread and overall target pool. This changed intent would have likely led to more government-connected entities and companies initially damaged by the attack, but a smaller infection total. These changes could have roughly doubled the impacted company count to around 350 companies, with around half estimated to file claims.

■ Wider Targeting

A change of target focus for the threat actor could have meant a much wider pool of potential victims and a higher impact rate for companies affected by the malicious update. To widen the target pool of the operation, threat actors would have needed to speed up the connection to infected systems to spread the operation beyond simple government-connected entities to the broader exposed user base. This would have increased the impacted company count to over 2,000 companies with roughly 1,000 estimated to file claims.

■ Indiscriminate Targeting

Had the threat actors moved to indiscriminate targeting of all affected users, the attack could have looked much different. With damage or financial focus from earlier changes in mind, attackers targeting all affected users would increase the impact rate to a much larger pool of companies. While resource constraints would mean not all the affected companies may get attention from threat actors, it is safe to assume many users would have had to deal with ransomware or wiper malware had they been deployed. This could have led to an estimated 4,100 companies impacted with roughly 2,200 estimated to file claims.

■ Larger SPoF

The final assumption is perhaps the most crucial of questions to ask. What if the attack had been spread through a much larger SPoF? Considering the pervasive access SolarWinds gave attackers, the logical large-scale SPoF to contemplate would be VMware. A virtualization and network management vendor with over 100,000 U.S.-based users, VMware would offer a similar level of network access and permissions to an attacker, and a similar foothold in top companies around the world. This would immediately widen the footprint of the attack by a factor of four. With counterfactuals #1 through #3 assumed to have still occurred, this widens the pool of impacted companies to around 17,000 with an estimated 8,200 filing claims.

Quantification

We then used our Portfolio Manager v5 catastrophe model in an unconventional way to estimate losses from these counterfactual variants that are not in our event catalog, or any event catalog. We applied these counterfactual variants, along with an estimate of the actual SolarWinds attack as a baseline, against our Industry Exposure Database for US Standalone Cyber insurance[3].

The compound effect of these changes can be seen in **Exhibit 1**. With each counterfactual change, a larger pool of companies is impacted, and a greater number of these companies are estimated to file claims.

Exhibit 1: SolarWinds Counterfactual Estimates

	Factor(S)	Baseline	Counterfactual #1 "Damage Focus"	Counterfactual #2 "Wider Targeting"	Counterfactual #3 "Indiscriminate Targeting"	Counterfactual #4 "Larger SPoF"
Footprint	Customers/Companies	24.8K	24.8K	24.8K	24.8K	101K
	Change	-	Lower Infection Rate	-	-	Larger SPoF
	Infected Companies	18.5K	16.7K	16.7K	16.7K	68.2K
	Change	-	Wider Deployment & Higher Impact Rate	Wider Deployment & Higher Impact Rate	Wider Deployment & Higher Impact Rate	-
	#Impacted Companies	144	349	2.1K	4.1K	16.9K
Severity	#Companies Filing	35	171	1.0K	2.0K	8.3K
	US Insured Loss					
	5th	$7.3M	$427M	$3.8B	$7.9B	$13.7B
	50th	$15.9M	$830M	$5.0B	$10.0B	$17.1B
	95th	$43.3M	$1.37B	$6.6B	$12.3B	$20.6B
	50th Percentile Loss as					
	Loss ratio %	0%	11%	66%	132%	224%
	Limits %	0.0%	0.1%	0.7%	1.5%	2.5%

All estimates shown are for the US Standalone Cyber market.
Numbers in blue indicate the result of an assumption change from the previous counterfactual scenario (or baseline) to the current one.

As the event footprint widens with each successful counterfactual scenario, so does the potential for losses. Counterfactual #1 illustrates how a simple change of threat actor intent could potentially have resulted in over $1 billion in claim costs from US insureds. As we progress from counterfactual #2 through #4, we begin to understand how wider targeting could have generated losses much larger than cyber insurers have seen to date – perhaps (in #4) also enabled by an even broader initial customer base for potential infection.

This counterfactual analysis allows us to understand modeled tail losses in a more tangible way. **Exhibit 2** shows the range of loss from each counterfactual plotted against the industry loss curve, or Occurrence Exceedance Probability (OEP) from CyberCube's Portfolio Manager v5. For example, counterfactual #3's median loss of $10.0 billion would be a 1-in-81 year industry occurrence, and counterfactual #4's median loss of $17.1 billion would approximately correspond with a 1-in-191 year event.[4]

[3] For this purpose, the Baseline footprint estimate is calibrated against the historical event data, but using current market take-up rates and loss cost assumptions.

[4] Note we are providing a range for the potential total loss amount. While each counterfactual scenario determines how many companies are affected by the attack, it is nontrivial to determine which specific companies are affected. This needs to be sampled, and results can vary meaningfully depending – for example – on how many Fortune 500 companies' insurance towers are pulled into the event. For this reason, we show the median estimated event loss as well as 5th and 95th percentiles.

Exhibit 2: SolarWinds Counterfactuals vs. US Standalone Cyber Industry Loss Curve

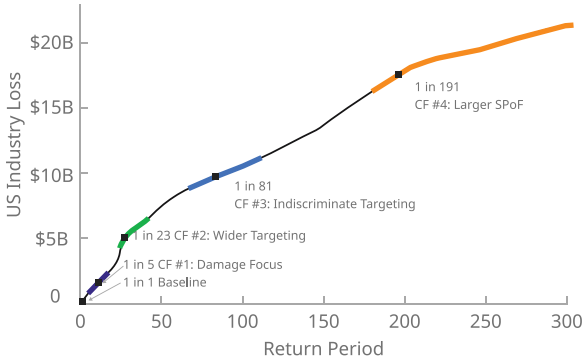

We present these findings not to provide a false sense of precision, but rather to give a tangible sense for what a roughly 1-in-200-year event could look like and the subtle (but important) ways it differs from events of recent history.

All estimates shown are for the US Standalone Cyber market.
CyberCube industry loss curve shown in black.
Confidence intervals presented in color for Baseline case and each counterfactual scenario.
Return periods shown are for median scenario loss.

Reconsidering WannaCry

Event	*WannaCry*
Occurrence	May 15th, 2017
Motivation	Financial
Event Type	Ransomware
Threat Actor	Lazarus Group aka APT 38 (North Korea)
Foot Print	230,000 systems, 150 countries, ~1,500 - 5,000 companies
Description	A widespread ransomware event perpetrated by a nation state linked group using an NSA-developed toolkit resulting in 230,000 global infections. The rapid spread was mitigated by previously released patches and a unified global response which identified the presence of a kill switch, poor coding logic, and errors in the encryption process.

Background

The *WannaCry* attack began with the creation of the Shadow Brokers toolkit – and its eventual leaking sometime in late 2016 or early 2017. Developed by the National Security Agency (NSA), these tools included a Server Message Block (SMB) protocol exploit commonly referred to as 'EternalBlue'. A group calling itself the Shadow Brokers somehow gained access and subsequently released this toolkit, which was later used in May of 2017 by a North Korea-linked group referred to as 'Lazarus' in a widespread ransomware attack. The malware often used another tool named 'DoublePulsar' to install and execute a version of itself as a backdoor. The attack spread through systems worldwide at an alarming pace, eventually infecting roughly 230,000 systems across 150 countries, including major government and corporate systems. Victim systems were encrypted and displayed a demand for bitcoin payment.

As the attack spread, a security researcher discovered that there was logic within the code of the malware to check for a non-existent domain before infecting each system. By registering the domain, the code would stop once it realized the domain was online. This "kill switch" logic meant the spread of the infection could be restrained, but it did not unencrypt systems that were already infected. Different variants of the code required differing kill switches to be enabled. New infections worldwide almost entirely ceased once a substantial portion of these switches were executed.

A coding error also prevented the threat actors from tracking which victims had paid, leading to most infected systems becoming unrecoverable. Eventually, errors in the encryption method were found, allowing researchers to develop decryption tools for many infected systems. Further post-mortem analysis indicates that the level of sophistication from the attack was rather low, with most researchers noting that simple IT security hygiene could have prevented many infections, including those on western government systems.

Considerations

WannaCry was quickly overshadowed when *NotPetya* was unleashed six weeks later, but the nature of the attack allows us to consider important counterfactuals. Whereas *NotPetya* started as a targeted attack against Ukraine and concentrated its payload against organizations operating there, *WannaCry* was indiscriminate and spread extremely quickly. Yet there were several important factors at the time of the attack that, in retrospect, contained the spread of *WannaCry* and prevented it from being as damaging as it might otherwise have become:

■ **Early warning:** Microsoft was purportedly tipped off about the Shadow Brokers tool release and the potential effect on vulnerable Windows systems, which likely provoked the release of Microsoft bulletin "S17-010 – Critical" on March 14, 2017. This led to a lower initial footprint for *WannaCry*; any company who patched or deployed the mitigations between March 14th and May 15th effectively removed themselves from the zone of effect or 'blast radius' of the event. It is widely assumed that Microsoft was informed either when the toolkit was stolen or sometime thereafter by the U.S. government or the NSA itself, leading to the security bulletin, mitigations and patch updates.

■ **Kill switch:** An important factor in curtailing the spread of *WannaCry* occurred hours after the first infection, when a security researcher discovered a kill switch in the *WannaCry* code. If this kill switch had not been discovered as quickly – or if no kill switch had existed at all – we could have seen a far more pervasive and lasting spread of the malware.

■ **Patching:** *WannaCry* propagated by exploiting the EternalBlue vulnerability. It is important to note that a patch and instructions for mitigating the vulnerabilities EternalBlue exploited had been released by Microsoft two months prior, on March 14th, 2017. While *WannaCry* was ultimately successful in infecting 230,000 systems, we can imagine that this number could have grown considerably if EternalBlue had been a true zero-day vulnerability with no available patch or advanced notice for mitigations.

■ **Sophistication:** A large hindrance to the spread of *WannaCry* was the mistakes within the code. This observation has been noted by numerous researchers: with most noting that the level of sophistication within the propagation, encryption and payment areas of the code were effective but not overly complex or adaptable. Had more sophisticated logic and tools been used, the malware could have spread at a faster and more efficient rate with payments flowing in. As it stands, this attack had one of the lowest profit margins of any major ransomware in history.

With *NotPetya* following just six weeks later, the scale at which *WannaCry* could have occurred is brought into perspective. *NotPetya* displayed a greater level of sophistication in propagation, supply chain usage, and superior coding logic for more permanent and complete damage on the affected systems. *NotPetya* was also targeted at Ukraine instead of the more widespread global impact felt by WannaCry. Had *WannaCry* used the sophistication of NotPetya and targeted western governments six weeks earlier, the damages could have grown considerably.

Counterfactuals

With these factors in mind, our team considered several counterfactual variants for *WannaCry*:[5]

■ **No Early Warning Tip-Off**

This change removes the early notification and patch release from Microsoft – represented by a higher external proliferation rate from network to network and an increase in the internal proliferation rate within a network. These changes result in higher counts of total infected systems and the number of impacted companies to roughly 3,300. The lack of an early patch and mitigation release would also cause a higher impact rate as firms have less guidance on remediation techniques to slow the spread of the attack or recover lost systems – leading to about 940 companies filing claims. Counterfactual #1 still assumes that Microsoft, likely through their own threat intelligence and research divisions, would have released a (delayed) patch before the attack started or spread at scale.

■ **No Kill Switch**

This stage considers a malware with no kill switch. The lack of a kill switch in the malware could have altered the global incident response capability. No kill switch means the infection rate would not have slowed as it had, with the only mitigations being patching or removing systems from harm's way. Combined with the change in counterfactual #1, the intercompany and intracompany proliferation rates again increase the number of impacted companies to an estimated 4,500. The resulting increase to incident response costs would increase the financial loss rate of the impacted companies, increasing the total number of companies filing claims to roughly 1,800, according to CyberCube analysis.

■ **Zero-Day (No Immediate Fix)**

This counterfactual addresses the question: ***what if this event had caught Microsoft completely off guard without any warning to the industry?*** Reimagining this event with a true zero-day exploit would mean there was no immediate help for victims as the event spread. Building on the previous counterfactuals, this would have resulted in higher numbers of impacted companies as the controls provided by Microsoft in the form of patches or mitigations would not have been available until days or weeks after the first known infection. Individual firms would also have had less capability to respond in the early hours and days of the attack, leading to nearly 7,500 companies estimated to be impacted. A higher incident response cost and longer downtime could have meant higher financial loss rates and an estimated 3,700 companies filing claims.

■ **More Sophisticated Attack**

Considering the ***WannaCry*** malware with more advanced propagation and encryption capabilities leads to an increase in external and internal proliferation rates as well as a decrease in the level of control that victims can maintain. Based on research into the ***NotPetya*** kill chain as well as additional propagation methods available at the time of the event, we can plausibly imagine a more sophisticated version of the ***WannaCry*** attack, leading to a larger number of vulnerable and initially infected systems. Based on our estimates this could have led to a pool of impacted companies of around 33,800 and a larger number of claims filed at around 16,500.

[5] These changes "compound" in this study, meaning that as each new change is described it is assumed the prior changes have also occurred.

Quantification

As done for SolarWinds, we then quantified the potential losses from each of these *WannaCry* variants. The results appear in **Exhibit 3**.

Exhibit 3: WannaCry Counterfactual Estimates

	Factor(S)	Baseline	Counterfactual #1 "No Tip Off"	Counterfactual #2 "No Kill Switch"	Counterfactual #3 "Zero-Day(no fix)"	Counterfactual #4 "Sophisticated Attack"
Footprint	Customers/Companies	251M	251M	251M	251M	251M
	Change	-	Higher Proliferation	Higher Proliferation	Higher Proliferation	Higher Proliferation
	Infected Companies	5.2K	7K	9.3K	12.5K	50K
	Change	-	Higher Impact Rate	Higher Financial Loss	Higher Financial Loss & Impact Rates	Higher Impact Rate
	#Impacted Companies	1.5K	3.3K	4.5K	7.5K	33.8K
Severity	#Companies Filing	267	941	1.8K	3.7K	16.5K
	US Insured Loss					
	5th	$469M	$2.1B	$4.5B	$9.2B	$28.6B
	50th	$803M	$2.9B	$5.9B	$11.5B	$34.8B
	95th	$1.3M	$3.8B	$7.4B	$14.0B	$41.3B
	50th Percentile Loss as					
	Loss ratio %	11%	38%	77%	151%	458%
	Limits %	0.1%	0.4%	0.9%	1.7%	5.2%

All estimates shown are for the US Standalone Cyber market.
Numbers in blue indicate the result of an assumption change from the previous counterfactual scenario (or baseline) to the current one.

With each successive counterfactual change, we can see the increasing financial consequences that could have unfolded. For example, the version contemplated in counterfactual #2 (No Kill Switch) could mean a U.S. industry loss in the range of $4.5 billion to $7.4 billion. This would exceed any loss seen by the cyber market to date – yet it is not difficult to imagine if the historical WannaCry had played out slightly differently.

As we did with *SolarWinds*, we can overlay these counterfactual scenarios against an industry loss curve to provide context to modeled losses at various points in the tail. **Exhibit 4** shows the range of loss from each *WannaCry* counterfactual plotted against the industry loss curve (OEP) from Portfolio Manager v5.

Exhibit 4: WannaCry Counterfactuals vs. US Standalone Cyber Industry Loss Curve

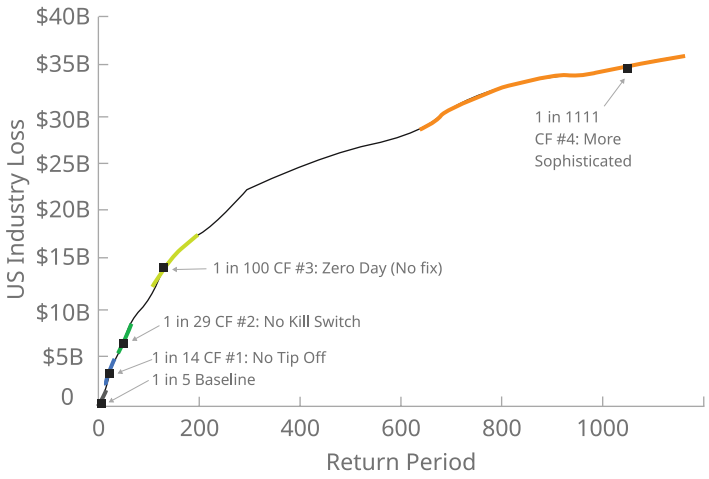

Here we see the counterfactual scenarios showing us three degrees of separation from historical *WannaCry* to the modeled tail, and four degrees to the extreme tail. Counterfactual #3's median loss of $11.5 billion would be a 1-in-100 industry occurrence in Portfolio Manager v5, and counterfactual #4's median loss of $34.8 billion would approximately correspond with a 1-in-1100. It's worth reiterating that counterfactual #4 departs from the actual events of *WannaCry* to a greater degree than counterfactual #3 or the previous steps. The differences between these

scenarios help to improve our understanding of what could distinguish a 1-in-100 event from 1-in-1000+ event. A lot goes badly at the 100-year level, but not everything. And it should surprise no one that one hypothetical 'worst case' cyber event would be a sophisticated attack deploying a destructive payload using a zero-day exploit in a widely used software or operating system.

Assumptions and limitations

This counterfactual paper summarizes what could have happened in the aftermath of the ***SolarWinds*** and ***WannaCry*** attacks, but its application requires an understanding of its limitations. This study omits several factors that are worth mentioning.

First, we have not quantified the applicability of war exclusions to these counterfactuals. Given that the historical ***SolarWinds*** and ***WannaCry*** attacks were both carried out by nation states, it is almost certain that if large losses accumulated from a similar event, (re)insurers would seek to apply policy language to sublimit or exclude claims. This topic is of critical importance to cyber insurers – it has received considerable attention and should continue to do so. However, quantifying the effect of such exclusions and the likelihood they would stand up to legal challenge is beyond the scope of this paper. That said, since counterfactuals allow us to examine the gray area between reality and fiction, we believe it's important to note that the line between nation states and other threat actor classes can be blurry[6]. Regardless of the insurance industry response, the potential economic damages to organizations from these scenarios would be significant.

Second, we have not factored for unusual remediation actions from the private or public sector. The 2024 ***Change Healthcare*** event provides an important example: United Healthcare's decision to introduce a temporary funding assistance program for its customers may have averted what otherwise could have been a sizable volume of claims for contingent business interruption. Such actions by an affected SPoF organization are as critical to the final event outcome as they are difficult to model for. While we might be tempted to assume that affected SPoF organizations will always provide some kind of relief, this also would be inappropriate – some SPoFs themselves do not outlive the contagion event that their technology made possible[7]. Similarly, we do not assume any financial relief provided by governments if such a cyber disaster were to occur.

It's worth noting that as more cyber events unfold, the entire cybersecurity industry continues to learn and adapt. For example, distributed denial-of-service attacks have been less impactful in recent years even as they rise in frequency and intensity. This has largely been due to cloud and critical infrastructure providers learning to defend and adapt with different load balancing and filtering techniques. As ransomware-as-a-service arose, incident responders and cybersecurity vendors adapted to the challenge of containing infections and recovering encrypted systems through advanced technology solutions (XDR, MDR, etc.) and advances in backup procedures.

[6] See "The Spectrum of State Responsibility", Jason Healey, 2011. Out of 10 categories of attack ranging from "state-prohibited" and "state-executed", five of them involve varying degrees of state and private actor engagement. Source: Atlantic Council, page 2.

[7] CloudNordic attack: a cloud service provider purportedly faced bankruptcy after a ransomware attack that disrupted service for its customers. Source: Cybernews.com.

Additionally, with the rise in publicity of cyber events in recent years, leading security firms are ready to act as first responders. These entities are often on retainer and given increasing latitude to mobilize quickly and with greater force to prevent events from reaching a larger scale. Our Portfolio Manager modeling does anticipate that better-resourced SPoF organizations will mobilize a faster and more effective response to an unfolding event – but we do not have confidence in modeling the complex dynamics of how the wider cybersecurity community participates in the response. With these limitations in mind, we can better understand the counterfactual analyses and what potential similar attacks mean for the industry.

Conclusions

This paper is intended to help the (re)insurance industry expand its understanding of cyber catastrophe risk. Many have expressed concerns that past events could have been worse, and counterfactual analysis provides a structure for us to investigate and confirm such concerns. Given the relatively short history of digital networks, we believe it's important to be able to glean as much insight as possible from the events that transpire. We hope that more (re)insurers will recognize the value of this approach.

When assessing the ***WannaCry*** event in particular, should one infer from this paper that we "nearly missed" an event estimated beyond the "1-in-1000" level? Absolutely not. "Degrees of separation" means exactly that – points of clear separation between these hypothetical scenarios and the events we have actually seen happen. But through this exercise, we have sought to illustrate that the potential for cyber tail risk is very real, and less difficult to imagine than one might think.

To this end, counterfactuals can help us better understand the variation across different parts of the modeled loss curve, allowing us to think more concretely about tail risk. Cyber tail events may look subtly different – not categorically different – from the events we all recognize. As Mark Twain reportedly said, "history doesn't repeat itself, but it often rhymes." While we cannot predict the exact impact of future events, comprehensive modeling tools as well as analyses on previous events, can help us better prepare and move forward together in our understanding of cyber risk.

Authors:
Jon Laux, VP of Analytics, j-on@cybcube.com

Josh Knapp, Principal Cyber Risk Modeler, joshk@cybcube.com

Contributors:
Doug Fullam, Principal Actuary, dougf@cybcube.com

Ethan Spangler, Lead Economist, ethans@cybcube.com

Ty Zeno, Senior Threat Intelligence Analyst, tyz@cybcube.com

Content Editors:
Rianna Mistry, Content and SEO Specialist, riannam@cybcube.com

Yvette Essen, Head of Content, Communications & Creative, yvettee@cybcube.com

Designer:
Fasen Zhao, Graphic Design Intern, fasenz@cybcube.com

About CyberCube

CyberCube is the leading provider of software-as-a-service cyber risk analytics to quantify cyber risk in financial terms. Driven by data and informed by insight, we have harnessed the power of artificial intelligence to supplement our multi-disciplinary team. Our clients rely on our solutions to make informed decisions about managing and transferring cyber risks. We unpack complex cyber threats into clear, actionable strategies, translating cyber risk into financial impact on businesses, markets, and society as a whole.

The CyberCube platform was established in 2015 within Symantec and now operates as a standalone company. Our models are built on an unparalleled ecosystem of data and validated by extensive model calibration, internally and externally. CyberCube is the leader in cyber risk quantification for the insurance industry, serving over 100 insurance institutions globally. The company's investors include Forgepoint Capital, HSCM Bermuda and Morgan Stanley Tactical Value.

For more information, please visit www.cybcube.com or email info@cybcube.com.

CyberCube

INFO@CYBCUBE.COM

CYBCUBE.COM