



SANS Institute Information Security Reading Room

A Question of Platinum Plus

Craig Wright

Copyright SANS Institute 2020. Author Retains Full Rights.

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

Of Black Swans, Platypii¹ and Bunyips. The outlier and normal incident in risk management.

GIAC (GSE) Gold Certification

Author: Craig S Wright, Craig.Wright@Information-Defense.com
Advisor: Stephen Northcutt

Accepted: December 21st 2010

Abstract

To act rationally requires that we forecast the future with inadequate information using the past as a guide for all its flaws. We make decisions in the absence of knowledge. We state that black swans and bunyips do not exist. From time to time, we find that we have decided in error and black swans are found. However, for every black swan, there is a unicorn, dragon and Bunyip that does not exist and of which we remain confident will never be found. Zero-day security vulnerabilities remain the fear of many security professionals. We present empirical evidence as to the rarity of these events as a source of system compromise. Instead, we demonstrate how common misconfigurations and old attacks are far more of a concern to the security professional. We show that predicting zero-day attacks is possible and that defending systems against common vulnerabilities significantly lowers the risk from the unexpected and “unpredictable”. The inherent psychological biases that have developed in the information security profession have centered on the outlier effect. This has led to a dangerously skewed perspective of reality and an increase in the economic costs of security. This paper demonstrates that producing resilient systems for known events also minimizes the risk from black swans without the wasted effort of chasing myths.

Keywords: Information Security, Economics, Risk, Black swans.

¹ Platypii is the plural for Platypus, an Australian marsupial once believed to be a hoax.

1. Introduction

The fallacy of the black swan in risk has come full circle in information systems. Just as the deductive fallacy, “*a dicto secundum quid ad dictum simpliciter*”² allowed false assertions that black swans could not exist when they do, we see assertions that risk cannot be modeled without knowing all of the ‘black swans’ that can exist. The falsity of the black swan argument derives from a deductive statement that “every swan I have seen is white, so it must be true that all swans are white”. The problem is that which one has seen is a subset of the entire set. One cannot have seen all swans.

Likewise, the argument that not enough weight applies to zero-day vulnerabilities and that these are a major cause of system intrusions relies on the same reasoning. The assertion that more compromises occur because of zero-day vulnerabilities comes from a predisposition to remember the occurrence of a zero-day attack more often than one remembers a more frequently occurring incident. Whereas near eidetic recall comes from events that are unusual, common events are more often forgotten (Roese & Olson, 2007). This leads to treating common events as if they matter less than they should.

Taleb (2007) formulated the Black Swan Theory with the assertion that unpredictable events are much more common than people think and are the norm and not the exception. In this, he has fallen into the logical fallacy and trap against which he rails. This fallacy of arguing from a particular case to a general rule, without recognizing qualifying factors lead people, before the exploration of Australia, to state that black swans could not exist instead of stating that it is unlikely that they exist. When Australia was finally explored, Platypii and Bunyips were reported along with black swans. At first, many refused to believe that a creature such as the platypus could be possible. The scientific discovery and examination of these creatures was unlikely, but far from impossible as their existence demonstrated. The discoveries of such an unexpected creature lead others to believe that Bunyips could also exist. They tried to assert that the discovery of other unlikely creatures made the discovery of the Bunyips more likely.

² The logical maxim, concerning where an acceptable exception is eliminated or simplified, a type of logical fallacy.

Though it is still possible that this is or at least was the case, the existence of Bunyips remains incredibly unlikely. In fact, it is so unlikely that we could state that Bunyips do not exist with a reasonable level of certainty. Many people have spent large amounts of money searching for mythical creatures. At times in the past, some have been discovered true. The fact remains, more monsters exist in our minds than could ever exist in the world.

For many years, information security and risk management has been an art rather than a science. This has been detrimental to the economy as a whole as well as to the operations of many organizations. The result has been a reliance on experts whose methodologies and results can vary widely and which have led to the growth of fear, uncertainty and doubt within the community. Although many true experts do exist, some of whom exhibit an insightful vision and ability, for each true expert, many inexperienced technicians and auditors abound.

This failure to be able to expend resources in securing systems has created a misalignment of controls and a waste of scarce resources with alternative uses. This paper aims to demonstrate that the common risk is the one against which to protect. Zero-day vulnerabilities and strange events are memorable, but this does not make them the target of an effective risk mitigation program. It also does not mean that they are the most likely event that will occur. Unusual and excepted events upset a number of models and methods that are common in many other areas of systems engineering, but which are only just starting to be used in the determination of information systems risk. This issue is not the one that should be considered. These processes can help both the inexperienced security professional as well as adding to the arsenal of tools available to the consummate expert. In place of searching for bunyips, we should implement systems that cover the majority of failures and prepare to act when unusual events emerge.

The standard systems reliability engineering processes³ are applicable to information systems risk. These formula and methods have been widely used in systems engineering, medicine and numerous other scientific fields for many years. The

³ For more detail on these processes see the Springer Series in Reliability Engineering, <http://www.springer.com/series/6917>

introduction of these methods into common use within risk and systems audit will allow the creation of more scientific processes that are repeatable and do not rely on the same individual for the delivery of the same results. Some failures will occur. A 99% confidence interval, though considered a good measure, brings a level of uncertainty, by definition. The issues are that it is unwise to discard the normal occurrences in favor of a black swan that may turn out to be something else again. By assuming that all black swans lead to catastrophic and unpredictable failure, we are again destroying the exception.

2. An investigation into the causes of system compromise

In order to test the causes of system compromises, we configured 640 Windows XP Professional systems that were on virtual hosts. The placement of each of the hosts was on an IP address external to a network firewall. Three separate tests formed the foundation of the experiment. For this, we set the baseline security of the system as a CIS (Centre for Internet Security) score. The CIS Windows XP Professional Benchmark v.2.0.1 (Shawgo, Whitney & Faber) formed the security test metric.

These are:

1. A base install of a Windows XP SP 2 system.
2. An increasing CIS score was configured on the hosts.
3. A snort IDS was used to separate worms and other automated malware from interactive attacks.

Network traffic monitors were used in order to determine if a system had been compromised. The hosts had no third party applications and initially had auto-updating disabled. A host, when compromised was reset and reconfigured for the next set of survival tests. The reset utilized the VMware 'snapshot' feature to take the system to a known good state (see Honeynet Project & Research Alliance).

In this paper, we have defined a zero-day attack as one that is generally unknown. Some authors (The SANS Institute, 2005) define a zero-day attack as a "*flaw in software*

Craig Wright, Author Name, email@address

code is discovered and code exploiting the flaw appears before a fix or patch is available.” For the purpose of this paper, we have defined a zero-day attack, as one that uses computer vulnerabilities that do not currently have a solution. This includes patching from the vendor, third party patches or workarounds. In this way, a vulnerability with a CVE number and third party protection (such as IPS filters or anti-malware updates that stop the attack) is not defined as a zero-day attack for the purpose of this paper.

This aligns with the definition of a “*zero-day exploit*” occurring “*when the exploit for the vulnerability is created before, or on the same day as the vulnerability is learned about by the vendor*” (Bradley). This is a superior definition of the term and should be used in place of the former. Many vulnerabilities remain unpatched for many months with user and vendor knowledge. These are commonly stopped using alternative approaches and work-arounds in place of vendor patches.

The reason for this lies in the ability to predict an attack. This paper seeks to measure the impact of controls that can be predicted and to compare these to attacks that have no known solution. A published attack with no official vendor patch may be mitigated and predicted. This type of an attack is not a ‘black swan’. The unpredictable requires an attack that unknown or unpublished. A select few experts could know of this type of vulnerability. This does not allow the public to have knowledge of this issue. As such, this limited knowledge would not lead to a generally deployed work around.

2.1. Modeling the impact of a single control

The first test process separated the hosts into two classes. Those with the Windows XP Firewall enabled, and those with the firewall disabled. No third party products (including anti-malware software) were used on either class of system. With the release of Windows Vista and Windows 7, the analysis of the impact of the inclusion of a firewall in Windows XP may seem a little dated. However, the same use and deployment of this control applies to both Windows 7 and Windows Vista. In addition, many organizations still use Windows XP.

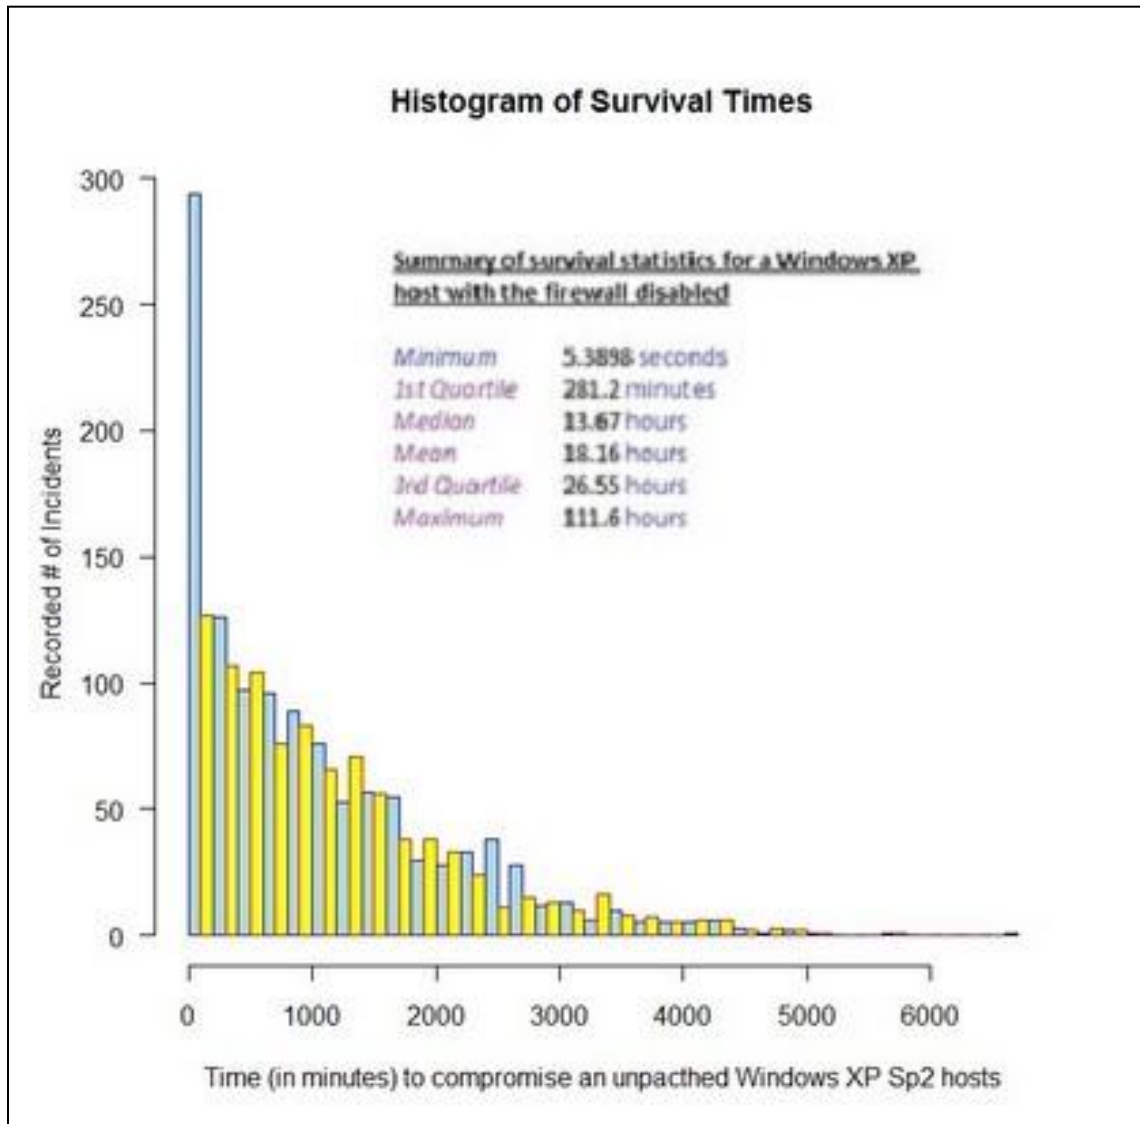


Figure 1. Survival times with the Windows Firewall Disabled.

The histogram in Figure 1 displays the distribution of survival times for the un-firewalled Windows XP hosts. The Conflicker worm that managed to compromise the un-firewalled hosts in quick succession skewed this result. The quickest time being 5.4 seconds from the network cable being connected to a scan occurred (this was in May 2009). This was an exception and hence an outlier. The mean time to compromise of the hosts was just over 18 hours, with only 25% of the sample compromised in less than 3 hours.

When the results of the firewalled and un-firewalled hosts are compared, we can confidently assert that the Windows host firewall is a control that has a statistically

Craig Wright, Author Name, email@address

significant effect when used. We say 'if', as this a control that is commonly overlooked or disabled. The results from enabling the Windows firewall are displayed in Figure 2 and a side-by-side box plot is displayed in Figure 3. With the firewall enabled, the mean survival time of the Windows XP SP2 systems increased to 336 days. No system with this control enabled was compromised in less than 108 days. With the maximum survival time for an unpatched and un-firewalled Windows XP system predominantly measured at less than 5 days and the minimum compromise time at 108 days with the enabling of the firewall and no additional patching, it is hard not to conclude that the Windows Firewall makes a statistically significant difference to the security of the system.

We used the Snort IDS for this exercise. This provided the details of the attacks allowing an analysis of worm (automated) compromises against manual (attacker, script kiddies etc). The IDS sat between the Internet connected router and the virtualized Windows XP hosts. Any outgoing traffic was investigated.

TABLE I. SUSVIVAL TIMES

	<i>Statistical Analysis of Survival times</i>	
	<i>Windows Firewall Enabled</i>	<i>Windows Firewall Disabled</i>
Mean	18.157 Hours	8,064.640 Hours
	t = -170.75 df = 2272 p-value = 2.2 Exp -16	

In the results of the 640 hosts that were used for this experiment, no system was compromised with a zero-day attack. Many new and novel attacks against known vulnerabilities did occur, but not a single compromise was due to an unreported vulnerability. Further, no attack without a patch was used to compromise any of the systems. This means that if the systems had been patched, none of the attacks would have succeeded.

In a simple test of a single control, the enabling of this control had a marked effect on the survivability of the system. We see that leaving the host running with the firewall enabled provided a good level of protection (without a user on the system). This does not reflect a true Windows XP system as any third party applications and user action have been introduced to confound the results. All connections to these hosts are from

external sources (such as a server model) to the host and no users are browsing malware-infected sites. In general, a Windows XP system will have a user and will act as a client. This introduces aspects of browsing and retrieving external files (e.g. email). These aspects of the host's security will change the survival rates of a system, but we can see that there is a significant advantage from even a simple control.

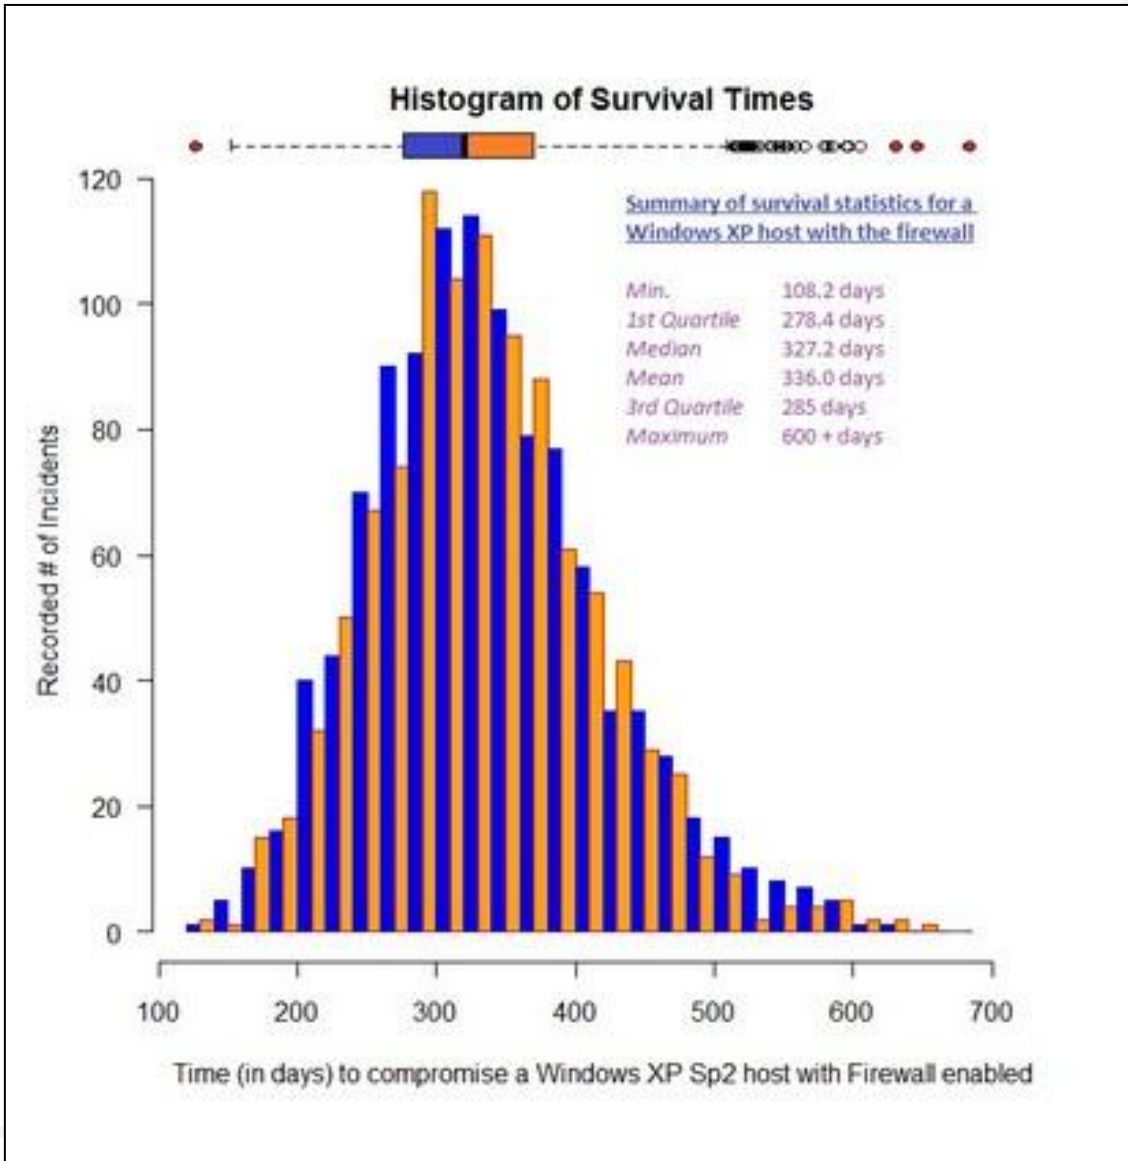


Figure 2. Survival time for Windows XP classified by interactive attacks and Automated malware (Worms).

The regret is that in a sample of 136 home systems from corporate computers that have been tested and a sample of 231 systems inside various corporate networks, few systems ran a firewall. Of the hosts tested, 31.28% (or 23 systems) had the Windows XP Firewall or a commercial equivalent installed and running. Of the internal systems tested

Craig Wright, Author Name, email@address

in this study, 6.1% had an internally (inside the corporate firewall) enabled firewall (14 hosts). The ability to enable IPsec and Group Policy within a corporate environment is a control that is generally overlooked or bypassed. The results of enabling (or rather not disabling) the Windows Firewall produce a pronounced benefit to the survivability of systems.

In this first experiment, we see marked benefits from a simple control without the worry of any black swan effect.

2.2. Modeling system survival by attack class

The results of a series of hazard modeling experiments on Windows XP that we presented in the last section were limited to a base Windows XP install. This was designed to test the effect of enabling or disabling the firewall. We next altered the experiment to investigate the attacks in classes. In particular, comparing the systems that have been compromised by an automated process (Worms etc) against those which have at least some level of interaction.

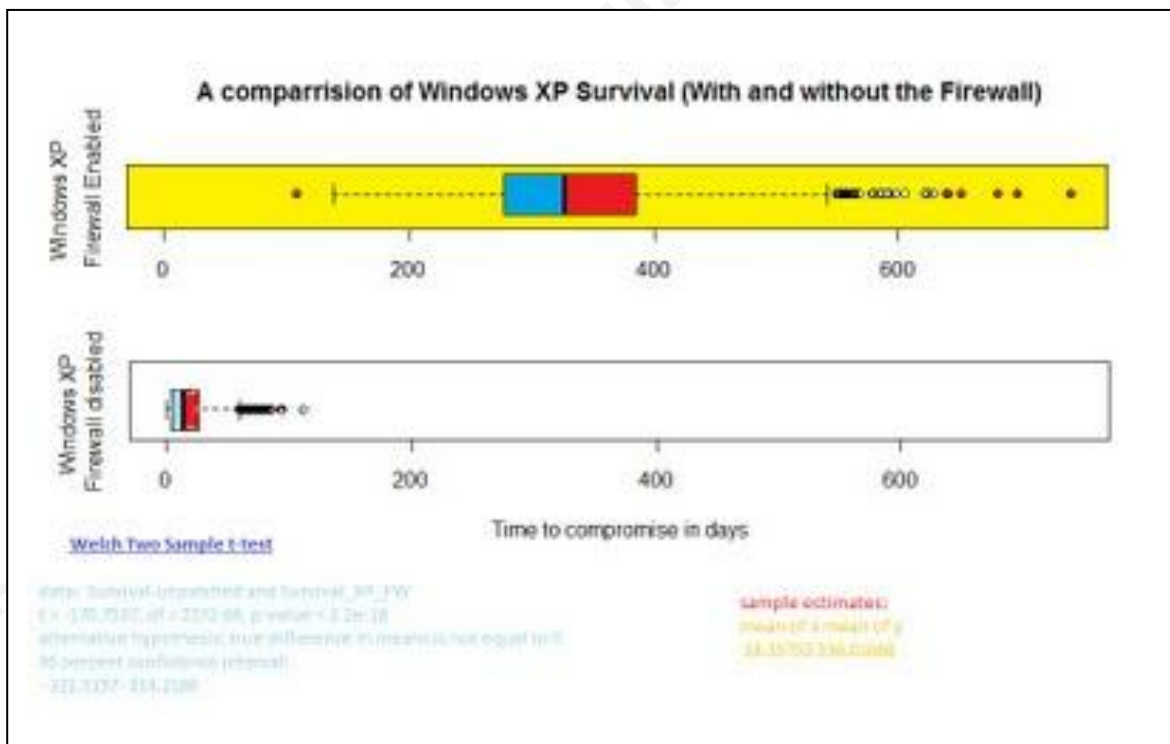


Figure 3. Comparing the use of the Firewall to an unprotected XP system.

The introduction with Windows XP SP2 of a firewall that is enabled by default is demonstrated to have had a significant impact on the overall security of networked systems.

Each of the systems was reset and configured with a varying level of controls. The CIS metrics were calculated using the automated tool. Systems were distributed evenly between the metrics in 5% intervals (that is, 32 systems were allocated to each 5% bracket). The systems have been made either more secure or less secure by enabling and disabling controls until a complete spread of scores was created.

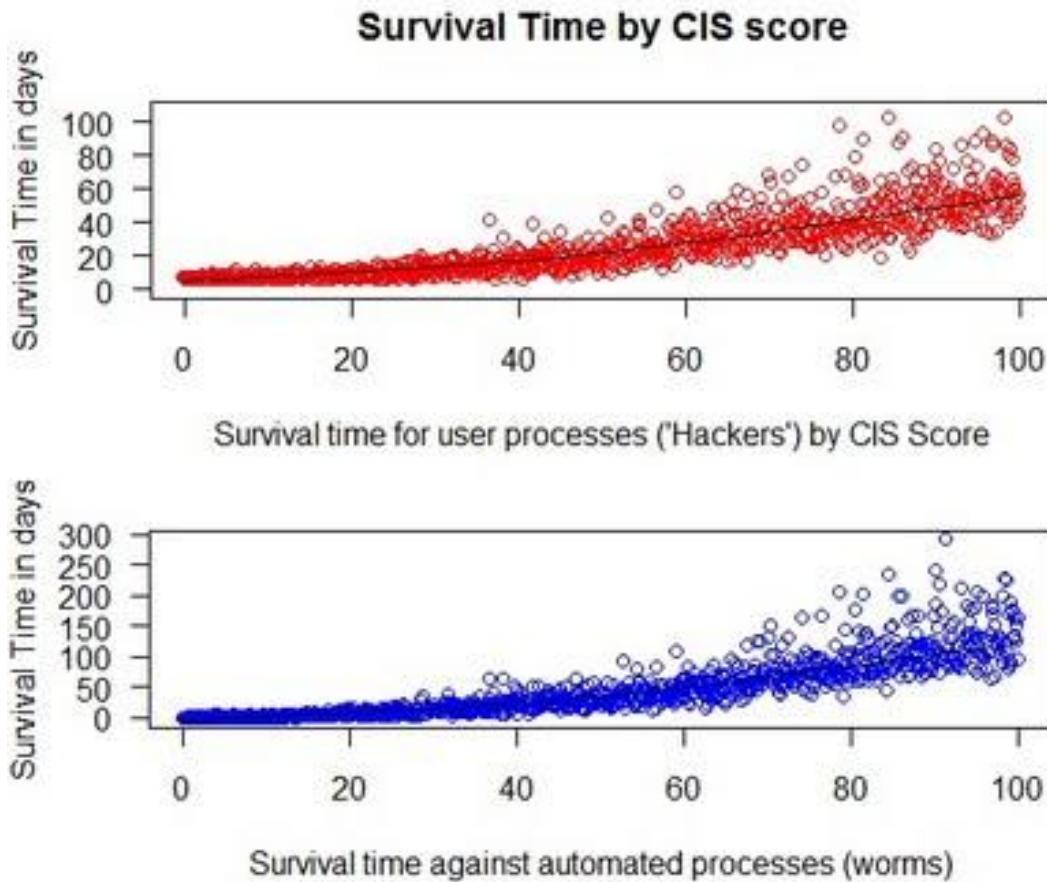


Figure 4. Survival time for Windows XP classified by interactive attacks and Automated malware (Worms).

Figure 4 and Figure 5 display a significant difference in the patterns of compromise due to automated and interactive attacks. We can see from the plots that worms act faster against vulnerable systems and that interactive users (attackers) are more capable at compromising more secure systems. This is more easily seen on an

Craig Wright, Author Name, email@address

overlay plot (Figure 5). This displays a plot of the survival time against automated processes (green) overlaid with that of manual processes (red). The Loess fit for each is also incorporated into the plot.

What we see from other results is that the more secure a system is (in this case patched of known vulnerabilities), the more likely that a compromise is manually initiated. Likewise, less secure (or patched and vulnerable) systems are exposed to more automated attacks (e.g. Worms).

2.3. System Modeling by CIS Metric

A selection of 48 Windows XP SP2 computers was used for a test that incorporated both 16 physical hosts and 32 virtual machines. This was conducted in order to examine the differences (if any) that may result with a virtualized host in place of a physical host. The tests were run over a 600 plus day period starting from November 2007. When a physical host was compromised, it was taken offline for 10 days. In this period, the host was rebuilt in a slightly different configuration. The 32 virtual hosts were built with differing levels of patching. These hosts have been reverted to a VM snapshot following a compromise. At this point, they would be re-patched and reassessed.

The same Snort IDS system used in the previous experiment was deployed to measure the attacks against the physical hosts. The 32 virtual hosts were configured on a single high-end Red Hat server running Snort. No filtering was conducted, but all attacks were logged. The survival time for the host is set as the time from when the host was placed as live on the network until a local compromise occurred. The 16 physical hosts were connected to a Cisco switch sitting behind a Redhat Linux host running Snort and acting as a forwarding router.

Each host in both the physical and virtual configuration was configured on a '/29' network. This was assigned an internal IP in the 10.x.y.z address ranges with the Redhat host being assigned the lower IP address and the upper to the host being tested. Static NAT was used to pass a real IP address to the host in the range of 203.X.Y.194 to 203.X.Y.242 with a netmask of 255.255.255.192. The full address is not reported at this time in order to minimize any impact on ongoing experiments.

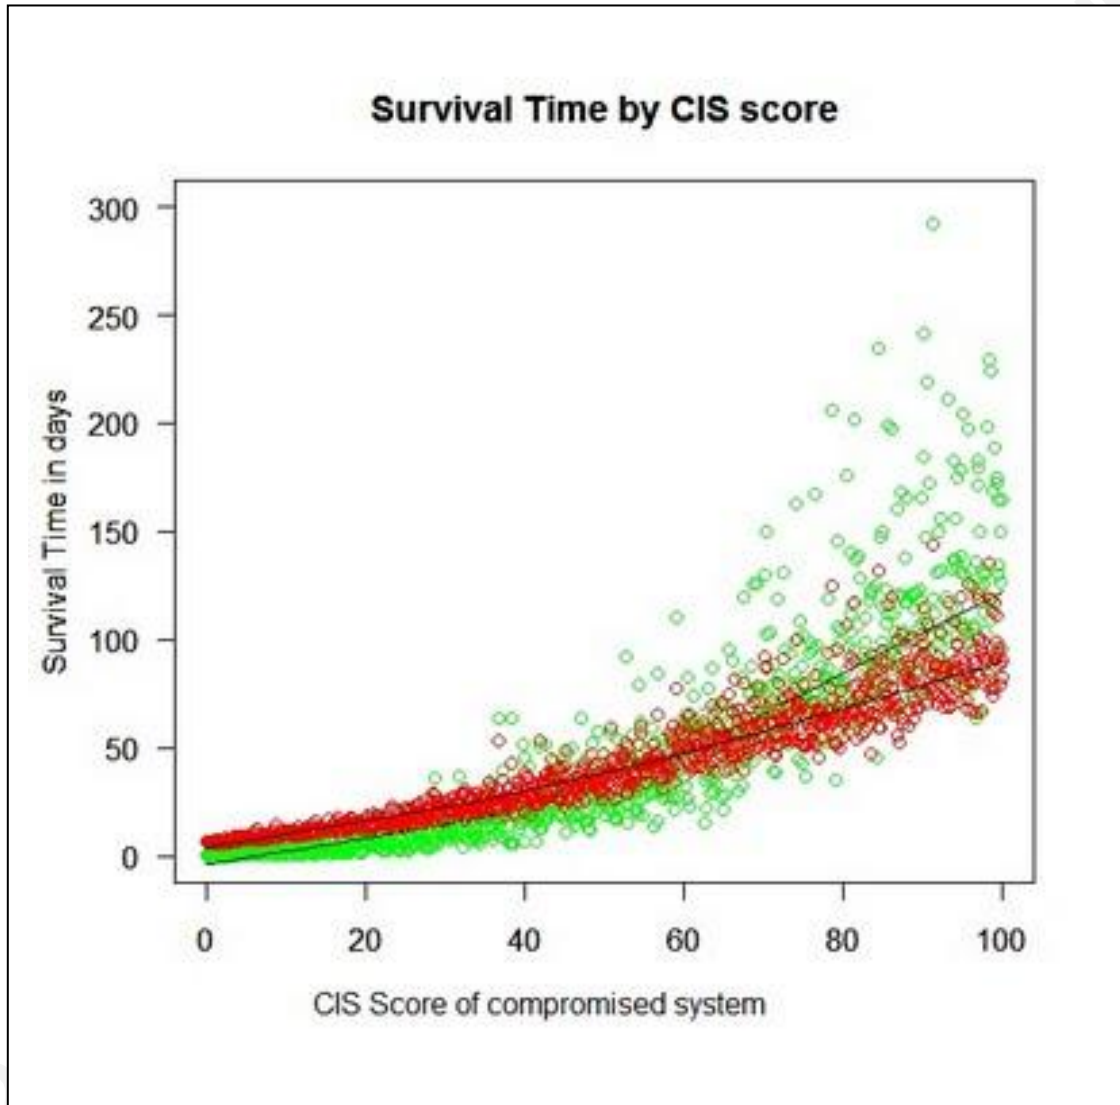


Figure 5. Automated vs Interactive attacks and survival times.

The iptables configuration on the two Redhat systems was configured to allow any IPv4 traffic from the Internet and to block any IPv6 traffic. The Redhat host did not have a publically routable IP address. Internet hosts were allowed to connect to any system on any port. The only restriction was designed to block traffic to and from the

Craig Wright, Author Name, email@address

Windows XP hosts to any other host on the same network. This allowed the host to be compromised from the Internet but a compromised host could not see another host on the same network. The Windows XP firewall was disabled for all CIS scores less than 90 and for some hosts with scores greater than 90 (although it is difficult to create a host with a score greater than 90 and the firewall disabled).

This was done to create a level of independence with attackers having to compromise systems from the same way and not being able to "hop" across systems (as occurs in real compromises). The goal of this experiment was to record initial compromises and not the subsequent process (that being the goal of a separate and ongoing experiment). The times and measures have all been recorded and analyzed. As before, no web browsing or other internal activity was conducted from the systems under test.

The scatterplot (Figure 6) is the plot of the measures score using the CIS scoring system against the time that it took to compromise the host. We see that there was a significant benefit in achieving a score of 80+. Any score of less than 40 was compromised relatively quickly. A score of 46 was compromised within 24 hours. All scores of 60+ remained uncompromised for at least a week. One host with a score of 59 on the CIS scale remained uncompromised for 98 days.

Similar results have been recorded for the hosts in the VM group (blue) and the physical group (red) in the scatter plot (Figure 6). A Loess best fit has been applied to this scatter plot marking the expected survival time by CIS scoring. As the score increases, the variance also increases, but this can be seen as a function of increasing survival times. No statistically significant differences in survival times have been noted because of the host being virtualized or physical.

From these results, we can assert that automated systems are more likely to compromise poorly configured systems than well-configured ones. This result is no more than common knowledge; however, we also see that an interactive attacker is more likely to succeed in compromising a well-configured system when compared to an automated process. We also see that even the best-configured system fails in time.

Craig Wright, Author Name, email@address

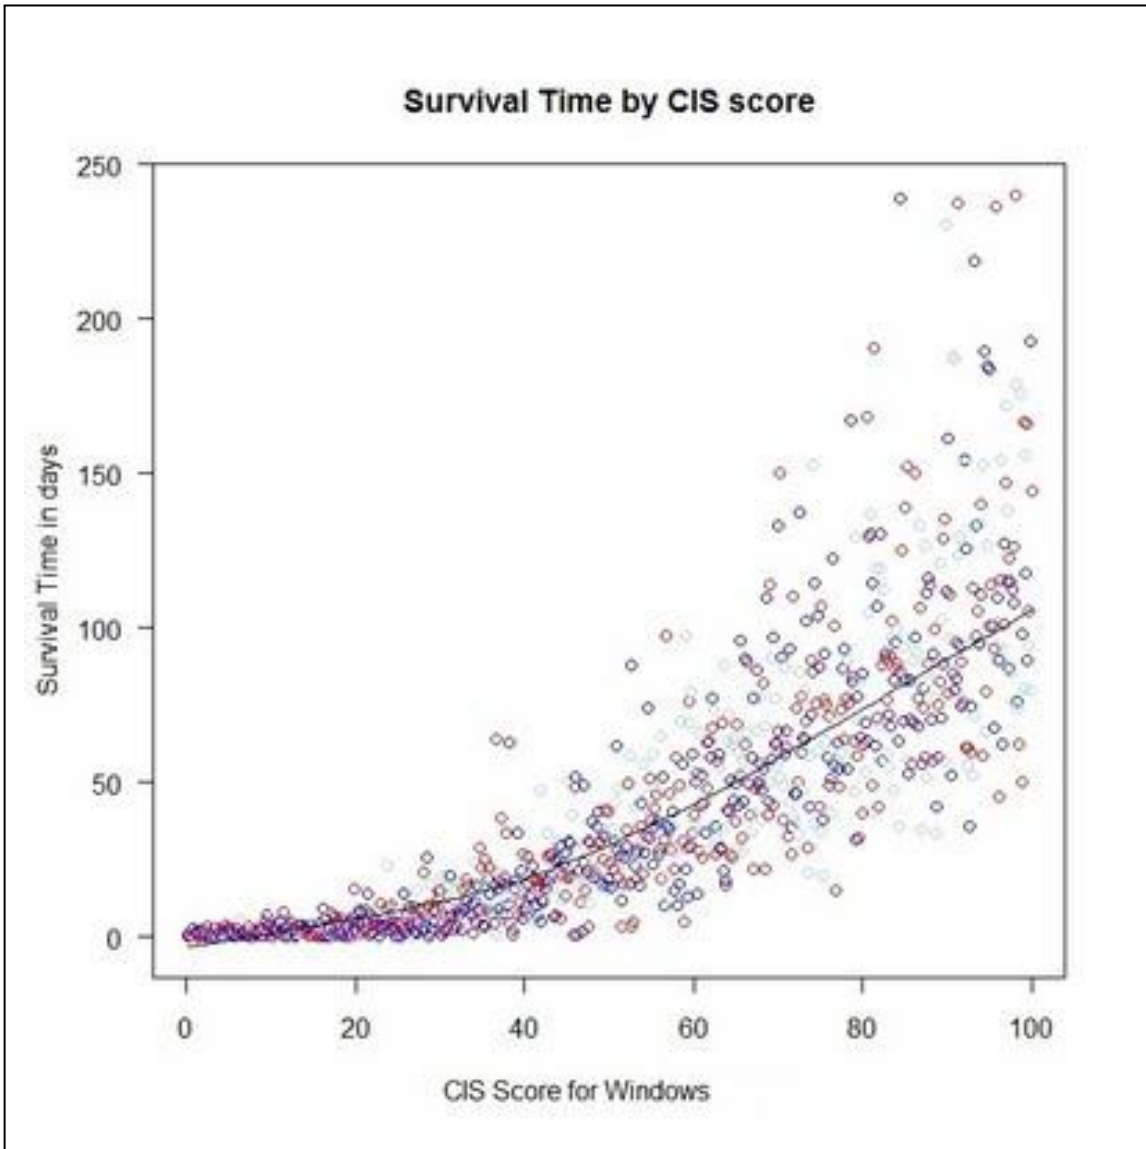


Figure 6. Survival for Physical vs Virtual hosts.

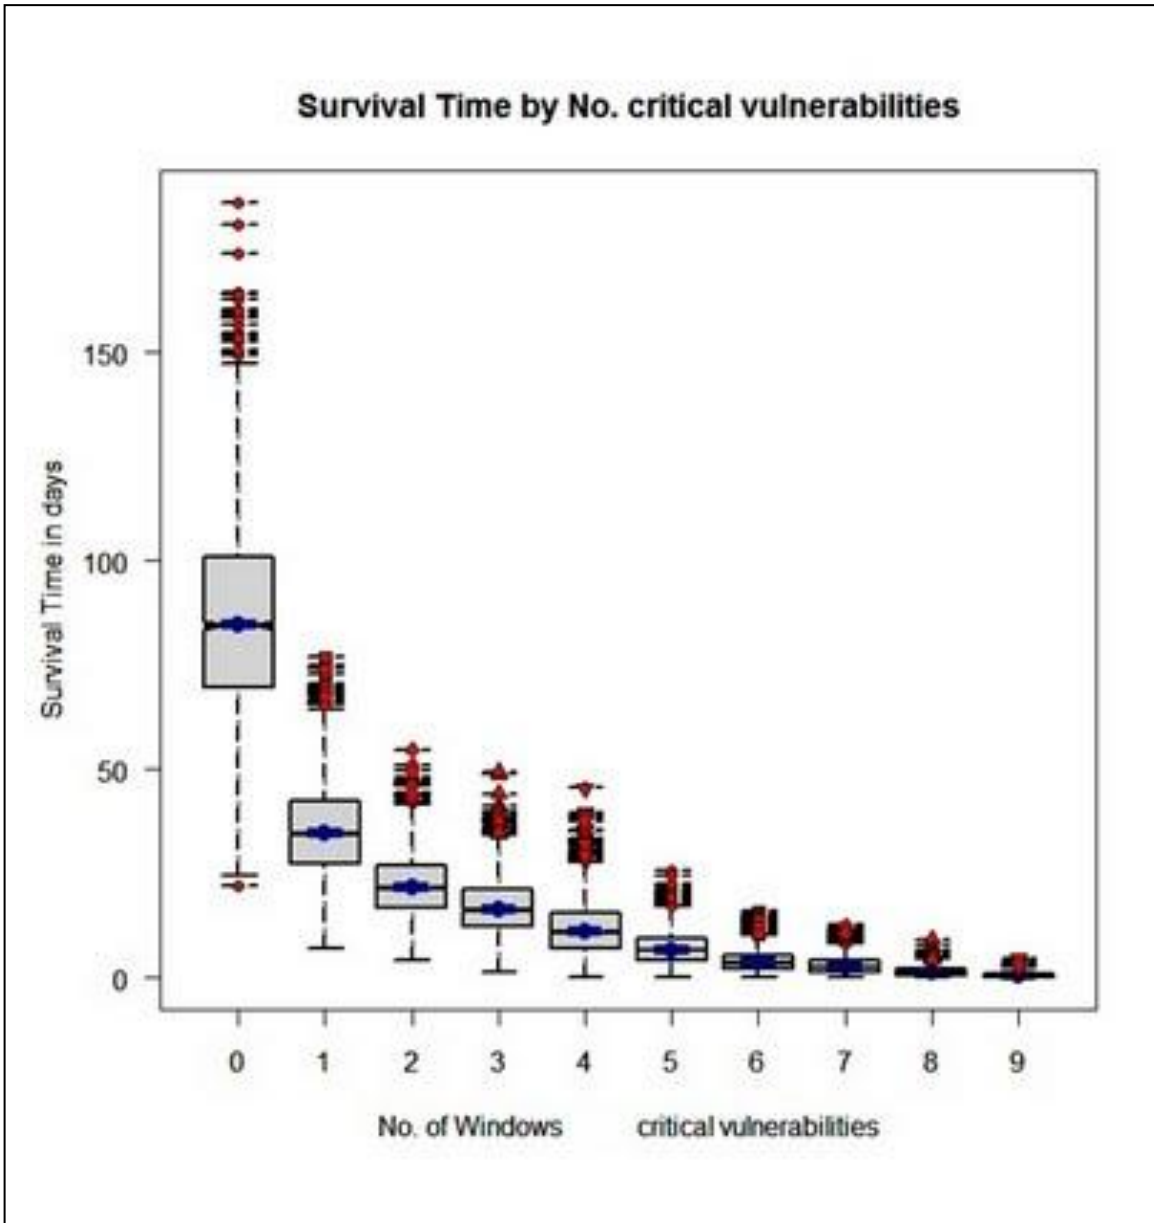


Figure 7. Mapping survival by critical vulnerabilities.

Again, no system failed from an unknown attack. Of note was that several systems were compromised using new but known attacks. In the majority of attacks against a system with a CIS score of greater than 60 and with the Windows firewall enabled, the system was compromised between patch cycles. This involved the attack occurring against a new vulnerability before the scheduled patch release was due. We further note, that in all instances, these attacks involved systems that are not interactively managed. Work-around existed for all of the incidents that lead to compromise of the

more secure systems. Further, more sophisticated anti-malware, firewall or other system security software would have stopped these attacks. This is why we have not classified these attacks as zero-days. The vendor did not have a public patch, but a work around or third party control existed in all instances.

The issue comes to economic allocation of scarce resources. Numerous solutions could have stopped all of the attacks against the secured hosts. Some of these solutions would have cost less than implementing the controls that gave the Windows system a greater CIS score.

2.4. Mapping survival time against vulnerabilities

The next part of the experiment involved the configuration of 16 Windows XP SP2 hosts with a set and measured number of critical vulnerabilities. We left these hosts unpatched (ranging from 1 to 10 unpatched vulnerabilities per host) for a selected set of vulnerabilities. The experiment involved applying all other patches for newer vulnerabilities as they became available. The particular vulnerability was randomly selected on each host. Each vulnerability was selected from the SANS Top 20 vulnerability list (SANS, 2007).

All of the hosts used virtualization with 'snapshots' enabled. A host that was compromised was reassigned with a new IP address and was reactivated 14 days later. Reactivation involved restoring the host to the snapshot and patching it. The host was left with the same number of critical vulnerabilities, but a different set of vulnerabilities was selected randomly from the SANS top 10 list.

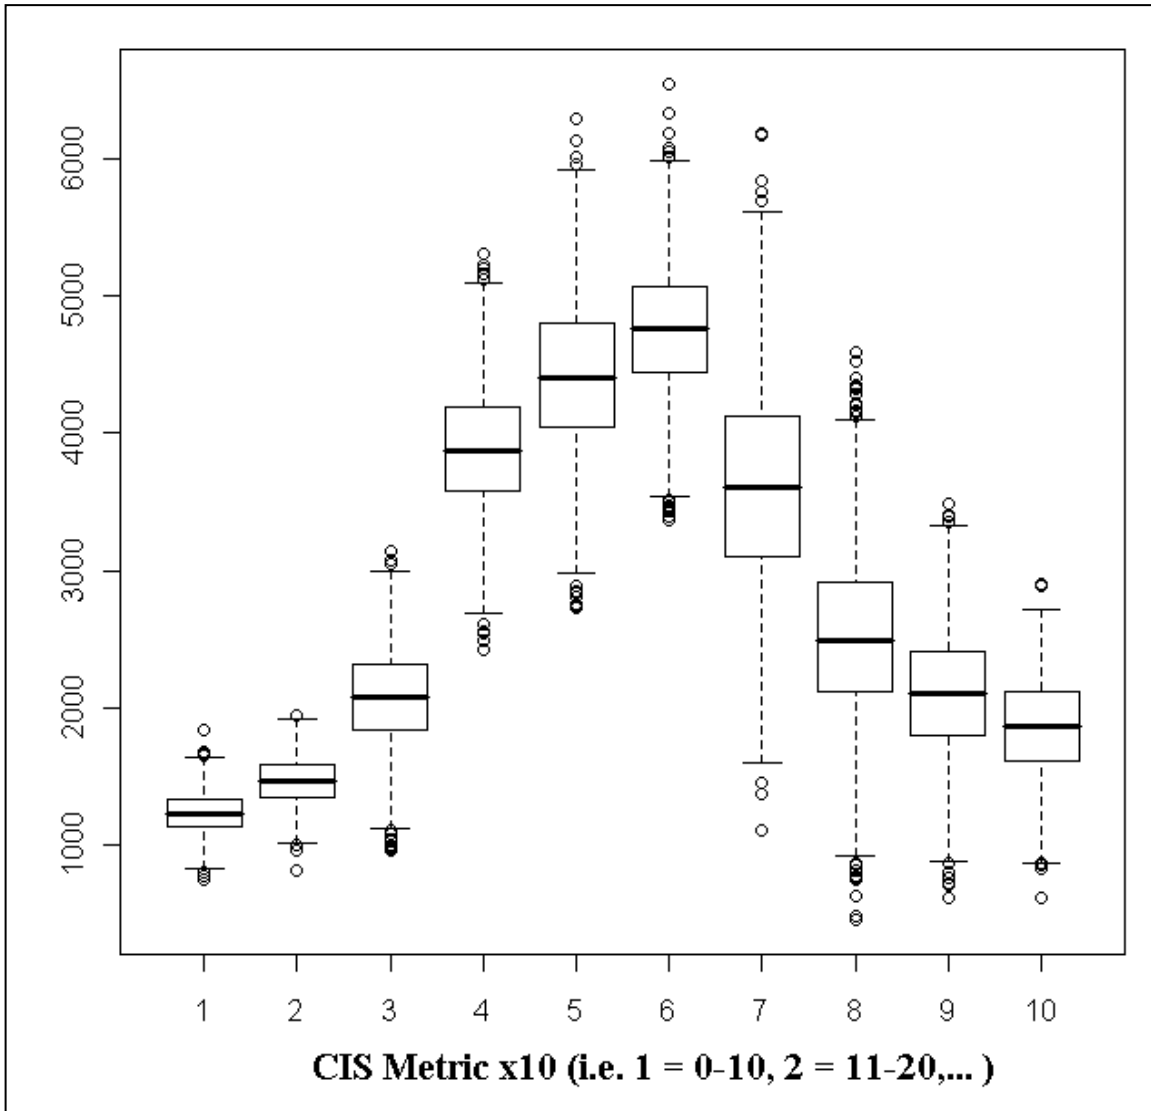


Figure 8. Attacker time by CIS metric.

The results of the experiment provided a good model for predicting system survival. A system with a greater number of vulnerabilities is compromised quicker. This is a negative exponential relationship. Additional vulnerabilities exposed on a host increase the likelihood of compromise significantly. We can hence assert that the greater the number of vulnerabilities that a system has, the faster it is compromised. No system with six (6) or more unpatched network accessible vulnerabilities remained uncompromised for more than 15 days. A compromise occurred in as little as four (4) days on systems with two (2) vulnerabilities. A system with no critical vulnerabilities can be expected to survive for several months even without administrative interaction. Again,

Craig Wright, Author Name, email@address

none of the attacks against these systems could be termed black swans. Each was known and predictable. In each case, known a work around existed.

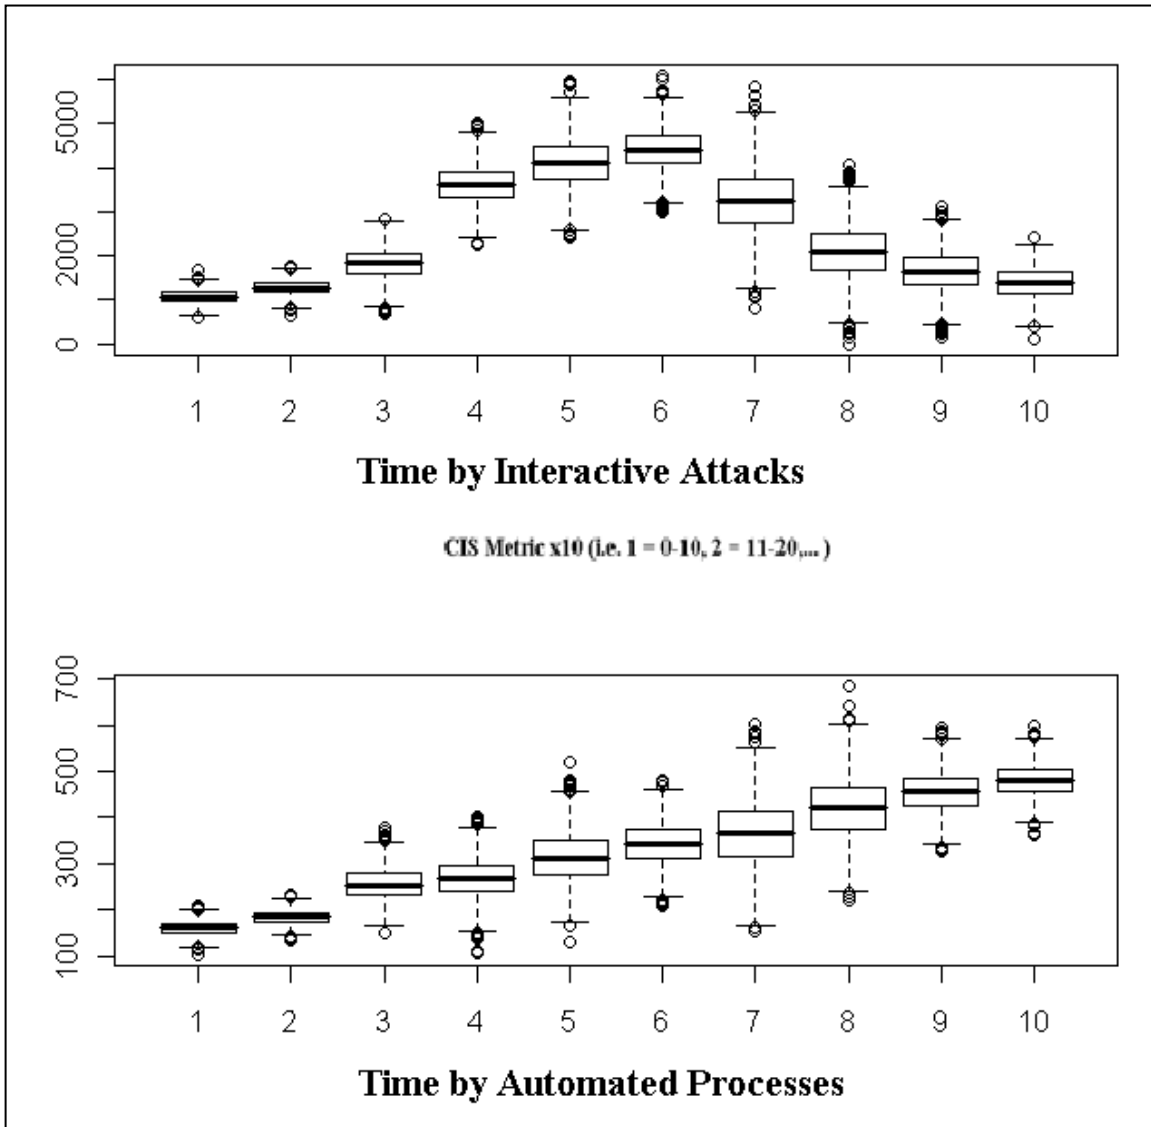


Figure 9. Attacker time by CIS metric and attack class.

2.5. Attack time by CIS Score

The time between the initial instigation of an attack until an attacker either moved on or compromised a host was analyzed for related systems. As we have no means to correlate between systems that an attacker may use, this value is lower in many cases than would be recorded if all the IP addresses used by a single attacker could be utilized. As such, this result is only indicative and does not take attacks from single attackers who use multiple addresses into account.

Craig Wright, Author Name, email@address

In Figure 8, we see that there is an inflection point on the amount of time spent attacking a system. More secure systems (a high CIS metric) would appear to discourage attackers where insecure systems (a low CIS metric) are quickly compromised. Some attackers are determined and will continue to attack a host for extended periods in time. Even when an attacker gets little or no positive feedback, many continue to test a system over time.

This holds result even more strongly when the attack class is separated by automated and interactive attacks (Figure 9). Automated attacks have a low amount of time on the system due to compromise. When a system is not compromised, they display little intelligence into the attack patterns deployed against a host. An interactive attack displays a marked drop in the time per host as the security metric increases. As the host exhibits fewer vulnerabilities, the attacker spends less time exploring the host. It must be noted that the attackers are random in this experiment. The distribution of attackers is unlikely to contain dedicated attackers who are targeting a particular site (such as many cyber criminals and ‘hactivists’ would do (Gordon & Ford, 2002)). The hosts appear as normal corporate and home systems. No information was provided to the attacker that would enable them to associate the hosts with any particular organization.

2.6. SANS Top 20 critical security controls⁴

As was seen above with the example of a simple control in the enabling of a Firewall by default in Windows XP, there are many easy to implement controls that make a significant positive benefit for little cost. Economically, these types of controls are the most beneficial returning far more than they cost to implement.

The Sans top 20 critical security controls methodology provides such a methodology. As can be seen from the other examples in this paper, a focus on known proven attacks and controls to mitigate those attack vectors is economically effective whilst at the same time proving an environment with significantly lowered risk of compromise. Even if the methodology is imperfect and covers far less than a more

⁴ <http://www.sans.org/critical-security-controls/>

comprehensive methodology (such as those from CIS, it gives the user a concrete method to significantly reduce risk with minimal costs.

3. Discussion

User interactions affect survival times in the real world. This will of course change the models produced by this experiment. The modeling of complete interactive systems was outside the scope of the experiment, but user actions can be modeled (Kumar, Pillai & Kumar, 2007) with more advanced techniques (such as clustering algorithms). The results of the experiments presented demonstrate that the focus on zero-day or black swan events is misplaced. These can cause damage, but they are no more likely to damage a system than an attack using a well-known vulnerability that has not been patched. As Anderson (2001) notes, this is hard. The economic costs (Arora, Nandkumar & Telang, 2006) of maintaining a system with all the required patches for all applications are frequently greater than the cost of purchasing and installing the software.

The problems with focusing on zero-day attacks are two-fold. First, the number of attacks that occur from true (Bednarski, Greg & Branson, 2004) zero-day events are fairly low. These incidents also cause the same damage as an incident that has resulted from a known vulnerability; a system compromise is a system compromise. A compromise will result in the same level of loss whether it was from an unmitigated known vulnerability or a true zero day attack.

Next, and more importantly, the total number of negatives is too large. There are simply too many black swans. For every attack that can succeed, there are a near infinite number of possible attack vectors, most of which never occur and will never occur. One such example is the oft-recurring argument as to the possibility of an attack against a fax server running on an analogue line⁵. Much effort that could have been better applied to securing known issues has been applied to such Bunyips. Even where zero day events occur as a platypus (that is a completely unexpected event that few would have believed

⁵ See the "Vulnerability testing in analog modem" thread on Security Basics (Securityfocus mailing list).

possible), the impact is rarely greater (Arora, Nandkumar & Telang, 2006 as well as others) than a compromise from an issue that was exposed but known.

As Carroll, (1872) noted when parodying Victorian inventions, we change little and often give little thought to the economic allocation of funds to mitigate risk.

"I was wondering what the mouse-trap was for." said Alice. "It isn't very likely there would be any mice on the horse's back."

"Not very likely, perhaps," said the Knight; "but, if they do come, I don't choose to have them running all about."

A focus on the unknown at the expense of the basics is foolhardy at best. We can expend effort on addressing all possible and even unknown issues like Carroll's knight, but this will divert expenditure from those events with the greatest impact. By focusing on the unknown, we fail to address the issues that have the greatest impact (Varian, 2004). The result of such an action is waste and loss (Friedman, 1953). By addressing the known issues, we also mitigate many of the unknown ones without trying.

Relative computer security can be measured using six factors (Aycock, 2009):

1. *What is the importance of the information or resource being protected?*
2. *What is the potential impact, if the security is breached?*
3. *Who is the attacker likely to be?*
4. *What are the skills and resources available to an attacker?*
5. *What constraints are imposed by legitimate usage?*
6. *What resources are available to implement security?*

In no event can we account for the unknown, but nor should we overly concern ourselves with it. Basic system hygiene and controls do more to counter black swan events in computer systems than does an effort to focus on the unknown. Of more concern is the limitation we place on responsibility. Focusing on software patches moves the responsibility from the user to the vendor. This makes it less likely (Katz & Shapiro,

1985) that the user will actively implement controls that can mitigate the issues that may occur through software bugs.

By limiting the scope of the user's responsibility, the user's incentive to protect their systems is also limited. That is the user does not have the requisite incentive to take the optimal level of precautions. Most breaches are not related to zero-day attacks (Cohen, 1976). Where patches have been created for known vulnerabilities that could lead to a breach, users will act in a manner (rational behavior) that they expect to minimize their costs. Whether risk seeking or risk adverse, the user aims to minimize the costs that they will experience. This leads to a wide range of behavior with risk adverse users taking additional precautions and risk neutral users can accept their risk by minimizing their upfront costs (which may lead to an increase in loss later). In any event, the software vendor as the cause of a breach is not liable for any consequential damages. This places the appropriate incentives on the user to mitigate the risk. As is noted below, the vendor has the incentive to minimize the risk to their reputation (Telang & Wattal, 2005).

The behavioral effect of *Loss Aversion* (defined as propensity of information security professionals to minimize the impact of loss even against risks that have expectation values of greater gain) should be explored in association with concepts of social capital and cognitive biases such as the endowment effect (for instance where an individual is “willing-to-reveal” at high price, “willing-to-protect” at low price). These issues will be appraised against psychological propensities for both anchoring and adjustment and the Status quo bias (the predisposition to resist changing an established behavior, unless incentive is overwhelmingly compelling). The open question is why are we more willing to blame vendors than to fix our systems and how can we align this to effect a more positive outcome?

The valence effect (as is associated with an individual's overestimation of the likelihood of favorable events being associated and impacting oneself) could be modeled in association to its impact and causal relationship with respect of information security and the feedback effect from rational ignorance and “Cold-Hot Empathy”. The failure to be able to expend resources effectively in securing systems has created a misalignment of

Craig Wright, Author Name, email@address

controls and a waste of scarce resources with alternative uses. The creation of models and methods that are common in many other areas of systems engineering, but which are only just starting to be used in the determination of information systems risk is feasible.

A follow-up paper covering the impact of new malware such as Stuxnet is being written and will address the issues associated with this type of attack in more detail.

4. Conclusion

The optimal distribution of economic resources allocated against risks expressed across information systems in general can only lead to a combination of more secure systems for a lower overall cost. The reality is that, as with all safety issues, information security derives from a set of competing trade-offs between economic constraints. The goal of any economically based quantitative process should be to minimize cost and hence minimize risk through the appropriate allocation of capital expenditure. To do this, the correct assignment of economic and legal liability to the parties best able to manage the risk (this is the lowest cost insurer) is essential. This allocation is what requires assessment. This will allow insurance firms to develop expert systems that can calculate risk management figures that can be associated with information risk. This will allow for the correct attribution of information security insurance products. These, when provided to businesses generally, will provide for the black swan incident.

It is rare to find that the quantification of an externality or the quantitative and qualitative effects on those parties affected by, but who are not directly involved in a transaction, has occurred. This is despite this calculation forming an integral component of any risk strategy. The costs (negative) or benefits (positive) that apply to third parties are an oft-overlooked feature of economics and risk calculations. For instance, network externality⁶ attributes positive costs to most organizations with little associated costs to themselves. In these calculations, the time-to-market and first-mover advantages are critical components of the overall economic function with security playing both positive and negative roles at all stages of the process.

⁶ Metcalf's law refers to the positive effect that can be related to the value of a network and is expressed as equaling 2x the network's number of users.

The processes that can enable the creation and release of actuarially sound threat-risk models that incorporate heterogeneous tendencies in variance across multidimensional determinants while maintaining parsimony already exist in rudimentary form. Extending these through a combination of Heteroscedastic predictors (GARCH/ARIMA⁷ etc) coupled with non-parametric survival models will make these tools more robust. The expenditure of further effort in the creation of models where the underlying hazard rate (rather than survival time) is a function of the independent variables (covariates) provides opportunities for the development of quantitative systems that aid in the development of derivative and insurance products designed to spread risk.

In spreading the risk from outlier or black swan events, organizations can concentrate their efforts into obtaining the best return from their scarce resources. There are far more bunyips than black swans. If we expend excessive resources looking for bunyips and black swans, we will find these from time to time but we will then miss the white swans. Focus on outlier risk incidents is unlikely to decrease the risk faced by an organization in mitigating the black swan event whether a consequence of a zero-day vulnerability, or a new form of attack. This approach will divert resources away from known risks and make these more likely. This lowers the level of security applied to an organization whilst still doing nothing to remove the discovery of an unexpected platypus from time to time. Conversely, good security practice, which leads to the minimization of risk through stopping known events, makes black swans incidents less likely. Good risk and security practice as expressed against known issues also minimizes the impact of zero-day and other outlier incidents.

5. References

Anderson. R. (2001) “Why information security is hard – an economic perspective”. In 17th Annual Computer Security Applications Conference, pp. 358–365.

⁷ GARCH and ARCH are statistical methods used in the analysis of heteroscedastic datasets. See <http://www.eecs.harvard.edu/~parkes/cs286r/spring08/reading3/ARCH.pdf> for more detail.

- Arora, A., Nandkumar, A., & Telang, R. (2006). "Does information security attack frequency increase with vulnerability disclosure? An empirical analysis". *Information Systems Frontiers*, (8:5), pp 350-362.
- Arora, A., Telang, R., & Xu, H. (2004). "Optimal policy for software vulnerability disclosure". The 3rd annual workshop on economics and information security (WEIS04). University of Minnesota.
- Aycock, J. "Computer Viruses and Malware" *Advances in Information Security*, Vol. 22, Springer US
- Bednarski, Greg M. and Branson, Jake; Carnegie Mellon University; "Information Warfare: Understanding Network Threats through Honeypot Deployment", March 2004
- Bradley, T. "Zero Day Exploits, Holy Grail Of The Malicious Hacker" *About.com Guide*
- Campbell, K., Gordon, L. A., Loeb M. P. & Zhou. L. (2003) "The economic cost of publicly announced information security breaches: empirical evidence from the stock market. In *J. Comput. Secur.* 11, 431.
- Carroll, L. (1871) "Through the Looking-Glass And What Alice Found There" Macmillan, USA
- Cohen, P. S. "Rational conduct and social life." *Rationality and the Social Sciences: Contributions to the Philosophy and Methodology of the Social Sciences* 1976
- Devost Matthew G. "Hackers as a National Resource. *Information Warfare – Cyberterrorism: Protecting Your Personal Security in the Electronic Age*". WinnSchwartau (Ed). Second Trade Paperback Edition. New York: Thunder's Mouth Press, 1996.
- Fowler., C. A. & Nesbit.R. F. "Tactical Deception in Air-Land Warfare" *Journal of Electronic Defense*. June 1995
- Friedman, Milton. "The Methodology of Positive Economics." In his *Essays in Positive Economics*. Chicago and London: Chicago University Press, 1953.

Craig Wright, Author Name, email@address

- Gordon, S. & Ford, R. "Cyberterrorism?" Symantec Security Response White Paper 2002.
- Halderman, J. (2010) "To Strengthen Security, Change Developers' Incentives," IEEE Security and Privacy, vol. 8, no. 2, pp. 79-82.
- Honeynet Project & Research Alliance, "Know your Enemy: Trend Analysis", 17th December 2004, <http://www.honeynet.org/papers/trends/life-linux.pdf>
- Honeynet Project & Research Alliance, "Know Your Enemy: Honeynets in Universities - Deploying a Honeynet at an Academic Institution", 26th April 2004, <http://www.honeynet.org/papers/edu/>
- Honeynet Project & Research Alliance, "Know your Enemy: Tracking Botnets - Using honeynets to learn more about Bots", 13th March 2005, <http://www.honeynet.org/papers/bots/>
- Katz, M. L. & Shapiro. C. (1985) "Network externalities, competition, and compatibility". In The American Economic Review 75, 424.
- Marti, K. (2008) "Computation of probabilities of survival/failure of technical, economic systems/structures by means of piecewise linearization of the performance function", Structural and Multidisciplinary Optimization, Vol135/3, Pp 225 - 244.
- Nassim Nicholas Taleb. The Black Swan: The impact of the highly improbable. Random House: New York, 2007
- Ozment, A.& Schechter. S. E. (2006) Bootstrapping the adoption of internet security protocols. In Fifth Workshop on the Economics of Information Security.
- Ramesh Kumar Goplala Pillai, P. Ramakanth Kumar, "Simulation of Human Criminal Behavior Using Clustering Algorithm," iccima, vol. 4, pp.105-109, International Conference on Computational Intelligence and Multimedia Applications (ICCIMA 2007), 2007
- Roese, N. J., & Olson, J. M. (2007). "Better, stronger, faster: Self-serving judgment, affect regulation, and the optimal vigilance hypothesis". Perspectives on Psychological Science, 2, 124-141.

Craig Wright, Author Name, email@address

Shawgo, J., Whitney, N., & Faber S., “*CIS Windows XP Professional Benchmark v.2.0.1*”
<http://cisecurity.org/en-us/?route=downloads.show.single.winxp.201>

The SANS Institute, SANS Top 20, 2007, <http://www.sans.org/top-cyber-security-risks/patching.php>

The SANS Institute, SANS “Top Cyber Security Risks”, <http://www.sans.org/top-cyber-security-risks/zero-day.php> The SANS Institute, “Survival Time History”, 2005,
The Internet Storm Centre, <http://isc.sans.org/survivalhistory.php>

Telang, R., & Wattal, S. (2005). “Impact of software vulnerability announcements on the market value of software vendors -an empirical investigation”. The 4th Annual Workshop on Economics of Information Security (WEIS05). Harvard University.

Varian. H. (2004) “System reliability and free riding. In Economics of Information Security”, L. J. Camp, S. Lewis, eds. (Kluwer Academic Publishers,), vol. 12 of Advances in Information Security, pp. 1-15



Upcoming SANS Training

[Click here to view a list of all SANS Courses](#)

SANS Sydney 2020	Sydney, AU	Nov 02, 2020 - Nov 14, 2020	Live Event
SANS Secure Thailand	Bangkok, TH	Nov 09, 2020 - Nov 14, 2020	Live Event
APAC ICS Summit & Training 2020	Singapore, SG	Nov 13, 2020 - Nov 28, 2020	Live Event
SANS Community CTF	,	Nov 19, 2020 - Nov 20, 2020	Self Paced
SANS Local: Oslo November 2020	Oslo, NO	Nov 23, 2020 - Nov 28, 2020	Live Event
SANS OnDemand	OnlineUS	Anytime	Self Paced
SANS SelfStudy	Books & MP3s OnlyUS	Anytime	Self Paced