



Sharing the Unverifiable: Prediction Exchange
Jason Gordon, infectionvectors.com
SANSFIRE 2006

Overview

In late 2005, one of the most widespread Internet “call-to-arms” predictions circulated among security and non-security groups: the danger of the Microsoft WMF (Windows Meta File) vulnerability. The damage from the resultant attacks (which can be considered relatively mild) was analyzed in depth; the predictions of danger were not. This is a common course for vulnerability concerns. Industry and technology experts take various pieces of data and craft forecasts of danger. With no standard way of formulating or evaluating predictions, especially with regard to applying them to specific environments, the community ends up with a “guru”-run discipline. That style of knowledge management fosters environments where Internet security is considered an imprecise or even near-magical practice. This examination is an attempt to make the exercise of security forecasting into a more scientific endeavor.

At its heart, a cyber attack of any kind is a criminal act. Predicting Internet crime may require the same types of skills and processes used by law enforcement to forecast crime in the physical world. The connection between cyber criminals and “real world” offenses has been established and continues to grow with the profile of the professional malware coder. What is lacking in the IT security discipline is an accepted means of crafting, sharing, and analyzing predictions. This paper looks at possible rationales for such an undertaking as well as the possibility of creating a model for making Internet attack predictions. Although a sample model will be presented, the use of a standard process is examined over making estimates simply based on an expert’s feeling.

Forecast

Previous forecasting research has been concerned with the number of attacks an enterprise may encounter or disseminating warnings for ongoing attacks.¹ The latter is the focus of the US Global Early Warning System.² However, the speed with which Internet attacks have crippled businesses (for example the Slammer/Sapphire worm)³ would indicate that once an automated assault is underway, it is too late to take appropriate action. In addition, to be more effective, a warning system needs to involve a prediction model that can quantify the dangerousness of particular vulnerabilities and proof-of-concept code. This analysis would be completed before attacks are launched, offering “predictions” of future harm. This section begins with a discussion of predictions and their roles in Internet security.

For every reader (or forecast “consumer”), a prediction should present fundamental concepts, allowing the reader to become more familiar with the discipline – so that the issue is better understood when it occurs. Whether or not the prediction is “correct” is not necessarily indicative of its value. That is, to be useful, a prediction about a big worm outbreak doesn’t have to be right about the exploit or timing if it outlines the reasons an outbreak may occur. To be of value to a researcher, the prediction only needs to present a technology, flaw, or trick thoughtfully (tempered without use of excessive fear mongering).

As information technology security matures, better-defined processes for evaluating systems are being developed. This has been seen in IDS technology, in patch management, and in overall practice improvement. Within the security community, new tools, processes, and fervent debate in particular have steadily bolstered one issue: full disclosure. Even without the passionate arguments of its supporters and detractors, we can evaluate claims of vulnerabilities (and related exploits); there is a scientific rigor that can be applied to the components of disclosure. As a tester, one can fairly simply review whether the vulnerability exists, if exploit code is functioning, and if the flaw has been previously reported. There are tools available for such examination: from exploit testers (Metasploit),⁴ to detection mechanisms (antivirus, intrusion detectors, etc.), to code analysis programs (disassemblers, debuggers, etc.). From the debate have come vendor alerts, patch management applications, and most importantly processes for managing the discourse: a well-established and considerate set of literature on “responsible disclosure” practices, Common Vulnerability Enumeration (CVE) identifiers, and Common Malware Enumeration (CME) identifiers.⁵ Within software development, standardized approaches like the Capability and Maturity Model (CMM)⁶ support improvements in secure coding. These efforts have provided for a strong foundation with which security managers can evaluate the bulk of vulnerability information and plan a strategy for their respective organizations.

Standardizing much of the approach to exploit/flaw disclosure has gone a long way to merging this research with the security practices of modern organizations, which are increasingly often process-driven, as opposed to technology-driven, systems. Strategic planning, aside from its value as a marketing term, is important to managers, as they are able to fit their own systems and processes to the requirements of Internet security. Predicting the dangerousness of problems with operating systems would help prioritize, evaluate, and schedule security investments.

Although vulnerability and exploit discussions are becoming better defined, the complementary prediction environment is not. Other disciplines have years of research and practice behind their forecast philosophies and corresponding models. Predictions are part of nearly every scientific endeavor – from earthquake research to criminology. IT security certainly has its share of predictions as well, as can be seen after every vulnerability release or New Year’s forecast column. Unfortunately, the forecasts and responses are not driven by standards-based analysis or presentation, as has been developed in other disciplines. To continue to develop the security community as a professional discipline, with common goals and constituent responsibilities, risk issues

will need to be standardized across the practice. That introduces the question, can security-related prediction research and disclosure be bound by a scientific framework?

On its face, the question is at least worth exploring because of two facets of IT security research as a professional discipline:

- Internet attacks are considered crimes and investigated by law enforcement, where forecasting has been shaped by professional standards with great success
- Responsibility among practitioners of IT security has been established and accepted for disclosure issues, which are similar in nature to prediction releases and would thereby set the stage for responsible forecasting

Predicting crime and dangerousness has been established within criminological, healthcare, geological, and environmental circles for decades and has endured (and continues to endure) a good deal of criticism. Agreement upon the correct way to shape predictions, the respective input values, and the correct way to interpret results has led to some frustrations within the sciences but also to great understanding of the practices. IT security, as its own discipline, will likely find the same trouble with respect within its community.

Responsibilities of Prognosticators

Answering the questions surrounding the responsibility of prediction makers cannot begin without establishing whether there is any reason for connecting responsibility with security forecasts. With respect to full disclosure of system vulnerabilities, a sense of responsibility has been constructed on the possible impacts of such announcements. This is especially true (for right or wrong) when someone releases the details of a vulnerability without coordinating with the vendor.⁷ Whether or not an “impact” alone establishes responsibility, the effects of such claims are the start of this analysis. Everyone that has an audience for their predictions has an impact. No matter the size or scope of the audience, every prognosticator finds they are influencing the choices of an IT owner. From the casual observer who may only affect his or her parents’ decisions about web security to the respected senior analysts planting seeds with business leaders, everyone has an impact. Internally, security managers have a great deal of influence over corporate spending (if they are effective managers) based on their own predictions of Internet security and insecurity. There is no lack of spectators for any prediction; however, there are certainly audiences of various sizes and financial scopes. Determining one’s responsibility based on a judgment over the type/size of audience affected would likely be a trying (if not impossible) exercise.

Determining how the prediction is intended to help the consumer may be more applicable to driving forecaster responsibility. Simply stated, what is the predictor’s analysis aiming to change? A vendor may have a much different view of a threat than an administrator, for better or for worse. Forecasters should incorporate specific actions with their recommendations, even if this action is simply to monitor a process/port. Moreover, there is a need for a prognosticator to assess whether the prediction is a general observation

about the security world or a call-to-arms in an effort to battle one precise threat. That establishes two types of prediction: the tactical and the strategic. Tactical predictions foretelling an impending (often very specific) attack generally ask for a security manager to turn his/her attention to certain vulnerabilities. The broader-based, strategic prediction may be looking for more basic changes in the way a manager views security or specific product classes (especially with regard to funding allocation).

Strategic predictions may appear over and over, in a very familiar cycle. For example, throughout 2004-2005 (and even continuing into 2006), the announcement of the “death of the mass mailing worm” has circulated throughout the IT security literature.⁸ Such a prediction is generally based on the old technology used in mass mailers and the ever-improving technology behind the SMTP scanners and filters. Evaluating such a claim, of course, can be approached in multiple ways. However, assuming that the reader accepts the premise of the prognosticator, then the effect may well be a decreased amount of spending on SMTP filters. With regard to strategic predictions in the last year:

Possible Goal/Impact of Specific Predictions

Strategic Prediction	Possible Goal/Impact
Mass Mailers are Dead	Eliminate SMTP Scanner Upgrade
Mobile/Phone Worm Imminent	Slow deployment of handheld devices
Intrusion Prevention Improvements	Reduce awareness training

It is necessary, even with broad-ranging strategic predictions to have a set of input criteria and decision rules (as well as the defined actions mentioned previously). Tactical predictions also have associated “calls to arms” for IT security administrators. With each prediction is an implicit or explicit directive for change to one or more layers in the reader’s “security in depth” approach. For example, with the prediction: “a network worm based on MS06-00x is imminent,” the writer may be pushing the reader to patch, make firewall changes, or add detection mechanisms to the network infrastructure.

In both cases, whether the prediction is considered strategic or tactical, the requisite changes may be justified, based on individual application of the prognosticator’s analysis to the respective network. Certainly there is no way for the predictor to make a case for each type of system that may be affected by the actions they suggest, however, there is still a need to consider the audience and the ramifications of such changes if there is any notion of responsibility applied to security futurists. The futurist, possibly working outside the realm of network administrator and removed from the “real” world of security administration, should be able to reconcile the theoretical (the prediction, no matter how certain he or she is of the impending attack) with the practical (making assessments of system worth and prioritizing changes).

Regardless of impact assessments, however, there is still a case for professional responsibility among security practitioners. Publishing information that hurts the discipline, due to such problems as exaggerated claims, could be considered irresponsible with or without the ability to point to a tangible impact on an organization. This claim will not be investigated further in this paper as it requires a separate (and likely lengthy)

discussion of professional ethics. However, it is no less important a factor, if the security professional analyzing the claim/claimant deems it to be a factor.

The primary consideration for both producers and consumers of predictions is the level of precision with which a prediction can be delivered. Unverifiable quantities are present a factor of security-related prognostication (as with any forecast model). For example, the majority of flaws and attacks are categorized in terms such as “high” or “severe” when presented to the public. The next section begins with problems of both sharing and consuming predictions.

Unverifiable Quantities

Immeasurability is not just a problem for specific predictions, but for the alert ratings these predictions create. Additionally, the scores used by security analysts are often neither chartable nor comparable. This is due in part to the fact that most alerts are raised by vendors, who are using the warnings in proprietary schemes – they did not intend for anyone to try and normalize scores, especially with their competitors’ alerts. Moreover, it is incumbent upon the security administrator to discern whether the alert is a predictive warning (something intended to warn of a future attack), or descriptive (intended to illustrate the current state of the Internet). In both cases, for a consumer, there are numerous outlets for someone interested in network security to poll; hence the reader is left with a myriad of uncorrelated tools.

Existing Alert Schemes

Alert System	Grading Scheme	Description of Alert
ThreatCon (Symantec)	1,2,3,4	Attack likely or in progress, malware with high risk
Radar Alert (F-Secure)	_,3,2,1	Malware outbreaks
INFOCon (ISC)	Green, Yellow, Orange, Red	Possibility of disruptions, changes in malicious traffic
Threat Watch (Panda)	Green, Yellow, Red	Probability of being affected by threat

Alerts such as these may end up feeding a forecast model or vice versa. These schemes could easily fit into a custom model – an organization would likely adopt the system that their respective vendor employed. Each group (and there are many others not noted in the table⁹) have their own means for initiating and elevating alerts. The examples are provided to give the reader a frame of reference for what types of alerting schemes exist. However, it is relevant to point out that without a verifiable, structured format to the alert and the supporting evidence there is a corresponding weakness in the surrounding debate and/or ability for a reader to refute the claim.

Using Caution

The way prediction and alert notices are released will define much of how security is consumed. For example, a general user may not be interested in looking to multiple sources for virus, spyware, and attack reports. A unified “attack” or dangerousness

warning may be the end result of research and analysis within the community. And certainly, disparate views of what constitutes a threat are not only common between casual Internet users and the security community; as seen above, there is some incongruity within the field as well. Moreover, security professionals no longer operate in a sociological vacuum where attacks and defense strategies are distinct from the rest of the world's crime. Over the last decade Internet crime has become a requisite part of the criminological landscape for any law enforcement body. As such, it is taken as a given that Internet attacks have bridged the law enforcement discipline with IT security.

Evidence for this has been found with increasing regularity over the last two years, which have been relatively fruitful for law enforcement agencies. This takes the form of arrests of malware coders and those that have been profiting from malware. The Beagle worm is reported as a professional piece of malware.¹⁰ This worm (and related Trojans) shows some of the most professional practices in software testing and distribution of any non-commercial product (and many commercial packages). Its ability to generate income for its author places it prominently in the scope of professional products. However, ask a security administrator how much he or she worries about Beagle, and one is likely to be told that the worm is not a major threat. Defining code as professionally written, although important for academic research, does not necessarily equate to tagging the code as "high risk."

It is just as important to point out, however, that like vandalism in the physical world, there are certainly crimes/worms that do not necessarily show a profit motive (consider threats such as Slammer/Sapphire).¹¹

Predicting Dangerousness

Reaching a consensus on what makes for a "significant" attack may be more time consuming than agreeing on a model with which to identify risks. The model could be proven to be accurate by historical modeling (seeing how well the model would have predicted past attacks). Nonetheless, the number of possible input factors for vulnerabilities and exploits is tremendous, deserving of extensive research.

Initial efforts in the field have shown a few recurring problems with categorization:

- When does the "dangerous" period start? Is it as soon as the vulnerability is discovered, the proof-of-concept, automated malware exploits, before the vulnerability is public, etc.
- What can be used for a subject to determine dangerousness? Is it a type of target (Windows machines), a class of malware, a specific sample of malware, a specific flaw in an application – and how will each of these be treated?

These are not different from the problems that clinical professionals working with criminals and health issues have difficulty answering. There is tremendous trouble reaching a consensus for what will constitute a valid model and set of inputs.

One approach to these problems is to simplify what one is trying to discern. If it is possible to take a very narrow set of the subjects and reach consensus on what would constitute a reliable risk model, then incorporating more of the questions raised above may be more palatable. For the purposes of setting an example, the research presented in the next section examines the 2005 official security bulletins released by Microsoft as they relate to possible malware attacks. In addition, a dangerous or serious attack is necessarily important to a model such as defined here. Simply showing the risk of an attack would be enough, allowing the organization to insert its own asset valuation and then devise a “dangerousness” score.

Forecasting in the security discipline is not unlike other sciences. The use of “guru”-like predictions has been a part of many fields, then yielding to standard approaches. Although there are a number of things that seem unique to information security, it is necessary to first examine the contributions of other disciplines to forecasting.

No matter the discipline, predicting how bad something will be, or how great an impact it will have in the future is a science that borders on a black art.¹² In most cases, however, there are models that have at least garnered enough backing to help the study avoid being called “pseudoscience,” as has been the fate of some forecasting systems. A brief survey of the predictive frameworks behind a few disciplines will help draw out some of the relevant correlations to IT security.

Interesting at the outset is the terminology used in other disciplines and how it has changed. Law enforcement doesn’t use “profiling” anymore, they use “criminal investigative analysis.”¹³ This implies more study and understanding than it does a determination. The health sector doesn’t “predict dangerousness” they assess risk and meteorologists make forecasts, not predictions.

- In the case of the health care field, biostatisticians frame the bulk of the hard-science (with the exception of the mental health care disciplines, which are discussed later). Biostatistics looks for predictors of risk through a myriad of screening tests and symptom evaluations. Of course, one will immediately note the logical correlation between malware/viruses and sickness in the physical world, however, this esoteric similarity is not all there is to learn from this field. Although there is great debate about how such things as predictive medicine will be employed, there is already a great deal of research available on genetic markers and other factors that may make a person more inclined to become ill with a specific disease. In addition, biostatisticians have researched the correlation between collected data and predicting future health risks. Electronic systems can be evaluated as being vulnerable and at risk of infection like a biological system.
- Examining health studies closer yields a very topical field: forecasting viral outbreaks. Through extensively detailed models, researchers attempt to project the reach and impact of various strains of virii.¹⁴ Factors such as average temperature in an area, tidal movements, rainfall, etc. are all analyzed with respect to how they

affect outbreaks. Researchers have shown that many viral outbreaks are not random, but can be traced back to local infections that move to a new host or location. Locating these local infections is similar to what many organizations do with firewall/IDS log amalgamation and honeypots. What is still required, however, is a means of taking these research projects and being able to accurately forecast widespread attacks. Environmental factors for the Internet are certainly applicable; many threat reports that already exist take these types of measurements (such as Panda software's threat meter).

In April of 2003, Joe Stewart and Steven Drew of Lurhq argued that a model for predicting worm outbreaks could be established with four factors: release of an advisory, the existence of exploit code, a suitable number of targets (to ensure propagation), and enough time to write and test the worm.¹⁵ This relies heavily on the expert opinion of the analyst, as "enough time" and "suitable targets" would be hard to pinpoint. But, it is a solid starting point for creating a more scientific model, as is outlined in the next section.

- Interestingly, one scientific discipline that does use the term "prediction" is earthquake forecasts.¹⁶ That may be a nod to the lack of scientific model that consistently and accurately predicts an earthquake strike. The reason: earthquakes are occurring all the time; we can see areas that are "likely spots" for quakes. What we care about are damaging earthquakes, and they rely on other events (i.e.: will a minor tremor continue and turn into a major quake). The landscape of the Internet does not always lend itself to "geographic" predictions. However, there could be a warning system that employed a similar type of alert for various technologies.
- Ozone Alerts are based on "unsafe conditions" outdoors for humans, especially those with breathing trouble already.¹⁷ They are founded on research showing that high levels of ozone are not healthy for people to inhale and are issued when our monitoring tools determine a specific threshold will be exceeded. These alerts closely mirror threat alerts for the Internet – times when it is relatively more dangerous to be connected to the Internet with a machine. However, to be especially useful for the security manager, we need some additional descriptors for the alerts. Things such as what types of machine are at a greater risk, what attacks are ongoing, and the level of impact associated with the attacks.
- Criminology is likely the most directly linked field of research, which already overlaps with IT security in terms of preventing and detecting crime. Many Internet crimes are quickly being accepted as part of illicit businesses, run by professional criminals. This allows for many ties to the physical world, where for-profit crime has long been part of the landscape. Criminologists working to predict crime are the closest group of researchers to those in IT security – as both groups attempt to decide where and when the climate is right for crimes to be committed.

If we attempt to predict Internet crime like physical world crime or any other scientific forecasting, attention must be paid to the types of inputs that our model will accept. Crime is often described in both qualitative and quantitative terms. With regards to qualitative information, the nature of the crime, how “bad” it seems are usual descriptors of criminal events. These measures are subjective, and are not generally part of forecasts. The crime projector is more interested in the rates, scope, and location of crime.

Furthermore, there is the question of using static and dynamic inputs for any possible model. These types of academic issues will persist and should be considered a priori if a scientific practice is to develop. The static factors (such as the number of vulnerable machines at any given time) may be difficult to measure on a system such as the Internet. Dynamic factors introduce even more difficulty, as important attributes such as a coder’s interest and time spent on a flaw may be important, yet impossible to measure. These problems point to the need for a very flexible model or possibly the impracticality of such an undertaking.

Static & Dynamic Factors in Threat Prediction

Static Factors	Dynamic Factors
Total Vulnerable Population	Difficulty of Coding Job
Number of Machines Targeted	Interest of Malware Coder
Number of Machines with a Firewall	Free Time of Malware Coder

In both areas, qualitative stats are generally discounted, but they are very important to IT security managers. Although they do need to know how “bad” and attack will be, it is really relative to asset value – something that is unique to each organization. Qualitative descriptions do appear in criminological studies, in a form that would be especially useful for security administrators as well: dangerousness.

Dangerousness is a powerful, yet very tough to define, quality.¹⁸ It is useful in situations such as evaluating inmates for potential release, as well as vulnerabilities analyzed for mitigation priority. If predictions of such a nature sound impossible to the reader at first glance, consider that the same protests were made in the mental health profession for years (decades in fact) before it was accepted as possible under the proper conditions.

Predicting dangerousness can be done by using statistical evidence or analyzing a single subject under controlled conditions. In IT security, when predictions are made it is usually done by considering the virtues of the flaw under review – a method similar to the latter approach. Although an analyst may reference one or two previous vulnerabilities that are similar and the respective outcomes from those reports, there is no mention of a model or statistical inference from those anecdotal reports. Without a foundation for judging what attacks are likely, or how bad an attack may be, we are left with the “warm fuzzy” feeling (or lack thereof) of the analyst. In the IT security world, “dangerousness” may be associated with a flaw that allows an attacker to run arbitrary remote code on a wide array of machines without user interaction. Widespread worms such as Sasser, Blaster, and Welchia were borne of such vulnerabilities.¹⁹

In clinical analysis, a defendant is often judged on their personal qualities (“offender characteristics”) relative to the local crime level – meaning a false positive may occur for an offender who happens to be in an area with a very low rate of crime. Moreover, there is a better chance for a false negative for an offender in a region with a very high rate of crime – as they don’t look so bad, relatively speaking. With the exceptional amount of work being completed on MS Windows vulnerabilities, this may hold true in the IT world as well, where a UNIX flaw may not get the attention it deserves because it doesn’t look quite so bad next to the monthly set of releases from Redmond.

Testing a Model

A model will need to incorporate static and dynamic factors as well as quantitative/qualitative data from the security community. The following section posits one such way a model could be constructed. To begin the analysis, let’s examine one set of possible input factors for analyzing vulnerabilities²⁰:

- Does the exploit require user intervention (can it be executed without needing a user to visit a web page, open an email, view a document, etc.?)
- Can the exploit be executed remotely (as opposed to requiring an attacker to have local/virtual local access)?
- Can the exploit execute arbitrary remote code (as opposed to a denial of service or information disclosure)?
- Is the pathway to the exploit filtered by most firewall policies?
- Has malware (or a proof-of-concept) code been released?
- Was the vulnerability unpatched at the time of the malware/proof-of-concept release?
- What was the vendor rating for the vulnerability/patch?
- Was the vulnerability in an operating system or specific application?

These eight possible descriptors do not even require a numeric score as an answer, but can be given yes or no in many cases. These are simply posited to engage the reader as to whether a model can have significance in evaluating vulnerabilities. Given these factors, the first course is to determine whether it is possible to extract a pattern of dangerous releases from past reports. At the time of this writing, a complete set of Microsoft bulletins was available for 2005, which became the start of the application.

Some of the factors obviously require knowledge that may come too late for administrators (i.e.: malware available at time of patch) – however, those should not be immediately discounted. Using any model such as this and reviewing historic vulnerabilities provides a tremendous understanding to threat researchers of what types of flaws are turned into malware. Also, there is no requirement to use all of the factors that were selected for the table. In addition to this, there are a number of cases where the knowledge of a publicly available malware generator, a widespread flaw, and the lack of a patch continue to produce warnings prior to what most would consider a major attack. The WMF vulnerability from December 2005 is one of the best examples of this. Proof-of-concept code, followed by malware and malware-generating tools existed (and were

publicly acknowledged) without a broadly distributed attack. There were, however, the most dire and broadly published predictions of such an attack in years.

Getting back to the set of possible factors, we need to adopt a set of known threat advisories. Taking the complete set of Microsoft security bulletins from 2005 and assigning a value for each of the factors above, there is a pattern to those that are called out as significant. To perform the analysis, there has to be an agreed upon measure of what constitutes a “dangerous” or “threatening” attack. Borrowing from the early stages of criminal predictions, it is necessary to “profile” the flaws in an efficient manner. For purposes of this exercise, a threat was considered most significant if it met the following:

- Required no user intervention (could occur automatically)
- Involved a remote exploit (as opposed to local access being required)
- Could execute arbitrary remote code (as opposed to DoS, info disclosure, etc.)
- Was not filtered by a firewall which allowed only TCP 80 & 25 outbound
- Publicly available proof-of-concept or malware was available (at any time)
- Malware/proof-of-concept was available prior to patch release
- Received a rating of Critical from the vendor
- Applicable to one or more operating systems (as opposed to applications)

None of the 55 vulnerability bulletins meet all of these factors; by the same token, all bulletins met at least one. It is necessary to determine which have either the most in common, or enough of those the analyst considers to be most important. Introducing the idea of what should be “most important” is a variable that would not yield an efficient return at this early stage of research. Therefore, the highest number of hits was used as match criteria. Given these attributes, the following flaws shake out from 2005 (using the attribute settings above, N for “user intervention required”, Y for “remote exploit”, Y for “arbitrary remote code execution”, N for “filtered by firewall”, Y for “public PoC or malware”, Y for “malware prior to patch”, Critical for its rating and, OS for the applied platform):

Vulnerability	Desc	Require User Intervention	Remote Exploit	Remote Code	Filtered	Malware/PoC	Unpatched at time of exploit	Vendor Rating	OS or App
MS05-001	HTML Help Flaw	Y	Y	Y	N	Y	Y	Critical	OS
MS05-010	License Logging Overrun	N	Y	Y	Y	Y	Y	Critical	OS
MS05-037	JVIEW Profiler	Y	Y	Y	N	Y	Y	Critical	OS

Each of the vulnerabilities above had 7 of the 8 criteria (shown in the order they were presented above, i.e.: the first “Y” or “N” indicates the flaw required user intervention). Two of these did, in fact, produce relatively widespread malware attacks. In December of

2004, a Trojan known as Phel was mass mailed²¹ – taking advantage of an unpatched vulnerability in MS Windows. The exploit involved sending a specially crafted HTML file, email, or a link to such a file to a user with the intention of simply having the user open the message/file. This flaw is interesting in the scope of recent predictions as it closely models the WMF vulnerability from December of 2005.²² In the latter case, however, the security community felt that much more attention was required, not to mention the requisite calls for action (which included a plea to install a 3rd party patch). The WMF flaw appeared on the top of security websites, the front of all types of IT magazines, and elicited predictions of doom from all corners of the Internet.

MS06-001	WMF Flaw	Y	Y	Y	N	Y	Y	Critical	OS
----------	----------	---	---	---	---	---	---	----------	----

The reasons for such a different prediction/reaction could be argued ad infinitum. That debate, however, would only go on to show the difficulty in creating a workable set of criteria for a prediction model.

Other flaws that would have appeared extremely dangerous based on the hypothetical model are the Plug & Play defect, JVIEW flaw, MSDTC, and an Exchange Server vulnerability. Each of these alerts would have matched 6 of the 8 criteria.

MS05-019	TCP/IP Vulns	N	Y	Y	N	N	N	Critical	OS
MS05-021	Exchange Server Vuln	N	Y	Y	N	Y	N	Critical	App
MS05-024	Explorer Web View	Y	Y	Y	N	Y	Y	Important	OS
MS05-026	HTML Help	Y	Y	Y	N	Y	N	Critical	OS
MS05-039	Plug and Play Flaw	N	Y	Y	Y	Y	N	Critical	OS
MS05-041	RDP Flaw	N	Y	N	N	Y	Y	Moderate	OS
MS05-048	CDO Object	N	Y	Y	N	Y	N	Important	OS
MS05-051	MSDTC and COM+ Flaw	N	Y	Y	Y	Y	N	Critical	OS
MS05-053	Graphics Rendering Engine	Y	Y	Y	N	Y	N	Critical	OS
MS05-054	IE Cumulative Update	Y	Y	Y	N	Y	Y	Critical	App

MS05-039 will look familiar as it was the flaw exploited by the Zotob worm²³, an attack called the worst of 2005. MS05-041, even though it was listed as a “Moderate” by Microsoft, turned out to be a rather widely discussed vulnerability due to the proof-of-concept code released for this denial of service flaw.

MS05-051 was the foundation of the Dasher worm²⁴; MS05-054 allowed Delf to compromise some machines in late 2005²⁵ (over 3 weeks before a patch was available). These pieces of malware, however, also did not result in widespread warning or emergency calls.

The JVIEW Profiler bulletin yields the same profile as the WMF flaw. The flaw was exploited by a Trojan named Jevprox²⁶, which came out just prior to the release of MS05-037.

The use of a model that attempts to identify the “most dangerous” may yield the starting point that is necessary for reaching a consensus within the security community. As such, the basis of the static input factors above were pulled from a study of previous malware attacks. In a historical context, some of the worst malware attacks have been predicated upon Microsoft bulletins:

Vuln	Desc	Require User Intervention	Remote Exploit	Remote Code	Filtered	Malware PoC	Unpatched at time of exploit	Vendor Rating	OS or App	Malware
MS03-026	DCOM RPC	N	Y	Y	Y	Y	N	Critical	OS	Blaster
MS02-039	SQL Resolution BO	N	Y	Y	N	Y	N	Critical	App	Slammer
MS04-011	Update for Windows	N	Y	Y	Y	Y	N	Critical	OS	Sasser
MS05-039	Plug & Play	N	Y	Y	Y	Y	N	Critical	OS	Zotob

Looking at the “dangerous” end of the spectrum is not the only way to judge modeling. On the other end of the table are those bulletins that are considered low risk items. Scores of 1 & 2 may not require nearly the investment that a 6 or 7 might. In fact, everything scored at less than a 5 saw no worm/Trojan activity associated with them. Since the existence of malware was a factor built into the table, it may seem artificial to then score the results in terms of how much malware was created for each flaw. To remedy this, the second wave of analysis involved removing those two criteria (did malware get created for the flaw and was the malware present before the patch). With these removed, one still gets very similar results to the first incarnation of the table. Although the TCP/IP Vulnerabilities bulletin (MS05-019) now comes to the top, along with a few position shifts among the top half of the model, the lower half of the table does not change. Those with four or more of the now six criteria are much more likely to be involved in a malware attack.

Microsoft warnings, of course, are not the only vulnerability announcements facing systems administrators and CIOs. In the summer of 2004, for example, there was a seemingly innocuous note concerning phpBB and a “highlight vulnerability.” Six months later, this flaw was used by Santy²⁷, the innovative worm that affected tens of thousands of web servers in December of that year. Predictions of such an attack were rare, as the vulnerability itself received little attention. A model that incorporates input factors such as those above must be careful not to exclude any particular technologies.

Vulnerability	Desc	Require User Intervention	Remote Exploit	Remote Code	Filtered	malware PoC	Unpatched at time of exploit	Vendor Rating	OS or App	Malware
CVE-2004-1315	phpBB Highlight	N	Y	Y	N	Y	N	Urgent	App	Santy

Danger to Society

When predicting the danger an incarcerated individual poses to society, it is a normal reaction to want to err on the side of caution. Most examiners would want to avoid releasing someone from custody just for that individual to go out and harm other innocent people. Given that no model for predicting how threatening any prisoner is has been shown to have 100% accuracy (or anywhere close to such a mark), criminal examiners are constantly faced with having to make a judgment about the subjects they encounter. Sometimes these predictions are wrong, of course, something each analyst is keenly aware of. Therefore, when drawing conclusions about someone's likelihood of committing additional crimes the criminal examiner will have to weigh the risks of being incorrect, i.e.: entering a false positive or false negative.

False positives identify a person to be a threat, unworthy of unsupervised interactions, who in fact is no longer a danger to society. Society goes along paying for the incarceration of these inmates, although they could be released without additional danger. A false negative is a misidentification of a dangerous individual as posing no additional threat. In the latter case, society is harmed by the threat of (or actual committal) of crimes that could have been avoided. The idea of a false negative is often seen as much worse than a false positive with regard to categorizing criminals. In fact, it is safe to say that a high percentage of people would rather the examiner have a high rate of false positives than false negatives. For the examiner, this is also true, as noted by Wayne Petherick in a "Crime Library" article²⁸:

It is often argued that it is better to be overcautious with predictions, stating more often that an individual will be violent when indeed they would not. The penalties for failing to correctly identify a dangerous individual have both social and professional implications. Consequently, most professionals will err on the side of caution and over predict dangerousness. Evidence suggests that even the most sophisticated methods yield a 60 to 70% rate of false positives.

With respect to IT security, sensor analysts are more than familiar with the concepts of false positives and false negatives. When a packet is identified as dangerous by an intrusion detection system (IDS), the analyst is forced to deal with the threat. Resources are expended for each investigation; false positives are gradually worked out of the system. False negatives, however, are the fear of most security administrators. The idea that their sensors are standing guard while attackers slip in and out of their networks is as frightening to the CIO as the newly released murderer is to a sociologist charged with predicting how threatening he/she is to the world.

To make a prediction about Internet security into a scientific venture, scientific questions need to be answered for each forecast: when the attack will occur, how it will occur, why it will occur, where it will occur, how damaging will the attack be?

Disclosure Practices & Goals of Predictions

Can making predictions help prevent an attack? In the mental health field, there is evidence that predicting someone is dangerous and putting them under scrutiny reduces the chances that those individuals will actually cause harm to anyone. As noted by Dr. William H. Reid²⁹:

As Norko and Baranoski pointed out in a recent paper, once someone has been placed in a high-risk group, his or her risk usually decreases. Suicide and other tragedies are often prevented by the closer monitoring, more intensive treatment, and greater attention such patients and groups receive. Many patients who appear to have been "false positives" would have experienced bad outcomes had they not been recognized.

This could certainly have a corollary in the IT security game; as the community scrutinizes a vulnerability and employs defensive tactics, the flaw that is being analyzed may be less attractive to an attacker. Certainly having the monitoring and prevention technologies active helps thwart an attack. However, there is ample evidence to the contrary as well. Although malware preceded 7 of the patches (as shown below), there was a malware (or proof-of-concept) release associated with 25 of the bulletins (28%).

MS05-001	HTML Help Flaw
MS05-010	License Logging Overrun
MS05-024	Explorer Web View Cumulative Update for ISA
MS05-034	Server
MS05-037	JVIEW Profiler
MS05-041	RDP Flaw
MS05-054	IE Cumulative Update

It could be argued that the bulletin release does not provide the type of "attention" that is required, that the regular advisories are not taken seriously on their own. This is possible, and a valid criticism, however, it further weakens the potential for finding consensus in prediction modeling as it would require that the community define levels of attention in addition to dangerousness.

It may appear that attacks are too random to be predicted in the first place. Evidence for this comes in many forms as has been seen in this paper:

- We don't fully understand an attacker's motives.
- We can't predict what vulnerabilities are "easy enough" to turn into exploits.
- We can't predict how easy it is to turn any exploit into a simple-to-use tool.
- We can't say what an attractive target looks like.

Is the question of how an Internet attack develops sociological or scientific? Study of a model and its requisite inputs will likely lead to a better understanding of attacks and exploit development, even if it does not yield an effective tool for accurate prediction.

However, there appears to be possible to identify important factors for malware attacks, even though it may turn out that none of the criteria used for this report are used in a model.

All of this exemplifies the problems with applying scientific rigor to forecasting. If it is accepted that predictions are completely unscientific, that is, there is no verifiable input/output and the methods used by the forecaster are not repeatable, then validation of such claims is a meaningless exercise. However, given that there are benefits of making and reading the prognostications of security professionals, there is much to be gained by reviewing and testing. As a CIO in charge of a large budget, there may be reason to care about the validity of predictions; as a casual reader, looking for entertaining reading, there may not be any interest in or reason to examine claims.

Although a model may be shown to have a certain degree of accuracy, validation of a future event is not possible until it occurs (or does not occur within the timeline provided). That, however, should not be the primary focus of prediction validation. There is an inherent value to providing forecasts that can be revealed during a validation. The research here is not meant to wholly prove or disprove any relationship between the example model and real-world attacks, only the possibility that a model can be constructed – even if it is built by individual security teams within each organization.

Many forecasts in all scientific areas are appended with the caveat, “I’m not saying this will happen, but that we need to be prepared if it does.” This goes without saying, but offers very little to the manager responsible for protecting their organization from the regular flurry of attacks. If there is no model for making accurate predictions in IT security, there is no need for such statements – we will all know that there is no guarantee with respect to a forecast, every local weather viewer will attest to that. If we had a refined, deterministic control over Internet attacks, it would be possible to know when an attack was going to occur, and the precise mechanisms used in the attack. We could then go as far as disconnect systems when the attack was taking place (similar to earthquake “evacuations”). What is proposed is to evaluate a range of factors for a vulnerability/attack and establish a scale by which we can make probabilistic determinations about attacks.³⁰ In essence, as a scientific endeavor, we could say how the resultant prediction was reached – not just why. This leads to a description of how likely the attack is, and the type of analysis that was done to make the prediction.

This paper has posited the idea that an efficient and effective threat profile is possible, provided the consumer is capable of molding the prediction to their own environment. Although reaching a consensus over what types of attributes are important enough to include in such a profile will require much additional research and discussion, the groundwork for how to build such a model has been completed in other scientific disciplines. How well Internet security practitioners learn from those efforts will determine how much additional work will be required for a cyber attack prediction scheme.

References

1. Alexander D. Korzyk, Sr. "A Forecasting Model for Internet Security Attacks." National Information System Security Conference 1998. Available at: <http://csrc.nist.gov/nissc/1998/proceedings/paperD5.pdf>.
2. Bob Brewin. "Fed planning early-warning system for Internet." Computerworld, October 18, 2002. <http://www.computerworld.com/securitytopics/security/hacking/story/0,10801,75248,00.html>
3. The UDP-based Slammer worm devastated Microsoft SQL machines Super Bowl Sunday 2003. One short, but powerful reference to the attack is found in: 86-992PS 2003 Cyber Security Research and Development, Committee on Science before the US House of Representatives, May 14, 2003. http://commdocs.house.gov/committees/science/hsy86992.000/hsy86992_0.HTM
4. MetaSploit: <http://www.metasploit.com/>.
5. CME information is available from Mitre, which maintains the naming database: <http://cme.mitre.org/>. CVE is also a Mitre standardization: <http://cve.mitre.org/>.
6. CMM and CMMI (Capability and Maturity Model Integrated) were developed and are administered by the Software Engineering Institute (SEI) at Carnegie Mellon: <http://www.sei.cmu.edu/cmm/>.
7. This paper deals with the forecasting aspect of publishing security information, not disclosure of vulnerabilities. The issue of responsible disclosure has been discussed by a number of its proponents and opponents: Frank Hecker's Mozilla Weblog (February 13, 2005) <http://www.hecker.org/mozilla/full-disclosure> and Bruce Schneier's Crypto-Gram (November 15, 2001) <http://www.schneier.com/crypto-gram-0111.html> are good starts.
8. A few of the many articles predicting the death of mass mailing worms can be found: John Leyden. "The strange death of the mass mailing worm." The Register. December 9, 2004. http://www.theregister.co.uk/2004/12/09/symantec_virus_forecast_2005/. Larry Seltzer. "The End of the Mass-Mailer Worm Era." eWeek. June 7, 2004. <http://www.eweek.com/article2/0,1759,1607743,00.asp>.
9. The Alert Systems in the table came from: Symantec's ThreatCon: <http://www.symantec.com/avcenter/threatcon/learnabout.html> F-Secure's Radar Alert: <http://www.f-secure.com/products/radar/> ISC's INFOCon: <http://isc.sans.org/infocon.php> Panda's Threat Watch: http://www.pandasoftware.com/virus_info/gtw/default.asp

10. Jason Gordon, “Years of the Beagle.” January 2006.
http://www.infectionvectors.com/library/years_of_the_beagle.pdf. More general evidence of the profitability/professionalism in malware: Jay Lyman, “Online Extortion Busts Highlights Profit, Problem.” TechNewsWorld. July 22, 2004.
<http://www.technewsworld.com/story/35288.html>.
11. A good review of how Slammer/Sapphire works is available from F-Secure: Mikko Hypponen, Gergely Erdelyi. “F-Secure Virus Descriptions: Slammer.” January 25, 2003. <http://www.f-secure.com/v-descs/mssqlm.shtml>.
12. The most notable example is currently earthquake prediction (as is noted later in the paper). One source discussing this opinion: Canadian Induced Seismicity Research Group, CISRG Discussions (undated web page): <http://www.telusplanet.net/public/retom/predict.htm>.
13. Information on the FBI’s Investigative Programs, which include Criminal Investigative Analysis: <http://www.fbi.gov/hq/isd/cirg/ncavc.htm>. Federal Bureau of Investigation – Investigative Programs.
14. One good review of how the bio-sciences forecast viral outbreaks can be found: <http://www.thepfizerjournal.com/default.asp?a=article&j=tpj38&t=Predicting%20The%20Next%20Outbreak>. “Predicting the Next Outbreak.” The Pfizer Journal.2003.
15. Stephen Drew and Joe Stewart, “Predicting the Next Outbreak.” SC Magazine. April 1, 2003. <http://www.scmagazine.com/asia/news/article/419704/predicting-next-outbreak/>.
16. Earthquake prediction is debated in a fascinating conversation documented at: http://www.nature.com/nature/debates/earthquake/equake_frameset.html. Introduced by Dr. Ian Main: “Is the reliable prediction of earthquakes a realistic goal?” Nature, February 25, 1999.
17. Information on Ozone Alerts from NIH (US) “Ozone Alert Values”:
<http://www.niehs.nih.gov/oc/factsheets/ozone/ozonevalu.htm>.
18. For a discussion of “dangerousness” from the clinical psychiatrist’s point of view, see: David A. Cohen. “Notes on the Clinical Assessment of Dangerousness in Offender Populations.” Psychiatry On-Line 1997 (originally published 1996).
<http://www.priory.com/psych/assessin.htm>.
19. Sasser: <http://www.microsoft.com/security/incident/sasser.msp>, Blaster:
<http://www.symantec.com/avcenter/venc/data/w32.blaster.worm.html>, Welchia/Nachi:
http://www.trendmicro.com/vinfo/virusencyclo/default5.asp?VName=WORM_NACHI_A.
20. These eight criteria were selected based on the author’s desire to find items that could be validated easily and compiled from a complete year’s worth of data (the Microsoft

2005 bulletins). They were the only items selected for the research (i.e.: they were not part of a larger set prior to beginning the research, nor were any items added after the results were documented).

21. Additional information on the Trojan known as Phel:

<http://www.symantec.com/avcenter/venc/data/trojan.phel.a.html>.

22. The WMF flaw, which precipitated one of the most active rounds of threat prediction, can be researched further at: Robert Lemos, "Zero-day WMF flaw underscores patch problems." SecurityFocus, January 12, 2006. <http://www.securityfocus.com/news/11368>, and the Microsoft bulletin for MS06-001, "Vulnerability in Graphics Rendering Engine Could Allow Remote Code Execution (912919):

<http://www.microsoft.com/technet/security/bulletin/MS06-001.msp>.

23. Zotob was called the worst outbreak of 2005, although it was on par with threats such as Blaster, Sasser, or Code Red: Mikko Hypponen from F-Secure is quoted in James Niccolai's "Experts see new variants of Windows 2000 worm." InfoWorld, August 17, 2005.

http://www.infoworld.com/article/05/08/17/HNwin2000wormvariants_1.html.

24. Dasher worm information: W32/Dasher.worm, McAfee:

http://vil.nai.com/vil/content/v_137567.htm.

25. Delf information: @Risk: The Consensus Security Alert, December 15, 2005. Volume 4, Issue #50.

<http://www.sans.org/newsletters/risk/display.php?v=4&i=50#widely1>.

26. Jevprox information: <http://www.sarc.com/avcenter/venc/data/trojan.jevprox.html>.

27. Santy information is available at: http://www.f-secure.com/v-descs/santy_a.shtml/

28. Wayne Petherick, "Predicting the Dangerousness of a Criminal." Crime Library (no date provided): http://www.crimelibrary.com/criminal_mind/profiling/danger/4.html.

29. William H. Reid, MD, MPH. Journal of Psychiatric Practice, Volume 9, January 2003, No.1. A copy of the journal article is available at:

<http://www.reidpsychiatry.com/columns/16ReidRisk0103.pdf>.

30. Related supporting documents can be found at <http://www.infectionvectors.com>.

Additional References

Martin Libicki, "The Future of Information Security." Institute for National Strategic Studies. No date provided (appears to have been published after September 2001 but before February 2003). <http://www.fas.org/irp/threat/cyber/docs/infosec.htm>.

Scott Berinato. "The Future Security: After the Storm, Reform." CIO Magazine, December 13, 2003. <http://www.cio.com/archive/121503/securityfuture.html>.

Craig Fosnock. "Computer Worms: Past, Present, and Future." July 27, 2005. Published at Infosec Writers:
http://www.infosecwriters.com/text_resources/pdf/Computer_Worms_Past_Present_and_Future.pdf.

National Institute of Justice. NIJ Journal Issue No. 253 "Predicting a Criminal's Journey to Crime." January 2006. <http://www.ojp.usdoj.gov/nij/journals/253/predicting.html>.

Jason Gordon. "Vector Report: 2005." and "Just in Time: Microsoft's Time to Exploit (Part III)." December 2005. Both available at infectionvectors.com:
http://www.infectionvectors.com/library/vector_report_2005_iv.pdf
http://www.infectionvectors.com/library/jit_ms_exploits_dec2005_iv.pdf

Yurcik, William, David Loomis, and Alexander Korzyk, Sr. "Predicting Internet Attacks: On Developing an Effective Measurement Methodology." September 2000. Published in the Proceedings of the 18th Annual International Communications Forecasting Conference.