Empirical Estimates and Observations of 0Day Vulnerabilities

Hawaii International Conference on System Sciences

Miles A. McQueen
Trevor A. McQueen
Wayne F. Boyer
May R. Chaffin

January 2009

Idaho National
Laboratory

Empirical Estimates and Observations of 0Day Vulnerabilities

Miles A. McQueen[1], Trevor A. McQueen[2] , Wayne F. Boyer[1], May R. Chaffin[1]
[1] Idaho National Laboratory, [2] Harvey Mudd College
[1] {miles.mcqueen, wayne.boyer, may.chaffin}@inl.gov , [2] trevor_mcqueen@hmc.edu

Abstract

We define a 0Day vulnerability to be any vulnerability, in deployed software, that has been discovered by at least one person but has not yet been publicly announced or patched. These 0Day vulnerabilities are of particular interest when assessing the risk to a system from exploit of vulnerabilities which are not generally known to the public or, most importantly, to the owners of the system.

Using the 0Day definition given above, we analyzed the 0Day lifespans of 491 vulnerabilities and conservatively estimated that in the worst year there were on average 2500 0Day vulnerabilities in existence on any given day.

Then using a small but intriguing set of 15 0Day vulnerability lifespans representing the time from actual discovery to public disclosure, we made a more aggressive estimate. In this case, we estimated that in the worst year there were, on average, 4500 0Day vulnerabilities in existence on any given day.

1. Introduction

Risk to a system is the probability of a negative event times the consequence of that event, summed over all possible events. Risk is defined in the National Infrastructure Protection Plan [1] as a function of consequence, vulnerability, and threat. One category of vulnerabilities within a system consists of those that reside in software. Within this category there are the questions of how many vulnerabilities exist, how well they are known, how easy they are to exploit, the resulting privileges that could be gained, and the damage that could ensue. This paper focuses specifically on the first two aspects of vulnerabilities and makes a first order estimate of the total number of 0Day software vulnerabilities which exist at any given time. We anticipate that with additional work this estimate will be able to be tailored for supporting an estimation of the exposure of individual systems from these relatively unknown vulnerabilities.

A vulnerability in software is "an instance of a mistake in the specification, development, or configuration of software such that its execution can violate the explicit or implicit security policy" [2]. The rate of vulnerability reporting for individual software applications has been studied for a number of years [3,4].

For software, in total, the number of publicly announced vulnerabilities has changed over the past 8 years from less than 3 per day in 2000 to over 16 per day in 2008 according to the National Vulnerability Database (NVD) [5]. This high rate of vulnerability announcements has focused attention on the very practical and immediate issues of patch management [6], vulnerability disclosure processes [7,8], and speed of patch generation [9], dissemination, and application.

Unfortunately, this focus has left unattended important issues which should be considered when pondering the security of a system consisting of software as well as hardware. One such issue is the question of how many software vulnerabilities are in existence which have been discovered by potential adversaries, but not yet publicly announced or patched (i.e. 0Day vulnerabilities). This paper's primary contribution is that it proposes and applies a novel method for making first order estimates of the number of 0Day vulnerabilities in existence on any given day.

There is no generally accepted formal definition for "0Day (also known as zero-day) vulnerability". The term has been used to refer to flaws in software that no one knows about except the attacker. Sometimes the term is used to mean a vulnerability for which no patch is yet available. For the purposes of this paper, we formally define a 0Day vulnerability as any vulnerability, in deployed software, that has been discovered by at least one person but has not yet been publicly announced or patched. These 0Day vulnerabilities are of particular interest in well managed systems which have effectively mitigated the publicly known vulnerabilities. In these well managed systems the risk contribution from 0Days will have proportionally increased. To aid understanding of how great a risk the 0Days may pose to a system, an estimate of how many are in existence is needed.

0Day Vulnerability Lifecycle Begins at Discovery

Figure 1. Vulnerability Lifecycle.

Using the 0Day definition given above, we developed and applied a method for estimating how many 0Day vulnerabilities are in existence on any given day. The estimate is made by: empirically characterizing the distribution of the lifespans, measured in days, of 0Day vulnerabilities; determining the number of vulnerabilities publicly announced each day; and applying a novel method for estimating the number of 0Day vulnerabilities in existence on any given day using the number of vulnerabilities publicly announced each day and the previously derived distribution of 0Day lifespans.

In this paper we first make use of 491 vulnerabilities, using the time they were privately reported to a vendor until their public announcement as a conservative stand-in for their lifespan as 0Days. After characterizing these lifespans we proceed to estimate how many 0Days existed on each given day in the past.

We then make a more aggressive estimate of 0Day lifespans by discarding the 491 vulnerabilities mentioned above, and using a small set of 15 0Day vulnerabilities for which we knew the actual date of discovery along with the date of public announcement.

As expected, the lifespans of these 15 vulnerabilities were, on average, quite a bit longer than those of the previous 491 vulnerabilities. Consequently, the new aggressive estimate of how many 0Days exist on each day in the past is significantly higher.

Finally, given the estimations for the number of 0Day vulnerabilities and their lifespans, we looked at whether the risk to a system from 0Days might be less than the raw number estimates would indicate. We tentatively discovered that the more serious the vulnerability the longer its lifespan tended to be and thus the risk to systems appears to actually be greater than the 0Day estimates would indicate.

The rest of this paper is organized as follows. Section 2 discusses a vulnerability's lifecycle and provides context for the discussion of 0Day vulnerabilities. Section 3 characterizes the 0Day lifespan of 491 vulnerabilities. Section 4 describes and applies a method for estimating the number of 0Day vulnerabilities in existence at each day in the past. Section 5 is a recalculation of the number of 0Day vulnerabilities in existence making use of a small set of 15 0Day vulnerabilities. Section 6 is a discussion of a potential relationship between 0Day lifespans and their

Common Vulnerability Scoring System (CVSS) base scores [10]. Sections 7 and 8 present our conclusions and discuss future work.

2. 0Day Vulnerability Lifecycle

The lifecycle of a vulnerability, from its creation until a patch is released and applied, or the software retired is illustrated in Figure 1. The red box in Figure 1 represents the start of a 0Day vulnerability lifespan while the blue boxes represent events that end the lifespan. A 0Day vulnerability only exists during the part of the lifecycle which starts with discovery and ends with it either being patched, the software retired, or it being publicly announced. Note that a vulnerability could, concurrently, remain at multiple stages of the lifecycle for an indefinite amount of time, not necessarily changing state.

The lifecycle represented in Figure 1 is described more fully in the next sections (sections 2.1 through 2.8).

2.1 Vulnerability Created and Deployed

A vulnerability's life begins when it is written into software by a developer using unsafe functions, not checking input data, etc. After a vulnerability is created and the software deployed, a vulnerability may be accidentally patched before it has been discovered (e.g. an undiscovered buffer overflow may no longer be reachable with a long enough string because a patch was developed and deployed to truncate the users input before writing to a database), stay hidden for the entire life of the software, or get discovered and become a 0Day.

2.2 Vulnerability Discovered

Once the vulnerability is deployed as part of the released software it may be discovered by a user of the software, "white hat" vulnerability researcher, or "black hat" hacker. After being discovered, the vulnerability, which is now 0Day, may or may not be exploited. It is possible that the discovery is made by the vendor and is then patched before it can be rediscovered and exploited.

2.3 Vulnerability Exploited

After being discovered a vulnerability may be exploited, depending on who discovered the vulnerability and for what purposes. If the discoverer chooses to keep it secret in order to exploit the vulnerability or sell it to others for exploitation, it may be rediscovered by multiple people.

The vulnerability may also be exploited by multiple groups, remaining a 0Day until it is reported to the vendor and patched, the software becomes obsolete, or it is publicly announced without vendor notification.

2.4 Vulnerability Reported to Vendor

When discovered, the vulnerability becomes a 0Day. Responsible disclosure involves reporting the vulnerability directly to the vendor so that they have an opportunity to develop a patch, or work around, before it is publicly announced. After a vendor has received notification, a vulnerability may be publicly announced before the vendor has developed a solution; the vendor may develop and release a patch; or the vendor may keep the vulnerability confidential and let obsolescence solve the problem.

2.5 Patch Developed for Vulnerability

If a vulnerability is discovered during testing, before the software is released, it may be fixed. However, the early discovery of a vulnerability does not necessarily lead to its being patched because sometimes it is determined that rewriting the code base to address a vulnerability is not cost effective. If the vulnerability is fixed during the testing stage then the vulnerability will never be deployed and its brief life is over. Note that this part of a vulnerability's lifecycle is not shown in Figure 1.

It is possible that a vulnerability is patched inadvertently during a software revision process before it has been discovered. At this point, its life is over as long as no vulnerable versions of the software are still in use.

If the vulnerability belongs to software that has been released and is still in use, it is a 0Day until it is publicly announced or a patch is released and applied.

2.6 Vulnerability Publicly Announced

Responsible disclosure allows the vendor to have a patch available to release as part of a public announcement, e.g. Microsoft's Patch Tuesday. Otherwise, the vendor is informed at the same time the vulnerability is publicly announced, allowing public exploitation of the vulnerability while a patch is developed and deployed.

For sections 3 and 4 of this paper, we define the 0Day lifespan to be the time between when the

vulnerability is reported to the vendor and when it is publicly announced. During this time period, the vulnerability is still a 0Day by the definition above even though the vendor is aware of it, because it has not been publicly announced.

2.7 Patch Released and Applied to Vulnerability

If the released and applied patch addresses a 0Day, the 0Day's life is over. Sometimes a patch is released and applied that fixes a vulnerability which has not been announced. Unfortunately, a new vulnerability may also be created during the patch process.

It is possible that upon patch release, the patch is reverse engineered and the vulnerability rediscovered by a different security research group. This process may even be automated in some cases to include automated generation of an exploit [11]. However, this possibility is not considered in this paper since we are, for the sake of exposition, making the simplifying assumption that once a patch is released that it has been uniformly applied and the vulnerability no longer exists. This simplifying assumption will be removed in our later 0Day research.

2.8 Software Obsolete

It is possible that a 0Day is never reported and therefore its lifetime only ends when the vulnerability software is no longer in use. This lifetime is not addressed in this research.

2.9 0Day Lifespans

By our definition, a 0Day vulnerability is one which has been discovered but has not yet been publicly announced or fixed. This represents the real lifespan of the 0Day. Unfortunately, the actual date of vulnerability discovery is not usually available. The NVD database has "discovery date" records for a small fraction of its entries, and for many of those entries the "discovery date" is the same or nearly the same date as the "publication date". For those cases it is highly unlikely that the recorded "discovery date" is a valid start date for a 0Day vulnerability lifespan. Consequently, in sections 3 and 4 of this paper, we use the time between reporting a 0Day vulnerability to the vendor and its public announcement as a conservative estimate of the actual 0Day's lifespan. In section 5 we make a potentially more realistic estimate of the lifespan using 15 vulnerability data points where the actual discovery date was known.

3. Estimate of 0Day Vulnerability Lifespans

In the 0Day vulnerability lifecycle discussed in the previous section, the time from vendor notification of the vulnerability to its public announcement is but a portion of the overall lifetime of a 0Day vulnerability and is referred to in section 3 and 4 of this paper as its lifespan. This section will describe our work in characterizing these 0Day lifespans.

3.1 Sources of 0Day Lifespan Information

In August, 2005 TippingPoint formed the Zero Day Initiative (ZDI). In TippingPoint's own words "The main goals of the ZDI are to:

- Extend our DVLabs research team by leveraging the methodologies, expertise, and time of others
- Encourage the reporting of zero day vulnerabilities responsibly to the affected vendors by financially rewarding researchers
- Protect our customers through the TippingPoint Intrusion Prevention Systems(IPS) while the affected vendor is working on a patch".

The second bullet is what is of most interest to us since it led ZDI to offering cash incentives to security researchers for the reporting of 0Day vulnerabilities to ZDI.

The process followed by ZDI when offered a 0Day vulnerability is to validate the vulnerability, attempt to negotiate a deal with the researcher, and, if a deal is reached, report the vulnerability to the vendor. When the vendor has developed a patch there is a coordinated public announcement about the vulnerability. The vulnerability details and the disclosure time line are then posted online. The difference between the report to vendor date and the public announcement date found in the disclosure time line may be used as an estimate of the lifespan of the 0Day vulnerability.

In May 1998 iDefense Labs (iDefense) was founded. Their business is to provide clients with leading edge intelligence on vulnerabilities and threats. iDefense efforts include the Vulnerability Contributor Program which is used to acquire 0Day vulnerabilities. They post online a list of all of their acquired 0Day vulnerabilities which have been made public. For each vulnerability the posting includes a description and analysis along with a disclosure timeline. The iDefense disclosure timeline includes the date the vulnerability was reported to the vendor and the date it was announced to the public. These two dates are the same

information captured by ZDI so may also be used as an estimate of the lifespan of the 0Day vulnerability. While other dates, such as the date in which the vulnerability was reported to ZDI or iDefense, are occasionally available for the 0Day vulnerabilities, we used the date reported to vendor and the date of announcement to the public in order to be consistent.

Using these two data sources, we collected vulnerability postings from January 5, 2006 through August 16, 2007 and analyzed the lifespans of the 309 0Day vulnerabilities. Due to higher priority commitments we were then forced to set this research aside until May 22 of 2008. At that time we reconstituted our work and collected vulnerability postings from the two data sources for August 17, 2007 through May 22, 2008, and then analyzed the lifespans of the 182 new 0Day vulnerabilities. This provided a total of 491 0Day vulnerabilities for analysis.

We were concerned about the possibility that the statistics of the 0Day vulnerabilities might have dramatically changed between the first collection of data and the second collection period, or that the statistics might be significantly different between the two sources of data due to some unknown differences in the underlying environment. Consequently, we initially kept the four sets of vulnerabilities separated for individual comparison. The four sets are "ZDI OLD", "ZDI NEW", "iDefense OLD", and "iDefense NEW". ZDI OLD and iDefense OLD consist of the 0Day vulnerability advisories, posted on ZDI and iDefense respectively, during the first data collection period. ZDI NEW and iDefense NEW consist of the 0Day vulnerability advisories, posted on ZDI and iDefense respectively, during the second collection period.

3.2 Characterizing 491 0Day Lifespans

When characterizing the 491 lifespans, four important questions were considered. The first was whether the mean and standard deviation of the lifespans from the four 0Day vulnerability data sets were similar; the second was whether, as a first order approximation, the lifespans could be reasonably characterized using a log-normal distribution; the third question was whether the underlying population distribution of vulnerability lifespans was stable over time; and the fourth question was whether the lifespans from ZDI and iDefense came from the same underlying population of vulnerability lifespans. The rest of this section answers those four questions.

3.2.1 Mean and Standard deviation of 0Day Lifespans. The calculated means and standard deviations of the four vulnerability data sets may be

seen in Table 1. The largest mean for lifespans was 169.81 days (ZDI NEW) and the smallest was 112.74 days (ZDI OLD). The largest standard deviation was 153.21 days (iDefense OLD). The mean and standard deviations of the data sets which were formed by combining the ZDI data into one set, the iDefense data into another set, and combining all of the data together may also be found in Table 1. The means for each of these three combined data sets are very close to each other, approximately 130 days.

**Table 1. Mean and Standard Deviation of 0Day Lifespans.**

Vuln. Data Set	Vuln. Count	Days: Mean Lifespan	Days: STD
ZDI OLD	100	112.74	76.13
ZDI NEW	63	169.81	145.57
iDefense OLD	209	125.11	153.21
iDefense NEW	119	139.47	138.33
ZDI COMBINED	163	134.80	111.49
iDefense COMBINED	328	130.32	147.93
All-Data COMBINED	491	131.81	136.81

For each data set, including the combined sets, the lifespans were placed in bins 25 days wide. We then plotted the vulnerability counts against lifespans. As an example, see Figure 2 for the plot of the four data sets combined. When visually inspecting the lifespan plot in Figure 2, and each of the other plots as well, it seemed possible that the lifespans might be log-normally distributed.

3.2.2 Modeling Lifespans Using a Log-normal Distribution. If the lifespan of a 0Day vulnerability is thought of as the outcome of the discoverer's unique set of attributes, each vendor's patch development process, the individual vendor's economic factors at any moment, the difficulty of patching, and the particular individuals involved in creating the patch then a log-normal distribution [12] may be an appropriate model. To determine how well a log-normal distribution would model the distribution of lifespans in the various 0Day vulnerability data sets, the natural log was taken of each lifespan and placed in bins 0.5 wide. Then the count of vulnerabilities were plotted against the log values. The plot for all of the data sets combined is shown in Figure 3. Visually, to varying degrees, each of the log plots seemed to have a Gaussian distribution.
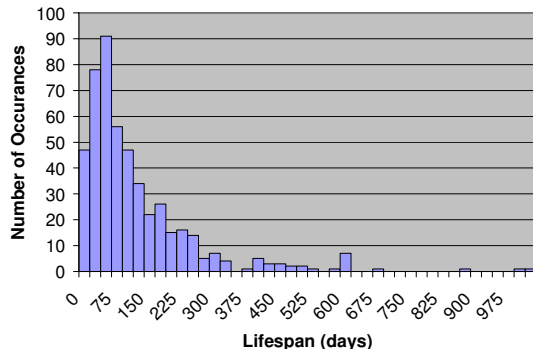
To quantitatively determine the goodness of fit of the log-normal distribution to the lifespan data the

mean of the log values and the corresponding standard deviations were calculated for each of the original four
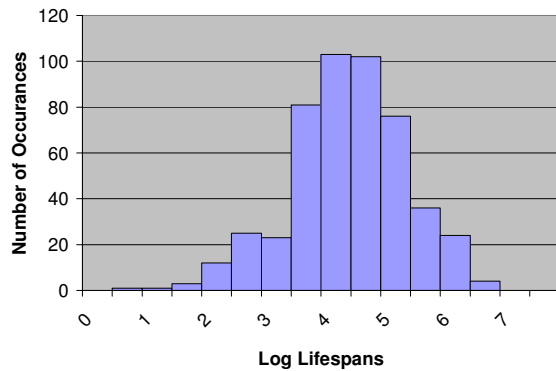
Figure 3. Log Scale Plot of the Four Combined Data Sets.

data sets, the combined ZDI data, the combined

Figure 2. Lifespan of the Four Combined Data Sets.

iDefense data, and all of the data combined. For each data set this information was then used to calculate the R^2 value which roughly represents how well the vulnerability lifespan distributions of the set are represented by a log-normal distribution. The closer R^2 is to 1 the better the fit.

The results of these calculations are shown in Table 2. From this we surmise that modeling the lifespans as a log-normal distribution is a reasonable first order estimate of the 0Day vulnerability distributions found in the data sets. We recognize that other distributions, such as the Poisson, might yield improved lifespan models and understanding of underlying causes, but those possibilities will be evaluated in later research.

Modeling the lifespans of each data set as a log-normal distribution leads to the question of whether the underlying population of 0Day lifespans stays constant over time and whether the ZDI data sets and the

iDefense data sets represent samples from the same population.

Table 2. Mean, Standard Deviation, and R^2 Values for 0Day Lifespans.

Vuln. Data Set	Vuln. Count	LN: Mean Lifespan	LN: STD	R^2
ZDI OLD	100	4.4802	0.7642	0.7547
ZDI NEW	63	4.7878	0.8662	0.6602
iDefense OLD	209	4.2740	1.0967	0.8942
iDefense NEW	119	4.5839	0.8665	0.8657
ZDI Combined	163	4.5991	0.8165	0.9067
iDefense Combined	328	4.3864	1.0287	0.921
AllData Combined	491	4.4570	0.9678	0.9157

3.2.3 Stability of Lifespan Distributions. In order to determine the likelihood that the means of the logs of the four vulnerability data sets came from the same underlying distribution the t-test was performed on each pair of vulnerability sets. The results are shown in Table 3. Corresponding to the t-value is the p-value which represents the chance that the actual measured means of the two data sets are due to them both actually having the same mean value (our null hypothesis). The table shows that when comparing ZDI OLD with ZDI NEW there is a well over 95% (1-0.0163) confidence that the mean of the underlying ZDI vulnerability distribution changed over time (or there is a changed bias in the sampling process used by ZDI). A similar statement may be made of the underlying iDefense 0Day vulnerability lifespan distribution by inspecting the p-value when comparing iDefense OLD with iDefense NEW. Thus, for ongoing estimates of 0Day lifespans, it may be important to continually collect the 0Day statistics as reported at ZDI and iDefense. Longer term 0Day vulnerability data collection, and analysis of the processes used by ZDI and iDefense are needed to more fully understand the apparent change in 0Day lifespans.

Table 3 also shows with high confidence that the ZDI 0Day vulnerability data sets are not being drawn from the same underlying population as the iDefense vulnerability data sets (or there are different selection biases being imposed by the two firms). Further exploration of the specific processes used by ZDI and iDefense are needed before drawing any firm conclusions.

Table 3. t-test On the Four 0Day Vulnerability Lifespan Data Sets.

t, p	ZDI OLD		ZDI NEW		iDefense OLD		iDefense NEW	
	t	p	t	p	t	p	t	p
ZDI OLD	0	1.00	2.4254	0.0163	1.7596	0.0794	-0.957	0.3396
ZDI NEW	-2.4254	0.0163	0	1.00	-3.4105	0.0007	1.5107	0.1326
iDefense OLD	-1.7596	0.0794	3.4105	0.0007	0	1.00	2.6472	0.0085
iDefense NEW	0.957	0.3396	-1.5107	0.1326	-2.6472	0.0085	0	1.00

4. Estimation of 0Day Vulnerabilities in Existence

The number of 0Day vulnerabilities that existed on a specific date in the past was estimated using the publication dates from the NVD and a model for 0Day lifespan as described above in section 3. The estimated number is simply a count of all the vulnerabilities with a vendor notification date less than the date of interest and a public announcement date greater than the date of interest. The vendor notification date of each announced vulnerability is calculated by projecting backward in time from the publication date based on a

```
M = number of Monte Carlo runs
N = number of calendar days in National
        Vulnerability Database (10/1/1988 to present)
for k = 1 to N
      count[k] = 0
      date[k] = calendar date for consecutive days
for m =  1 to M
      for i = 1 to N
            P = number of vulnerabilities published on
                  day i
            for j = 1 to P
                  r  = sample from 0Day lifespan
                        distribution model (in days)
                  x = maximum( 1, i - r)
                  for k = x  to (i -1)
                        count[k]  = count[k] +1
for i = 1 to N
      avg[i] = count[i] /M
      print  date[i] , avg[i]
```

Figure 4. Pseudo code for 0Day vulnerability estimation algorithm.

sample lifespan selected from one of the 0Day lifespan log-normal distributions. This process is repeated multiple times and the result is averaged to obtain the statistical mean. Figure 4 is the pseudo code for the initial method used to obtain estimates for all dates since the earliest NVD vulnerability was published (October 1988). Notice that this method tends to underestimate the number of 0Days in existence for dates near the present day because the number of published vulnerabilities for future dates is unknown.

Figure 5 shows the initial method's results using three different lifespan models from Table 1 (ZDI NEW, ZDI OLD and AllData COMBINED). The results are the average of 1000 runs. The shapes of the graphs are similar for the three different lifespan models. The graphs show several sharp drops in the estimated number of 0Days. The sharp drops are caused by the irregular nature of the publication dates from NVD. For example, there were a large number of vulnerabilities publically announced on December 31 for each of the years 2002, 2003 and 2004. (798 on Dec. 31 2002, 441 on Dec. 31 2003 and 1113 on Dec. 31 2004). On those dates the estimated number of 0Day vulnerabilities has a large reduction because the estimation method treats a large number of publications as an immediate drop in the number of 0Days in the pipeline. Based on the Combined results shown in Figure 5, it is reasonable to estimate the number of 0Day vulnerabilities in existence on any given day during 2006 to be around 2500.

We compared the estimates obtained for log-normal vulnerability lifespan distribution models to estimates obtained using a simpler model where the 0Day lifespan is assumed to be a constant value. Figure 6 shows a comparison between the ZDI NEW log-normal distribution model versus a constant 169 day lifespan model (169 days is the average lifespan for the ZDI NEW data set). The comparison shows that this simpler model produces estimates that are a reasonably good approximation to the results obtained from the log-normal distribution model.

The simplest method we used for estimating the number of existing 0Day vulnerabilities is the average lifespan of the chosen model multiplied by the average daily vulnerability public announcement rate. Figure 7 is a comparison of this method with the initial method described earlier. The ZDI NEW log-normal lifespan distribution model using the initial estimation method is plotted along with this simple method. The daily average public announcement rate is calculated once per year and is averaged over the succeeding 365 days. The estimate is the daily average multiplied by 169 days (the average lifespan for the ZDI NEW data set). This comparison indicates that the simple method

**estimated number of zero-day vulnerabilities**
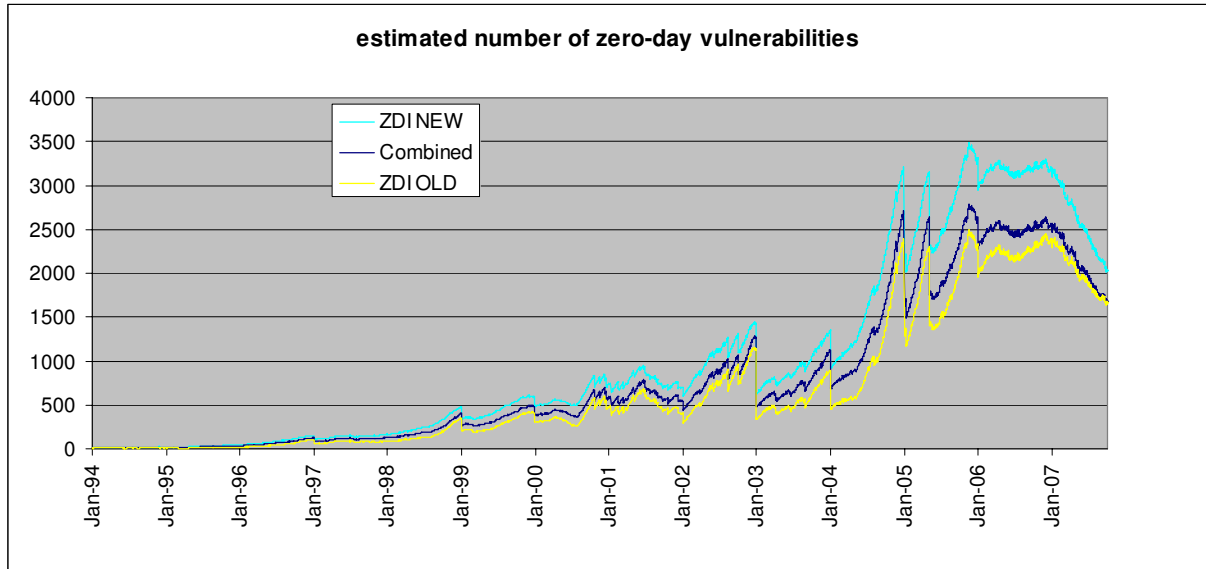
Figure 5. Estimated number of 0Day vulnerabilities using NVD public announcement dates and three cases of log-normal distribution for lifespan.

provides a first order approximation that is comparable to the other more complex methods and models.
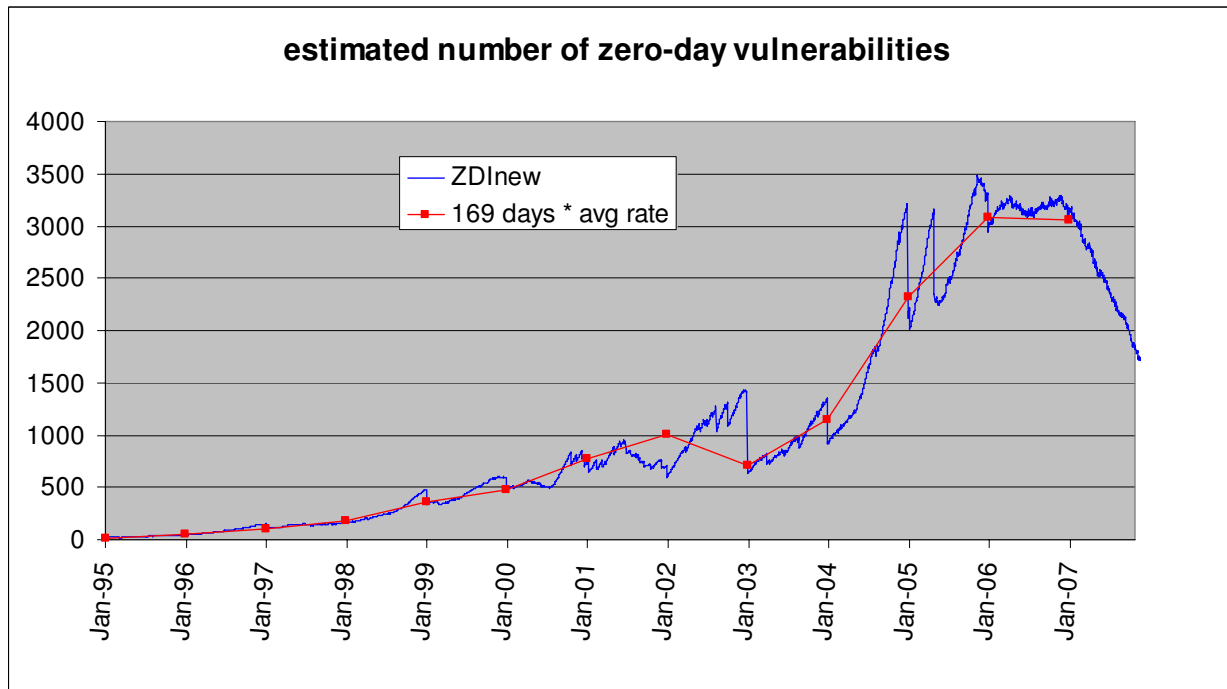
5. Modified Estimation of 0Day vulnerabilities in Existence

The 0Day estimations discussed in the previous section were based on a lifespan that begins with vendor notification and ends with a public announcement. However, this represents only a portion of the true lifespan which actually begins with vulnerability discovery. Unfortunately, the date of initial discovery is usually not known nor as well documented as the vendor notification date. As explained in section 2.9, the NVD database has "discovery date" records for a small fraction of its entries, and for the cases where there is a recorded

estimated number of zero-day vulnerabilities

Figure 6. Estimated number of 0Day vulnerabilities using NVD public announcement dates and log-normal distribution for lifespan (ZDI NEW) compared to constant 169-days lifespan.

estimated number of zero-day vulnerabilities

Figure 7. Estimated number of 0Day vulnerabilities using NVD public announcement dates and log-normal distribution of lifespan (ZDI NEW) compared to average rate times 169-days. Average rate is the number of announced vulnerabilities per day averaged over the succeeding 365 days.

discovery date, the validity of the data is suspect, therefore "discovery date" data from the NVD database could not be used to estimate 0Day lifespan.

A small but intriguing set of 15 vulnerability lifespans was supplied to us by a research group that does vulnerability discovery. We cannot disclose the identity of the research group, and there is no evidence that these 15 lifespans are representative of all 0Day vulnerabilities. However, these lifespans were included in our study because they are the elapsed time from actual discovery date to public announcement date. These same 15 vulnerabilities were also discovered independently and publicly disclosed by a different research group, which provides some added assurance that the lifespan data is valid. Figure 8 is a histogram that shows the lifespan data for these 15 vulnerabilities. The average 0Day lifespan for these 15 vulnerabilities is 256 days and the median lifespan is 200 days. As expected, the average lifespan is larger than the lifespans discussed in the previous sections and would be expected to result in larger estimates for the number of 0Day vulnerabilities in existence.

The 15 lifespans described above were used to estimate the number of 0Day vulnerabilities in existence. The estimates were calculated using the

Figure 8. Fifteen 0Day life spans from discovery to public announcement.

NVD public announcement dates and the method in Figure 4 but with lifespan samples obtained by a uniform random selection from the 15 lifespans. The results are shown in Figure 9. The results show that the new estimate for the worst year (2006) is about 4500 0Day vulnerabilities in existence on any given day. This is a less conservative estimate than the results from the previous section and since the lifespan dataset is small it should not be considered a high confidence

estimate. However, since these 15 lifespans begin at the moment of discovery rather than the vendor notification date, this estimate could very well be more realistic.

Figure 9. Estimated number of 0Day vulnerabilities using NVD public announcement dates and statistics of 15 discovery to announcement.

6. CVSS Base Score and 0Day Lifespans

Given the estimations of the number of 0Day vulnerabilities and their lifespans it is important to ask whether the estimates may be misleading about the risk 0Days pose. If the lifespans are very long for low impact, difficult to exploit, vulnerabilities but very short for high impact, easy to exploit, vulnerabilities then the risk would be lower than the raw numbers lead us to believe. To evaluate this possibility we used the 405 0Day vulnerabilities from ZDI and iDefense which had a CVSS Base score assignment, and broke them into categories based on their score. CVSS is an industry standard for assessing the severity of computer system security vulnerabilities.

The CVSS Base score consists of a weighting of the impact and exploitability subscores. The impact subscore is an evaluation of the potential impact to confidentiality, integrity, and availability from exploitation of the vulnerability. The exploitability subscore is an assessment of the complexity of actual exploitation. This includes measures such as whether the attacker may exploit the vulnerability remotely, how complex the actual attack process is expected to be, and the required number of authentications during attack execution. Details of the CVSS Base score may be found in the CVSS scoring guide [10].

We expected to find that the vulnerabilities which are easily exploitable and have a high impact (very high CVSS Base scores) would have significantly shorter lifespans since it would seem to be in the software owner's and vendor's interest to have the vendor devote their limited resources to fixing them quickly.

Table 4. Mean Lifespans Using Two CVSS Base Score Categories.

	LOW CVSS 0.0 – 7.9	
	Vuln. COUNTS	**MEAN LIFESPAN**
ZDI OLD	45	90.0889
ZDI NEW	22	181.9545
iDefense	102	100.4216
iDefense	57	118.7719
ALL DATA Combined	226	110.9292
	HIGH CVSS 8.0 – 10.0	
	Vuln. COUNTS	**MEAN LIFESPAN**
ZDI OLD	53	130.8491
ZDI NEW	34	164.8824
iDefense	31	145.9032
iDefense	60	142.5500
ALL DATA Combined	178	143.9157

The initial analysis results may be seen in Table 4. We were surprised to see that the mean lifespans of the most severe vulnerabilities (CVSS Base scores 8-10) were actually longer than for the less serious vulnerabilities in three out of the four vulnerability data sets. Upon investigation of the exception we found in ZDI NEW, we discovered that the higher average lifespan value in the low severity category was due to a single vulnerability with a CVSS Base score of 7.8 and a very long lifespan. It is interesting to note that with all of the 0Day vulnerability data sets combined, the mean lifespan of the high severity vulnerabilities are about 30% longer than the low severity cases. This is just the opposite of what would be hoped for from a security perspective.

Rather than use the two, somewhat ad-hoc, CVSS Base score categories found in Table 4, we decided that for the sake of validation and completeness we would do the same analysis as before but make use of the low, medium, and high severity categories as formally defined in the CVSS documentation. This change lead to the categories and analysis results seen in Table 5. In this case all four data sets show a longer lifespan for the high severity vulnerabilities, and the combined data set shows that the medium severity vulnerabilities have a lifespan approximately 20% longer than the low severity vulnerabilities and the

high severity vulnerabilities have a lifespan approximately 20% longer than those of medium severity.

Table 5. Mean Lifespans by CVSS Standard Base Score Categories.

	Low CVSS Base Score (CVSS 0.0 - 3.9)	
	VULN. COUNTS	MEAN LIFESPAN
ZDI OLD	0	0.0000
ZDI NEW	0	0.0000
iDefense OLD	7	101.8571
iDefense NEW	1	13.0000
ALL DATA Combined	8	90.7500
	Medium CVSS Base Score (CVSS 4.0 - 6.9)	
	VULN. COUNTS	MEAN LIFESPAN
ZDI OLD	16	106.5000
ZDI NEW	16	146.6875
iDefense OLD	57	95.2982
iDefense NEW	26	123.1538
ALL DATA Combined	115	110.3043
	High CVSS Base Score (CVSS 7.0 - 10.0)	
	VULN. COUNTS	MEAN LIFESPAN
ZDI OLD	82	113.2317
ZDI NEW	40	181.5500
iDefense OLD	69	124.9420
iDefense NEW	90	134.5333
ALL DATA Combined	281	132.6548

The tentative conclusion from this analysis is that the previous estimates for the number and lifespans of 0Day vulnerabilities may actually underestimate the risk since the estimates don't account for the extended lifespans of the more severe vulnerabilities.

However, more investigation is needed. The question of 0Days with high CVSS base scores having longer lifespans may simply be an artifact of the data we used, or it may be real and be caused by some hidden attribute of the process such as a greater difficulty in developing patches for the more serious vulnerabilities.

7. Conclusions

We demonstrated a method for estimating the number of past 0Day vulnerabilities. In the worst year (2006) we conservatively estimated that there was an average of 2500 0Days in existence on any given day. Using a much smaller vulnerability data set, but one where the calculated vulnerability lifespans ranged from the moment of discovery to the date of public announcement, we more aggressively estimated that there was an average of 4500 0Day vulnerabilities in existence on any given day during the worst year (2006). These estimates are first order approximations that are subject to change as more data becomes available.

We also provided preliminary evidence that the most serious of these 0Day vulnerabilities have longer lifespans than lower severity vulnerabilities. Consequently, 0Day vulnerabilities appear to represent a greater risk to our systems than even the estimated number of 0Day vulnerabilities would indicate.

8. Future Work

We are pursuing investigation into a variety of issues related to estimating both the past and current number of 0Day vulnerabilities, the risk they may pose to a variety of systems, and potential mitigations. The work includes the modification and application of the research described in this paper to individual software programs. This will then be followed by an investigation in applying the attack surface metric concept to systems.

Further, critical infrastructure control systems make use of many programs which are not pervasive and thus they have not undergone the vigorous assault of the larger security research community. Consequently the vulnerability data related to control system software is expected to require a modified approach for characterizing the lifespans and for estimating the number of 0Day vulnerabilities.

Of course the risk posed to systems from 0Day attacks rest not just on the number of 0Day vulnerabilities but also on how easy it is to acquire the desired vulnerability. The ease of acquisition is impacted not just by how many potential attackers are aware of it, but also by the markets available for both buying and selling the 0Day vulnerabilities. This impact will be investigated.

Also, the number of 0Day vulnerabilities and their lifespans may impact the form of an optimal disclosure process which minimizes the risk window, particularly to our most sensitive critical infrastructure

systems. Thus, the relation of 0Days to control system vulnerability disclosure will be investigated.

The question of 0Days with high CVSS Base scores having longer lifespans will also be investigated further to determine whether it is simply an artifact of the data we used, whether it relates to the difficulty in developing patches for the more serious vulnerabilities, or whether the more significant vulnerabilities are more closely guarded by the discoverer.

We are in the process of acquiring many more 0Day vulnerability discovery and public announcement dates from firms and individuals who make the discovery themselves. This data will be used to create an improved estimate of the lifespan of a 0Day vulnerability from initial discovery to public announcement.

9. Acknowledgments

10. References

[1] National Infrastructure Protection Plan, June 30 2006.

[2] A. Ozment, "Improving Vulnerability Discovery Models", ACM Workshop on Quality of Protection, Alexandria, Virginia, October 2007.

[3] E. Rescorla, "Is Finding Security Holes a Good Idea", Security & Privacy, IEEE, Jan-Feb 2005, 14-19.

[4] A. Ozment,,S. Schechter, "Milk or Wine: Does Software Security Improve with Age?", Proceedings of the Fifteenth Usenix Security Symposium. Vancouver, BC, Canada, July 31 - August 4 2006.

[5] NIST, National Vulnerability database, http://nvd.nist.gov.

[6] H. Cavusoglu, J. Zhang, "Economics of Security Patch Management", The Fifth Workshop in the Economics of Information Security, University of Cambridge, England, June 2006.

[7] A. Arora, R. Telang, "Economics of Software Vulnerability Disclosure", Security & Privacy, IEEE, Jan-Feb 2005, 20-25.

[8] D, McKinney, "New Hurdles for Vulnerability Disclosure", Security & Privacy, IEEE, March-April, 2008, 76-78.

[9] S. Frei, B. Tellenbach, B. Plattner, "0-Day Patch-Exposing Vendors (In)security Performance", BlackHat Europe, Amsterdam, NL, March 2008.

[10] P. Mell, K. Scarfone, S. Romanosky, "CVSS- A Complete Guide to the Common Vulnerability Scoring System Version 2.0", On the Forum for Incident Response and Security Team, June 2007.

[11] D. Brumley, P. Poosankam, D. Song, J. Zheng, "Automatic Patch-Based Exploit Generation is Possible: Techniques and Implications", Proceedings of the IEEE Security and Privacy Symposium, May, 2008.

[12] J. Aitchison, J. Brown, "The Lognormal Distribution", Cambridge University Press, 1957.