

**Project Report
DCO-1**

Exploit Probabilities Conditioned on CVSS Scores

**A.B. Wollaber
P.C. Trepagnier**

4 November 2016

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS



This material is based on work supported by
the U.S. Army Program Executive Office, Enterprise Information Systems
under Air Force Contract No. FA8721-05-C-0002 and/or FA8702-15-D-0001.

Approved for public release: distribution unlimited.

This report is the result of studies performed at Lincoln Laboratory, a federally funded research and development center operated by Massachusetts Institute of Technology. This material is based on work supported by the U.S. Army Program Executive Office, Enterprise Information Systems (PEO EIS) under Air Force Contract No. FA8721-05-C-0002 and/or FA8702-15-D-0001. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of PEO EIS.

© 2016 MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Delivered to the U.S. Government with Unlimited Rights, as defined in DFARS Part 252.227-7013 or 7014 (Feb 2014). Notwithstanding any copyright notice, U.S. Government rights in this work are defined by DFARS 252.227-7013 or DFARS 252.227-7014 as detailed above. Use of this work other than as specifically authorized by the U.S. Government may violate any copyrights that exist in this work.

**Massachusetts Institute of Technology
Lincoln Laboratory**

Exploit Probabilities Conditioned on CVSS Scores

*A.B. Wollaber
P.C. Trepagnier
Group 51*

Project Report DCO-1

4 November 2016

Approved for public release: distribution unlimited.

Lexington

Massachusetts

ABSTRACT

We present data-driven results that probe the relationship between the existence of cyber exploits – a stand-in for threat assessment – with CVSS scores, which are a particular metric of cyber vulnerability. Initial results indicate a roughly power-law relationship with an exponent of 7.5 ± 1.4 , rising to a maximum of about 9% for vulnerabilities with the most severe (highest) CVSS scores.

TABLE OF CONTENTS

	Page
Abstract	ii
List of Figures	iv
List of Tables	v
1. BACKGROUND	1
2. ASSUMPTIONS	2
3. RESULTS	5
4. CONCLUSIONS	7
References	8

LIST OF FIGURES

Figure No.		Page
1	Some vulnerabilities have many exploits. These are the numbers of exploits for the top 20 CVE entries in the Symantec malware database.	2
2	Probability function representing the days elapsed between appearance in the NVD and Symantec's malware database. Some incidences are negative, indicating that Symantec published a patch before the vulnerability was entered into the NVD.	3
3	Conditional probability of an exploit existing with a CVSS score v falling in $[v, v + 1]$.	5

LIST OF TABLES

Table No.		Page
1	Curve Fitting Coefficients and Standard Deviations for Exploit Probabilities Conditioned on CVSS Scores	6
2	Conditional Probabilities of an Exploit Occurring for CVSS Severity Rankings	6

1. BACKGROUND

In the absence of emergent probabilities of compromise conditioned on known vulnerabilities, Lippmann et al. devised a simple heuristic to convert a known vulnerability’s Common Vulnerability Scoring System (CVSS) score [1] to a probability of compromise given its observed CVSS score,

$$P(\text{Compromised}|\text{Observed}) = \left(\frac{v}{10}\right)^2, \quad (1)$$

where v is the CVSS score, ranging from 0 to 10, that is associated with a vulnerability enumerated by a Common Vulnerabilities and Exposures (CVE) identification (ID) number [2–4]. Lippman and Riordan remark that “This simple approach could be improved if more detailed information on vulnerabilities and exploits were available.” The purpose of this summary is to better probe this relationship using data made available by Allodi and Massacci [5], who have provided digestible databases of CVE IDs, CVSS scores, and known threats in the form of exploit kits and Symantec inventories of malware and network attacks.

2. ASSUMPTIONS

To compute the conditional probabilities, we use the definition

$$P(A|B) = \frac{P(A \cap B)}{P(B)}, \quad (2)$$

where A is the event that an exploit exists for a particular CVE ID, and B is the event that a CVE ID has an associated CVSS score $v \in [V_{Lo}, V_{Hi}]$. Then $P(B)$ is computed by accumulating the total number of CVE IDs that have CVSS scores in that range, and $P(A \cap B)$ is computed by accumulating the total number of scores in that range that also have known exploits, each of which is divided by the total number of CVE IDs.

Tacit in the construction of these probabilities are several assumptions. Namely, we assume that the National Vulnerability Database (NVD) information provided by Allodi and Massacci enumerates the entire “universe” of vulnerabilities and that the exploit kit and Symantec databases account for all known exploits (at the time of database construction). Additionally, we only query the existence of an exploit’s association with a CVE ID. However, some vulnerabilities have many associated exploits; presumably, the threat presented to these vulnerabilities is higher; see Figure 1.




Figure 1. Some vulnerabilities have many exploits. These are the numbers of exploits for the top 20 CVE entries in the Symantec malware database.

Additionally, many exploits exist without associated CVE IDs. Only about 9% of the Symantec data, for instance, is associated with a CVE ID. These unassociated exploits are unaccounted for in our discovered probabilities. We also attempt to consider the time at which the data was taken by examining the date-stamp in Symantec’s malware database and comparing it with the publication date of the vulnerability in NVD. The latest date in Symantec’s data is from November 16, 2012, implying that NVD entries beyond 2012 are superfluous to this analysis. We also attempt to allow time for a patch to emerge by considering the days elapsed between an NVD entry and a Symantec database entry; a probability function describing that difference is provided as Figure 2. The median of this data was 13 days, so we pruned the NVD data to end at November 3, 2012.




Figure 2. Probability function representing the days elapsed between appearance in the NVD and Symantec’s malware database. Some incidences are negative, indicating that Symantec published a patch before the vulnerability was entered into the NVD.

Other caveats are that:

- The CVSS scores in the provided database are integers, but CVSS version 2 scores should be floating point. We have addressed this by re-computing the base scores from the score vectors, which were also provided in the database. It appears that the data were truncated (“floored”) when they were ingested.
- The data are static. Therefore, these probabilities would change as a function of the date range for the databases if we were to update them (also, CVSS scores are revised over time). As future work, it may be of interest to time-window these probabilities.

- We did not perform statistical analyses to measure uncertainties in these probability estimates.
- Probability estimates only make sense in the context of this specific universe of samples, which are likely undirected attacks against many hosts. A determined attacker may systematically develop their own exploits against discovered vulnerabilities for a high-value target, which would substantially increase the risk to that target beyond any probability associations discovered here.

3. RESULTS

About 2.7% of the CVE entries in this dataset have an associated exploit, regardless of their CVSS score, implying that there a great deal of known vulnerabilities that have no documented exploits against them. Only 9.2% of vulnerabilities with CVSS scores $v \in [9, 10]$ have known exploits in this dataset. Figure 3 presents the conditional probabilities of an exploit existing given a CVSS score $v \in [V, V+1]$, where V ranges from 0 to 9. The top subfigure of Figure 3 provides a comparison to the quadratic heuristic introduced in [4]; the bottom subfigure presents the data alongside a curve-fit to a generic function of the form $f(v) = c + av^b$. Under this restriction, the functional probability fit

$$p(v) = 0.01 + 0.18 \left(\frac{v}{10} \right)^{7.51} \text{ where } v \in [V, V+1] \text{ and } V \in \{0, 1, \dots, 9\}, \quad (3)$$

emerges between the exploit probability and the CVSS score falling into a unit score window.




Figure 3. Conditional probability of an exploit existing with a CVSS score v falling in $[v, v+1]$.

Table 1 provides the values and standard deviations for the curve fit parameters, indicating that the relative standard deviations are in the range of 17% to 24% of their stated estimates. In particular, a one-standard-deviation range of the exponent is $b \in [6.12, 8.9]$, so the value of 7.5 should not be considered a “gold standard.”

Consider the functional form to which we are fitting: $p(v) = c + av^b$. The fit error of the constant, score-independent term $c = 0.01$, is ± 0.0024 , so it is significant at more than 4σ . Thus, the term invites speculation as to its interpretation. One way of interpretation is to assign it to errors in the CVSS assessment process. That is, vulnerabilities that were initially assessed to have a negligible impact turned out to be significant in practice. Another approach, which might be

considered more aggressive, is to consider it as stemming from a low, ineradicable, background exploit probability for any software in the NVD, regardless of its score. This approach agrees with the intuition that no software is “unhackable.” Extending that train of thought a little further, one might consider all software, not just software in the NVD, to have some irreducible residuum of vulnerability that is represented by the constant term. It would be interesting to further categorize the data to determine if statistically significant classes of software appear more or less secure than others.

TABLE 1
Curve Fitting Coefficients and Standard Deviations for Exploit Probabilities
Conditioned on CVSS Scores

$p(v) = c + av^b$	Value	Standard Deviation	Fractional Standard Deviation
c	0.01	0.0024	0.24
a	0.18	0.031	0.17
b	7.51	1.39	0.18

Table 2 presents conditional probabilities that emerge from expanding the score window into the traditional CVSS severity rankings of low, medium, and high. It indicates that medium severity scores are about twice as likely to have an exploit as low severity, and high severity scores are about 3.5 times as likely to have exploits than medium scores. However, they are all low probabilities since so many vulnerabilities have no coincident exploit.

TABLE 2
Conditional Probabilities of an Exploit Occurring for CVSS Severity Rankings

Score Range	Severity Ranking	Exploit Probability
$v \in [0, 3.9]$	Low	0.8%
$v \in [4, 6.9]$	Medium	1.4%
$v \in [7, 10]$	High	4.5%

We also queried the distribution of CVSS scores for CVE IDs that have known exploits and found that 48% of them have scores in the range [9,10], indicating that about half of the known exploits targeted the highest scored vulnerabilities.

4. CONCLUSIONS

Using conditional probabilities, we have made some progress into ascertaining a relationship between vulnerability scores and the existence of an exploit. Initial results indicate that the exploit probability increases with the CVSS score according to a power-law function with an exponent of 7.5 ± 1.4 , with a maximum probability just over 9%. Higher severity threats generally show a trend of having a higher probability of exploits, although exceptions occur for isolated scores. All of the computed probabilities are relatively low since many CVE IDs have no known associated exploits.

As future work, we may consider alternative data sources, accounting for time-windowing the exploit emergence with vulnerability reporting (truncating very old NVD entries), and probing for other trends in the vulnerability types, for instance, amongst Common Attack Pattern Enumeration and Classification (CAPEC) [6] categories.

REFERENCES

- [1] P. Mell, K. Scarfone, and S. Romanosky, “Common Vulnerability Scoring System,” *IEEE Security Privacy* 4(6), 85–89 (2006).
- [2] R.A. Martin, “Managing vulnerabilities in networked systems,” *Computer* 34(11), 32–38 (2001).
- [3] R.P. Lippmann and J.F. Riordan, “Threat-Based Risk Assessment for Enterprise Networks,” *Lincoln Laboratory Journal* 22(1), 33–45 (2016).
- [4] R.P. Lippmann, J.F. Riordan, T.H. Yu, and K.K. Watson, *Continuous Security Metrics for Prevalent Network Threats: Introduction and First Four Metrics* (2012).
- [5] L. Allodi and F. Massacci, “Comparing Vulnerability Severity and Exploits Using Case-Control Studies,” *ACM Trans. Inf. Syst. Secur.* 17(1), 1:1–1:20 (2014), URL <http://doi.acm.org/10.1145/2630069>.
- [6] S. Barnum and A. Sethi, “Attack patterns as a knowledge resource for building secure software,” in *OMG Software Assurance Workshop: Cigital* (2007).