INL/CON-11-22125
PREPRINT

Are Vulnerability Disclosure Deadlines Justified?

International Workshop on Security Measurements and Metrics

Miles McQueen
Jason L. Wright
Lawrence Wellman

September 2011

The INL is a
U.S. Department of Energy
National Laboratory
operated by
Battelle Energy Alliance

Are Vulnerability Disclosure Deadlines Justified?

Miles McQueen, Jason L. Wright, Lawrence Wellman
Cyber Security R&D
Idaho National Laboratory
Idaho Falls, Idaho 83415
Email: {miles.mcqueen,jason.wright,lawrence.wellman}@inl.gov

Abstract

Vulnerability research organizations Rapid7, Google Security team, and Zero Day Initiative recently imposed grace periods for public disclosure of vulnerabilities. The grace periods ranged from 45 to 182 days, after which disclosure might occur with or without an effective mitigation from the affected software vendor. At this time there is indirect evidence that the shorter grace periods of 45 and 60 days may not be practical. However, there is strong evidence that the recently announced Zero Day Initiative grace period of 182 days yields benefit in speeding up the patch creation process, and may be practical for many software products. Unfortunately, there is also evidence that the 182 day grace period results in more vulnerability announcements without an available patch.

1. Introduction

For this paper we define grace period and vulnerability lifespan as follows:

- *Grace period* is the amount of time the security researcher allots to the vendor for providing a fix, after which the researcher may independently announce the vulnerability.
- *Vulnerability Lifespan* is the amount of time a vulnerability has spent in a vendors queue. It starts when the vulnerability is reported or discovered by the vendor, and ends when the vendor provides a patch.

From 2002 through 2011, CERT/CC stated that they allow vendors a 45 day grace period [1]. In 2005 Phil Zimmerman, of PGP renown, was quoted as stating that the vendor should be allowed 30 days to fix a vulnerability [2]. In late 2010 three security organizations that perform vulnerability research, among other business functions, very publicly announced new grace periods they would give vendors. Rapid7 insisted on 15 days followed by a report to CERT/CC. Thus, Rapid7 effectively allows for a 60 day grace period [3]. Google announced they would allow a 60 day grace period [4]. The one notable outlier during late 2010 was the Zero Day Initiative (ZDI) which announced that they would allow vendors a 6 month, approximately 182 day, grace period [5].

The explanations and justifications from the vulnerability researchers emphasized the intent to protect end-users. Aaron Portnoy, a representative of ZDI, was quoted as saying "For every day a vulnerability goes unpatched, end users are susceptible, vendors are being a little bit irresponsible by not patching them." [6]. Google Security team members posted on the Google Online Security Blog an article titled "Rebooting Responsible Disclosure: a focus on protecting end users." [4]. And in August 2010, HD. Moore, CSO of Rapid7, stated "The core issue is that the product has a security flaw; debating about the correct disclosure process shouldn't take away from the fact that the vendor is indeed responsible for anyone exploiting a problem in their product ... The argument for disclosure is simple; the more the end user knows about the problem, the better they can defend against it." [7]. Unfortunately, none of these vulnerability researchers provided verifiable quantitative evidence supporting their chosen grace periods.

Thus, the announced grace periods once again raise questions about the appropriate disclosure timelines. The most important events in the vulnerability lifecycle are shown in Fig. 1. Ideally, each event could be independently observed and validated. Unfortunately,

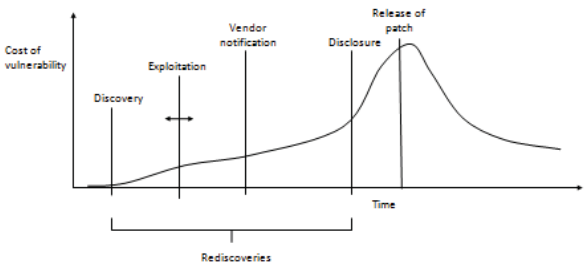

Figure 1. Vulnerability lifecycle events.

most events are not credibly and verifiably known.

The initial discovery of a vulnerability cannot be firmly known by anyone, even a discoverer, since there is always the possibility it has been previously discovered. The date when exploitation begins is after initial discovery but otherwise unknown since exploitation detection mechanisms for previously unseen exploits have questionable detection rates. The rediscovery of a vulnerability is also problematic since the rediscoverer may not take action that makes that fact directly observable. The total cost of a vulnerability, which would include the total losses from its exploitation, the cost to mitigate by end-users, and the cost for the vendor to create a patch, is not dependably and verifiably known for similar reasons.

This leaves us with the events shown in Fig. 2 which occur for each vulnerability within the product which is reported to and, eventually, publicly announced by the vendor. The date reported to the vendor can in principle be known and verified, through oversight of the security researcher who reports the vulnerability. Of course, the disclosure date, independent of patch availability, is publicly known. And patch release dates for a vulnerability are usually publicly disclosed but may be determined through patch reverse engineering if necessary (e.g. [8]). These three events represent

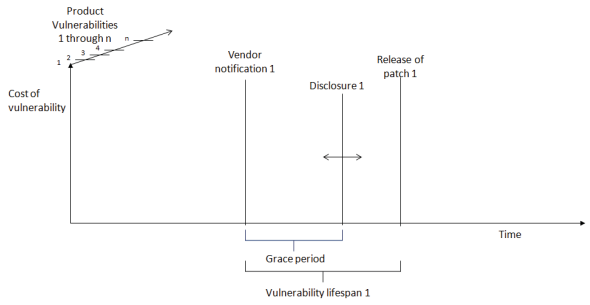

Figure 2. Vulnerability lifecycle events which may be dependably and verifiably known.

the dependably known and measurable aspects of the vulnerability disclosure process debate, and thus they should become the solid foundation on which the grace period and other vulnerability disclosure questions should be discussed.

We focus our analyses on comparing and using the grace periods, and the vulnerability lifespans. The analyses in this paper do not focus on individual products nor on the rate of vulnerability reporting.

1.1. Summary of Results and Contributions

For assessing the value of grace periods specified by vulnerability researchers we first provide some indirect evidence that the shorter grace periods of 45 and 60 days do not appear as realistic as the 182 day grace period if applied to most vulnerabilities.

We then provide strong evidence that vendors still do respond, as previous work suggests [9], to the threat of disclosure. An analysis of the imposed 182 day grace period demonstrates that vendors do modify their patch process so that they are more likely to have a patch available within 182 days. Unfortunately, the 182 grace periods also results in a significant increase in the number of vulnerability announcements made without a patch being available.

1.2. Organization of Paper

The rest of this paper is organized as follows. In Section 2 we provide an overview of the two primary disclosure processes. In Section 3 we compare the announced grace periods to current vulnerability lifespans. In Section 4 we assess the impact of the ZDI, 182 day, grace period on the speed of patch creation. Then, in Section 5, we discuss our conclusions and introduce our planned future work.

2. Overview of Vulnerability Disclosure

Software products have vulnerabilities. The absolute number of vulnerabilities within any given software product is currently unmeasurable with any degree of confidence [10], [11]. What can be determined, and what software vendors must confront, is the number of vulnerabilities being reported and how long it takes to produce a patch. The length of time it takes to produce a patch is directly under the control of the vendor and can be directly influenced by the quality and quantity of resources devoted to the task. It is a business decision, and each vendor (perhaps each vendor's product line) has their own unique costs and benefits to consider.

2.1. Current Disclosure Processes

The current vulnerability disclosure process has two primary forms. The first form is usually referred to as full disclosure and in effect means that upon discovery the vulnerability researcher may publicly announce full details of the vulnerability. The vendor is given no forewarning. The second form, responsible disclosure, generally means that the vulnerability researcher reports the vulnerability to the vendor and gives the vendor time to create a patch. A coordinated public disclosure of the vulnerability is then often made when a patch has been created by the vendor and is ready to be released.

Responsible disclosure has been the topic of heated debate. Some argue that vendors are much to slow at patch development unless they are threatened with the potential of public announcement of the vulnerability, independent of whether a patch is available (supporting evidence may be found in [9]). Consequently, some vulnerability researchers and firms allot a specific amount of time, the grace period, for vendors to create and release a patch. At the end of the grace period these researchers feel free to partially or fully disclose the vulnerability with the idea that end-users may find ways to mitigate the problem even without a patch.

This raises the question of whether these vulnerability researcher specified grace periods are sensible from an end-user vulnerability exposure perspective. We begin to answer this question by analyzing publicly accessible lifespan data of vulnerabilities used in the Pwn2Own competition, and then analyze the lifespans of a much more general set of vulnerabilities. All lifespan data was collected from ZDI.

3. Grace Periods Compared to Vulnerability Lifespans

3.1. Pwn2Own Vulnerability Lifespans

Pwn2Own is a well known computer hacking competition held every year since 2007 as part of the CanSecWest security conference. The contest has high visibility and the vulnerabilities exploited by contestants gain a fair amount of attention from the vendors of the exploited products. For example, Daniel Veditz the Security Group Moderator for Mozilla made this comment about a Firefox vulnerability exploited at Pwn2Own in 2009: "...Since this is a high profile bug (Firefox cracked during a public hacking contest) we need to focus on it. If we had a fix I'd like to shoehorn it into 1.9.0.8 even though we're past code freeze (April release) but May's 1.9.0.9 is more realistic....".

Table 1. Pwn2Own Vulnerability Lifespans

Lifespan (days)		Product	Year	CVE
< 45	8	Apple QuickTime	2007	CVE-2007-2175
	10	Firefox	2010	CVE-2010-1121
	11	Firefox	2009	CVE-2009-1044
	19	Safari	2010	CVE-2010-1120
	20	Safari (WebKit)	2008	CVE-2008-1026
< 60	55	Safari (WebKit)	2009	CVE-2009-0945
	55	Mac OS X	2009	CVE-2009-0154
< 182	61	Adobe Flash Player	2008	CVE-2007-6019
	72	Safari (WebKit)	2010	CVE-2010-1119
	83	Internet Explorer 8	2009	CVE-2009-1532
> 182	310+	Internet Explorer 8	2010	CVE-2010-1118
	310+	Internet Explorer 8	2010	CVE-2010-1117
	676+	Safari	2009	CVE-2009-1060
	676+	Safari	2009	CVE-2009-1042
	676+	Internet Explorer 8	2009	CVE-2009-1043

Because of the increased vendor attention we guessed that the lifespans of Pwn2Own vulnerabilities would be shorter than most other vulnerabilities, and we wanted to see how Pwn2Own vulnerability lifespans compared to the announced grace periods. The results for vulnerabilities exploited at Pwn2Own from 2007 through 2010 are shown in Table 1.

There are a total of 15 previously undisclosed vulnerabilities that could be identified as being exploited at Pwn2Own. Of these 15, 10 have patches available and have lifespans that range from 8 days to 83 days. 50% of the lifespans were 45 days or less, 70% were 60 days or less, and all of them had patches available in 182 days or less. These vulnerability lifespans seem to be in line with the shorter grace periods allotted by vulnerability researchers.

However, looking at the complete data in Table 1 the picture becomes more muddled. There are 5 Pwn2Own vulnerabilities which, for unknown reasons to us, have yet to be fixed. When these five are included in the analysis then we find that only 33% of Pwn2Own vulnerabilities are fixed in 45 days, 47% are fixed in 60 days, and 67% are fixed in 180 days. We decided to look at a larger set of vulnerabilities, which perforce, meant vulnerabilities that were less highly publicized.

3.2. ZDI Vulnerability Lifespans

We collected 473 vulnerability lifespans from ZDI. The lifespans included all vulnerabilities which were initially sold to ZDI and then disclosed sometime between November 7, 2009 and April 30, 2011. Around 15% of the lifespans were less than 45 days, 17% less

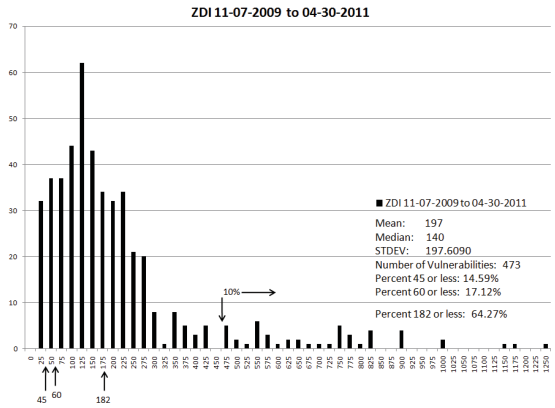

Figure 3. ZDI vulnerability lifespans: Nov, 2009-April, 2011.

than 60 days, and fully 64% of lifespans were less than 182 days. These lifespans can be seen in Fig. 3.

These lifespans do raise the question of whether the 45 and 60 day grace periods are practical at this time. With a median of 140 days, a mean of 197 days, and a maximum of over 3 years we wonder if it is reasonable to expect vendors to be able to meet the grace periods. Even if 10% of the vulnerabilities are granted an exception by the vulnerability researchers, the range of lifespans is 2 days up to 421 days. The 182 day grace period may be more hopeful since it would require a maximum vendor patch creation speedup of 57%. Not easy, but not as difficult as the shorter grace periods.

Of course the adoption of grace periods assumes that vendors actually will speed their patch creation process when confronted with the possibility of an independent disclosure before they have a patch available. Work by Telang and others in 2005 provides some indication that vendors do indeed speed up in these circumstances, in order to protect business value [9]. But we decided to see if there was evidence in the collected ZDI data which indicated that is still the case.

4. Did Vendors Speed Up Their Patch Creation?

On August 4, 2010, ZDI announced their intention of imposing a grace period of 6 months on vendors. ZDI indicated that the new policy would begin immediately.

4.1. Impact On Initial Pool of Vulnerabilities

The initial pool of vulnerabilities (InitPool) is defined to be those ZDI acquired vulnerabilities which, as of August 4, 2011, had been reported to the vendor but were yet to be publicly disclosed. ZDI stated that each of the vulnerabilities in the InitPool would be treated as newly acquired. Thus February 4, 2011 would be the grace period deadline for all InitPool vulnerabilities.

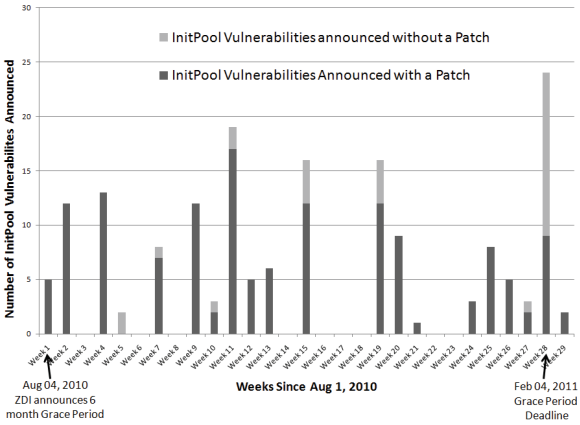

Figure 4. Number of InitPool vulnerability public announcements (each week after ZDI grace period announcement).

When publicly announced, there are three different states for a vulnerability. At announcement, the vulnerability might have a patch available, it might have a vendor specified fix other than a patch, or it might not have any mitigation at all. Since some of the vendor specified fixes did not, in our opinion, credibly address the vulnerability we decided to conservatively group together all vulnerabilities announced without a patch and treat them as unmitigated.

The InitPool started with vulnerabilities which had ages ranging from 14 to 1170 days. The mean age was 183 days and the median age 64 days. There were a total of 172 InitPool vulnerabilities, of which 29 (16.9%) were eventually announced without patches being available. All InitPool vulnerabilities were publicly announced no later than 2 weeks after the grace period deadline. The week in which InitPool vulnerabilities were publicly announced are shown in Fig. 4. The number of patched and unpatched vulnerability announcements are given for each week. It is interesting to note that 16 vulnerability announcements for which there were no patches occurred just days before and after the grace period deadline of February 4, 2011. ZDI enforced their 6 month grace period.

The 16.9% of InitPool vulnerabilities which were announced without a patch, can be thought of as having exceeded the maximum 6 month lifespan imposed by the grace period. In the 87 days immediately preceeding the ZDI announcement of a grace period,

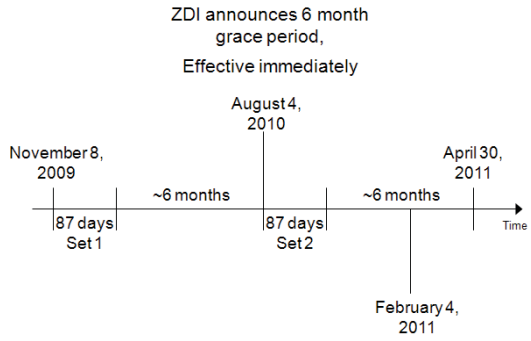

Figure 5. Two sets of ZDI acquired vulnerabilities for lifespan comparison.

over 32.8% of announced vulnerabilities had lifespans exceeding 6 months. So it does appear that some patch speedup was attained after ZDI imposed their grace period on vendors.

However, the InitPool vulnerabilities might have been biased towards more difficult to patch vulnerabilities or vendors more susceptible to speeding up their patch creation when confronted with a threat of disclosure. Since InitPool vulnerabilities might not be representative of all vulnerabilities, we compared two other sets of vulnerabilities.

4.2. Lifespan Comparison Before and After ZDI Announcement

In Fig. 5 the ZDI lifespan data is partitioned into two equal length periods. Period 1 is the 269 days immediately before the ZDI grace period announcement on August 4, while Period 2 is the 269 days immediately after the announcement. Each period has an 87 day time frame at the beginning. Vulnerability Set 1 represents the group of vulnerabilities acquired by ZDI and reported to the vendor in the first 87 days of Period 1. Vulnerability Set 2 is the set of vulnerabilities acquired by ZDI and reported to the vendor in the first 87 days of Period 2. The expectation was that with Set 2 vulnerabilities the vendors would be more likely to have produced patches before the 6 month grace period elapsed than with the Set 1 vulnerabilities for which there had not been a grace period deadline. We compared the vulnerability lifespans of Set 1 and Set 2 to determine if this effect did in fact occur.

The lifespan distributions for both sets of vulnerabilities are shown in Fig. 6. The lifespan statistics for Set 1 and Set 2 can be seen in Table 2. In Set 2, the vulnerability lifespans should not exceed 6 months, but 9 vulnerabilities were announced with patches less than one month after their 6 month grace periods, and

Table 2. Did ZDI announcement have an impact?

Set (i)	1	2
Total Vulnerabilities (A_i)	56	81
range of lifespans	11 - 416	21 - 210
mean lifespan	189	122
median lifespan	190	118
vulnerabilities announced without a patch	1	12
lifespans > 6 months (B_i)	29	21
lifespans < 6 months (C_i)	27	60
proportion > 182 ($P_i = {}^{B_i}/_{A_i}$)	0.518	0.259
Null hypothesis (H_0)	$P_1 = P_2$	
Alt. hypothesis (H_a)	$P_1 > P_2$	
Standard normal (Z)	3.10	
p-value	0.0013	
Conclusion	Reject H_0	

12 vulnerabilities were announced without patches. Together these were counted as 21 vulnerabilities which exceeded the grace period of 6 months. Of note is that Set 1 had 51.8% of lifespans exceed 6 months while Set 2, those vulnerabilities acquired immediately after imposition of the grace period, had only 25.9% of vulnerability lifespans exceed 6 months.

Choosing a null hypothesis of no difference between the percentages of vulnerabilities with lifespans greater than 6 months in the two data sets, and using an upper tailed two sample z-test (e.g. [12]) on the difference in population proportions yielded a P-value of 0.001. For any reasonable test of significance the null hypothesis can be rejected. This hypothesis test is summarized in Table 2.

We should note that the ZDI grace period deadline is not entirely firm. Exceptions may be made if ZDI,

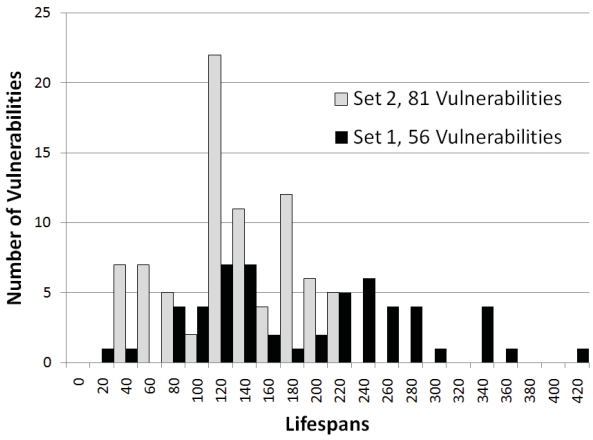

Figure 6. Distribution of vulnerability lifespans for Set1 and Set2.

through discussion with the vendor, decides there are mitigating circumstances. Consequently, a fairer analysis of the impact of the ZDI grace period might focus entirely on the publicly announced vulnerabilities for which no patch was available. Rather than 21 vulnerabilities exceeding the grace period in Set 2, we would have just 12 vulnerabilities, 14.8%, announced without a patch since 9 of the 21 vulnerabilities exceeding the 6 month grace period were announced with a patch in the seventh month. Those 9 vulnerabilities reflect some flexibility in the grace period when ZDI believes the vendor is responding to fix the problems.

To make a corresponding change to the 29 Set 1 vulnerabilities which exceeded the 6 month grace period, only 2 were removed since they were also announced in the seventh month. Thus there were still 27 vulnerabilities in Set 1, 48.2%, which had lifespans longer than the extended grace period of 7 months. Following the same analysis as above this leads to a new P-value of 0.00001. Once again the null hypothesis of equivalent lifespans for both Set 1 and Set 2 can be rejected for any reasonable level of significance.

This is strong evidence that in the general pool of ZDI acquired vulnerabilities, the 6 month grace period did result in vendors speeding up their patch process. There is also evidence that the grace period results in more vulnerability announcements without a patch being available.

5. Conclusions and Future Work

Vulnerability research organizations such as Rapid7, Google Security team, and ZDI have imposed grace periods for public disclosure of vulnerabilities with or without an effective mitigation from the affected software vendor. At this time we were not able to find data which either firmly support or refute the usefulness of the shorter grace periods of 45 and 60 days. There is evidence that the ZDI grace period of 182 days yields some benefit in speeding up the patch creation process. From a risk perspective it is important to note that even after the new grace period there were still 25.9% of ZDI reported vulnerabilities which did not have patches available in the specified time frame.

Whether or not the ZDI grace period of 182 days has a lasting effect will be tested after we can collect another 87 days of lifespan data from ZDI. This will occur on July 26, 2011.

Acknowledgment

The authors would like to thank Debbie McQueen for her assistance in gathering the ZDI data.

References

[1] CERT Coordination Center. (2008, May) CERT/CC vulnerability disclosure policy. [Online]. Available: http://www.cert.org/kb/vul_disclosure.html

[2] J. E. Dunn. (2005, Jan) 'Serious' Microsoft Office encryption flaw discovered. IDG News. [Online]. Available: http://www.pcworld.com/article/119483/serious_microsoft_office_encryption_flaw_uncovered.html

[3] Rapid7. (2010, June) Vulnerability disclosure policy. [Online]. Available: http://www.rapid7.com/disclosure.jsp

[4] C. Evans, E. Grosse, N. Mehta, M. Moore, T. Ormandy, J. Tinnes, and M. Zalewski. (2010, July) Rebooting responsible disclosure: a focus on protecting end users. Google Online Security Blog. Google Security Team. [Online]. Available: http://googleonlinesecurity.blogspot.com/2010/07/rebooting-responsible-disclosure-focus.html

[5] A. Portnoy. (2010, Aug) ZDI disclosure process changes. Zero Day Initiative. [Online]. Available: http://dvlabs.tippingpoint.com/blog/2010/08/03/zdi-disclosure-changes

[6] K. McLaughlin. (2010, August) HP's zero day initiative gives vendors patching deadline. CRN Technology News. [Online]. Available: http://www.crn.com/news/security/226500302/hps-zero-day-initiative-gives-vendors-patching-deadline.htm

[7] S. Ragan. (2010, August) The new era of vulnerability disclosure – a brief chat with hd moore. The Tech Herald. [Online]. Available: http://www.thetechherald.com/article.php/201033/6025/The-new-era-of-vulnerability-disclosure-a-brief-chat-with-HD-Moore?page=2

[8] D. Brumley, P. Poosankam, D. Song, and J. Zheng, "Automatic patch-based exploit generation is possible: Techniques and implications," in *IEEE Symposium on Security and Privacy*. IEEE Computer Society, May 2008, pp. 143–157.

[9] R. Telang and S. Wattal, "An empirical analysis of the impact of software vulnerability announcements on firm stock price," *IEEE Transactions on Software Engineering*, vol. 33, no. 8, pp. 544–557, June 2007.

[10] A. Ozment and S. E. Schechter, "Milk or wine: Does software security improve with age?" in *15th USENIX Security Symposium*. USENIX, July 2006, pp. 93–104.

[11] E. Rescorla, "Is finding security holes a good idea?" *IEEE Security and Privacy*, vol. 3, no. 1, pp. 14–19, January 2005.

[12] J. L. Devore, *Probability and Statistics for Engineering and the Sciences*, 7th ed. Brooks / Cole, January 2008, ch. Inferences Based on Two Samples, p. 354.