The dark side of vulnerability exploitation: a proposal for a research analysis. [*]

Ph.D. Student: Luca Allodi
Advisor: Prof. Fabio Massacci

Università degli studi di Trento
Trento, Italy

Abstract. Software security research has put much effort in evaluating security as a function of the expected number of vulnerabilities and their criticality. As hackers become more sophisticated and economically-driven, I argue that exploitation activities are a much more interesting index of risk than the number of vulnerabilities: the economics of the black market can shed light on attacking processes and trends, and can be very useful in better assessing security and re-thinking patching behavior and patches priority.

1 The problem is in the approach

Security is not easy to quantify. The usual approach [1–4] is to evaluate security in a semi-static way: the researcher takes into account the number of vulnerabilities that affect a system in some time-frame and their respective exploitation easiness; the vendor has to choose which vulnerabilities' patches should be prioritized and typically uses the CVSS Impact Score [5] as an index to make that decision. The main claim of my research proposal is that both the researcher and the vendor should not only be concerned with the volume and criticality of vulnerabilities, but rather with the *effective risk factor* that those would introduce in the operational system. I argue indeed that a vulnerability represents a risk only if it is not latent and is efficiently exploited.

I show in Section 3 that, while on first look this could seem an optimistic and naive claim, and actually an '*under*-semplification' of the problem, preliminary evidence exists that attackers' intentions are more predictable than considered in the literature: precious information can be inferred by more carefully monitoring criminal activities and exploitation trends. This type of information might seem hardly reliable and trustworthy, but the fact that criminal underground activity is becoming, as I show in section 3, more and more structured and economically-driven makes it easier and more meaningful to evaluate risk as related to actual criminal behavior and trends rather than simply as an unproven 'exposure to potential attacks'. My Ph.D. research goal is to identify, by means of auditing and understanding criminal underground activities, schemes or trends that would help in better defining security metrics and would be useful for: (a) the user,

[*] Work partially supported by EU-funded project FP7- 256980 NESSOS.

that could better choose which software configuration is more secure in that particular time-frame; (b) the vendor, that could better allocate human and economic resources by means of a more knowledgeable understanding of actual exploitation risks, thus increasing product security and monetization of effort.

In the next section I give some brief definitions helpful in describing vulnerability exploitation. In section 3 I introduce the different markets involved in the process and describe the structure supporting it. Finally in section 4 I draw my conclusions and Ph.D. proposals, formulating three hypothesis that, if hold true, could help improve software security and patches scheduling.

2 Some quick definitions

Vulnerabilities. I use Ozment's definition of vulnerability [6], according to which vulnerabilities are mistakes in the code or in the configuration of a software that can cause violations in its security policy. These mistakes can be exploited by an attacker to get access to the vulnerable system.

Exploits. One could identify different levels of maturity of an exploit as they usually are born as simple proofs-of-concept, are then scripted and eventually automated [7]. Frei et al. in [8] analyzed more than 14 thousand vulnerabilities. Out of these only about 3400 were exploited, most of which within a month from the disclosure of the vulnerability.

Attacks. An attacker needs to exploit (at least) one vulnerability in the system to reach his goal. The relation between exploitation time and vulnerability disclosure date is shown in [9] by Arora et. al: attacks increase at the time of the vulnerability disclosure. There also seems to be a correlation between random-wide information scans and attack probes, evidencing that untargeted attacks are common practice [10, 11].

3 The Markets

Vulnerability and exploit markets are distinct but related: while the former is divided between legitimate and illegitimate markets [12], the latter is mostly an underground activity usually labeled as 'black market'. On the other hand, the financial consequences of vulnerability exploitation have been shown to go far from solely their market value [13–15].

The market of vulnerabilities. No extended study exists, to the best of my knowledge, on the value of vulnerabilities in the black market. In [12] a very interesting insight on the legitimate vulnerability market is given; there are many difficulties in the legitimate selling of vulnerabilities to vendors, because of the 'secretive' nature of the good. The relationship between the software vendor and the security researcher, especially if independent and external to the company, can be trouble[1]: the vendor may indeed not appreciate the bad publicity that the disclosure of a vulnerability earns him [16].

[1] http://news.cnet.com/8301-27076_3-57320190-248/apple- boots- security-guru-who-exposed-iphone-exploit/

The market of attacks. The economical value for the attacker seems significant: in [13] Franklin et al. investigate the amount of financial and economical-related information that circulate in the market; they calculate the value of the market *of the credit cards only* to be about 37million USD; if one considers bank data theft and identity theft, their estimation increases up to 93million USD. The magnitude of these estimates is also confirmed in [17]. Motoyama et al. in [14] study the dynamics of underground forums, in which these data are actually traded, and show the high interest the criminals put in online payment accounts and stolen financial information.

The market of exploits. Lately spam has become a way to diffuse malicious links that drive the user toward domains controlled by the attackers, that can then try (and often succeed) to exploit their systems; very diffuse vectors for such an activity are porn sites [18], botnets [15,19,20], and social networks [21]. Once the attacker gets access to the victim's machine, he can install keyloggers or any kind of malware that will provide him with the victim's private data or 'permanent' access to the machine, at his will.

The profile of the coders that write exploits may vary a lot, ranging from security enthusiasts to professionals. Some coders put a lot of effort in efficiently exploiting vulnerabilities; these 'efficient' exploits are featured in web applications with a MySql backend; the community calls them *Exploit Kits*.

It is my opinion that the Exploit Kits phenomenon can shed some light on the exploiting economics underlying the whole data-theft market, and due to some of its peculiarities can perhaps be of great value in better evaluating *effective risk*. Moreover, it provides preliminary evidence that exploiting activities are governed by an economical process not yet investigated by the scientific community.

There are many different Exploit Kits on the market, very often advertised in underground forums such as *exploit.in* and *vendors.pro*. Examples of these are Phoenix, Eleonore, Blackhole, Crimepack. Exploit kits are rented to the interested attacker for different periods of time, usually up to an year; a one-year license would cost from 1000USD to 2500USD[2]. Perhaps the most popular Exploit Kit around is now Blackhole[3], but Phoenix and others have a significant market share too. Their coders seem to put a lot of effort in code obfuscation and encryption[4]. Even more importantly and perhaps counter-intuitively, but supporting the hypothesis that exploitation is driven by economical processes, the number of exploited vulnerabilities in these packs is in the order of ten or less, and many of them are very old ones.

As an example, these are the softwares exploited in Eleonore v1.6.5, released in March 2011[5] featuring only 10 exploits, most of which are at least 1-year old and two of which are 5+ years old: MDAC(2006), WMI Object Broke (2006),

[2] http://malwareint.blogspot.com/2010/01/state-of-art-in-eleonore-exploit-pack.html

[3] http://dvlabs.tippingpoint.com/blog/2011/04/26/blackhole-exploit-kit

[4] http://research.zscaler.com/2011/02/blackhole-exploits-kit-attack-growing.html

[5] http://exploit.in/forum/index.php?showtopic=46653 (account required to access the page; the reader might want to use a TOR network or a secure proxy to access the page, depending to whom belongs the IP used)

Snapshot (2008), IEpeers (2010), HCP (2010), PDF libtiff mod v1.0 (2010), Flash <10.2 (2011), Flash < 10.2.159 (2011), Java Invoke (2010), Java trust (2010). Analogous are Blackhole's[6] and Phoenix's[7] offerings, as many others' too[8]. The vulnerabilities in those Exploit Kits concern a small set of widely diffused and exposed softwares such as Java, Flash, or Adobe Reader plugins; while at the time of writing Java seems to be the main target in the most diffused exploit kits[9] (Blackhole, Phoenix), in the past were mainly targeted Office Plugins and Flash[10], suggesting there might be additional, software-related trends in the process. Exploit kits are advertised by screenshots and exploiting success rates[11].

The actual exploitation takes place when the victim requests the, say, 'exploit.php' page on the attacker's domain[12]. The attacker must fool the user in requesting that web-page: apart from social engineering and direct link spam techniques, the attacker usually compromises one or more hosts (often by means of SQL Injection) and insert an iframe in the domain's homepage that redirects connections towards the attacker's 'exploit.php' page; this is a very common practice, as evidenced by sites such as Malware Domain List[13] that serve as a database of hosts that have been compromised. Once the victim reaches the attacker's host, a set of exploiting scripts is run; as a consequence, the successful attacker can often execute code on the target machine: install keyloggers, steal data, download malware and/or make the machine part of his botnet. In order to increase the hit rate, the compromised sites might be acquired by somebody else; the attacker could (and this may not be an exhaustive list):

- buy a set of hosts compromised by somebody else
- rent connections to compromised hosts from whom acquired them
- rent connections from traffic brokers[14,15] that buy traffic from some third party (2-6USD per 1k connections).

In particular, the second approach is made easier by the existence of traffic dispatchers (e.g. SimpleTDS[16]), and often augmented by botherders themselves [22]; the third is widely diffused in pay-per-click(-install) scenarios such as porn networks [18] and others [22]: the traffic from a compromised host is sold by the 'compromiser' to the traffic broker, which will then receive a certain amount of connections from victims that accessed the compromised host. These

[6] http://exploit.in/forum/index.php?showtopic=41662

[7] http://exploit.in/forum/index.php?showtopic=37627

[8] http://vil.nai.com/images/FP_BLOG_100527_1.jpg

[9] http://www.kaspersky.com/about/news/virus/2011/Ja va_the_Target_of_Choice_for_Exploit_Kits_in_2011

[10] https://threatpost.com/en_us/blogs/carberp-and-black-hole-exploit- kit-wreaking-havoc-120511

[11] http://malwareint.blogspot.com/2010/09/black-hole-exploits-kit-another.html

[12] http://blog.imperva.com/2011/12/deconstructing-the-black-hole- exploit-kit.html

[13] http://www.malwaredomainlist.com/

[14] http://www.trafficshop.com/

[15] http://www.trafficholder.com/

[16] http://www.simpletds.com/

very connections are redirected by the traffic broker to his clients, that in turn have bought a certain amount of 'traffic' that will be directed straight to the 'exploit.php' page under their control [18, 22].

4 A research plan

From this preliminary analysis, the exploitation market results far from being simply driven by enthusiasts, unorganized hackers or groups of hackers: there is a whole infrastructure supporting both the exploitation of vulnerabilities and the economic investment that the attacker must (and apparently actually do) sustain. This gives preliminary evidence and, to my opinion, a very good reason to further investigate the dynamics of exploitation and the attackers' goals, in order to provide insights about actual security and perhaps, eventually, to better evaluate security metrics, countermeasures, risk assessment and to support vendors' patching behavior.

My research goal is to find a novel, more precise way to describe vulnerability exploitation, and thus to evaluate the *effective risk factor* affecting a system. In order to accomplish that, I formulate the following three hypotheses:

- **Hypothesis (1).** Attackers are economically rational.
- **Hypothesis (2).** There is a substantial difference in success rates between public and commercial exploits.
- **Hypothesis (3).** Commercial exploits are not redundant (i.e. not many exploits exist in the same time-frame for the same system configuration).

Therefore by (2) higher risk would come from those vulnerabilities for which a commercial exploit exists; if (1) holds, then the most dangerous vulnerabilities will be those that are efficiently exploitable, because those would optimize the exploitation success rate and thus maximize attackers' return on investment. Following (3) vulnerabilities that provide access to a certain system configuration for which other, easier or more efficiently exploitable vulnerabilities exist would represent a lower risk because of less interest to the attacker.

I'm planning to investigate those hypothesis during my Ph.D. program here at the University of Trento. **Hypothesis 3** can be validated by analyzing hackers' exploitation resources; I'm planning to further understand how much diffused those tools are as attack vectors. I'm also willing to understand who is behind their development and how profitable this activity is. **Hypothesis 2** will involve testing the efficacy of publicly released exploits against the ones featured in exploitation tools from (3). *Dulcis in fundo*, **Hypothesis 1** will be the toughest one to investigate: to collect evidence of the importance of the economic aspects in the attacking process may not be sufficient; I'm planning to conduct interviews with (professional) hackers and to design and deploy a social experiment with the purpose of better understanding how much effort the attackers are willing to put into the exploitation of a system.

The validity of those hypotheses could smooth the way toward a more precise and realistic risk assessment process, and significantly improve security metrics's reliability, patching priorities, and system hardening efficiency and efficacy.

References

1. M. Howard, J. Pincus, and J. Wing, "Measuring relative attack surfaces," *Comp. Sec. in the 21st Century*, pp. 109–137, 2005.
2. P. K. Manadhata and J. M. Wing, "An attack surface metric," *TSE*, vol. 37, pp. 371–386, 2011.
3. I. Kotenko and M. Stepashkin, "Attack graph based evaluation of network security," in *Proc. of CMS'06*, ser. LNCS. Springer, 2006, vol. 4237, pp. 216–227.
4. L. Wang, A. Singhal, and S. Jajodia, "Measuring the overall security of network configurations using attack graphs," in *Proc. of DAS'07*, 2007, pp. 98–112.
5. P. Mell and K. Scarfone, *A Complete Guide to the Common Vulnerability Scoring System Version 2.0*. CMU, 2007.
6. A. Ozment, "Improving vulnerability discovery models," in *Proc. of QoP'07*, ser. QoP '07. New York, NY, USA: ACM, 2007, pp. 6–11.
7. W. Arbaugh, W. Fithen, and J. McHugh, "Windows of vulnerability: a case study analysis," *Computer*, vol. 33, no. 12, pp. 52 – 59, 2000.
8. S. Frei, M. May, U. Fiedler, and B. Plattner, "Large-scale vulnerability analysis," in *Proc. of LSAD'06*. ACM, 2006, pp. 131–138.
9. A. Arora, R. Krishnan, A. Nandkumar, R. Telang, and Y. Yang, "Impact of vulnerability disclosure and patch availability-an empirical analysis," in *Proc. of WEIS'04*, 2004.
10. S. Ransbotham and S. Mitra, "Choice and chance: A conceptual model of paths to information security compromise," *ISR*, vol. 20, 2009.
11. B. W., H. M., H. A., and H. C.David, "2011 data breach investigation report," Verizon, Tech. Rep., 2011.
12. C. Miller, "The legitimate vulnerability market: Inside the secretive world of 0-day exploit sales," in *Proc. of WEIS'07*, 2007.
13. J. Franklin, V. Paxson, A. Perrig, and S. Savage, "An inquiry into the nature and causes of the wealth of internet miscreants," in *Proc. of CCS'07*, ser. CCS '07, 2007, pp. 375–388.
14. M. Motoyama, D. McCoy, S. Savage, and G. M. Voelker, "An analysis of underground forums," in *Proc. of IMC'11*, 2011.
15. C. Kanich, C. Kreibich, K. Levchenko, B. Enright, G. M. Voelker, V. Paxson, and S. Savage, "Spamalytics: an empirical analysis of spam marketing conversion," in *Proc. of CCS'08*, ser. CCS '08. ACM, 2008, pp. 3–14.
16. R. Anderson and T. Moore, "The economics of information security," *Science*, vol. 314, p. 610, 2006.
17. A. Cárdenas, S. Radosavac, J. Grossklags, J. Chuang, and C. Hoofnagle, "An economic map of cybercrime," in *Proc. of TPRC'09*, 2009.
18. G. Wondracek, T. Holz, C. Platzer, E. Kirda, and C. Kruegel, "Is the internet for porn? an insight into the online adult industry," in *Proc. of WEIS'10*, 2010.
19. B. Stone-Gross, M. Cova, L. Cavallaro, B. Gilbert, M. Szydlowski, R. Kemmerer, C. Kruegel, and G. Vigna, "Your botnet is my botnet: analysis of a botnet takeover," in *Proc. of CCS'09*. ACM, 2009.
20. J. B. Grizzard, V. Sharma, C. Nunnery, B. B. Kang, and D. Dagon, "Peer-to-peer botnets: overview and case study," in *Proc. of HOTBOTS'07*, 2007.
21. T. Kurt, G. Chris, P. Vern, and S. Dawn, "Suspended accounts in retrospect:an analysis of twitter spam," in *Proc. of IMC'11*. ACM, 2011.
22. J. Baltazar, "More traffic, more money: Koobface draws more blood," TrendLabs, Tech. Rep., 2011.