

Using a Mixed Data Collection Strategy to Uncover Vulnerability Black Markets

Jaziar Radianti

University of Agder
jaziar.radianti@uia.no

Eliot Rich

University at Albany
e.rich@albany.edu

Jose. J. Gonzalez

University of Agder
jose.j.gonzalez@uia.no

Abstract

Information security researchers hypothesize that black markets exist for the trading of software vulnerabilities and zero-day exploits. Such markets would encourage the development and exploitation of vulnerabilities through direct attack, malware spread or extortion. It is hard to assess the presence of vulnerability black markets and their associated transactions, as they are naturally hidden from general view, with only insiders knowing where and how to locate potential buyers and sellers. Nevertheless, the importance of the possible presence of such markets to software security makes their study requisite.

After reviewing evidence for the notion of the vulnerability black markets, we conclude that there seems to be a credible case for their existence. Next, we present data collection and modeling methods that may be used to explore this area further. In the absence of observable black market transactions, we define a “closed anchor” approach that assumes that transactions exist, and look for indirect data of their influence on observable data. To this end we have begun an analysis of online vulnerability trading sites that apparently support underground activities. Archival analysis of existing Internet sources may be used to create a common vocabulary and semantic structure for conversations about vulnerability risks and markets. In addition, we posit that interviews with participants in these sites would be useful to identify causal models that support or refute the ongoing viability of black markets. Our ultimate goal is the integration of the survey and interview data into a simulation model that may shine light onto the effects of these hidden activities.

Keywords: Mixed Data Collection Strategies, System Dynamics, Software Vulnerability, Vulnerability Black Market, Simulation

Introduction

The possibility of illicit software vulnerability trading is generating interest among security researchers. Successful vulnerability black markets (VBM) would likely lead to increased threats by creating economic and social incentives to discover new vulnerabilities and exploits, which in turn would increase the exposure of computer systems to malicious activities. Understanding the motivations for participating in these markets and collecting statistics about such activities are important for threat assessment and mitigation.

It is apparent that malicious attacks are growing, apparently due in part to the transfer of vulnerability information. IBM (2007) links underground exploit sales and markets for Web-browser exploits to the apparent growth in targeted attacks against specific customers and sites. PandaLab (2007) describes the trading of malware kits, with price quotes attached. These data indicate indirectly that there are software developers and black hat attackers exchanging information about targets and tools. Such information exchange is the core function of a VBM.

VBM activities are elusive and it is difficult to get accurate information about black market transactions. Individuals engaged in these activities are reluctant to be identified and interviewed. The extent of their presence and economic viability is largely speculative. Accordingly, an effort to investigate the VBM systematically cannot rely on direct market observation. This paper proposes an indirect approach using archival information and textual analysis to identify the broad dynamic characteristics of these markets. We demonstrate this approach by applying it to an online source that purports to trade vulnerability zero-day exploits. These efforts advance our long-term research goal of establishing the presence of VBMs and, in turn, creating policies and activities that counter the effects of such a market.

Looking for evidence of vulnerability black markets

The term “black market” originally described underground economic transactions during the Second World War (Clinard 1969). While individual transactions may be hidden from view, their aggregate role in an economy may be quite significant. Accordingly, macroeconomists estimate the size of illegal activities within a nation’s GDP as the discrepancy between estimated expenditures and reported income (Bajada et al. 2005). The *Global Index of Illicit Markets* (Havocscope 2007) provides estimates of commodity black markets and their effect on the world economy. These studies start with a theory of how legal markets should perform, and then attribute gaps between various collected statistics to extra-legal activities. We will return to this principle – which we term the “closed anchor approach” – when we propose an analogous procedure for a model-based assessment of the VBM (p. 4).

Security practitioners are confident that VBMs exist. Miller (2007) notes that “there has long been a black market for computer exploits and...4 security researcher may choose to sell the vulnerability information on the black market, but faces potential criminal prosecution for such action...”. Naraine (2006) observes that zero-day exploits are becoming more accessible and interesting to spammers and criminals, perhaps as means to disrupt and extort specific targets. There are also reports of attacks motivated by political disputes. While it is arguable that few individuals have the skills needed to develop exploits, VBMs turn specialized knowledge into a procurable commodity.

Theoretical work on vulnerability black markets has so far been limited largely to scholars in the economics of information security. These authors focus on the mechanisms that should be in

place to support a legal market, to establish price mechanisms, and to counter market failures. In response to the apparent need to create an open market, Böhme (2006) advocates the development of an auction strategy, where “an adversary would have an incentive to report the bug instead of exploiting it or selling it on the black market.” Ozment (2004) proposes a Dutch Auction model, among others, for vulnerability white markets, where sellers expose their goods for sale with an initial offering price that can only be lowered, ensuring that vulnerabilities are put into markets quickly. One unintended consequence of institutionalizing vulnerability markets is the creation of economic incentives for vulnerability detection and distribution. Ozment (ibid) finds the possible resale of unpublicized vulnerabilities particularly troubling, concluding that there are few solutions to this dilemma; he is optimistic that having auctions for vulnerabilities would reduce the incentives for resale.

Sutton and Nagle, both from iDefense Labs, present two models for underground vulnerabilities, focusing on revenue streams rather than transactions (2006). In the *contracted model*, a malicious actor hires a hacker to find vulnerabilities in specific software targets. However, Sutton and Nagle emphasize that there is little public information on active uses of the contracted model. In the *purchase model*, a hacker starts by identifying vulnerability, creates an exploit, and sells it to malicious actors. Sutton and Nagle emphasize that all parties have to broker the deal, involving some potentially risky contracts, while making sure that they are not caught by law enforcement. Naraine (2006) reports on a purchase model transaction, where the Microsoft Windows WMF vulnerability was discovered by a vulnerability researcher and sold on the underground market to malicious actors.

Security companies have created their own vulnerability marketplaces. An example is the Zero Day Initiative (ZDI) launched by TippingPoint. This program is designed to pay researchers and would-be hackers for data on product vulnerabilities. This provides a counter-argument to those who claim that security researchers are not properly compensated (Schechter 2002). Brokered markets have several advantages over direct contacts; for example they can prevent resale of vulnerability information to malicious individuals. The opportunity for a trusted market decreases the chance that overt vulnerability research will be marginalized and moved further underground. An active and sustained marketplace, such as ZDI, where vulnerability exploits are traded, is also a good indication of the viability of a parallel VBM (Stone 2007).

Evidence of markets and Prices

Our observations of Internet sources provide additional, though speculative, support for interest in VBMs. The dialogue on electronic news surrounding vulnerabilities often includes estimated prices for hacks. Several authors have also speculated about prices for specific vulnerabilities (Table 1). Later in this paper, we will briefly discuss some initial results from the observations

Case	Price	Source
A Vista flaw on Romanian Web forum	\$ 50,000	Trend Micro in Naraine (2006)
WMF selling underground	\$ 4,000	Higgins (2006); Naraine (2006).
Offers for zero-day flaws in the gray market	\$5,000 and \$20,000	Raimund Genes, Trend Micro from Lemos (2007)
A ‘weaponized’ exploit in the black market	\$20,000 to \$30,000	David Maynor, SecureWorks from Higgins (2006)
Zero-day vulnerabilities on the Internet black market	\$25,000	Landesman (2007)

Table 1: Estimated black market vulnerability prices

on one website used for conducting vulnerability trading. If we assume that the price speculations below are accurate, then the price on the black market is relatively high compared with the price in the legitimate market. Legal markets, such as ZDI, offer \$2,000-\$10,000 for a vulnerability purchase.

There are also discussions about how black hat hackers meet buyers. Public bulletin boards and private online chat rooms are available to link criminals with technologists. Stone (2007) reports that buyers appear to solicit developers directly. There is a substantial amount of evidence that Internet resources are used for executing illicit acts, including financial fraud and identity theft.

News reports about VBMs

Our review of the literature shows that much of the public reporting about black markets is grounded in news reports or interviews with expert informants. The informants include security researchers, both named and unnamed (Higgins 2006); personnel from well-known security companies (Whipp 2006); a virus analyst (Naraine 2006); a founder of a small security company (Greenemeier 2007) and a pseudonymous hacker who claims to engage in illegal activities. Some texts refer the VBM issue to the reports made by some well-known security companies (ibid). Clearly these works need to be regarded with some skepticism, as the potential conflicts of interest are apparent. Sometimes the information is incomplete, partial, or based merely on expert opinion.

We conclude that software VBMs are likely to exist, although very few of the specific claims about them are based on formal research and rigorous data collection. The occult nature of the problem motivates us to investigate the problem in an indirect but scientifically defensible manner by combining structural modeling with the available data.

A systems perspective on VBMs

We hypothesize that VBMs arise not from one single origin, but instead emerge as the consequence of flawed software and from the complex interplay of individuals who want to fix or exploit those flaws. Radianti and Gonzalez (2006; 2007) propose a System Dynamics model that links the discovery and trading of vulnerabilities among hackers, attackers and software vendors to changes in the quality and security of software products over time

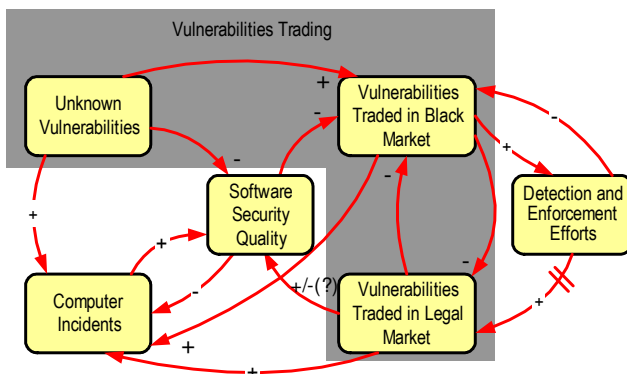


Figure 1: VBM Model Overview

(Figure 1). The shaded area of the diagram denotes the activities of software vulnerability researchers (both black and white) to discover and exploit their findings. The exposure of vulnerabilities in the black market is expected to increase the number of computer incidents, security breaches, and problems facing vendors. The exposure of vulnerabilities in the white market also contributes to a growth of computer incidents. Unlike black market activities, however, white market exposures may also directly stimulate vendor efforts towards improving the quality of its software

before incidents occur in the marketplace. This asymmetry is important to the effectiveness of legal vulnerability markets. Another important effect may be the presence of law enforcement efforts to locate and prosecute black market participants. This effort would contribute to an increasingly risky environment for black market researchers, stimulating them to participate in the legal market instead. The detailed simulation of this model raises concern about secondary effects of the establishment of white markets to combat black markets. For example, an active marketplace for vulnerability discovery creates additional pressure on software vendors to develop patches. If these vendors do not have the capacity to complete this work, then discovered vulnerabilities may remain exposed. Further discussion of the simulation is found in Radianti and Gonzalez (2007).

Developing a dataset

We propose two linked approaches to constructing theory and collecting data that capture the little we know about black markets. The first is the use of a “closed anchor” paradigm. As it may not be possible to directly observe black market activities, we employ an anchoring assumption that these markets exist based in part on the literature cited earlier. Once the anchor is set, we then consider the implications of such a market on computer security. In particular, we can attempt to estimate the scale of unobservable activities by considering how multiple observable indicators behave compared with behaviors predicted by theory. This approach, while certainly fraught with uncertainty, takes the discussion beyond a debate of the presence of VBMs. Instead, we may consider whether the problems VBMs create are important enough to attempt to control, and if these controls may create other unintended consequences.

The closed anchor approach requires both a formal statement of theory and the incorporation of whatever credible evidence exists, supplemented with reasonable assumptions. The formal statement of theory is presented through the simulation model discussed above. The collection of evidence is based on established methods employing multiple qualitative sources. As the bulk of the available data comes from text, we propose to perform archival study, anonymous surveys, and interviews with selected security researchers to build our dataset. The archival analysis develops a repository of raw material for analysis. We propose to search Internet security sites and the general text of the Internet for keywords related to the problem domain. This material will likely contain a great deal of chaff, but remains a starting point for discussion. In parallel, we propose to interview some researchers who have written on this topic as well as overt contributors to vulnerability analysis. The objective is to identify various possible causal mechanisms that may support or negate the existence of vulnerability markets. We will have to rely on formal and defensible mechanisms to link information in the text repository. We would also attempt to contact researchers who actively contribute to the various public vulnerability databases, with the hope of developing a pool of informants through snowball techniques.

We are aware of the weaknesses this proposal faces. The current set of experts face similar problems to ours, and may have formed their opinions based on hearsay. The sample frame may or may not be representative. Participants in Internet-based data collection may provide misleading information, choosing to mask their work or create entire personas for their own reasons. We concur with Axinn & Pearce (2006) that mixed method strategies and multiple sources of information will reduce error and counteract biases inherent in every type of data collection. In this case, though, there may well be so much noise that it will be hard to find consistent insights.

Some Preliminary Results

The existence of a VBM presupposes that an economic marketplace exists for software vulnerabilities. Such a market requires sellers and buyers, a mechanism to bring these actors together, and some medium of exchange. In addition, the presence of a black market means that the white market is unable to satisfy the needs of all buyers and sellers. We began our data collection by observing the contributions and discussions on hacker websites that features an explicit Black Market (BM), marketplace or trading forum. Over the last year, we identified nine Internet websites that purport to trade exploits and different types of malware (Table 2). Some of these forums have stopped accepting new postings. Others were taken offline before basic characteristics were read.

Code	No. of Threads (as of)	Earliest Post in	Hosting Country	Forum Name	Primary Language
W1	686 (24 Aug 07)	2006-04-20	Cayman Islands	Black Market	English
W2	560 (Nov 9 2007)	2007-06-18	USA	Marketplace	English
W3	45 (Nov 7 2007)	2007-08-02	USA	Advertise	English
W4a	Used to have a BM forum, now unavailable				
W4b			Australia	No BM	English
W5	82 (N/a)	N/A	USA	Forum down	English
W6	387 (Nov 7 2007)	N/A	Germany	Buy-Sell-Trade	English
W7	5 (Nov 7 2007)	2007-10-31	USA	Black Market	English
W8	17 (Nov 7 2007)	2006-07-14	India	Money	English

Table 2: BM Forums

Our initial analysis focuses on the site we code as W1. This site has an explicit black market forum with almost 700 discussion threads over the period of April 2006 to July 2007. This forum has the largest number of postings available for analysis and the longest history of postings, though the forum appears to have become dormant. Data was extracted from the site's public interface, without other access to the server functions.

We have begun categorizing thread contents by activity (e.g., buy, sell, or trade), posting date, membership rank, goods offered, contact methods, price offered, and payment method, if any. Our initial analysis covers the first six months of the forum, October 2006 through March 2007. During this period we observe increasing membership, growing varieties of offered or solicited exploits. In addition, the text indicates the emergence of social norms among participants. The analysis does not yet include activities from April through July 2007.

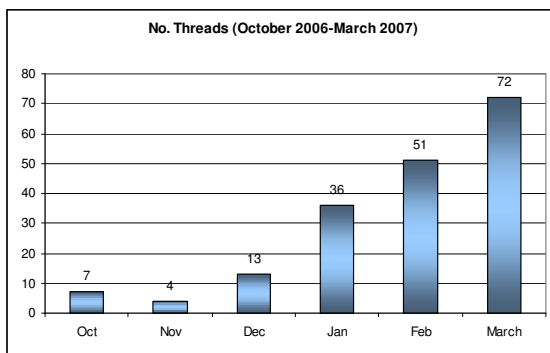


Figure 2:
No. of new threads in the W1 BM Forum
Source: Observation of a hacker forum

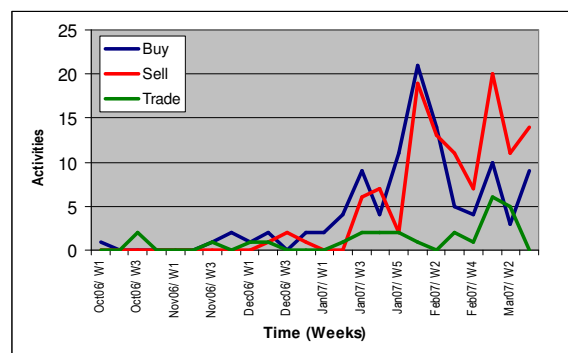


Figure 3:
Buy, Sell and Trade trends in the W1 BM forum
Source: Observation of a hacker forum

These first months of the forum included 235 threads. We eliminated threads that dealt with general questions or instructions from the moderator, leaving 183 threads for textual analysis. From April 2006 through September 2006, this forum contained a general discussion of black market activities. Starting in October, however, posts to the forum started to solicit or offer explicit vulnerability trading. In the following six months of data the number of threads increased at an exponential rate (Figure 2). This may be attributed to increased interest in the black market, the idea of vulnerability trading, or may have resulted from changes in moderator policies.

Specific offers or other transaction-related activities in the threads also showed a marked increase in the early months of 2007 (Figure 3). Perhaps not coincidentally, Microsoft introduced its Vista operating system in late 2006 to businesses and to customers in late January 2007.

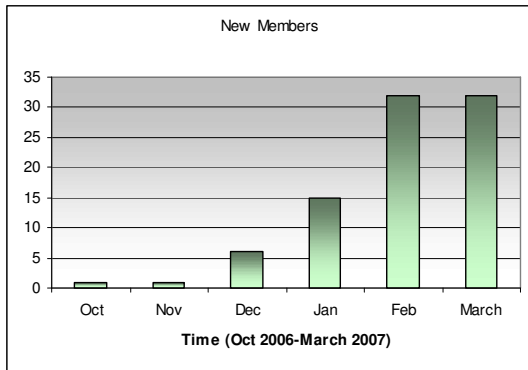


Figure 4
New Members Joining the BM Forum
 Source: Observation of a hacker forum

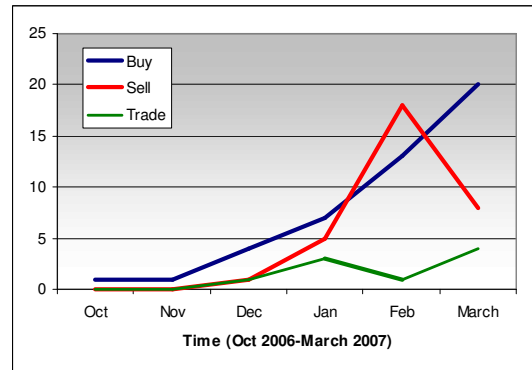


Figure 5
New Member Activities in BM
 Source: Observation of a hacker forum

The number of new members that actively purport to sell, to buy or to trade in the forum rises in lockstep with the number of threads (Figure 4). A report from one well-known newspaper citing this site as potentially performing illegal activities around the same time seems to have triggered interest in the forum. It is more difficult to track exiting members, so we do not know what the number of current and active members may be. In response to concerns about the credibility of new participants, senior members of the forum have recently applied entry barriers to enrollment. When new member activity is disaggregated, it appears that putative buyers have increased steadily, while sellers have a less consistent pattern (Figure 5). The data indicate that the February new members were attempting to sell multiple exploits during that short period. Though this data must be taken with some degree of skepticism, being over a short period and from a single forum, it does imply that there is increasing interest in the area.

Once interest in an exchange is established, many participants go off the public areas of the forum. In many instances, real-time contact information for use with Internet Relay Chat and instant messaging information is shared. The W1 forum also supports a private message feature for discreet conversations. It is difficult to know which proposed exchanges are actually consummated. There is also some evidence of a viable medium of exchange. Most messages that solicit vulnerability exchanges describe the offering and leave pricing to the offline discussion. Payments are commonly offered through e-Gold (www.e-gold.com), an online currency backed by precious metal reserves. While the details of the final transactions are not in the public view, these messages are evidence of interaction between buyers and sellers.

Analysis of the social interactions among the participants provides additional insight into the world of VBMs. It appears that the reputation of buyers and sellers is an important consideration. A buyer or seller with a blemished history will be identified by other users when they return to the forum. It is not easy to establish a reputation as a legitimate buyer or seller. In this forum, seller legitimacy is established by sharing the purported exploit with a respected and neutral party who assesses the exploit's quality. Some sellers are reluctant to share their wares with these neutral parties, fearing that their intellectual property will be distributed without compensation. Written and unwritten rules and group ethics have developed among the members concerning verification, credibility, membership hierarchy, authorities of higher-rank membership, and social contacts. Users build credibility by regularly sending useful posts that contribute to the general discussions, and by buying, selling, or trading genuine tools and exploits in the forum. Verification procedures are introduced to ensure the quality of goods. Access to some forums is limited to users with a high number of postings in order to keep out newbies and casual visitors. The forum moderators actively monitor postings, and have banned users that break established rules.

These findings are incomplete and are based on examination of a single forum. In this forum it appears that the three criteria for a potentially viable VBM have been met. We have observed the presence of buyers and sellers, identified mechanisms for them to meet, and noted a viable medium of exchange. We have not observed the completion of a transaction, nor should we expect to, as such visibility might well lead to prosecution of the author of a maliciously-applied vulnerability. The recent dormancy of W1 may be due to some enforcement activity or the desire for the moderators to maintain their occult status. Efforts to collect data need to include a continually expanding number of sites, as well as their content, in the event that some sites are being cut off.

Conclusions and Future Research

Some scientists consider the evidence of a VBM as insufficient and deem its existence as rumor. Indeed, our review of the existing sources shows that we do not yet have data that would fully satisfy scientific standards. However, a legal vulnerability market does exist, and it seems to owe its existence to a firm belief in some minds that the VBM is a real threat. We surmise that their belief is grounded in actual perceptions of the black market and that actors in the legal market may wish to disguise their participation in black market transactions out of commercial interest. Although data about a potential black market for software vulnerabilities is scarce and not yet trustworthy, we join the ranks of researchers who believe in the existence of this trading forum and share their concerns.

We have delineated two complementary approaches to develop better theory and elicit better data about VBMs. The use of a closed anchor model gives us a mechanism for inferring the behavior of vulnerability markets. The careful examination of postings from websites purporting to host discussions about VBMs provides additional insight into how illicit transactions may occur. Over the coming months, we plan to supplement this data with a round of expert interviews with security researchers and practitioners in order to review and criticize the model structures (Luna-Reyes et al. 2003) and create a common vocabulary and semantic structure concerning vulnerability risks and markets. Directed surveys can follow to develop better knowledge regarding market structures and participation. In addition, we are building elaborated system dynamics models with a black market superimposed on a legal market of software vulnerabilities. Simulations of scenarios will result in trends of observable parameters. Comparison with empirical data would allow inferences about the attributes of the

black market. In this way we will employ both simulation and grounded data to learn more about this challenging problem.

References

- Axinn, W.G., and Pearce, L.D. (2006) *Mixed Method Data Collection Strategies* Cambridge University Press, Cambridge.
- Bajada, C., and Schneider, F. (2005) *Size, Causes and Consequences of the Underground Economy: An International Perspective* Aldershot.
- Böhme, R. (2006) "A Comparison of Market Approaches to Software Vulnerability Disclosure," International Conference, ETRICS 2006, LNCS 3995 Springer-Verlag Berlin Heidelberg, Freiburg, Germany, pp. 298-311.
- Clinard, M.B. (1969) *The Black Market: A Study of White Collar Crime* Patterson Smith, Montclair, New Jersey.
- Greenemeier, L. (2007) "A Security Researcher Gets Offered The Big Score," <http://www.informationweek.com/story/showArticle.jhtml?articleID=197004915> Retrieved February 12, 2007.
- Havocscope (2007) "Global Index of Illicit Markets," <http://www.havocscope.com/>, Retrieved September 10, 2007.
- Higgins, K.J. (2006) "Bucks for Bugs," http://www.darkreading.com/document.asp?doc_id=99518, Retrieved July 20, 2007.
- IBM (2007) "IBM Internet Security Systems X-Force 2006 Trend Statistics," Report XF-XFORCEEXECBRIEF-0107, IBM Internet Security Systems, Atlanta, GA http://www.iss.net/documents/whitepapers/X_Force_Exec_Brief.pdf.
- Luna-Reyes, L.F., and Andersen, D.L. (2003) "Collecting and analyzing qualitative data for system dynamics: Methods and models," *System Dynamics Review* (19:4), pp 271-296.
- Miller, C. (2007) "The Legitimate Vulnerability Market: the Secretive World of 0-day Exploit Sales," Workshop on Economics of Information Security, Pittsburgh, USA.
- Naraine, R. (2006) "Researcher: WMF Exploit Sold Underground for \$4,000 " <http://www.eweek.com/article2/0,1895,1918198,00.asp>, Retrieved September 15, 2007.
- Ozment, A. (2004) "Bugs Auctions: Vulnerability Market Reconsidered," Workshop of Economics and Information Security (WEIS), Minneapolis, MN.
- PandaLabs (2007) "Quarterly Report PandaLabs," (April-June 2007) <http://www.pandasecurity.com/>.
- Radianti, J., and Gonzalez, J.J. (2006) "Toward a Dynamic Modeling of the Vulnerability Black Market," Workshop of Economic of Securing Information Infrastructures, Washington, D.C.
- Radianti, J., and Gonzalez, J.J. (2007) "A Preliminary Model of The Vulnerability Black Market," the 25th International System Dynamics Conference, Boston, USA.
- Schechter, S. (2002) "How to Buy Better Testing: Using Competition to Get The Most Security and Robustness for Your Dollar," Infrastructures Security Conference.
- Stone, B. (2007) "Moscow Company Scrutinizes Computer Code for Flaws," <http://www.iht.com/articles/2007/01/29/business/bugs.php>, Retrieved April 28, 2007.
- Sutton, M., and Nagle, F. (2006) "Emerging Economic Models for Vulnerability Research," The Fifth Workshop on the Economics of Information Security (WEIS), Robinson College, University of Cambridge, England.
- Whipp, M. (2006) "Black Market Thrives on Vulnerability Trading " <http://www.pcpro.co.uk/news/84523/black-market-thrives-on-vulnerability-trading.html>, Retrieved March 2, 2007.