DRAFT

The legitimate vulnerability market:

the secretive world of 0-day exploit sales

Charles Miller, Ph.D, CISSP

Independent Security Evaluators

cmiller@securityevaluators.com

Trading of 0-day computer exploits between hackers has been taking place for as long as computer exploits have existed. A black market for these exploits has developed around their illegal use. Recently, a trend has developed toward buying and selling these exploits as a source of legitimate income for security researchers. However, this emerging "0-day market" has some unique aspects that make this particularly difficult to accomplish in a fair manner. These problems, along with possible solutions will be discussed. These issues will be illustrated by following two case studies of attempted sales of 0-day exploits.

Introduction

There has long been a black market for computer exploits. For a long time, hackers were content to trade or sell exploits amongst themselves, mostly for prestige. Computer security researchers normally followed "responsible" disclosure which entails contacting the vendor and usually receiving acknowledgment when the vulnerability was announced along with the supplied patch. In the last few years, the market for 0-day exploits, those for which there is no available patch, has begun to migrate into the commercial space.

A few companies openly pay researchers for this information [1] [2]. For example, the Zero Day Initiative program, run by TippingPoint, a division of 3Com, has 532 registered security researchers [3]. In 2006, they purchased 82 vulnerabilities from security researchers and disclosed 57 [3]. Likewise, iDefense has a large number of participants and has released many advisories. While these companies openly pay researchers for their discoveries, their primary business is not in the buying and selling of vulnerabilities and so they have the least incentive to offer a large sum for a particular vulnerability.

Certain companies sell tools or packages which contain zero day exploits [4] [5] [6], while some others purport to be able to broker deals between researchers and U.S government agencies [7]. Likewise, the illegal market for these exclusive tools has begun to become more economically based as spammers and criminals become interested in the use of 0-day exploits for use in illegal activity [8].

As a computer security researcher, there are many options available after discovering a vulnerability in a high-profile application or operating system. She may choose to report the vulnerability to the vendor, or simply announce it publicly without vendor notification. Such a choice may be made in order to increase her reputation or add to her resume. She may choose to sell the information on the black market, but faces potential criminal prosecution for such an action. Finally, she may choose to attempt to sell this information to a legitimate buyer. Such legal buyers may include government agencies, commercial tool suppliers, large penetration testing and consulting firms, intrusion detection companies, and subscription services. This paper documents the problems such a researcher will face when attempting to sell this vulnerability information or exploit to a legitimate buyer. It will then discuss possible solutions to some of these fundamental issues. Finally, these problems and solutions will be addressed in the context of two actual attempts at selling 0-day exploits by the author, one which ended somewhat successfully and one that did not.

Inherent obstacles

Due to the nature of vulnerability information and 0-day exploits, there are many obstacles that traditional businesses and services do not have to face. Some of these problems are outlined in this section.

Vulnerability information is a time sensitive commodity

An interesting problem concerning the commodity of vulnerability information is that it's value can go from extremely high to almost zero instantaneously. This is due to the fact the information is only valuable when it is not widely known. As soon as the vulnerability is announced or a patch is released, the vulnerability information becomes worthless. Worse yet, other factors may completely reduce a vulnerability's value, such as the introduction of a new technology. For example the target binary may be recompiled with the /GS flag, SElinux may become turned on by default, or patches for unrelated issues may change the binary in some manner which makes the vulnerability impossible to exploit. These events are usually outside of the researchers control and visibility. There is usually no way to know when another researcher or the vendor will announce the same vulnerability. Therefore, the researcher must always assume today is the last day for their issue. Literally, hundreds of thousands of dollars could be lost by waiting an extra day to complete a sale.

The implication of this problem is that any transactions involving vulnerability information must take place as quickly and discreetly as possible. Once a researcher has found a vulnerability and/or written an exploit, they must be able to quickly identify a buyer, negotiate the price, and complete the sale. In the current market, none of this is possible, as will be seen below.

No transparency in pricing

There is no publicly available information regarding the prices of different vulnerabilities for different applications on different platforms. The value of a vulnerability depends on many factors, most of which are difficult to measure. Some of these factors include

- How widespread is the use of the application that is vulnerable?
- Does the application come by default with the operating system?
- Is the application turned on by default?
- Is authentication required to exploit the application?
- How well do typical firewall configurations block access to the application?
- What versions of operating systems/application are vulnerable?
- Is the vulnerability in a server or client application?
- Is user interaction required to exploit the vulnerability?
- How difficult it is to find the vulnerability (which is a proxy measurement of how long it will be before it is discovered by someone else)?
- How many people know about the vulnerability?
- How reliable is the exploit?
- Does a single exploit work against many versions?

Research reveals some examples or vague guidelines as to the value of vulnerabilities, but not nearly enough for a seller or buyer to know the fair market value for a vulnerability or exploit. Without this information, it is impossible for the two to come to an agreement in a fair and informed manner. One side will always lose out. Below is a list of estimated values for various vulnerabilities or exploits. Some of these numbers represent "exclusive rights", while others do not.

Vulnerability/Exploit	Value	Source
"Some exploits"	$200,000 - $250,000	A government official referring to what "some people" pay [9]
a "real good" exploit	over $100,000	Official from SNOsoft research team [10]
Vista exploit	$50,000	Raimund Genes, Trend Micro [8]
"Weaponized exploit"	$20,000-$30,000	David Maynor, SecureWorks [11]

Vulnerability/Exploit	Value	Source
ZDI, iDefense purchases	$2,000-$10,000	David Maynor, SecureWorks [11]
WMF exploit	$4000	Alexander Gostev, Kaspersky [12]
Microsoft Excel	> $1200	Ebay auction site [13], [14]
Mozilla	$500	Mozilla bug bounty program [15]

With this incredible range of prices, if a researcher discovers a flaw in Internet Explorer, for example, what price should they seek? The above numbers indicate a fair price may be anywhere between $5000 and $250,000. With such vague pricing information available, there is no real way to place an accurate price to the value of such a vulnerability.

More pricing information will soon be available as the ZDI program plans to begin publicizing their sale prices in March 2007 [3]. This will at least reveal a lower bound for pricing.

Difficulty finding buyers and sellers

The current market for the legitimate sale of 0-day exploits is not openly accessible. When a researcher discovers a vulnerability and wishes to sell it, there is no centralized way to locate a buyer. The researcher is forced to "cold-call" any contacts they may have, or any companies they think might be interested. This approach has many problems. The first is that sometimes it is difficult to accurately describe a vulnerability without making the vulnerability easier to find (see below). Therefore, the researcher does not wish to inform people who otherwise wouldn't be interested. Another problem is that this process is time consuming, as even with companies that would be interested, the researcher is not likely to initially contact the right person. This time delay could ultimately make the information worthless.

Conversely, with the exception of iDefense, Tippingpoint, and Netragard, companies do not typically advertise the fact they purchase vulnerability information. Additionally, vendors do not normally pay for vulnerabilities in their own product. Perhaps if they did offer such a bounty, the security of their products would increase -- such as when Netscape offered a $1000 bug bounty program back in 1995 [16].

Checking the buyer

Due to the fact there is no centralized way to locate a buyer of vulnerability information, the researcher is often forced to try to tell many individuals about the discovery in an attempt to find a buyer. A consequence of this is a buyer may emerge with whom the researcher is not familiar. Assuming the researcher wants to sell only to a legitimate buyer, it can be difficult for the researcher to verify the buyer's intentions and avoid a trip to "Gitmo". This is complicated further by the time constraints the researcher faces.

Value can not be demonstrated without loss

One of the more fascinating problems a researcher attempting to sell vulnerability information or a 0-day exploit may face is proving the validity of the information without disclosing the information itself. The only way to prove the validity of the information is to either reveal it or demonstrate it in some fashion. Obviously, revealing the information before the sale is undesirable as it leaves the researcher exposed to losing the intellectual property of the information without compensation. Demonstrating the vulnerability via an exploit is no better. It is not possible to exploit a system in the possession of the researcher because the seller will not be able to verify that the system has not been altered in some fashion. Exploiting a system under the buyer's control is also problematic as the researcher has no way to know that the buyer is not recording the working of the exploit, by say, capturing the packets involved or monitoring the application under attack. Without the use of some trusted third party (TTP) -- which at this time does not exist in a generic fashion -- at some point in the process of the sale, one party must implicitly trust the other. Either the researcher will have to reveal the vulnerability before payment or the buyer will have to pay the researcher before they receive the information. For example, in both the VCP or ZDI programs, the exploit or vulnerability information must be given to the respective programs *before* an offer is made.

What is worse is that even a vague description of the vulnerability may be too much to reveal. As was discussed in a previous section, two factors that play into the value of a vulnerability are its reliability and whether authentication is required. However, suppose a researcher were to reveal that they had found a vulnerability in WU-FTPD, that didn't require authentication to exploit and was reliable because it was a stack overflow. Any potential buyer could quickly "rediscover" this vulnerability by using static analysis and looking at the small amount of code that is run pre-authentication and focusing on stack buffer manipulation. In this scenario, the potential buyer has no reason to pay for the information, and even if the researcher managed to sell the information to someone, it would likely quickly become public since all the people who the researcher tried to sell it to would know about the vulnerability, or could find the vulnerability with minimal effort.

Even supplying versioning information can sometimes lead to a vulnerability being revealed. For example, suppose a researcher were to reveal that a particular vulnerability affected an application for all versions more recent than version 3.21a. A potential buyer would only have to look at differences introduced in this version of the application as compared to previous versions. If there are few enough changes, they will quickly find the vulnerability information without having to pay for it. Therefore, the researcher is in the awkward position of wanting to reveal as little information as possible about the vulnerability to prevent the loss of their discovery, but needing to provide as much as possible to ensure they are paid sufficiently for it. Combined with the fact that they must try a large number of potential buyers and there is no good idea what a fair price is, the researcher is in a very poor position.

Ensuring claim to vulnerability

Worse yet, as the commodity in question is information which is not widely known, it is difficult to reveal the information without risking that the other party will claim the infor-

mation as their own. For example, if the researcher supplies the vulnerability information to a potential buyer, she has no protection from the potential buyer rejecting the sale and then attempting to sell the information as their own. This is complicated by the fact that many of these sales may be international, thus limiting the enforceability of potential contracts.

Exclusivity of rights

The final hurdle involves the idea of exclusive rights of the information. In order to receive the largest payoffs, the researcher must be willing to sell all rights to the information to the buyer. However, the buyer has no way to protect themselves from the researcher selling the information to numerous parties, or even disclosing the information publicly, after the sale. The problem is fundamental since, unlike physical commodities, the researcher cannot truly give up the information. They still possess knowledge about the vulnerability in question, even if they destroy all traces of it from their computers. In this case, they are merely agreeing not to share it.

A contract of the sale may include language which would force the researcher to return the funds in the case where the researcher reveals the information openly. However, it could prove extremely difficult to prove that the researcher had sold it to other parties or revealed it to the vendor under a different name. This leaves the buyer at the mercy of the seller to not reveal the information to others. As Dave Aitel of Immunity stated with regards to his company's policy of buying vulnerabilities, "Sometimes we get burnt, sometimes not." [11]

Possible solutions

While there is no easy solution to the problems faced by researchers trying to sell vulnerability or exploit information, there are some steps they can take to protect themselves. Most of the problems outlined above are due to the secretive nature of the marketplace. Adding transparency and organization can help to alleviate some of these problems. Organizations and companies need to step into this space and help address these issues. Below are some possible solutions.

Researcher actions

Most of the consequences of this opaque marketplace outlined in the previous section require industry wide changes in order to be addressed. There are, however, a few things that a researcher can do to attempt to protect their discovery.

In order to "prove" that the discovery is theirs, a cryptographic hash of the information can be obtained and posted to some public place. This idea has recently been discussed on mailing lists [17], [18] and was carried out by the author in the first case study. Commercial services also exist which log hashes of information in order to verify the date the information was known [19]. Such an action protects the researcher from a potential buyer attempting to take credit for the discovery without paying. This is only limited protection, however, as the potential buyer still has the information and can use it or release it at their discretion. The researcher can only prove that they were aware of

the information before the release by the buyer. This doesn't provide any income to the researcher.

Another problem the researcher faces is proving they possess a specific vulnerability. It is incredibly difficult to prove to a potential buyer that the researcher has a given vulnerability without giving all of the information away. One such possible method is as follows:

1. Buyer and seller meet in a physical location.

2. Buyer brings physical media (CD's, DVD's, etc) consisting of the operating system and any required applications.

3. Seller brings the exploit or demonstration of the vulnerability and the necessary hardware.

4. Under the seller's observation, the buyer installs and upgrades the operating system and any necessary applications using their media.

5. Using a cross-over cable, the seller exploits the vulnerability.

6. The seller retains the hardware after the demonstration.

There are many drawbacks to this method including that it is time consuming, expensive, requires physical contact, and can be circumvented with some very advanced hardware implants by the seller. But, given the constraints, this may serve as a way to prove the existence of the vulnerability without giving it completely away.

The final tool in the researcher's arsenal is the concept of "mutually assured destruction". In the case where the researcher believes a potential buyer has taken the vulnerability information and not paid for it, they can announce the vulnerability in a public manner. Such an announcement makes the stolen vulnerability worthless, thus negating the value of the taken information. Obviously, this scenario does not have the desired consequence for the researcher either. Hopefully, the mere threat will help to keep buyers honest.

Market place solutions

In his paper "Vulnerability Markets", Bohme suggests five different market types for the sale of vulnerability information; bug challenges, bug auctions, vulnerability brokers, exploit derivatives, and cyber-insurance [20]. Of these market types, the first two can benefit individual security researchers, but must be initiated by the vendor. Therefore, only vulnerability information regarding vendors participating in such a solution would be valuable under such a mechanism. Depending on the level of participation, this could deeply hamper a security researcher.

Vulnerability brokers and cyber-insurance don't have an immediate incentive for the individual researcher. It would be difficult to leverage a found vulnerability into income under these two systems.

The final market type, exploit derivatives does provide a way for a researcher to profit from vulnerability information. In this market type, a mechanism is built around contracts that pay out a defined sum in the case of a security event. Different parties would freely trade these contracts based on whether they thought such an event would occur or not. Such a system would benefit an individual researcher by allowing her to purchase a large amount of contracts that pay off for vulnerabilities found. Then she could release the vulnerability and profit from the information.

Using an exploit derivatives market mechanism avoids many of the problems from the last section. It is not entirely clear how well using exploit derivatives would compare to the current system from a pure pricing perspective, though. It may be impossible to make the amount of money possible in the currently system using this approach. One problem that might limit the amount of money a researcher could receive is that traders may observe what is occurring and attempt to profit from it. For example, if David Litchfield starts to buy derivatives that state Oracle will have a vulnerability, many people will also attempt to buy these derivatives. This will limit the financial gain that Mr. Litchfield can make from a given vulnerability. However, the biggest problem would be establishing the market and having enough participants to provide the necessary liquidity.

Direct auctions

In contrast to Bohme's concept of an auction instigated by the vendor, there is also the possibility of one started by the researcher. In 2005, a security researcher known as fearwall attempted to sell a vulnerability he uncovered in Microsoft Excel on Ebay. The highest price obtained when the auction was removed was approximately $1200 [13]. This begs the question of exactly what would have transpired had the auction finished. How does fearwall prove that he has a vulnerability to the auction winner without revealing the vulnerability information to them? How does the auction winner know that fearwall is not selling the information to many other individuals?

The idea of security researchers using auctions to attempt to get a fair market price for their discovery goes back to at least 2003 [13]. Using this approach, researchers could post the fact they found vulnerabilities or had functional exploits on an auction site and allow interested parties to bid on it. Some sort of reputational system could be used to establish whether researchers and buyers had successfully used the site which could help respected buyers and sellers to trust one another. The auction site could also offer third-party verification and escrow services that specialized in vulnerability information and exploits.

Such an approach would solve many of the problems faced in the current system. It does suffer some drawbacks, however. First, the reputational system would be limited as compared to a typical auction site, such as Ebay, due to the limited number of transactions that would take place. An individual can sell more Pez containers than find 0-

day exploits in a given year. Furthermore, it could prove difficult to verify the intentions of the buyer. Likewise, exclusivity would still remain problematic.

The biggest drawback to this system is in its questionable legal status. Noted information security attorney Jennifer Granick stated that while running such an auction site is probably legal, it would certainly be risky [21]. She stated that if someone using the site went on to commit a crime, it is possible that a particularly aggressive prosecutor could attempt to prosecute the site owner. Similar legal advice is what made Greg Hoglund decide not to pursue this route in 2005 [13].

Need for trusted third parties

The need for trusted third parties in the legitimate sale of vulnerabilities is crucial. Third parties provide a method for the seller to ensure they will be paid and the buyer to ensure they are getting the information the seller claims to have. It could also serve a role in bringing interested buyers and sellers together.

As will be seen in the case studies, without a third party, there comes a time when either the buyer or the seller must completely trust the other and give the other valuable information or money without knowing whether they will receive what is due to them.

Again, legal issues of this service would have to be addressed.

Case studies in the failures of the current system

A "successful" sale

In the summer of 2005, I discovered a remote vulnerability in a common Linux daemon one evening. Like most researchers in this position, I wasn't entirely sure what to do with it. Having worked for the National Security Agency for five years, I had some government contacts that I thought might be interested in it. However, I became involved in a protracted legal dispute as to the classification of this vulnerability with my former employer. When this was finally resolved in my favor, I contacted a friend who ran a company, Transversal Technologies, which had more contacts with the government. For a 10% cut in the sale, he offered it to many different agencies and I received an offer of $10,000 from one organization and asked for $80,000 from another. The second organization agreed (too quickly - which likely means I had probably not asked for enough) to this amount, provided it would work against a particular flavor of Linux, which it currently didn't. I handed over my intellectual property to them so they could try to get the exploit working against this platform. After over two weeks they were still determining whether it was viable or not. As I became increasingly nervous, I decided to stake a claim to my discovery, so I posted a hash of the exploit in a bogus email to the full disclosure mailing list [22], see below. A few days later, I renegotiated a deal for $50,000 regardless of its effectiveness against the platform in question. A little over a week later, I got a check in the mail.

[Full-disclosure] Security researcher

DRAFT

From: asdfasf (zerodayinithotmail.com)
Date: Fri Aug 25 2006 - 09:01:39 CDT

I'm looking for a security researcher named "Gobbles". If anyone could send me his contact information I would appreciate it.

sadf
e9a4f234e0f5d3e587c3d27e709b7eda

Figure 1: The final line of this email contains a hash of the exploit.

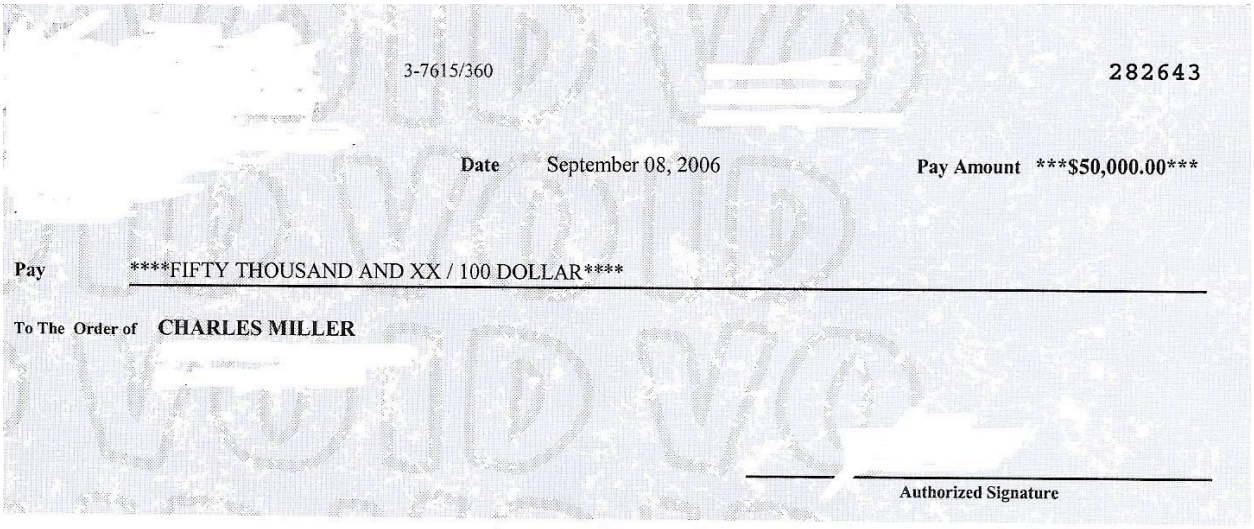

Figure 2: Finally, the Benjamins.

The timeline for these events follows:

Date	Action
Summer 2005	Vulnerability discovered.
11/7/05	Submitted to prepublication review at NSA.
7/27/06	Approved for release by prepublication review.
7/27/06	Offered to government via Transversal Technologies.
8/10/06	Verbally agreed to $80,000 conditional deal.
8/11/06	Exploit given for evaluation (at this point I have no leverage).
8/25/06	Hash of exploit published.

Date	Action
8/28/06	Agreed to lesser amount.
9/8/06	Paid.

While I was paid, it wasn't a full success. First, I had no way to know what a fair market value for this exploit was. I may have been off by a factor of ten or more. Also, the only reason the sale happened at all was because of personal contacts I had, which should not be necessary for a security researcher who wants to make a living. Not only did these contacts allow me to get in touch with the right people, but it served to give the buyer a certain level of trust in the vulnerability I had and that I wouldn't resell it to someone else. So while the sale did happen, it was in spite of the market mechanisms in place, not because of it.

A spoiled sale

Very recently, an acquaintance of mine who knew of my interest in the sales of exploits approached me and asked me to help him sell a vulnerability he had discovered in Microsoft Powerpoint XP and 2003. I helped him get the vulnerability into a form I thought people might be interested in purchasing and made some calls and sent some emails. It was frustrating not to know who to contact or have any idea of the value of this exploit. The only numbers I could base it on were the Excel vulnerability that fearwall tried to offer ($1200+) and numbers I had heard regarding iDefense and Tipping Point ($1000-$3000???). I guessed it might be worth $50,000 to the government or $20,000 to a company. I received a few offers and had settled on $12,000 from a security company - until I learned that at some point in the process it had been patched. The timeline for these actions is below:

Date	Action
1/20/07	Vulnerability discovered
1/25/07	Offered to government via Transversal Technologies
1/28/07	Exploit finished
2/10/07	Offered to computer security companies
2/13/07	Patched - KB929064

This experience perfectly illustrates the failure of the legitimate vulnerability market. There was difficulty finding an interested buyer. (In fact, once the right company was located, the sale would have been completed in a matter of a day or two, however, it is difficult to know which companies to contact) There was difficulty in establishing a price. I felt it was worth $20,000, the original offer by the final company was $8,000. I received offers as low as $5,000. With no industry figures to compare this with, it is im-

possible to come to a fair price. Furthermore, the lack of speed of the entire process caused the final sale to not be possible. Also, personal contacts became important, again, as the prospective buyer wanted to speak to people that we knew in common. Finally, I was prepared to send the exploit to the company before I had received any payment. I was completely exposed to losing the value of the information with no recourse (the company was not a U.S. company).

Conclusions

From the perspective of a security researcher, selling vulnerability information or 0-day exploits is a very risky ordeal. Due to the secretive nature of the market at the present time, it is difficult for them to find a buyer, determine a price for the information, prove the value of the vulnerability, and exchange the goods for money. On top of this, at any point in this process, the vulnerability may be announced by someone else, making the discovery worthless.

Some solutions exist which help to alleviate some of these problems, however their actual implementation remains far off in the future.

References

[1] Zero Day Initiative I 3Com I TippingPoint, a division of 3Com, http://www.zerodayinitiative.com/

[2] Vulnerability Contributor Program // iDefense Labs, http://labs.idefense.com/vcp/

[3] Email from David Endler, Director of Security Research, TippingPoint

[4] Argeniss - Information Security, http://www.argeniss.com/products.html

[5] GLEG, http://www.gleg.net/products.html

[6] IMMUNITY: Knowing You're Secure
http://www.immunityinc.com/products-canvas.shtml

[7] SNOsoft Research Team: Exploit Acquisition Program
http://snosoft.blogspot.com/2007/01/exploit-acquisition-program.html

[8] *Hackers Selling Vista Zero-Day Exploit*, eWeek, December 15, 2006
http://www.eweek.com/article2/0,1895,2073611,00.asp

[9] Conversation, Summer 2006.

[10] Instant message conversation with "greybrimstone", February 20, 2007.

[11] *Bucks for Bugs*, Dark Reading - Risky Business
http://www.darkreading.com/document.asp?doc_id=99518

[12] Researcher: WMF Exploit Sold Underground for $4,000, eWeek, February 2, 2006,
http://www.eweek.com/article2/0,1895,1918198,00.asp

[13] *eBay Pulls Bidding for MS Excel Vulnerability*, eWeek, December 9, 2005
http://www.eweek.com/article2/0,1759,1899697,00.asp?kc=EWRSS03129TX1K0000614

[14] *Researchers: Flaw auctions would improve security*, SecurityFocus, December 15, 2005, http://www.securityfocus.com/news/11364/1

[15] Mozilla Security Bug Bounty Program,
mhttp://www.mozilla.org/security/bug-bounty.html

[16] *Pay-off Time for Bug-busters, Netscape Pledges "Dogfight"*, Investor's Business Daily 11 Dec 95 A7

[17] ADD / XOR / ROL, February 5, 2007,
http://addxorrol.blogspot.com/2007/02/i-would-like-to-use-this-blog-to-make.html

[18] Dailydave: Some Sums, February 5, 2007,
http://lists.immunitysec.com/pipermail/dailydave/2007-February/004045.html

[19] Surety Integrity Advantage, http://surety.com/solutions/integrityadvantage.htm

[20] *Vulnerability Markets: What is the Economic Value of a Zero-Day Exploit?*, Bohme, 22 C3. 2005. Berlin, Germany,
http://events.ccc.de/congress/2005/fahrplan/attachments/542-Boehme2005_22C3_VulnerabilityMarkets.pdf

[21] Phone conversation, February 2006.

[22] Full-disclosure: Security researcher, August 25, 2006,
http://archives.neohapsis.com/archives/fulldisclosure/2006-08/0653.html