PHISHING AND MONEY MULES

Dinei Florêncio and Cormac Herley

Microsoft Research, One Microsoft Way, Redmond, WA, USA

dinei@microsoft.com, cormac@microsoft.com

ABSTRACT

Data breaches, phishing and spyware have compromised millions of end-user records and credentials. Mules are the preferred means for draining compromised accounts. These are unwitting accomplices who provide a stepping stone between the victim account and the attacker. The key role they play is to turn reversible traceable transactions into irreversible untraceable ones. This, together with the fraud protections enjoyed by US banking customers, generates some surprising findings. First, it is the mule's money not the victim's or the bank's money that the attacker steals. Second, mule recruitment and not credential theft appears the true bottleneck in online fraud. Third, this suggests an explanation of why stolen credentials sell so cheaply: there is a shortage of mules.

I. INTRODUCTION

Attacks on end-user accounts have been increasing in recent years. Conventional wisdom suggests that hackers, having concentrated on vandalism for many years have finally become serious about making money. This is borne out by the enormous effort devoted to phishing, and the growth of keyloggers and Man-In-The-Middle attacks. The popular press contains many accounts of how severe the problem is. Tens of millions of customer details have been compromised and sold on the black market economy. This raises an interesting question: how much fraud do banks actually detect? That is, when someone is phished or keylogged, or otherwise has their credential stolen, what percent of attempted fraud actually succeeds? The question is important as it speaks directly to the balance between front and back end security investments. Front-end investments might be stronger authentication technologies and user education, while back-end investments are the efforts at fraud detection and bank-to-bank cooperation. If no attempted fraud was ever detected at the back-end, then clearly great efforts have to made at the front, as this would be the last line of defence. If, on the other hand, a bank detects 90% of attempted fraud at the back-end, then it can risk weaker authentication knowing that the overall losses will still be low.

The dominant form of account draining appears to be the use of a mule. We show why the US consumer protections against fraud ensure that irreversible untraceable transactions are hard, and hence mules are necessary. Mules essentially receive bad transfers and initiate good ones. A surprising consequence is that in the series of transfers between victim, mule and attacker it is really the mule's rather than the bank's or the victim's money that is stolen. This means that the size of the online fraud business is determined not by the number of credentials that can be stolen, but by the number of mules who can be recruited, and how much they can send. This suggests an explanation for the fact that stolen credentials sell for small fractions of the underlying account value: there is shortage of mules. It also suggests that banks find investments in back-end fraud detection provides greater return on investment than front-end schemes such as stronger authentication.

II. RELATED WORK

Anderson [1] first suggests examination of security questions from an economic standpoint. Thomas and Martin [2] first demonstrated the enormous activity in the underground economy and observe that those who drain accounts are in high demand. Franklin *et al.* [3] followed up on this work with more detailed measurements and estimates of the size of the market. Although the money volume may be questioned [4], the activity is real. A number of other papers point to the *money mules* as the main path to draining stolen accounts [5], [6], [7], [8], [9], [10]. In most of these studies, however, the authors seem to assume the process ends when the criminal gets his hands on the money. Moore *et al.* [9] is one of the few to point out that *"the mule becomes personally liable for the funds already sent."* They stop, however, short of investigating the full implications

of this. Krebs [11] and a Verisign report [12] details many of the recruitment techniques for mules. Curiously, a lot of work has been devoted to investigating phishing, while, as we argue in this paper, the main crime being committed may actually be in the mule recruitment. In fact our own research has long concentrated in the several aspects of phishing, from prevention [13], [14], [15], to economic analysis [16], [4], to statistics and indirect implications [17], [18]. In a previous study [4], we analyze the underground markets, and offer an explanation for why stolen credentials sell for pennies on the dollar. The conclusion is that rippers, who cheat other participants on underground IRC markets create a lemon market which depresses the value of all sales. That does not, however, implies a reduction in the intrinsic value of the stolen credential itself, which is what we show in the current paper. Our study on password usage[17], includes an estimate of the rate of phishing at 0.4% a year. That rate indicates stolen credentials should be easily available for a large number of accounts.

III. HOW TO ROB A BANK (OR A MULE)

When it comes to fraud detection one size does not fit all. Different financial institutions have very different fraud detection success rates depending on their customer base. This can be seen from the inverse correlation between assets under management and the degree to which banks are targeted by phishers. For example, the four largest US financial firms by assets under management are State Street Investors ($1.4 trillion), Fidelity ($1.3 trillion), Capital Group ($1.2 trillion), and Vanguard ($0.85 trillion). Collectively these account for less than 1% of phishing attacks in 2007. Paypal, on the other hand, had transaction volume of $12 billion in 2007 but was targeted by fully 40% of phishing sites [19]. Why do phishers devote so much more energy to Paypal when there's $100\times$ more money at Fidelity? It's not that Fidelity employs better front-end security: Fidelity customers access their accounts online using a username and password. Paypal offers two factor securID access to customers who desire it, while this is not available to Fidelity customers (as of June 2010). The reason why Fidelity is less attacked by phishers would appear to be that it's so much harder to get the money out of Fidelity. Fidelity handles customers savings; while wire transfers and check writing is possible, the average number of such outbound transactions per account is likely very low. Paypal, by contrast, developed as a business to make it easy for strangers to wire money to strangers. The main purpose of most Paypal accounts is to make payments outside of the conventional check-based banking system.

Once an attacker obtains access to an account he must drain the funds. We will focus on the draining process without regard to whether access to the account has been gained by phishing, keylogging or even realtime session hijacking. The fact that credentials sell for pennies on the dollar [3] makes clear that draining is either very difficult or fails frequently. Suppose, for example, that draining an account were as simple as using online billpay to remit payment to the attacker. In this case there would be no reason to ever sell credentials for less than face value. Since numerous accounts indicate that credentials sell for only 5% or so of face value [3], [2], [4] there are clearly factors reducing the overall return.

III-A. Federal Reserve Regulations and Consumer Protections

In the US consumers are protected against unauthorized transfers from financial accounts by Regulation E of the Federal Reserve Board [20]. This covers all transfers except by check and credit card, and limits the user's liability to $50 if the loss is reported within two days of discovery. Interestingly, even in cases involving negligence the user's liability is limited: "Negligence by the consumer cannot be used as the basis for imposing greater liability than is permissible under Regulation E. Thus, consumer behavior that may constitute negligence under state law, such as writing the PIN on a debit card or on a piece of paper kept with the card, does not affect the consumer's liability for unauthorized transfers. [20]" Thus the victim does not lose money in most online fraud cases: the bank must make him whole so long as he fulfills his burden of timely disclosure.

The procedures for clearing transactions in the US are governed by Federal Reserve Regulation CC [21]. To summarize, this requires banks to make available after two business days the following types of deposits: cash, electronic payments, US federal or local government checks and postal orders, cashier's or tellers checks, and checks drawn on an account at the same bank. The bank can delay availability of funds in the following cases: deposits greater than $5000, redeposited checks (*e.g.* where the check was not properly endorsed), deposits to accounts that are repeatedly overdrawn, deposits where there is reason to doubt the collectibility of the check, and deposits into accounts that are less than 30 days old.

III-B. Untraceable Irreversible Transactions are Hard

We start from the observation that any business that cannot protect itself against dishonest customers faces the possibility of serious fraud. Regulation E [20] appears to favor the customer greatly, almost to the point of making fraud in the form of self-theft a very profitable possibiity. Suppose a dishonest bank customer wishes to double his savings. To do so he need merely transfer his money to

another account that he controls, and then claim fraud. Under Regulation E his money will be "restored" unless the bank can show the customer's complicity in the crime. From this we infer that back-end protections are at least as important as front-end ones for the bank. In addition to credential-stealing attackers, the bank potentially faces an adversary who knows not merely the username and password, but also the answers to any backup authentication questions, the user's mother's maiden name, SSN, address and recent transaction history. Only if the bank can prove culpability on the user's part can it avoid paying out. Observe that better front-end security and stronger authentication does not help. If setting up an untraceable irreversible channel for transactions were easy, then Regulation E creates an enormous loophole by which self-theft becomes very profitable. The reason that this fails is that creating an untraceable irreversible transfer of money without appearing in person is hard. This makes Regulation E inspired self-theft very hard, but it also makes transfer of money from compromised accounts a difficult matter. Hence the need for mules.

III-C. Mules: Turning Reversible Transactions into Irreversible Ones

Mules are the preferred means of draining accounts [11], [8]. The procedure goes as follows. When the attacker compromises an account he sends the money from the victim account to the mule, and the mule forwards to the attacker (minus commission). The transfer from the victim account to the mule can be by check or online bill pay. Thus it is covered by Regulation CC [21] and the funds must be made available within a few days. While fraud of this type may be hard to detect the transaction is still traceable and reversible, since the money is still in the US banking system. The transfer from mule to attacker, by contrast, will be untraceable and irreversible. This can be a wire transfer to a foreign account, for example (so long as there is no chance of reversal once the fraud is discovered). Because of their irreversible nature most banks require that these transactions be carried out in person, or that the customer confirm the transaction by phone. Several banks now require that the customer sign a statement saying that they understand the risks and dangers of scams (*e.g.* Chase requires customers to sign a form indicating that they are not acting as a money forwarding agent for a third party). This cannot be done from the victim account of course, since the victim can't be asked to sign such a form.

Thus the role of the mule is to turn a bad check into a good transfer. The bad check doesn't bounce at first, since the victim has the funds to cover it, but a reversal will be initiated once the fraud is discovered. The situation with a sample transaction of $100 and 10% mule commission is summarized in Table I. Observe in the end it is the mule and not the victim who loses money. In essence the attacker "borrows" $100 from the victim and convinces the mule to exchange this for $90 in cash or untraceable instruments. The funds from the victim account are key: unless that check cleared the mule might not be willing or able to forward the money to the attacker. However, the mule is the center of the whole operation, and recruiting and managing mules becomes a limiting factor.

	Before Discovery	After Discovery
Victim	-$100	$0
Bank	$0	$0
Mule	+$10	-$90
Attacker	+$90	+$90

Table I. *Gains and losses of the various parties for a $100 fraudulent transfer via a mule. Before discovery the victim is down the full amount and the mule receives 10%. After discovery the bank makes the victim whole (as required by Regulation E), and reverses the payment to the mule. The attacker is in effect stealing from the mule and not from the account he has compromised. If the mule has insufficient funds to cover the reversal, the bank is left with a (perhaps uncollectible) debt.*

III-D. The Feeding and Care of Mules

Clearly the size of the opportunity for online fraud is governed not just by the number of victim accounts but by mule recruitment. Access to 1000 banking accounts is of little use unless the mules can be found to drain the funds. In the end, the mule is the one robbed, not the compromised banking account. This suggests an explanation for the fact that credentials sell for fractions of a penny on the dollar: mules are in short supply and without them draining accounts is hard and risky.

What makes a good mule? Ideally, the mule can forward large sums quickly. Table I is a very simplified analysis, showing the attacker up and the mule down by equal amounts. In practice, especially if the amounts are large, the mule may be unable to repay the money that has been forwarded to the attacker. For example, if the mule has cleared $7000 worth of victim money, kept $700 for himself and forwarded $6300 to the attacker he may simply not have the funds to repay when reversal is initiated. In this case the mule still owes the money but the bank has a possibly uncollectible debt (the victim is still made whole on the strength of Regulation E). In a report to the US congress in 2007 it was reported that banks recover 30% of fraud checks [22], which is probably a reasonable figure for the amount that banks recover from mules.

Effectively, when allowing the mule to transfer money irreversibly the bank is making a judgment. They must verify not just that the funds are currently available, but also judge that there is little likelihood of reversal. If they get it wrong they will end up an uncollectible debt which is the difference between what the victim lost and what they recover from the mule. Collection is simple when dealing with people who have significant net worth or assets. For example, if someone with significant assets is foolish enough to act as a mule and forward $30k to an attacker, the bank can probably recover the entire amount. If a student, or someone with little financial history and low net worth wishes to act as mule his bank has far greater exposure. Thus the decision when allowing an irreversible transfer is similar to the decision when issuing a loan: the ability of the customer to pay is key. That is, the bank of a would-be mule knows that they are risking exposure to uncollectible debt when he requests an irreversible transfer. While he may have the funds to cover the transaction the provenance of those funds matters, and the bank will take a loss if anything goes wrong. The best way to reduce this risk is to simply contact each bank from which the mule recently received funds and ask them to check if any of the transactions appear suspicious. This reduces the useful lifetime (to the attacker) of a mule greatly. It might ordinarily be weeks or months before victims notice missing funds and contact their banks. However, contacting customers to check on the validity of transfers greatly speeds this process up. Banks that fear the prospect of uncollectible debt have both the inclination and the means to detect likely mule accounts. Customers who initiate irreversible transfers and who act as stepping stones (minus commission) for money that is passing through are good candidates for further scrutiny.

Recruitment of mules is obviously a problem for attackers, but so also is their care. Remember the attacker is stealing from the mule, not the bank. However, the typical mule is a low-income individual with few assets and little to steal. Thus, the attacker does best by driving the mule into debt. This happens by having the mule forward more than he is worth. The mule is then effectively borrowing from his bank (though this will only become clear when the victim demands reversal under Regulation E). Thus the main job of the attacker is to raise the mule's credit-worthiness, in order to cause the bank to lend. Transferring money from the compromised accounts to the mule's bank account is just a quick way of raising the mule's credit with the bank. Raising credit-worthiness takes time, but each mule has only a certain lifetime before one of the accounts that he helps drain sounds the alarm and the fraud is discovered. Thus attackers have a hard problem in determining the profit-maximizing strategy for a mule once recruited. Channeling as much money as possible through the mule seems optimal, since discovery is only a matter of time. However too much money at once might present too great a temptation for a mule; he might forgo future commissions and decide to keep the money (he will find out only later that he will keep nothing). Equally, too much money channeled too quickly makes the stepping stone nature of the account all too obvious [23], and lessens the banks willingness to carry out successive irreversible transactions. Thus, managing mules and optimizing them as a resource seems like a genuinely hard problem and perhaps more challenging than the task of stealing passwords.

IV. CONCLUSIONS

Our examination of mules has interesting implications. First, it suggests that it is not the victims of phishing and keylogging that lose money but the mules. They receive bad checks and write good ones and as (albeit perhaps innocent) co-conspirators are not protected by the Federal Reserve consumer protections. Second, this implies that mule recruitment is probably a major bottleneck in the fraud pipeline: without them stolen credentials are worth little. Third, this suggests a simple explanation for the fact that credentials sell for small fractions of their face value; *i.e.* there is an insufficient supply of mules to drain the number of compromised accounts. Fourth, banks find it better to invest in back-end fraud protections than front-end improvements such as two-factor authentication. Finally, it shows there is no shortage of compromised accounts. Thus, a small reduction in the rate of account compromise will not reduce fraud at all, at least until account compromise is at a level small enough that it becomes the bottleneck. The only effective way to reduce online fraud is by making mule recruitment even harder.

V. REFERENCES

[1] R. Anderson, "Why Information Security is Hard," in *Proc. ACSAC*, 2001.

[2] R. Thomas and J. Martin, "The Underground Economy: Priceless," *Usenix ;login:*, 2006.

[3] J. Franklin and V. Paxson and A. Perrig and S. Savage, "An Inquiry into the Nature and Causes of the Wealth of Internet Miscreants," *Proc. CCS*, 2007.

[4] C. Herley and D. Florêncio, "Nobody Sells Gold for the Price of Silver: Dishonesty, Uncertainty and the Underground Economy," *WEIS 2009, London.*

[5] Kim-Kwang Choo and Russell Smith, "Criminal exploitation of online systems by organised crime groups," *Asian Journal of Criminology*, vol. 3, pp. 37–59.

[6] M. Aston, S. McCombie, B. Reardon, and P. Watters, "A preliminary profiling of internet money mules: An australian perspective," in *Symposia and Workshops on Ubiquitous, Autonomic and Trusted Computing. UIC-ATC '09.*, jul. 2009, pp. 482 –487.

[7] Shujun Li and R. Schmitz, "A novel anti-phishing framework based on honeypots," in *eCrime Researchers Summit, 2009. eCRIME '09.*, sep. 2009, pp. 1 –13.

[8] Lorrie F. Cranor, "Can Phishing be Foiled?," *Scientific American*, pp. 104–110, December 2008.

[9] Tyler Moore, Richard Clayton, and Ross Anderson, "The economics of online crime," *Journal of Economic Perspectives*, vol. 23, no. 3, pp. 3–20, 2009.

[10] Oscar Delgado, Amparo Sabater, and J. M. Sierra, "Analysis of new threats to online banking authentication schemes," in *X Spanish Meeting on Cryptology and Information Security - RECSI*, sep. 2008, pp. 337–344.

[11] Brian Krebs, "Money Mules Help HaulCyber Criminals Loot," *Washington Post. Jan. 25. http://www.washingtonpost.com/wp-dyn/content/story/2008/01/25/ST2008012501460.html*, 2008.

[12] Verisign, "Money Mules: Sophisticated Global Cyber Criminal Operations," http://labs.idefense.com/intelligence/researchpapers.php.

[13] Dinei Florêncio and Cormac Herley, "Stopping Phishing Attacks Even when the Victims Ignore Warnings," *MSR Tech. Report*, 2005.

[14] Dinei Florêncio and Cormac Herley, "KLASSP: Entering Passwords on a Spyware Infected Machine," *ACSAC*, 2006.

[15] D. Florêncio and C. Herley, "One-Time Password Access to Any Server Without Changing the Server," *ISC 2008, Taipei.*

[16] C. Herley and D. Florêncio, "A Profitless Endeavor: Phishing as Tragedy of the Commons," *NSPW 2008, Lake Tahoe, CA.*

[17] D. Florêncio and C. Herley, "A Large-Scale Study of Web Password Habits," *WWW 2007, Banff.*

[18] Dinei Florêncio, Cormac Herley, and Baris Coskun, "Do Strong Web Passwords Accomplish Anything?," *Proc. Usenix Hot Topics in Security*, 2007.

[19] Anti-Phishing Working Group, "http://www.antiphishing.org," .

[20] "Regulation E of the Federal Reserve Board," http://ecfr.gpoaccess.gov/cgi/t/text/text-idx?c=ecfr&sid=0283a311c8b13f29f284816d4dc5aeb7&rgn=div9&view=text&node=12:2.0.1.1.6.0.3.19.14&idno=12.

[21] "Federal Reserve Board: Compliance with Regulation CC," http://www.federalreserve.gov/Pubs/regcc/regcc.htm#determin.

[22] "Federal Reserve Board: Report to the Congress on the Check Clearing for the 21st Century Act of 2003," http://www.federalreserve.gov/boarddocs/RptCongress/check21/check21.pdf.

[23] Yin Zhang and Vern Paxson, "Detecting stepping stones," in *SSYM'00: Proceedings of the 9th conference on USENIX Security Symposium*, Berkeley, CA, USA, 2000, pp. 13–13, USENIX Association.