Reducing legitimate

Online fraud continues to be a signi□cant and growing cost for merchants of all sizes, not only in operational cost, but also in lost business revenues due to false positives. David Jones from ThreatMetrix makes some valid points on the loss of revenues from legitimate customers and how this can be alleviated through true device intelligence.

A leading industry report estimates that in 2006 US$3 billion in online revenues was lost to online fraud. Merchants have been combating fraud using various techniques and technologies, but even so, 81% of merchants are engaging in manual order review and, on average, one third of all orders are reviewed, resulting in high transaction costs.

Various methods of transaction verification, both automatic and manual, are used in relation to verification of credit card data itself, such as:

- Real-time authorisation from credit companies;
- The credit card Address Verification System (AVS);
- Card verification codes (CVV2 for Visa, CVC2 for MasterCard, CID for Amex, etc);
- Scrutinising orders that are unusually large or request overnight shipping;
- In-house evidence data collected from previous fraudulent activity ('negative files', etc).

All these methods are effective some of the time, but are all subject to both false positives and false negatives. On the whole, the threat of online fraud has led merchants to over-compensate, spending large sums on looking for and blocking suspicious transactions. While they have probably succeeded in minimising the number of fraudulent transactions that get through, there has been insufficient attention paid to the false positives that lead to loss of revenue from legitimate customers. It may well be that considerably more revenue is lost from spurned customers than is lost from fraud. In fact, according to CyberSource, the share of incoming orders that merchants declined to accept in 2006 due to suspicion of payment fraud was 4.1%. If only 20% of these turned out to be valid, then as much as US$1.6 billion may have been lost from loss of valid sales. It has been estimated that for every dollar lost to direct fraud, about four dollars worth of valid orders are declined.

eCommerce fraud

While the various technologies and protections built into the credit card 'system' are helpful and prevent fraud by amateurs, they were originally designed to prevent fraud attempted by means other than the internet. Professional internet fraudsters are using much more sophisticated measures and are constantly finding better ways to circumvent detection.

Simply put, the nature of the problem is that:

- Personal details can be lost and stolen — by keystroke logging, phishing, blog scraping, card theft, etc;
- Fraudsters can use stolen credit card details to perform online transactions — often the stolen

e-commerce losses

details can include information such as address and card verification codes that can circumvent the credit card authorisation systems;
- Fraudulent transactions can be performed quickly using automation from anywhere in the world.

Of course, there are various transaction tracking and authentication technologies that can be used to help mitigate the problem, but they have their own drawbacks:
- Improved authentication methods, such as multi-factor authentication using PIN-code tokens, installed client software or call-back methods such as SMS are more practical for repeat visits to a single merchant, such as in online banking, and suffer from increased support overheads.
- Transaction profiling systems (also called 'fingerprinting' by their proponents) that also take some machine identifiers and turn them into 'fingerprint' hashes are systems of intelligence local to the merchant (not shared with other merchants) that would need repeat customer visits to be reliable, so are less effective for the majority of commerce on the internet done today.

IP geo-location

One large US financial institution recommends that if the location of the ordering computer is more than 500 miles from the ship-to address, then this can be considered suspicious. Others will treat all transactions from overseas as suspicious. This raises two issues:
- ThreatMetrix has statistics that show that the US has the largest number of zombied nodes, which is not surprising, considering that North America accounts for nearly 50% of internet nodes. China and the combined countries of the European Union are also large, but still less than the USA. American merchants using IP geo-location as a criterion based on international location exclude an unacceptably high number of infected nodes in the USA from suspicion.
- Travelling customers can be excluded inadvertently by filtering based on IP geo-location. In 2008, many people will travel to China for the Olympic Games — using IP geo-location as a filter may cause a significant increase in customer complaints.

The use of IP geo-location for risk mitigation essentially presents a large false-positive risk to merchants who rely on it as a filtering criterion and will at least place a larger burden on manual verification processes.

Malware and proxies

It is estimated by some that over 16% of the world's computers (connected to the internet) are infected with zombie malware and that the average time an unprotected computer can stay uninfected is six minutes. In June 2007, the FBI reported that over one million internet users might have been the victims of compromised computers used to steal passwords and identities.

"Merchants need to be able to 'see through' the proxy that is hiding the location and more accurately identify the machine making the transaction and its level of trust."

The fact is that professional eCommerce fraud perpetrators use networks of infected machines (botnets) to perform fraudulent transactions. In most cases, normal internet users have no idea that their machine is being used for fraudulent activities.

These fraudulent transactions are also commonly performed through an intermediate server, or so-called 'proxy' server, which hides the originating IP address. For example, a machine

in China may use a proxy in the US to hide its real IP address, making it look to an American merchant as if the transaction is occurring domestically. It is also difficult to separate legitimate proxy services from others, with AOL's proxy services being a classic example. A recent investigation by ThreatMetrix of 500 suspicious transactions with a single site indicated that 80% of the transactions were via open proxies.

The need for individual device intelligence

Without some intelligence about the specific node that is performing the transaction, merchants are no better off than if they limit their methods to the traditional services provided by the credit companies and to manual verification. Merchants need to be able to 'see through' the proxy that is hiding the location and more accurately identify the machine making the transaction and its level of trust. Lacking the granularity of individual device intelligence, existing methods are overly false positive prone and lead to lost revenue. In addition, device intelligence and reputation enables a greater use of automated detection and decision filtering, so that manual verification costs can be reduced and revenue can be increased.

A proxy server acts on behalf of another node in the internet. In the case of eCommerce transactions, the transacting node appears to have the IP address of the proxy server.

A node performing fraudulent transactions may be indicated when multiple transactions occur from the same IP address with different customer or credit card details in a short period of time (account anomalies), or there is simply a high transaction velocity from that address. If a proxy is used, it would serve to hide the true location of the perpetrator.

"Travelling customers can be excluded inadvertently by filtering based on IP geo-location. In 2008 many people will travel to China for the Olympic Games — using IP geo-location as a filter may cause a significant increase in customer complaints."

A more sophisticated perpetrator will rotate his transactions through different proxies. This will have the effect of passing any test for account/IP anomalies and transaction velocity. If the proxies are located in the merchant's country, then they will also pass tests based on IP geo-location.

ThreatMetrix has a malicious host database (ThreatIndex) tracking about one million activities per day and is tracking more than 10 million malicious PCs.

A first generation open proxy is a standard proxy server listening on a well-known proxy port, which allows anonymous connections. Some are deliberately configured, while others are due to misconfigurations of proxy servers that should be closed to anonymous connections. The ThreatIndex maintains intelligence on currently active open proxies, adding about 35,000 per day. Of course, these proxy servers are also used for purposes other than fraud that are not always illegal or undesirable.

The second generation of proxies involves using a purpose-built closed proxy in the form of malware, with a non-standard port for the private use of the perpetrator.

Until recently, these zombied nodes have used techniques such as IRC connections to a command-and-control host (CCH) to announce their availability and receive commands. These techniques are traceable and once understood could be blocked by firewalls.

We are now seeing more subversive methods being employed. Each time the infected machine is rebooted a new proxy port is selected and the malware code 'phones home' to its controller using a single encrypted UDP packet to announce its availability and port. Malware is also available that does the same thing over HTTP and 'pulls' its work commands from the CCH, allowing it to work transparently through corporate firewalls.

Device intelligence

More complete and accurate assessment of device risk involves identifying the true identity of the perpetrating node, unmasking the anomalies hidden by proxies and sharing the node information through a global threat intelligence network.

The use of a global identity database based on device parameters independent of the apparent IP address (TrueIP) allows the tracking of a computer device behind proxies (including AOL) and NAT firewalls. This then allows for true IP geo-location (TrueGeo) to be possible. The ThreatMetrix fingerprinting and fingerprint verification processes can be performed without affecting the user's website experience. Because ThreatMetrix stores the fingerprint data, multiple merchants can benefit from the accrued device intelligence.

Even if a perpetrator acts through multiple proxies, their fingerprint allows their activities to be tracked and anomalies detected. For example, a perpetrator who uses multiple stolen credit card identities via different proxies will be exposed as originating from the same device, thus exposing the fraudulent activity.

Conclusion

Online merchants are potentially losing billions of dollars every year by turning away legitimate business as a result of fraud detection processes that are overly conservative. In addition, manual transaction verification is resource intensive and costly.

True device intelligence through device fingerprinting, and automated decision processes through the use of a global threat database such as ThreatIndex, can allow merchants to 'see through' the commonly used internet technologies that can obfuscate the activities of the fraud perpetrators, exposing their activities and allowing merchants to share fraud intelligence in a global intelligence network. Such a network will allow merchants to lower false positive risk by identifying suspicious nodes more accurately, reducing the need for manual verification and reducing the loss of legitimate business.

David Jones co-founded ThreatMetrix in 2004 and has been the CEO since inception. He has more than 15 years' experience in technical and management roles within the software industry. Prior to co-founding ThreatMetrix, Jones served as VP of Global Research and technical director of SurfControl, an internet content filtering company. Among his major achievements at SurfControl was the development of the Email Content Filtering and Riskfilter products, which evolved to become SurfControl's flagship anti-spam solution. Jones arrived at his position at SurfControl when EmUTech, another company he founded, was acquired by SurfControl in 2001. He holds a BSc in electrical engineering from the University of Technology, Sydney.