Controlling fraud by reducing online anonymity using a centralized Device Reputation Authority™

Overview

In theory, conducting ecommerce can be efficient and cost effective. In practice, the damage caused by the use of stolen financial data, identity theft, spam, phishing, hackers, and other fraudulent activities can be enormously expensive and difficult to manage. These realities have significantly increased the business risks and real costs associated with doing business on the Internet.

A key enabler of fraud is the fact that the online environment is highly anonymous, with little concrete connection between the "account" and the user. While e-commerce networks regularly identify "bad" accounts after they have misbehaved, it has been difficult to identify such high-risk accounts proactively, before the damage has been done. Moreover, preventing the same individuals from repeating their offenses within a given network has proven to be a stubborn issue (much less preventing those individuals from repeating their offenses elsewhere on the Internet).

At iovation, we believe that the resolution of these issues will be achieved to a great extent when online businesses can pierce the veil cloaking online fraudsters and miscreants, and establish better controls over the connections and relationships established with their sites.

Businesses need to be able to:

- Better understand and use the relationship between computers and end-user accounts to proactively identify potential offenders;
- Gain intelligence from other networks, allowing the business to be proactive when encountering high-risk relationships, while still maintaining an extremely high level of privacy.
- Control access from any device that has been associated with undesirable behaviors.

In this paper, we'll explore a solution for these issues: iovation's ieSnare™ system, utilizing the company's proprietary Device Reputation Authority™.

Implications of fraud for the Internet economy

Most online businesses are all too familiar with the bottom line costs associated with fraud. In fact, the problem has become so pervasive and significant that most businesses are loath to publicly expose the extent of the damage to their operations. Conservative estimates place online chargeback rates at more than 10 times the chargeback rate experienced with traditional card-present transactions. In fact, Gartner group cites online chargeback rates 19 times higher than brick-and-mortar transactions. It is our experience that many online businesses operating in high-risk categories suffer chargeback rates that can exceed 10-percent of revenues. Without significant changes to the way we do business on the Internet, these losses are going to pose a significant hurdle for many e-commerce operations.

Perhaps more significant is the damage that fraudulent and illicit behavior is having on the ecommerce sector as a whole. The success of the online economy is highly dependent on the public's perception of the security of their interactions. In late 2002, Gartner Group and Harris Interactive released the results of a survey showing that 7% of online adult customers had already reported being victims of credit card fraud by October of 2002.* The rapid growth of online fraud threatens to undermine customer confidence throughout the Internet economy, limiting its ultimate potential.

Some types of online fraud and other negative behaviors

The anonymity and immediacy of the Internet has helped create a fertile environment for the invention and spread of many new types of fraud and otherwise delinquent behavior. The motives of the perpetrators are varied and range from profit-making scams to the simple challenge of proving that they can "beat the system." Whatever the motive, the costs are enormous.

Some of the most common forms of fraud and delinquent behavior that currently pervade the Internet include:

- Friendly charge backs
- Stolen credit card information, leading to charge backs (often operated in rings)
- Phishing
- Spam, spam zombies ans SPIM
- Use of Trojan applications to gather information or control processes
- Behavioral issues in online communities
- Reputation spoofing in online auctions
- Unauthorized network access and activities
- Fraudulent use of services

Common fraud management and network security methodology

Because of the prevalence and impact of fraud, there are a significant number of tools and methodologies that have been used to try to combat the problem. There is no shortage of vendors in the marketplace with proposed solutions for various aspects of online fraud. Some of these solutions are strong; some have limited real world application for online businesses that need to remain competitive. And it is extremely unlikely that any one of them will ever be the "silver bullet." Effective fraud control requires a collection of tools, each with their own specific application. When correctly paired and used together, each tool will usually enhance the others.

Some examples of commonly used fraud management and network security tools include:

- Dual authentication
- Digital certificates
- Geolocation services
- Screening services
- Negative card databases
- Heuristic analysis and pattern recognition
- Biometric identification

The ieSnare™ System

iovation's ieSnare system adds a new layer of intelligence and control for online enterprises – a layer that has never been readily available before now.

As with virtually every activity on the Internet, fraudulent or inappropriate behavior requires that a device connect to an online business's network. At it's most fundamental level, ieSnare provides real-time protection to networks by identifying computers and other network devices being used for illegal, fraudulent, malicious, inappropriate, or otherwise unwanted activities. Once these devices are identified, ieSnare can share this information with all networks protected by ieSnare. This allows subscribed networks to make business decisions about individual connections, and allow, limit, or prevent access based on the reputation of the devices involved.

The ieSnare system is comprised of three primary elements:

i) DevicePrint™ (a system that uniquely identifies devices)
ii) The centralized Device Reputation Authority™ (DRA™), and;
iii) Customized business process changes.

> DevicePrint™

ieSnare utilizes proprietary methods to uniquely identify various types of devices connected to the Internet, essentially creating a fingerprint for each device. That fingerprint remains constant across all subscribing networks. For example, a personal computer connecting to one e-commerce site protected by ieSnare is assigned a device identifier by the identical method used to identify personal computers connecting to other e-commerce sites protected by the system. The fact that the same device identifier is carried across multiple networks is a key to providing enhanced, gateway protection for subscribers to the ieSnare system.

> Device Reputation Authority™ (DRA™)

The Device Reputation Authority centralizes device identifiers, the reputation of device identifiers by network, and the relationships between device fingerprints and network specific end-user account identifiers. Identifiers contain no real customer information, and therefore limit privacy concerns for both the subscribing networks and end-users.

Each subscribing network establishes their own rule sets for specific end-user interaction. ieSnare follows these rules to return simple ‘proceed’ or ‘stop’ responses to network queries at such touch points as log-in or at the time of a transaction. If suspicious activity worthy of investigation is encountered, ieSnare can also notify audit groups within the subscribing network. Furthermore, subscribing networks are provided with an interface into the DRA for more robust research into the relationship between devices and account activity. They can then maintain their own rule sets, update their own device reputation information, run queries and generate reports through simple HTML and SOAP interfaces.

> Business process changes

Practical use of the intelligence provided from the Device Reputation Authority is generally implemented through simple and unobtrusive business process changes within the subscribing network's system. The principal functions implemented within the network include the following:

- Networks activate the DevicePrint application either by including a small code set in their own downloadable application (where an enterprise-specific application is used to connect to the network), or through ActiveX controls;
- Networks gather device identifiers at login, append a unique account identifier, and pass this combination of information to the Device Reputation Authority;
- At key end-user interaction points, such as account creation, login, purchase, access to confidential information, and other touch points, the subscribed network will reach out to the DRA to determine, according to the subscriber's own rules, whether to proceed;
- Customer support and audit groups will be trained on using the intelligence provided by ieSnare to make business decisions regarding flagged relationships.

ieSnare in Action

The best way to illustrate the ieSnare system's benefits for your business are through examples of how the system can be used by subscribers.

Example 1

The problem
Regardless of the lengths to which networks go to try verifying the identity of end-users logging into their systems, the devices involved generally remain anonymous. If, instead, networks have a clear understanding of the relationship between accounts and specific devices, potential problems could be identified before significant damage is done.

For instance, 'SUPER EMAIL NETWORK' might decide that an excessive number of accounts coming from the same device is a strong indicator of malicious behavior. They already limit the number of accounts that can use the same email address, and deploy cookies to try to track use, but both of these techniques are easily defeated by creating more email addresses and disregarding cookies. How does 'SUPER EMAIL NETWORK' manage this problem?

The ieSnare solution
'SUPER EMAIL NETWORK' has established thresholds in the Device Reputation Authority regarding accounts per device. Every time an end-user provides a valid username and password, as a final step in the log in process, the network sends the Device Reputation Authority a unique identifier associated with the account, and the unique device identifier from the device used to connect. The Device Reputation Authority logs this association and in real-time follows the pre-established rules created by 'SUPER EMAIL NETWORK' to determine whether or not the network should grant or terminate the login request. In this example, the 'SUPER EMAIL NETWORK' may receive a <proceed; no additional action> response if there are three or fewer accounts associated with a particular device. They may receive a <proceed; audit suggested> response when there are between four and six accounts associated with a particular device, at which time the Device Reputation Authority automatically notifies an audit group per 'SUPER EMAIL NETWORK'S' instructions. And finally, the 'SUPER EMAIL NETWORK' may receive and act on a <deny access> response from the DRA if seven or more accounts are associated with a particular device.

In addition to device white lists, black lists, thresholds for devices per account, and thresholds for accounts per device, networks may establish any number of rules for various end-user interactions. Each subscriber network establishes its own rule sets.

Example 2

The problem
Even after an online business knows that the activity of an account has caused damage, there is no reliable way to limit access to the network.

After a fraudulent attack or other misbehavior, Individual accounts can be shut down if they are registered to the same name, email address, physical address or other like information. Unfortunately all of this information can be completely false, or is simply so transitory and easy to sidestep that it is valueless. New accounts with new information can be set up at any time, and may include new stolen credit card information.

How can a network control access by known high-risk individuals, when there is no way to truly reach out and touch the individuals involved?

The ieSnare solution
'SUPER BOOK NETWORK' is notified by a transaction processing partner that they will not receive payment for a particular transaction because the credit card number had been stolen. After closing the account that created the fraudulent transaction, the audit team at 'SUPER BOOK NETWORK' queries the Device Reputation Authority for a list of all network devices which have been used to access the effected account. In this case, 'SUPER BOOK NETWORK' receives a list of ten device identifiers. The audit team then queries the Device Reputation Authority for a list of all account identifiers associated with these ten device identifiers, resulting in a list of six additional accounts. After research by the audit team, a total of seven accounts are closed on 'SUPER BOOK NETWORK'S' system, and ten devices are marked as 'bad' in the Device Reputation Authority.

In the future, any network device that ever connects to one of these seven bad accounts will be automatically marked in the Device Reputation Authority. In addition, any other account that attempts to log in from one of these devices will be automatically marked as bad in 'SUPER BOOK NETWORK'S' system.

Understanding relationships between accounts and specific network devices allow networks to connect bad accounts that might otherwise appear unrelated.

Once the devices used by bad accounts are identified, network access can be denied at the device level.

Example 3

The problem

Negative card databases can tell you if a card has been reported stolen, but most of the damage with these cards is done well before they are reported. Screening services can tell you if the card information (and, in some cases, the account information) is high-risk. But there is no way to know whether the person who is accessing your network has any history of fraudulent or high-risk behavior. Mostly, networks just have to take the risk and hope that, if the individual intends to do harm to the site, the network can contain any bad behavior quickly enough to limit damage.

The ieSnare solution

'INTERNET MUSIC NETWORK,' operating out of New York, finds a customer who has used a stolen card and left them with chargebacks. 'INTERNET MUSIC NETWORK'S' audit team marks the account as bad in the Device Reputation Authority™ and checks to make sure that no other devices or accounts are related. By marking the device in the Device Reputation Authority, 'INTERNET MUSIC NETWORK' can be assured that that device will not connect to them again.

Minutes later, in London, the same device tries to connect to 'SUPER BOOK NETWORK' who receives a response from the Device Reputation Authority that the device has a negative reputation from another retailer. Depending on the rules previously established by 'SUPER BOOK NETWORK', the trust level network four has assigned to 'INTERNET MUSIC NETWORK'S' information, the reason code associated with the bad reputation, and other factors, 'SUPER BOOK NETWORK' can make an immediate decision to either grant access with no additional action, grant limited access with notification to an audit group, or deny access altogether.

Effectively, the ieSnare system shares information learned through that device's actual behavior, enabling each subscriber to make informed decisions about the risk associated with allowing a connection to their network.

By sharing information about specific devices, networks can share intelligence without sharing any private customer information. In fact, networks only share information about bad devices.

More importantly, this information is far more valuable, as bad end-users may not use any of the same account information from network to network.

Summary conclusion

Some of the Internet's core strengths, speed and the ability to conduct anonymous transactions, also present unique challenges for businesses and organizations that use the Internet to interact with end-users.

By definition, the Internet requires a unique device as an access point for users. By identifying the devices, and associating them with known activity within participating networks, ieSnare is capable of removing a layer of anonymity without compromising the privacy of merchant and customer data.

ieSnare has proven effectiveness in controlling online credit card fraud. It also has significant implications for online communities who wish to manage behavior in their community, for networks who wish to limit the number of devices used by clients, and for internal networks who wish to establish an additional layer of trust for certain devices.

Additional Information:

For more information about ieSnare or the centralized Device Reputation Authority, contact iovation Inc.

Jon Martin Karl
VP of Marketing and Business Development / Founder
jon.karl@iovation.com
+1 503 943 6702 Direct

[*] 3dec02 *Fraud will cost online retailers $500 million during the holidays;* Avivah Litan, Gartner Group.