VeriSign®

WHITE PAPER

VeriSign® Identity Protection Fraud Detection Service

An Overview

Where it all comes together.™

CONTENTS

VeriSign® Identity Protection Fraud Detection Service

An Overview

+ Introduction

Identity theft and fraud are growing problems for Internet businesses, affecting the cost of doing business, heightening consumer concern, and inviting government regulation. In a 2003 survey, the Federal Trade Commission (FTC) estimated that identity theft and account fraud cost businesses an average of $10,200 per incident.[1] In 2005, the FTC found that 55% of all fraud originated from web sites or email.[2] A recent survey of US households by Forrester Research showed that 36% of consumers have scaled back their purchase of goods and services online because of security concerns.[3] Government regulations, such as the recent FFIEC guidance on Authentication in an Internet Banking Environment, which is aimed at US financial services companies, have put even more urgency around evaluating and adopting stronger authentication.[4]

The best way to prevent identity theft and fraud is through a layered approach. A critical layer in this type of approach includes fraud detection – risk-based authentication.

+ Fraud Detection – Risk-Based Authentication

A fraud detection system examines logins and other sensitive transactions, evaluating the risk of each transaction based on all contextual and historical information available for that user and transaction. Based on the inherent risk of a transaction, fraud detection systems can trigger intervention for additional authentication. This information typically includes who the user is, from where they are logging in, what kind of device they are using and what they are doing. A fraud detection system addresses the typical shortcomings with second factor authentication systems because it does not require a significant change to the end users web experience and it does not force them to carry an authentication device. Because of that, fraud detection systems can be deployed in a much shorter time frame and with a much lower total cost of ownership. When used in conjunction with second factor authentication devices, a fraud detection system can help detect second factor credentials (such as OTP) that were phished by man-in-the-middle attacks or Trojans.

In order to detect fraud, it is important to understand the normal and legitimate behaviors that users exhibit when they access web applications. As is the case with other types of human interactions, most of the time, a user's actions will fit a pattern. For example, a user might log into his bank account from his office in Boston during the work day, and log in from home in Newton over the weekend. He might check his account balances, pay his utility bills online, and log off. Now, suppose that someone logs into this user's account from a computer in Russia at dawn on a Thursday, and tries to transfer all the money from the account to a Swiss bank. This transaction is an

1 From http://www.consumer.gov/idtheft/pdf/synovate_report.pdf
2 From http://www.consumer.gov/sentinel/pubs/Top10Fraud2005.pdf
3 From "Keeping Financial Transactions Online", Forrester Research, January 12, 2005
4 The recent guidance on Authentication in an Internet Banking Environment [FFIEC-AUTH] issued by the Federal Financial Institutions Examination Council (FFIEC), states: "The agencies consider single-factor authentication, as the only control mechanism, to be inadequate in the case of high-risk transactions involving access to customer information or the movement of funds to other parties."

anomaly that does not fit the normal pattern. Using sophisticated algorithms, a fraud detection solution can "learn" normal usage patterns over time, and detect if a transaction does not match the pattern. This is called anomaly detection or behavior monitoring, and can be used to spot potentially fraudulent transactions.

Additionally, most companies have learned from experience that some transactions are inherently risky. Even if the machine learning system does not classify them as anomalous, some transactions are so suspicious that a company will want to double-check them. For example, an e-commerce site in the US might want to double-check all logins from Asia or all purchases greater than $10,000. A complete fraud detection solution should allow companies to easily create, test and publish rules like this to flag suspicious transactions. However, not every suspicious transaction is really fraud. Sometimes, an end user will do something unexpected like logging in from a different location, or making a large money transfer. A web site does not want to stop a legitimate user from logging into a service or conducting a transaction, but does want to double-check the user's identity to make sure that the transaction is not fraudulent. Thus, a fraud detection system should be complemented by an identity confirmation or intervention system. The identity confirmation system is an automated system for confirming an end user's identity, ideally through a variety of channels such as web, phone, e-mail or SMS. The system should allow an end user to confirm their identity quickly and easily, without contacting customer support. This helps minimize the cost of identity confirmation to the web site, and minimize the inconvenience to the end user. After the user's identity is confirmed, the system should be able to learn automatically that the new behavior is legitimate, and not repeat the intervention if the user shows up again on a transaction with the same characteristics.

+ Detecting Fraud in Web Transactions

A fraud detection system uses four categories of information to detect unusual transactions:

- **Computer.** A fraud detection system uses characteristics about the user's computer: operating system, language, browser, and other characteristics that help make each computer unique.
- **Clock.** The fraud detection software can also use information about when each transaction occurred: hour of the day, day of the week, frequency of login, and other information.
- **Connection.** A fraud detection engine should use information about the user's connection to the Internet: IP address, Geolocation, Connection Speed, Proxy types, and other information.
- **Category.** Finally, a fraud detection system looks at the transaction type (such as initial login, balance transfer, or payment) and the user type (such as student, high net worth individual, or retiree).

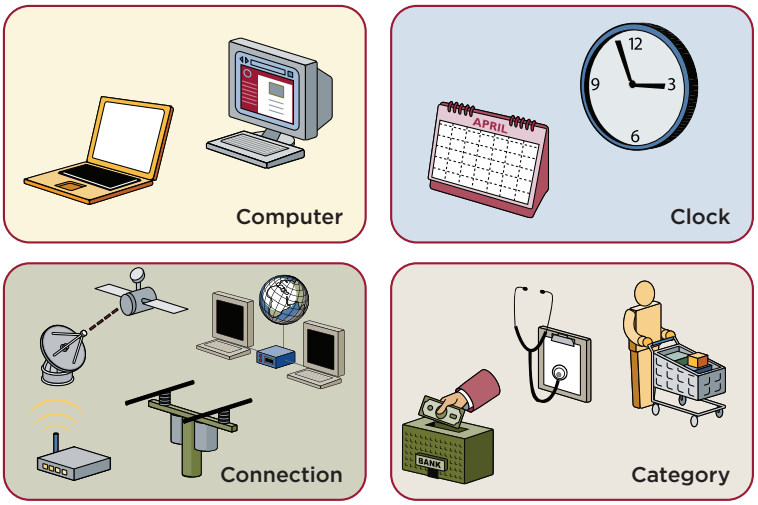

Figure 1 – Four types of transaction characteristics

Most importantly, since each web application environment is unique, a fraud detection system must be able to leverage any relevant user or transaction information available. It should also be able to operate without requiring the installation of components such as cookies, flash objects or ActiveX on the user's machines.

+ Rules Engines

Most commercial fraud detection systems include a rules engine, allowing the customer to code rules for common patterns of fraud. Each transaction is fed into the rules engine. The rules engine checks each transaction to see if it matches any pre-determined pattern for fraudulent or high-risk transactions. The rule system should facilitate the creation of complex rules, and allow enterprises to describe many subtle relationships. For example:

- If a user logs in from two different locations that are more than 1000 miles apart within one hour, mark the transaction as suspicious;
- If the user's IP address is on a list of IP addresses from which confirmed cases of fraud have been found, mark the transaction as suspicious; or
- If a user with very little history is attempting to access an account at an American Bank from an Eastern European location between 1 AM and 5 AM EST, mark the transaction as suspicious.

Rule-based systems can be very powerful and very effective if they include the right rules. Unfortunately, rule-based systems can only protect you for known types of attacks and they may be difficult to maintain. As banks and web sites learn to spot each type of attack, attackers change methods. Because of that, it is particularly dangerous to rely solely on lists of machine addresses that are sources of fraud or sources of good transactions since attackers change computers frequently to avoid detection. So, over time, attackers learn to execute attacks that are not detected by existing rules. Many times, it only takes one new coordinated attack to inflict serious damage and losses to an online business.

+ Anomaly Detection

A machine learning, anomaly-detection system can be used to address the shortcomings of rule-based systems. Rather than having to wait for a new attack to be detected and for a new rule to be written by an expert, these systems automatically and immediately detect unusual behavior for each user and for groups of users. Behavioral systems are inherently future proof—they can spot new types of attacks the first time that they are executed.

An effective anomaly detection system relies on clustering algorithms. A clustering algorithm groups similar transactions into a small number of clusters. Each cluster represents a common pattern of activity. Each time a new transaction is processed by the anomaly detection system, the system tries to fit it into an existing cluster. If a transaction does not fit into any cluster, it is classified as an anomaly. The business can then investigate the anomaly, to see if it appears fraudulent.

Clustering-based Anomaly Detection: An Intuitive Example

For example, suppose that an end user were to log onto his bank's web site 10 times with the following characteristics*:

Day of week	Time of day	IP Address	Location	Browser
Monday	10:30 AM	10.10.10.10	Providence, RI	Mozilla Firefox 1.5
Tuesday	1:00 PM	10.10.10.10	Providence, RI	Mozilla Firefox 1.5
Thursday	12:00 PM	10.10.10.10	Providence, RI	Mozilla Firefox 1.5
Thursday	9:00 PM	192.168.10.10	Boston, MA	IE 5.5
Saturday	2:00 PM	192.168.10.10	Boston, MA	IE 5.5
Sunday	11:00 AM	192.168.10.10	Boston, MA	IE 5.5
Monday	3:30 PM	10.10.10.10	Providence, RI	Mozilla Firefox 1.5
Monday	7:30 PM	192.168.10.10	Boston, MA	IE 5.5
Tuesday	6:30 AM	192.168.10.10	Boston, MA	IE 5.5
Wednesday	9:00 AM	10.10.10.10	Providence, RI	Mozilla Firefox 1.5

** This is based on a real example from users within VeriSign (IP addresses have been masked to protect user's privacy). We have an office in Providence, RI. These users often live in Boston suburbs, and sometimes travel to our headquarters in Silicon Valley.*

This user logs in from two locations: Providence, RI and Boston, MA. (The user logs in from Providence during normal business hours and from Boston on evenings and weekends). In this example, the fraud detection system will automatically build two clusters from this data: one corresponding to each usage pattern.

Suppose that the user now logs in with the following characteristics:

Day of week	Time of day	IP Address	Location	Browser
Tuesday	1:00 PM	172.16.17.18	Mountain View, CA	IE 5.5

This transaction record could correspond to a trip to California by Grandma Putterman who never travels or it could represent a fraudster from California trying to impersonate the user. Since a transaction like this does not match, Grandma Putterman's typical usage pattern, an anomaly may be triggered.

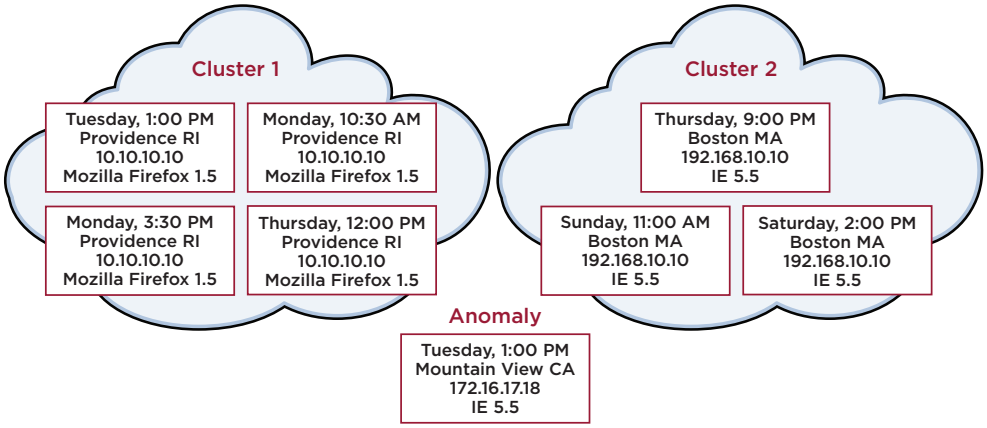

Figure 2 - Cluster data of user usage patterns

+ Overview of VeriSign® Identity Protection

VeriSign Identity Protection (VIP) is the VeriSign response to online consumer identity theft. VIP is a comprehensive suite of identity protection and authentication services that enable consumer-facing applications to provide a secure online experience for end users at a reasonable cost. VIP includes a combination of both in-premise and VeriSign-hosted components which can be accessed through standard network protocols for easy integration into existing Internet applications. VIP enables both invisible security through VIP Fraud Detection Service as well as more visible security through VIP Authentication Service. To minimize costs and maximize security by sharing intelligence and resources, VIP services are backed by the power of a network. More particularly, the VIP Authentication Network enables the sharing of a single authentication credential across a network of online service providers and enterprises. In the second half of 2006, it is anticipated that customers will be able to leverage the VIP Fraud Intelligence Network, which will allow a sharing of fraud intelligence data between members of the network. This white paper explores the VIP Fraud Detection Service and how it helps organizations mitigate issues related to identity theft and insider threats.

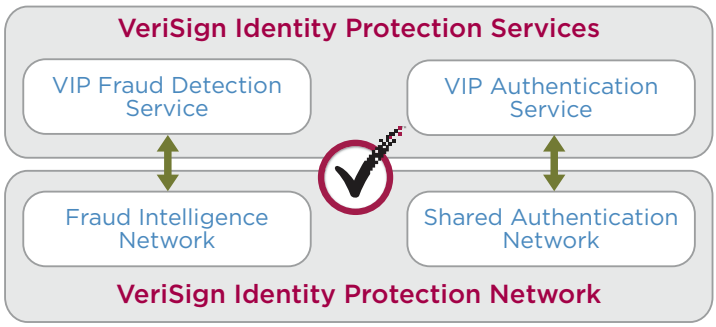

Figure 3 - VIP Services and the VIP Network

Risk-based authentication and fraud detection will enable Internet businesses to assess security risks and use out-of-band challenge and response second factor authentication only when necessary. This type of authentication gives businesses the flexibility to be able to provide additional authentication as necessary, while not being cost and resource prohibitive to the business.

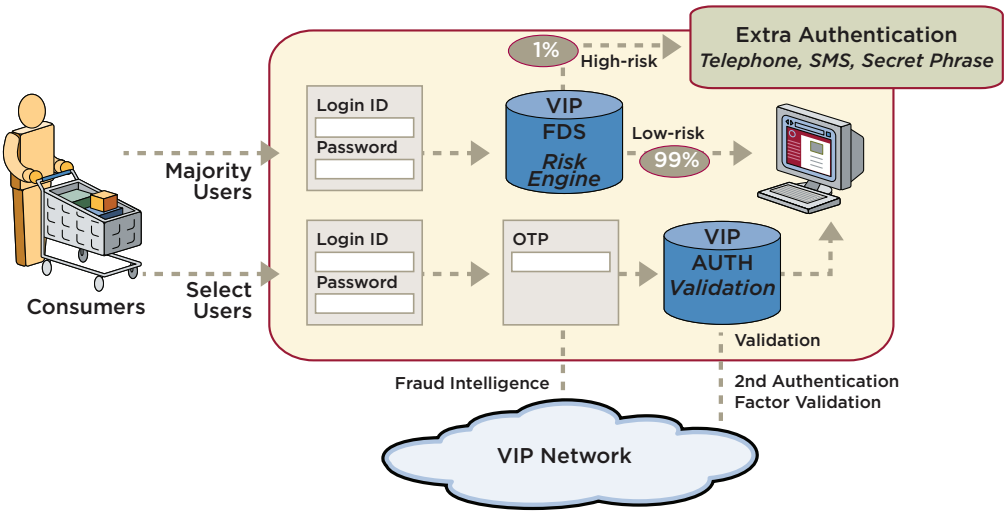

Figure 4 - One Integrated Platform for All Users

+ VIP Fraud Detection Service

The invisible part of VIP is the Fraud Detection Service (FDS). VIP FDS is a service that works in real time to detect and prevent identity theft and transaction fraud. It includes both a rule-based system and a unique behavioral heuristics engine. The service is designed to be simple and unobtrusive for both web sites and for end users. If the system detects a suspicious transaction, end users can quickly confirm their identities using an automated system. This automated system may query the user to identify themselves further with any of the following types of credentials: one time password, unique question and answer, e-mail, SMS, automated call or a customer service call.

VIP Fraud Detection Services: A More Intelligent Solution

Invisible & in-premise

VIP FDS is implemented as in-premise software that runs in your data center (instead of an ASP model that requires you to send all information outside the bank). The solution is totally invisible to the end user. It does not require any client, cookie or flash object to be installed on the user machine.

Superior anomaly detection technology

Our analytics software includes both a rule engine and a behavioral engine. The rule engine is designed for scalability and speed. Out-of-the-box rules for login are included. The behavioral engine is based on unsupervised clustering algorithms that are far superior to Bayesian logic or neural networks for these types of applications.

Protects both against known fraud and zero-day attacks

The rule engine is used to detect known fraud patterns or implement pre-defined risk policies that the bank defines. In addition to known fraud pattern, the behavioral engine can defeat zero-day attacks by flagging user activity that is inconsistent with their past behavior whether or not the attack has been seen before.

Increased robustness
Most solutions tend to privilege a specific parameter (e.g. a cookie or artifact to identify a known device or IP address). Because our system is built on true clustering technology, the anomaly engine does not privilege any specific parameter but detects behavior changes by looking across multiple variables with no pre-defined or arbitrary weighing. This approach drives lower false positive rates and increased robustness over time.

True transaction fraud detection, not just login fraud detection
Because it is not limited to a fixed or pre-determined set of parameters or a pre-determined set of rules, our solution is applicable to transaction fraud. In particular, it can be deployed to detect login fraud as well as higher risk transactions such as money transfers (for each user, the solution will automatically build behavioral models based on parameters specific to each transaction type).

Reduced maintenance
Solutions that are based on a rules engine only approach (e.g. credit card fraud engines) are maintenance intensive. They require the ongoing development of rules to model the normal behavior of a specific segment of users that exhibit similar usage pattern. Because user behavior is not only very diverse, and can rapidly change over time, such an approach applied to online banking login and high-risk transactions can be extremely time-consuming and maintenance heavy. A large rule set is also harder to maintain. On the other hand, for each transaction type, our behavioral engine automatically models the observed normal behavior of each user through unsupervised clustering, therefore eliminating the process of developing and maintaining rules to model non-fraudulent activity across distinct user groups and transaction types.

Our Fraud Intelligence Network will provide unique fraud intelligence data to enhance the effectiveness of the fraud detection process. VeriSign has proven experience dealing with fraud on a global scale. At VeriSign, we operate intelligent infrastructure services that enable and protect billions of interactions every day across the world's voice and data networks. Every day, we process as many as 18 billion Internet interactions and support over 100 million phone calls. Over 93% of Fortune 500, the world's 40 top banks use VeriSign Authentication.

Simpler integration – Faster deployment
One of the challenges for banks to meet the FFIEC guidance is speed of deployment. Many solutions require complex integration. At a minimum, custom code must be added to the application to invoke the fraud engine custom APIs. Our solution does not require any new code to be written by the bank. It does not require the application to be changed. In fact, our solution can be deployed in two modes: the "Zero Integration" mode and the "Real-Time Intervention" mode. In the first mode, the fraud detection engine feeds from the application logs, hence does not integrate with the application directly. In the second mode, only the configuration of the application server needs to be modified (config file + jar file in J2EE environment, config file + DLL in Windows environment). The second and more integrated mode also implements real-time intervention (e.g. challenge-response, out-of-band calls, etc). In both modes, no custom code needs to be written by the bank. The application code remains untouched.

Support distributed deployment model
Our solution has a parallel architecture. This means that the anomaly engine can be deployed across multiple clusters. Clusters can also be deployed across multiple data centers. In fact, the solution leverages standard relational databases as its persistent store (behaviors and configuration). Therefore, database replication tools, including cross-site replication can be used to support a multi-site deployment model.

+ The Fraud Detection Process

The VIP FDS is designed to be part of a company's business process. We anticipate that most institutions will adopt a process like the one shown below in Figure 5. First, the end user enters his login information (with or without second factor credentials). Next, the site validates the credentials and forwards the results to the fraud engine. The fraud engine then inspects the transaction. If the fraud engine determines that the transaction is low risk, the end user is logged into the web site. If the fraud engine determines that the transaction is risky, the web site will then attempt to validate the user's identity before allowing the user to login. This may include additional security questions, out-of-band automated messages (through voice, email or SMS messages), or phone calls to customer service. If the transaction turns out to be legitimate, then the FDS system can be configured to automatically ignore future similar anomalies.

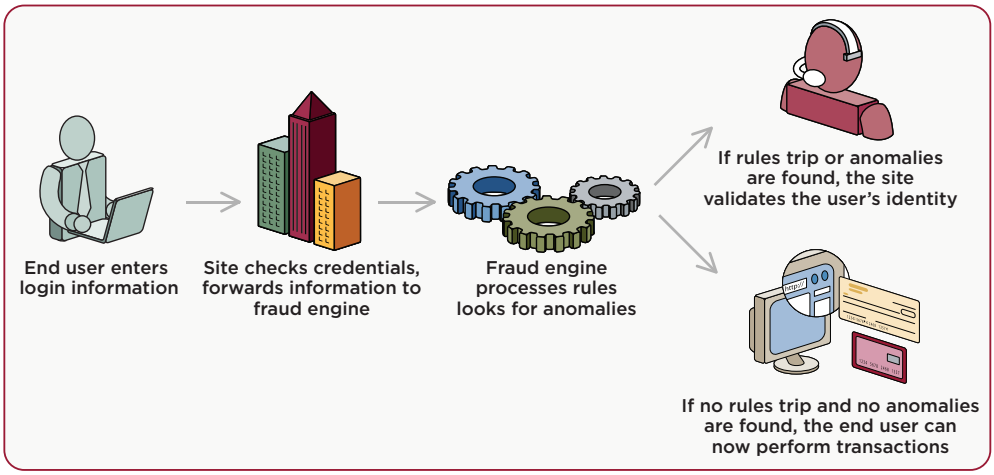

Figure 5 - The Fraud Detection Process

The VIP FDS Rules Engine

The rules are completely customizable for each enterprise. VeriSign provides an out-of-the-box set of rules to identify common patterns of fraud and ensure FDS can protect your users on day one.

As we continue to learn more about new fraudulent behavior, we will update the default set of rules. Enterprises are free to use the default rules alone, to mix standard and proprietary rules, or to strictly use their own rules. The FDS system includes an intuitive graphical user interface for creating, modifying, and monitoring rules. Rules are stored in XML files and can also be manipulated programmatically.

The VIP FDS Behavior Engine

After processing the transaction through the rules engine, the FDS system uses a machine-learning system to analyze the history of transactions for each user and look for anomalous behavior. The system uses a patent-pending algorithm to cluster categorical transactions based on similarity. Whenever a transaction is found that does not meet an existing historical pattern for that given user, it will be marked as suspicious. FDS learns typical transaction patterns on four levels.

All VIP FDS systems allow patterns to be detected on a single web site:

- **One user on one site.** Does the current transaction fit the typical usage pattern for this user on a single application (e.g. web banking)?
- **Many users on one site.** Does the current transaction fit the typical usage pattern for this type of user (e.g. high net worth individuals) on a single application (e.g. web banking)?
- **One user across many sites.** Does the current transaction fit the typical usage pattern for this user (e.g. high net worth individuals) across multiple applications (e.g. web banking, brokerage and insurance)?
- **Many users across many sites.** Does the current transaction fit the typical usage patterns for all users across multiple applications (e.g. web banking, brokerage and insurance)?

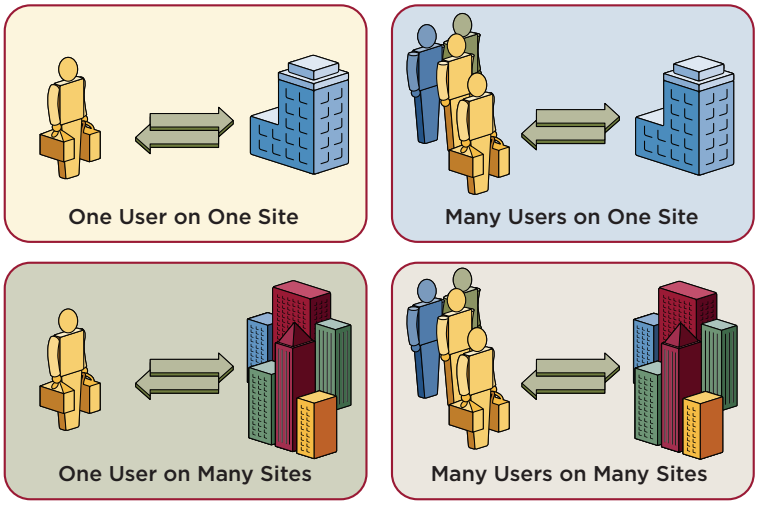

Figure 6 - Four levels of learning

The FDS system uses a Partitioned Categorical Clustering Algorithm to find anomalies. Over time, as an end user logs into applications, the system builds a statistical representation of all transactions. The system compares transactions with each other, grouping similar transactions into "clusters." Each time the user logs into the site, the latest transaction is compared with the statistical representation of that user's transactions in the system. If the new transaction is similar to other transactions in an existing cluster, then the user is allowed to log in. If, however, the new transaction resembles no previous transactions, then it is marked as an anomaly. FDS automatically adjusts the weight of each transaction parameter to reflect the relevancy of that parameter for a given user. For example, FDS would increase the weight of the geolocation parameter for a user that always does transactions from home, but would decrease the weight of that parameter for a user that travels all over the globe and does transactions from multiple locations.

In addition, the system administrator can decide and configure which data is used for clustering and how transactions are compared. The administrator can also adjust the sensitivity of the system to minimize false positives and maximize the number of fraudulent transactions detected.

Identity Confirmation Services

If the VIP FDS system detects a suspicious transaction, it passes the transaction to an identity confirmation or intervention system before authenticating the user. The system asks the user for additional confirmation of his identity, depending upon the degree of risk as determined by the organization's policy. The system may ask the user additional questions to confirm his identity or send a message using an out-of-band mechanism such as a telephone call, SMS message, or email. If the user successfully confirms his identity, he will be logged in as usual. If the user cannot confirm his identity, he will be asked to call customer service. If desired, FDS can call an external intervention system and wait for a response to determine whether or not the intervention was successful.

+ The VeriSign® Identity Protection Fraud Intelligence Network

Both the VIP Fraud Detection Service and VIP Authentication Service offer a compelling value proposition — a comprehensive, easy to deploy, and economical suite of services for reducing fraud, improving security, and achieving compliance. However, VIP offers much more than that.

The VIP Network is a set of shared services that builds on the VIP Fraud Detection Service and VIP Authentication Service. The VIP Network consists of two components: the Shared Authentication Network and the Fraud Intelligence Network, the latter of which will be available in the second half of 2006. Each of the VIP services is enhanced by network effects. The VIP Network provides lower costs, better and standardized end user experience and higher security through network services. The VIP Authentication Service helps businesses share authentication resources to reduce costs and improve the experience for end users. The Fraud Intelligence Network service will help businesses share intelligence about online identity fraud to improve security.

Customers may choose to subscribe to either service alone, or to maximize the benefits of network membership by using both services. The VIP Network combines the two core VIP Services in a unique business framework that facilitates sharing costs, data, and resources.

Criminals on the Internet use many different mechanisms to capture personal information including phishing web sites, key loggers, false store fronts, and database theft. Often, criminals will try to use the same information on multiple web sites, testing login information by trial and error, establishing multiple fraudulent accounts, or other malicious activities. In the offline world, banks and credit card companies know that attackers will often re-use stolen identity information or copy-cat other fraudster processes, and have established data sharing consortia to identify fraudulent applications and account usage. The same approach can be used to stop identity theft and account fraud on the Internet.

+ VeriSign: A Trusted Partner

VIP is backed by VeriSign, the premier Internet security services company. Today, over 400,000 web sites display the VeriSign Secured Seal, allowing customers to confirm the identity of e-commerce sites. VeriSign is among the most trusted consumer brands for Internet Security.

VeriSign acts as a trusted third-party service provider for a diverse set of applications, ranging from issuing SSL Certificates to e-commerce sites, to providing Inter-Carrier SMS Messaging, to operating the Electronic Product Code Global information services. We will bring the same experience and expertise to a shared authentication network.

VeriSign operates many core services for the Internet and telecommunications networks, including the .com and .net domain registries, SSL Certificate Authorities, and SS7 Signaling Networks. The VeriSign operations staff has decades of experience in running critical infrastructure, keeping it secure, and keeping it available. VeriSign monitors, manages, and protects the networks of many other financial institutions, utilities, government agencies, and other companies through our Managed Security Services. VIP Services are run by VeriSign, so you can be sure that they will be secure and reliable.

Finally, VeriSign researches emerging threats and finds new vulnerabilities through our iDefense division. Additionally, iDefense researchers monitor hacker forums in English, Russian, Chinese, Portuguese and other languages. We will use original research from iDefense to identify new methods of committing fraud and improve the VIP Fraud Detection Services.

+ Learn More

VeriSign Security Services protect online interactions, enabling companies to manage reputational, operational, and compliance risks in the simplest and most cost-effective way possible. For more information about VeriSign Identity Protection, please call 650-426-5310 or email enterprise_security@verisign.com.

Visit us at www.VeriSign.com for more information.

07-25-06