



W h i t e P a p e r

Fraud Prevention Guide

Abstract

This *Fraud Prevention Guide* provides guidelines for online merchants on how to design and implement a best practice risk management process and in particular how to leverage the fraud management functionality provided by the ClearCommerce Engine.

Contents

Overview	2
Internet Fraud: A Realistic Sizing of the Problem	2
Cost of Fraud for Merchants	3
Need for Fraud Prevention	4
How Fraud Starts	5
The Key Components of Risk Management	7
Authentication	8
Screening	11
Case Management	15
Analysis and Tuning	21
Traits of Fraudulent Orders	24
International Fraud	25
The Aftermath of Fraud: Reporting and Prosecuting Fraudsters	29
More Information	31
About ClearCommerce	31
References	31
Document History	31
Notices	31

Overview

The rapid growth of online commerce presents new opportunities and new challenges for both traditional retailers and “pure-play” Web merchants. With business-to-consumer online commerce projected to grow to 150 billion dollars by 2004, the Internet is set to become a major selling channel for the retail sector; however, doing business over the Internet exposes merchants to a much greater risk of losses due to fraud. Because of the liability policies for *card-not-present* transactions and because of the anonymity, reach and speed that the Internet provides for fraudsters, risk management and fraud prevention have become necessary components of every e-commerce infrastructure.

This *Fraud Prevention Guide* provides guidelines for online merchants on how to design and implement a best practice risk management process and in particular how to leverage the fraud management functionality provided by the ClearCommerce Engine.

Internet Fraud: A Realistic Sizing of the Problem

The incidence of fraud in the online industry has been a subject of speculative and contradictory reports that have contributed to a widespread concern among both merchants and consumers. Even traditionally reliable research firms, like the Gartner Group, have published conflicting reports on chargeback rates experienced by online merchants; with figures ranging from 15%¹ to 1%². Often these figures fail to distinguish between actual credit card fraud and other forms of cyber-crime. Online auction fraud, for example, accounts for nearly half of all incidents of fraud on the Internet³.

Recent statistics provided by the card associations put online fraud rates between 0.8% and 0.9%⁴. The latest statistics obtained from the ClearCommerce Data Consortium⁵ show that fraudulent chargebacks represent 0.6% of all completed transactions, and 0.8% of the entire dollars transacted. Although these rates are significantly higher than those reported for face-to-face transactions (for which Visa reports 0.06%), they are indeed comparable to fraud rates that issuers have experienced in the past with Mail Order / Telephone Order (MO/TO) transactions. In practice, credit card transactions over the Internet appear to be only slightly more risky than other *card-not-present* transactions.

¹ Gartner Group “Limiting Credit Card Fraud and Chargebacks on the Internet”, June, 1999

² Gartner Group “E-Tailers Squeezed by Higher Credit Card Fraud and Rates”, July, 2000

³ According to the Internet Fraud Complaint Center, an initiative by Federal Trade Commission.

⁴ Bank Technology News, July, 2000

⁵ The ClearCommerce Data Consortium is a multi-merchant database containing millions of historical Internet transactions and chargeback data.

Thus, at the moment, the incidence of fraud for Internet merchants appears to be higher than for brick-and-mortar merchants, but not of the catastrophic proportions initially speculated. Nonetheless, since online merchants are more directly exposed to losses due to credit card fraud, fraud prevention remains a must-have for every merchant conducting business via the Internet.

Cost of Fraud for Merchants

Since card issuers classify purchases completed via the Internet as *card-not-present* transactions, online merchants have to bear not only higher interchange rates⁶, but also the full liability for losses due to fraud. Whenever the legitimate cardholder disputes a credit card charge, the card-issuing bank will send a chargeback to the merchant, reversing the credit for the transaction. In most cases it is very difficult for the merchant to reverse the chargeback, because there usually isn't any physical evidence (e.g. delivery signature) available to challenge the chargeback. Therefore most of the time the merchant will absorb the cost of the fraudulent order, which unfortunately includes several "line items":

- Cost of goods sold: since it is very unlikely that the merchandise will be recovered in a case of fraud, the merchant will have to write off the value of the goods involved in a fraudulent order. The impact of this loss will be highest for low-margin merchants.
- Shipping cost: since the shipping cost is usually bundled in the value of the order, the merchant will also absorb the cost of the shipping for fraudulent orders. Furthermore, fraudsters typically request high-priority shipping for their orders, because it allows a rapid completion of the fraud. Therefore fraudulent orders also carry high shipping costs.
- Card association fees: Visa and MasterCard have put in place fairly strict programs that penalize merchants generating excessive chargebacks. Typically, if a merchant exceeds established chargeback rates for any 3-month period (e.g. 1% of all transactions or 2.5% of the total dollar volume) the card issuers will penalize the merchant with a \$25 fee for every chargeback. The penalty will increase up to \$100 per chargeback if the merchant is unable to control the chargeback rate in the following months. Furthermore, issuers will also apply monthly fines charges ranging from \$5,500 to \$100,000 per month for merchants with excessive chargeback rates. Finally, in extreme cases, the issuer can decide to terminate the merchant service agreement; in practice, which would make the merchant unable to conduct business online.

⁶ Interchange rate for online merchants are, on average, 2.5% of the transaction value, plus an additional 20 to 30 cents per transaction, while brick-and-mortar merchants typically pay 1.5% of the transaction value and 10 to 20 cents per transaction.

- Merchant bank fees: in addition to the penalties charges from the card associations, the merchant will typically pay an additional processing fee to the acquiring bank for every chargeback. These fees vary between \$10 and \$25 dollars per incident.
- Administrative cost: every transaction that generates a chargeback requires significant administrative costs for the merchant. On average, each chargeback requires between 1 and 2 hours to process. This is because processing a chargeback requires the merchant to receive and research the claim, contact the consumer, and respond to the acquiring bank or issuer with adequate documentation. Large organizations often dedicate several people exclusively to the chargeback handling process.

The line items described above quickly add up to a per-case cost that can easily reach several hundred dollars, even for merchants selling relatively low-cost merchandise. The cost of fraud is obviously much higher for merchants that sell high-ticket items, like computers, consumer electronics and jewelry.

Need for Fraud Prevention

A recent Jupiter report suggested that one of the most important factors in determining the risk profile of a merchant is the overall traffic experienced by the site.⁷ Survey data shows that the fraud rates tend to grow with the number of unique visitors as a site grows its recognition and visibility on the Internet. This means that online merchants that are just beginning to grow their businesses will experience fewer fraud cases initially, and therefore may underestimate the threat of online fraud. However, not having a preventative risk management process in place may expose these merchants to unpleasant surprises.

One of the main challenges with fraud prevention is the long time lag between the time a fraudulent transaction occurs and the time when the merchant receives notice (chargeback notification). ClearCommerce internal analysis shows that the average lag between the transaction date and the chargeback notification is 72 days. However, over 20% of all fraudulent charges may become apparent only 100 days after the transaction or even longer. This means that, if no fraud prevention is in place, one or more fraudsters could easily generate significant damage to a business before the merchant even realizes the problem. When chargebacks begin to flow in, it is typically too late to take defensive measures.

In summary, the unavoidable time gap between the actual fraud event and the chargeback notification requires the online merchants to put in place fraud prevention programs right from the start. Without these preventative measures, merchants expose themselves to the risk of unexpected significant losses due to fraud and severe penalties from card issuers.

⁷ Jupiter “Merchant’s Risk Management”, November, 2000

How Fraud Starts

In most cases, the only thing that a fraudster needs to place a fraudulent order is a valid (e.g. issued and active) credit card number. In fact, although card processors request additional data for expiration date, AVS (Address Verification Service) and CVM (Card Verification Method) checks, these fields are not mandatory, and do not necessarily result in declined transactions.

Fraudsters can gain access to valid credit card numbers in a number of ways. In spite of dominant consumer concerns, credit card numbers are rarely stolen in cyberspace. Today nearly all online merchants use secure communication channels (such as secure socket layer, or SSL) when sensitive data is transmitted between the consumer browser and the web site, therefore the likelihood of a fraudster intercepting card numbers during a transaction is extremely small. A greater risk is theft of credit card data from storage on the merchant's web site. Break-ins resulting in theft of credit card data have occurred at sites with insufficient security, and media coverage of these thefts has heightened consumer alarm at the potential risks of e-commerce. These risks, however, can and are being addressed by implementation of appropriate site security.⁸ While the Internet provides a *card-not-present* environment in which to use stolen cards, the original source of the data is most often traditional.

Stolen or lost credit cards obviously provide fraudsters with full access to account information, expiration data and billing name. However, these card numbers typically provide a short window of opportunity to the fraudster since the legitimate cardholder will usually report the accident to the issuer and the account will be blocked. Credit card numbers collected from card imprints, receipts or monthly statements collected in dumpsters give fraudsters a wider window of opportunity, since the cardholder is unaware that the card number has been compromised until he or she receives a statement from the issuer that includes unauthorized transactions. Newer Point of Sale systems protect the cardholder by printing out receipts that only include the first and last four digits of the card.

New high-tech tools now commonly used to steal credit card information are hand-held credit card *skimmers*. These devices can read the card information encoded in the magnetic stripe and store thousands of card numbers that are later uploaded to a PC. Since these devices are easily concealed, an unethical waiter can easily swipe the card while walking between the cash register and the table. As with credit card numbers stolen from imprints and receipts, the cardholder is typically unaware of the event for weeks or even months.

⁸ For more information on securing your e-commerce site, see the *Visa Cardholder Information Security Program* (www.visabrc.com).

While the methods described above require some form of access to the physical card or paper receipts, others can give fraudsters access to a card number without even leaving home. Card number generator programs are software tools that can produce hundreds of “valid” card numbers, that the fraudster can easily generate and then “test” online. The software generates numbers that are valid with respect to the coding scheme used by card issuers (the so-called MOD10 check), however the vast majority of these numbers will not correspond to active accounts. Nonetheless, it is relatively easy for the fraudster to test large sets of generated numbers against a target web site. Once the fraudster finds a “hit” the card is charged up to its limit with fraudulent orders.

The Key Components of Risk Management

Managing the risk of fraud effectively and economically can be a challenging task for online businesses. While minimizing the financial losses caused by charged back orders is the main objective, the shopping experience of legitimate customers should not be impacted by restrictive order acceptance policies, which ultimately would result in lower sales. This means that merchants need to carefully evaluate suspicious order before eventually rejecting them, which also has a cost and an impact on customer satisfaction, as it may cause delays in the fulfillment and delivery of orders. Achieving an optimal balance between these various constraints requires not only sophisticated detection technologies, but also an efficient case management process and data analytics.



Figure 1: The four key components of online risk management

A comprehensive online risk management solution comprises of four key components:

- Authentication
- Screening
- Case Management
- Analysis and Tuning.

Authentication mechanisms attempt to verify that the person placing an order is indeed the legitimate account owner. This is done by requesting some type of verification information at the time the order is placed, information that should only be available to the legitimate cardholder. Authentication represents the first level of defence against fraud and, in some cases, it poses a sufficient deterrent to fraudsters.

Screening utilizes various order profiling techniques to automatically identify transactions that could indeed be fraudulent. Order profiling allows merchants to minimize the number of orders that require human intervention, focusing resources only on the truly suspicious ones.

Case management is the process and the infrastructure to support order verification for those orders that are deemed suspicious or risky. This component of a risk management solution is very important because it should allow a rapid and cost effective resolution of orders held for verification.

Analysis and tuning leverages the data that flows into the system to identify emerging patterns, monitor the performance of the screening algorithms and improve the detection abilities of the overall system.

Authentication

There are currently three mechanisms available to online merchants to authenticate consumer transactions: payer authentication, address verification system (AVS) and card verification methods (CVM). While AVS and CVM were designed by issuers well before online commerce began, payer authentication is a new technology specifically designed for Internet transactions.

Authentication methods, especially CVM and payer authentication, can be very effective in deterring fraud. However, the concern that they could have some negative impact on customer experience and shopping carts abandonment has led to only moderate adoption rates among online merchants.

Payer Authentication

Payer authentication promises to bring a new level of security to business-to-consumer Internet commerce by providing assurance to the seller that the payer is authentic and the purchase is authorized.

The first implementation of this type of service is the Verified by Visa (VbV) or Visa Payer Authentication Service (VPAS) program, launched worldwide by Visa in 2002. The program is based on a Personal Identification Number (PIN) associated with the card, similar to those used with ATM cards, and a secure direct authentication channel between the consumer and the issuing bank. The PIN is issued by the bank when the cardholder enrolls the card with the program, and will be used exclusively to authorize online transactions.

When registered card-holders check out at a participating merchant's site, they will be prompted by their issuing bank to provide their password. Once the password is verified, the issuing bank notifies the merchant that the payer has been authenticated, and the merchant may complete the transaction and send the verification information on to their acquirer.

This program provide benefits to both parties involved in the transaction. The cardholder benefits from enhanced protection from unauthorized use of their card number on line. This, Visa believes, will eliminate one of the major barriers to the growth of business to consumer ecommerce, namely the consumer's concerns about security. The benefits for merchants are even greater. Merchants participating in the VPAS program will benefit from protection from fraudulent chargeback claims. This liability shift on fraudulent orders is expected to be a major driver for VAPS adoption among online merchants.

The ClearCommerce Payer Authenticator (CCPA) is a Visa-certified plug-in module that allows merchants to enable their secure site for VPAS-authenticated transactions. The CCPA module resides on the merchant's sites and communicates directly with the Visa Directory Server (DS) and the Access Control Server (ACS) via the Internet. When a consumer enters a card number, the CCPA first interrogates the DS to verify that the card is VPAS-enabled. Then, it redirects the consumer to the appropriate ACS to authenticate the card. Finally, the CCPA receives a special authorization number from the ACS, which is then submitted by the merchant to the acquirer during the authorization.

Address Verification System (AVS)

Card issuers first established the Address Verification System (AVS) as a security mechanism for *card-not-present* transactions. AVS validates the billing address information provided by the consumer (via a web form or over the phone) against the billing address information that the issuer has on record for the account. Specifically, AVS checks the ZIP code and the numeric part of the street address and returns a match/mismatch response. Notably, the AVS response provides additional information for the merchant, but AVS match is not required for approval, nor is a transaction that obtained an AVS match response guaranteed against chargebacks. The decision on whether to accept an order based on the AVS response is completely left to the merchant.

Although the AVS response provides useful information for the merchant to determine the risk level of an order, it suffers from several limitations that make it a weak screening tool if used alone. First, AVS has a relatively high failure rate: typically, less than 60% of the transactions will obtain a full match on AVS (for one reason, AVS is not available for credit cards not issued in the United States). Second, the vast majority of orders that completely fail AVS are valid: typically over 98% of the transactions with failure on both zip and street address are legitimate. Finally, AVS only validates the *billing* address: fraudsters who obtain billing addresses can still pass the AVS check, and ship fraudulent orders to a different address.

In spite of these limitations, the AVS response will be part of any comprehensive order evaluation scheme. Often the merchant will combine the AVS response code with other order attributes, for example the amount and a “shipped to different address flag” to determine if an order is deemed a review.

FraudShield allows merchants to create rules based on the AVS response, alone or in combination with any number of checks on order attributes. For example, the merchant can easily create a rule like:

```
IF AVS-Response="NN" AND Amount>$100 AND Email-Domain  
IN "Free-Email-Domains-List" THEN Review
```

Card Verification Methods (CVM)

The Card Verification Methods (CVM⁹) consists of a 3 or 4-digit numeric code that is printed on the card, but not embossed on the card nor available in the magnetic stripe. The merchant can request the consumer to provide this numeric code with the order and submit it with the authorization. The card processor will then validate the code supplied with the number on record for the specific card, and return a match/no-match response. As with AVS, the CVM response is only provided as additional information for the merchant, however a match will not protect the merchant from possible chargebacks on the transaction.

The purpose of CVM is to ensure that the person submitting the transaction is in possession of the actual card, since the code cannot be copied from receipts or skimmed from the magnetic stripe. Although CVM provides some protection for the merchant, it doesn't protect from orders placed on physically stolen cards. Furthermore, a fraudster who had temporary possession of a card could, in principle, read and copy the CVM code.

⁹ The various card issuers use different names to indicate this security feature; CVV2 for VISA, CVC2 for Master Card and CID for American Express.

Nonetheless, fraud rates on transactions with verified CVM codes have historically been significantly lower than those for transactions without CVM (fraud rates on CVM-validated transaction are reportedly 80% lower than those for non-CVM transactions).

The ClearCommerce Engine fully supports authorization with CVM verification. If the merchant requires the CVM code for every order, the Payment Engine will return the CVM response. The response code can then be validated in a FraudShield rule, alone or in combination with any number of additional checks.

Screening

Screening includes a variety of techniques that allow merchants to automatically profile incoming orders to identify indicators of possible fraud. The screening process often interacts with the results of authentication, as orders that have been successfully authenticated via a strong authentication mechanism could be excluded from screening. On the other hand, even if a merchant has implemented upfront authentication methods, like payer authentication or CVM, not all transaction can be authenticated, because, for example, not all credit cards have CVM codes or are enrolled in payer authentication programs.

Each screening technique has its strengths and limitations. Therefore, it is typical for merchants to “layer” and combine screening methods to achieve effective fraud protection, and thus it is important to rely on an environment, like the ClearCommerce Engine, that offers multiple methods and that allows to combine them easily.

Fraud Rules

Rule-based systems rely on a set of expert rules designed to identify specific types of high-risk transactions. The rules, typically expressed using an if-then logic, are created using the knowledge of what characterizes fraudulent transactions. For example, a fraud rule may be designed to flag all orders over \$500 with multiple units of the same product. Another rule may look for orders that failed the AVS check and are shipped to an address different from the billing address. “Positive” rules can also be designed to profile low-risk orders and then skip any subsequent fraud check. For example, orders that passed the AVS check, are below \$50 and are shipped to the billing address may be deemed safe and therefore approved right away.

Fraud rules enable the merchant to automate the screening processes leveraging the knowledge gained over time regarding the characteristics of both fraudulent and legitimate orders. Rules can also be used in combination with negative and positive files and processor-provided verification mechanisms like AVS and CVM. Typically, the effectiveness of a rule-based system will increase over time, as more rules are added to the system and negative and positive files grow in size. It should be clear, however, that ultimately the effectiveness of the system depends on the knowledge and expertise of the person designing the rules.

Rule-based systems can also be used to implement policies on order acceptance that may or may not be driven by risk-related considerations. For example, a merchant may want to restrict the number of units that can be ordered in a single transaction, or may not be able to ship certain products to certain countries. With a rule-based system these policies can be easily formulated and maintained as part of risk management infrastructure.

The FraudShield Rule Editor provides a powerful web-based user interface to create, edit and manage a library of fraud rules. The FraudShield engine will apply the rule base created via the Editor to incoming orders, in real-time, flagging orders that triggered review or decline conditions.

Negative and Positive Lists

A negative list is a database used to identify high risk orders based on specific data fields. An example of a negative list would be a file containing all the card numbers that have produced chargebacks in the past, used to avoid further fraud from repeat offenders. Similarly a merchant can build negative lists based on billing names, street addresses, emails and IPs that have resulted in fraud or attempted fraud, effectively blocking any further attempts.

Negative files can also be used to flag, rather than decline, orders that may be particularly risky. For example, certain foreign countries, in particular in East Europe, have historically been associated with online fraud. Therefore a merchant could create and maintain a list of high-risk countries and decide to review or restrict orders originating from those countries.

Positive files are typically used to recognize trusted customers, perhaps by their card number or email address, and therefore bypass any fraud check for their orders. Since customer satisfaction is key to any online operation (where customers are just a click away from a competitor) positive files represent an important tool to ensure unnecessary delays in processing valid orders.

FraudShield provides a complete environment to maintain and utilize negative and positive lists. Items can be added or deleted from individual lists via an easy-to-use user interface. The lists are then used in FraudShield rules to determine the appropriate action to take when a match is found. For example:

```
IF CardNumber IN "Chargeback-Cards-List" or Email IN  
"Fraud-Emails-List" THEN Decline
```

Lockout Mechanisms

Automatic card number generators represent one of the new technological tools frequently utilized by fraudsters. These programs, easily downloadable from the Web, are able to generate thousands of “valid” credit card numbers. A valid credit card number is a 14 or 16 digit sequence that begins with an assigned BIN (Bank Identification Number) and is compliant with the so-called MOD10 check, a standard consistency check used by issuers. Basically, these programs are able to generate numbers that a bank *might have* issued; however, the vast majority of these numbers correspond to nonexistent account numbers. The fraudster would typically target a site and submit same-amount transactions just to test out a sequence of numbers. Once the processor authorizes a valid number corresponding to an active account, the fraudster will begin charging the card to the limit on the same site or on other sites.

The traits of a card number generator attack are the following:

- Multiple transactions with similar card numbers (e.g. same BIN)
- A large number of declines
- All transactions will fail an AVS check (because the fraudster typically does not have access to the actual cardholder billing address data)

Merchants can put in place prevention mechanisms specifically designed to detect number generator attacks. When an attack is detected, the merchant should attempt to respond immediately by blocking further orders originating from the suspected fraudster.

FraudShield provides automatic lockout mechanisms specifically designed to defend merchants from attacks perpetrated using card number generator programs.

The system detects sequences of declines occurring within a user-configurable period of time. A sequence is detected based on same card number, same source IP or customer ID. When a sequence is detected, the card, the IP or the user ID is then locked out of the system for a user-defined period of time.

Risk Scoring

Risk scoring tools are based on statistical models designed to recognize fraudulent transactions, based on a number of indicators derived from the order characteristics. Typically these tools generate a numeric score indicating the likelihood of an order being fraudulent; the higher the score the more suspicious the order.

Neural networks, a class of non-linear statistical models, have been widely used in the past by issuers and fraud detection vendors to develop systems that helped reduce fraud in the brick-and-mortar world. Today the same technology is being applied to Internet fraud and made directly available to merchants.

Risk scoring systems provide one of the most effective fraud prevention tools available. The primary advantage of risk scoring is the comprehensive evaluation of an order being captured by a single number. While individual fraud rules typically evaluate a few simultaneous conditions on an order, a risk-scoring system arrives at the final score by weighting several dozens of fraud indicators, derived from the current order attributes as well as card number historical activities.

The second advantage of risk scoring is that, while a fraud rule would either flag or not flag a transaction, the actual score indicates the degree of suspiciousness of each transaction. Thus, transactions can be prioritized based on the risk score and, given a limited capacity for order review, only those with the highest score would be reviewed. Furthermore, the risk score allows a finer control over order review decisions and the ability to take different actions based on the risk level of an order.

Finally, risk-scoring systems deliver the “statistical knowledge” contained in extensive databases of historical transactions, and fraudulent ones in particular. These databases are essential in the design and tuning of the neural network models that are behind risk-scoring systems. These models are basically “trained” by using examples of both legitimate and fraudulent transactions and are able to correlate and weight the various fraud indicators (e.g. unusual transaction amount, card history, etc) to the occurrence of fraud. Thus, the implicit information contained in these historical databases is made available in the form of a score.

The ClearCommerce Engine provides risk-scoring capabilities with the FraudAnalyzer component. FraudAnalyzer utilizes a neural network model to generate, in real time, a normalized risk score between 0 and 100. The neural network model that is the core of FraudAnalyzer was designed and trained using a very large, multi-merchant, database that contains exclusively Internet transactions. The FraudAnalyzer risk score can be used to create FraudShield fraud rules that flag orders based on the value of the score, or a combination of score value and other checks based on order attributes.

Risk scores can be used independently or in combination with other transaction attributes. For example, a merchant may want to create rules that flag orders at different risk score levels based on the type of goods being purchased in the order, with lower thresholds for more risky items, such as DVD players, and higher thresholds for low-risk items, such as books.

Case Management

Successful and cost effective fraud prevention cannot be achieved by exclusively implementing one of the several fraud detection tools currently available. Effective fraud prevention is achieved by implementing an overall *risk management process*, which is enabled by both fraud detection algorithms and case management tools.

As it was mentioned earlier, today’s best practice approach to fraud prevention is based on an *arsenal* of automated fraud screening techniques to maximize the efficiency of the *order verification* process. The first key point is that every fraud detection technique has its limitation; no tool can identify all types of fraud and every tool is best at identifying one particular type of fraud. Only by compounding the benefits of an array of tools can merchants effectively defend themselves from the threat of fraud.

The second important observation is that fraud detection tools should be used for automated screening, but only in limited circumstances for automated decision-making (e.g. rejecting an order). The reason for this is that every detection tool will inevitably flag some legitimate transactions as potentially fraudulent. This is inherent with the nature of the problem itself, since fraudulent orders represents (in most cases) a tiny percentage of all orders. This makes it very difficult to identify a significant percentage of the fraudulent orders without also flagging a large number of valid orders. Too often, perfectly legitimate orders look just like suspicious ones. The conclusion is that automatic rejection based on fraud screens is not a viable approach for most Internet merchants, since the risk of declining valid orders is unacceptably high and potentially more damaging than the risk due to fraud itself.

Fraud Prevention Workflow

To better understand the role of human review in an online risk management implementation, it is helpful to first describe the overall order lifecycle, or workflow, through the system.

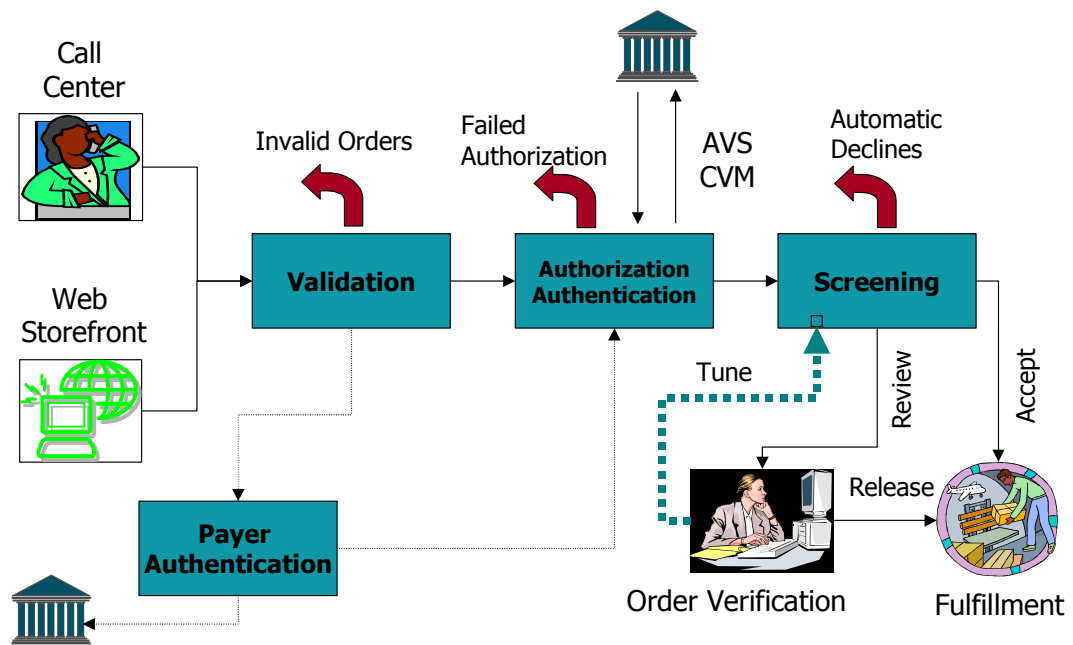


Figure 2: Fraud prevention workflow

The first step of the process is order Validation; at this stage various checks are applied to the order, primarily to discard orders that contain invalid payment information and therefore would be certainly rejected if submitted for authorization. Further checks may include the presence and validity of the billing information, required for AVS validation, or compliance with merchant business policies (e.g. a merchant may decide not to ship certain products abroad). Negative files can also be used at this stage to ensure that the card number is not in a “hot cards” list.

FraudShield provides a variety of built-in checks for order validation, including:

- Card number validity (MOD10 check)
- Zip code verification
- Duplicate orders
- Check for well-formed email addresses
- User-defined rules

After an order is validated, if a merchant is implementing payer authentication schemes such as VPAS, the user will interact directly with the bank to authenticate the transaction. If the order is successfully authenticated a special authorization code will be returned to the merchant web storefront

After an order is validated, a real-time authorization will be obtained by communicating with the card processor. Real-time authorization is preferable to batch authorization because it provides critical information for fraud screening when the order is received. Some of the orders may fail to obtain authorization and will be declined by the processor. This eliminates the need to further evaluate the order for fraud purposes since the transaction could not be completed. Real-time authorization also provides access to processor level security mechanisms (AVS and CVM response codes) that can then be used to scrutinize approved transactions.

Since FraudShield is integrated with the ClearCommerce payment engine, real-time authorization results are readily available for fraud screening. This is only one of the advantages derived from a solution that provides payment services and fraud screening within the same integrated platform.

The next stage of the management process is screening. This is the part of the process where the “arsenal” of automated fraud screening tools is put into action. Fraud rules, block files, neural network scoring, lockout mechanisms and other screening tools are simultaneously applied to every incoming order for a comprehensive risk profiling. The vast majority of the orders examined by the screening tools will pass all checks and will be forwarded to order processing and fulfillment. A certain percentage of the orders will trigger one or more fraud checks and will be routed to order verification for further examination. The percentage of orders that will require manual intervention because of suspected fraud will vary from merchant to merchant, depending on the amount of fraud prior to the implementation of the risk management process, the amount of fraud that is deemed tolerable by the merchant, and the tuning of the fraud screens. Typically a merchant in a high-risk industry for online commerce (e.g. consumer electronics, jewelry, travel, etc) will review 10 to 20 percent of the approved orders.

Orders flagged by the fraud screens will be inspected by members of a fraud prevention team; typically customer service representatives trained specifically to identify fraudulent orders. The billing and shipping information will be validated and the customer may be contacted by email or phone for further verification. Some of the orders reviewed will be quickly recognized as legitimate orders and approved. However, the majority of the fraud attempts will be uncovered during the verification and stopped before the order is processed.

The final element of the risk management process described in Figure 2 is the feedback loop from the order review to screening. As lessons are learned from both detected fraud attempts and chargebacks due to undetected fraud, more knowledge and data will flow back into the automated fraud screens, improving the effectiveness of the system. Fraud is a moving target, with new schemes emerging every day while established one are understood and prevented. Therefore this feedback loop is critical to maintain effectiveness of the the entire risk management process.

Case Study – Cost of Fraud

Risk management involves some costs in order to avoid fraud-related losses. Therefore it is important to understand how these various costs and savings interact in a return on investment (ROI) analysis for a risk management solution.

Merchant ABC processes an average of 25,000 orders per month and in the last 6-months has incurred a sustained 0.6% chargeback rate due to fraudulent orders.

The merchant incurs other costs due to fraudulent orders. With an average fraudulent order size of \$100 and a profit margin of 10%, the merchant loses, on average, \$90 for the cost of goods sold, plus an average of \$12 for shipping costs. Furthermore, since each chargeback requires on average an hour of administrative work to be settled, the merchant is also incurring an additional cost of \$25 per chargeback. Finally, the acquiring bank charges the merchant a surcharge for every chargeback incident; the merchant incurs an additional \$10 cost for every fraudulent order. Overall, each fraudulent transaction is costing the merchant an average of \$137, or an annual total of \$246,600.

To address the problem, the merchant decided to set up a small review team to screen suspicious orders. At an average loaded cost of \$25/hour, a reviewer can typically screen 10 transactions per hour, leading to a review cost of \$2.50 per order. If the merchant could not rely on a screening tool, the only option to reduce fraud would be to randomly select orders for review. However, since the overall per-transaction cost of fraud is, in this scenario, 82 cents per transaction (total cost/total orders), spending \$2.50 to review an order will guarantee a negative financial result. However, if the merchant is able to selectively apply the review process only to a relatively small segment of high-risk orders, the review process can deliver a positive financial outcome.




Figure 3: Minimizing the Cost of Fraud

Risk scoring, rules and other detection techniques can help merchants identify “high-risk” orders. The optimal percentage of orders to be reviewed depends on several factors, including the accuracy of the detection tools. The key to maximizing the return on investment is to identify the point of diminishing return, namely the level of risk that equals the cost of review. Figure 3 shows the total cost of fraud (losses due to fraud plus cost of order review) as a function of the percent of orders reviewed. In this particular example, the total cost is maximized when reviewing 5.5% of the orders, which also corresponds to a detection of 40% of the fraud. The chart also shows the corresponding total cost if random selection is used. Notice that any strategy based on random selection fails to decrease the total cost.

By implementing a review process based on fraud screen, the merchant will spend an average of \$3,437 a month for order reviews, but at the same time it will cut the cost due to fraudulent orders by \$8,220 per month. On an annual basis, the implementation of a risk management process based on targeted fraud screens will reduce the overall cost of fraud from \$246,600 to \$189,000. In practice, the merchant is able to take control over chargebacks, reducing the overall fraud rate to 0.36%, and save a significant amount of dollars by implementing an efficient process.

SCENARIO	No Review	Random Review	Score-Based Review
Avg. Monthly Sales	\$2.5M	\$2.5M	\$2.5M
Net Revenue	\$250K	\$250K	\$250K
Cost of Fraud			
Goods and Shipment	(\$15,300)	(\$9,180)	(\$9,180)
Penalties and Administrative	(\$5,250)	(\$3,150)	(\$3,150)
Total Cost of Fraud	(\$20,550)	(\$12,330)	(\$12,330)
Risk Management Cost	\$0	(\$25,000)	(\$3,437)
Fraud Losses Reduction	\$0	\$8,220	\$8,220
Savings	N/A	(32%)	239%

Table 1: ROI analysis

Analysis and Tuning

The most successful risk management organizations are those who regularly analyze and tune their screening methods based on historical data. While substantial knowledge can be gathered by fraud analysts as they deal with new fraudulent schemes, more subtle fraud patterns can be detected by deeply analyzing historical transaction data and fraudulent orders. Leveraging information about both fraudulent orders prevented by early detection and perpetrated fraud that resulted in chargebacks is key to identify trends and patterns and then to using them to fine tune the detection algorithms.

Analytics can help improve the fraud detection process through:

- tuning and discovery of fraud rules;
- tuning of risk scoring models; and
- review process improvement.

Tuning Rules

Fraud analysis can be beneficial for merchants implementing rule-based detection because it allows them to fine tune existing rules and discover new ones that can stop additional fraud. The tuning of rules is typically done by using specific performance metrics to assess how well each of the rules in production has performed in the past. For example, one may utilize metrics such as *selectivity* and *detection*. Selectivity indicates what percentage of orders are flagged by the rule, while detection indicates what percentage of those orders turn out to be actually fraudulent. By retrospective analysis one can assess how each rule has performed in the past and project the effect of modifications to the rule in terms of performance metrics.

New rules can also be designed and tested using historical data. Rule design can be driven by analysts knowledge of recent fraud cases, but they can also be extrapolated directly from the actual data by using sophisticated data mining algorithms.

Tuning Models

It is well known that predictive fraud models deliver best results when they are trained on merchant-specific data and are kept current with emerging fraud patterns by periodic retraining. Current data on known fraudulent orders provides an essential feedback for predictive models based on neural networks, which indeed can only “learn” and improve performance over time when data on the actual outcome of previously scored orders is available.

When a merchant has collected a sufficient amount of historical data and the corresponding chargeback records, it is often best to use the data to completely retrain the model. This is because the optimal weighting of certain fraud indicators may be quite different from merchant to merchant, depending on the type of business, order type, volume and fraud schemes. In these cases, a completely retuned model could perform significantly better than a model trained on data not specific to the merchant. However, to maintain such performance it is usually appropriate to retrain the model 2 to 4 times a year, using the latest available data.

Tuning Order Review

When a merchant implements an order review process for suspicious transactions, it is important to monitor the process using appropriate metrics, to identify aspects of the process that could be improved. For example, one typical measure is the percentage of “leakage”: namely fraudulent orders that were held for review but then, inappropriately, released for approval. This often represents a significant percentage of the chargeback orders that cannot be addressed by simply improving the automatic screening algorithms. In this case merchant should be focusing on training, establishing guidelines for the order verification process and providing risk analysts with appropriate research and investigation tools. Other metrics may focus order resolution time, throughput per person, etc.

ClearCommerce offers a variety of risk management services to help customers maximize the benefits of the ClearCommerce Risk Management solution by analyzing historical data.

- The *Risk Management Assessment* package is based on quantitative data analysis of historical data to assess the performance of any existing rules and generate tuning recommendations. The service also includes the design of new rules, designed to both reduce the number of false positives and increase the detection of fraudulent orders. This is done via an established analytical process to ensure timely delivery of high quality results.

- The *FraudAnalyzer Tuning* package offers a complete retraining of the standard FraudAnalyzer neural network model, while still leveraging the extensive Data Consortium database. This service allows merchants to maximize the benefits from their investment in FraudAnalyzer and it also includes ongoing maintenance of the model.
- The *FraudAnalyzer Custom Model* package, provides the most sophisticated risk scoring solution, which includes the design of additional, customer-specific, fraud indicators for FraudAnalyzer and the implementation of a custom model. Custom models may include special indicators, like product-specific flags, external authentication results, identity verification responses, etc.

Traits of Fraudulent Orders

Although new online fraud schemes continue to emerge every day, merchants should pay attention to certain characteristics of an online order that have consistently been associated with cases of fraud. Although each of these “red-flags” per se is not sufficient evidence of a fraud attempt, an order for which multiple of these conditions are met should definitively raise an alert. The following “Decalogue” of fraud prevention lists some of the most common indicators used in online fraud prevention.

- **Larger than normal orders:** a fraudster is always going to maximize the size of the bounty, because a stolen credit card number may only be active for a short amount of time. Watch for suspiciously large orders that look too good (for the merchant) to be true. If possible, limit the dollar amount of an individual order.
- **Orders containing several units of the same item:** an unusual number on units of the same item being ordered at once may indicate an attempt to “stock up” on goods that could be sold on the black market or Internet auctions. Consider what type of order is being purchased: why would a legitimate customer want to order 25 copies of the same DVD? On the other hand, ordering 25 different DVDs may simply indicate that you have a new movie-lover customer.
- **Orders shipped overnight:** since ultimately they are not going to pay for shipping charges, fraudsters would typically request the most expensive and expedite delivery method. This also gives the merchant a smaller time window to investigate and possibly put the order on hold.
- **Orders shipped to an address other than the billing address:** even if the fraudster is able to obtain complete cardholder information (identity theft), in order to pass the AVS check, the goods will not be shipped to the billing address. The further an order is shipped from the billing location the more attention should be paid to the order. Nowadays, it is not unusual for someone to have orders shipped, for example, to their office because no one is at home to pick up packages during working hours. However, if an order is shipped to a different state, or even country, the inherent risk level increases.
- **Change of destination:** at times, in order to avoid notice, a fraudster would first place orders using the billing address as the shipping address. After the merchandise leaves the merchant, the fraudster will call the shipping company and reroute the order to a different address. If possible, have the shipping company notify any destination changes to you.
- **Orders that failed AVS verification:** although a failure to validate an AVS verification is not, by itself, sufficient evidence of fraud, orders on US-issued credit cards that fail AVS should be closely scrutinized. In most cases the AVS fails because the consumer has mistyped the billing information or has not updated the billing address after a change of address.

- **Anonymous email addresses:** the majority of free email services allow users to create an account without providing verifiable personal information. Furthermore, most of these services typically will not disclose logging information that may help trace a fraudster that was using the service. Therefore merchants should be more suspicious of orders where an anonymous email address is provided and attempt to verify both billing and shipping information.
- **Multiple orders on the same card, in a short amount of time:** whenever a fraudster has access to a valid and active credit card number, he or she will try to charge the card to the limit in a short amount of time, before the account is deactivated. Card numbers with unusually high activity (specially during a slow season) should raise an alarm.
- **Multiple card numbers from the same IP:** a sequence of orders originating from the same IP address in a short period of time, especially if placed on multiple credit card numbers, may indicate an ongoing fraudster attack. However, merchants should be aware that an hyperactive IP address may also indicate a proxy server of a large organization or service provider.
- **Multiple orders to the same shipping address, on multiple cards:** although it is possible that a legitimate consumer may utilize multiple credit cards to place a series of orders, this situation should raise an alarm if the number of cards goes beyond 3 or 4. Merchants should also be aware that fraudsters will slightly change the format of the address to avoid notice (e.g. “1204 Main St.” in one order and “1204 main street” for the next order).

International Fraud

International fraud probably represents the fastest growing problem for Internet merchants today. Three main factors contribute to the problem: first, the unlimited reach of the Internet, which allows fraudsters and hackers to reach commerce-enabled sites from any location in the world. Second, because some countries do not have the resources to investigate and prosecute online fraud cases, organized criminal organizations find an ideal haven in these countries to pursue their scams undisturbed. Thirdly, since standard cardholder validation tools, such as AVS, are only supported for US-issued credit cards, online merchants have a limited ability to detect suspicious transaction.

Even merchants that do not ship products abroad nor accept non-US credit cards are not immune to the problem, as international fraudsters devise very complex schemes to circumvent these restrictions.

- A very recent “triangulation” scheme used online auctions to perpetrate fraud. The fraudster auctioned dozens of consumer electronics products, at deeply discounted prices, on a popular online auction site. When a person (unaware of the scam) placed a winning bid for the product, the fraudster placed an order at the manufacturer’s web site, using a compromised card number, requesting the goods to be delivered directly to the winner of the auction. The auction winner paid the fraudster via a person-to-person money-transfer service that routed the funds to a bank account in an Eastern European country. The manufacturer shipped the goods to the winner of the action but obviously several weeks later the received a chargeback for the transaction. In this case the fraudster, probably located in Eastern Europe, never touched the goods as he or she successfully orchestrated the scam from thousands of miles away.
- Another recent case of international fraud involved an organization that was probably based in Nigeria. The fraudsters placed dozens of relatively small-size orders using US credit cards. All the orders had a US delivery address (the merchant did not ship internationally), specifically three different addresses which corresponded to three international freight forwarders based in different US locations. When the scheme was uncovered and the shipment companies contacted, the merchant discovered that the products had been shipped to Nigeria.

Statistics on International Fraud

ClearCommerce recently conducted a study to quantify the problem of international online fraud and investigate the use of IP (browser network address) geographic location as a tool for detecting these fraud cases. The analysis leveraged the extensive historical database of the ClearCommerce Data Consortium and a IP database that maps network addresses to the country to which each IP sub-network is assigned.

The first staggering result was that, although transactions originating outside the US represented only 17% of the sample, they accounted for over 42% of all fraudulent orders. International fraud also represents a growing percentage of all the online fraud cases: in the last quarter of 1999 international transactions (based on the IP originating country) represented only 10% of all transactions that resulted in chargebacks. In the last quarter of 2000, international transactions represented nearly 50% of all chargebacks.

International fraud is highly localized: the majority of fraudulent international orders originate from a relatively short list of countries. Figure 4 shows the top high-risk countries, ranked by the percentage of orders originating from that country that resulted in chargebacks. For reference, the last bar in the chart shows the fraud rate for orders originating from IP addresses located in the US. From this chart it is clear that the risk associated with orders originating from certain foreign countries is of different magnitude compared to domestic orders. It is also important to notice that the top ten high-risk countries account for over 35% of all the international fraud, although less than 5% of all the international orders originate from these countries.

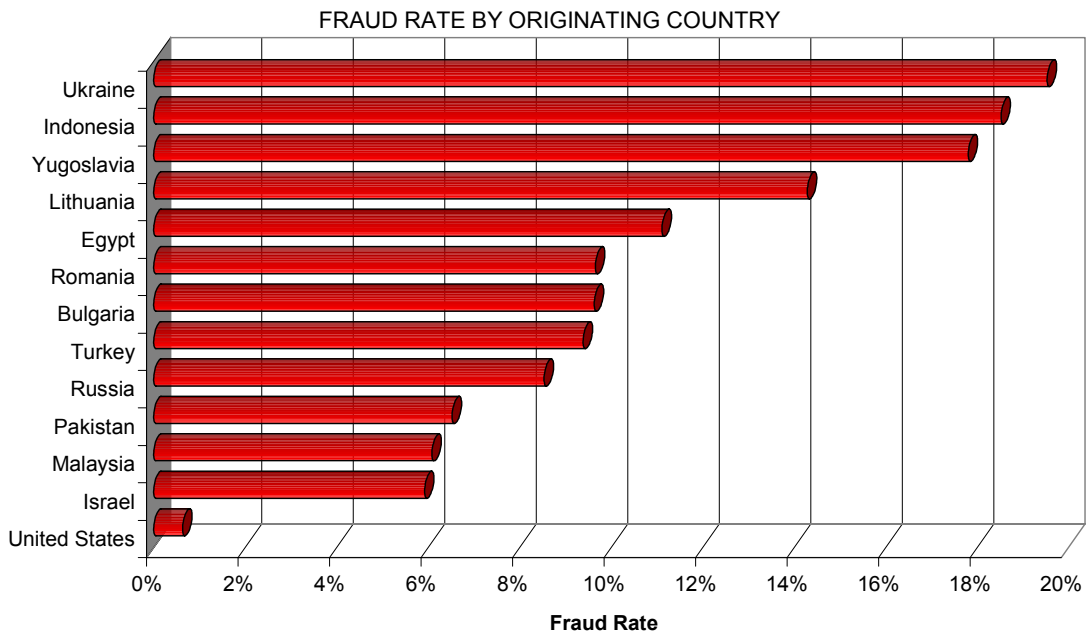


Figure 4: Top High-Risk Countries

Another fact that the analysis demonstrated, confirming evidence from the two reported cases, is that international fraudsters are often able to completely mask their transactions as legitimate domestic orders. Over 25% of all fraudulent orders originating outside the US used credit card number issued in the US. Furthermore, 23% of all the international orders had US shipping addresses. The good news is that this type of inconsistency between the billing addresses, shipping addresses and the geographic location on the IP address can be exploited to detect suspicious orders and therefore prevent international fraud.

Uncovering International Fraud with FraudShield GeoLocator

The examples of international fraud schemes discussed above demonstrate how difficult it is for Internet merchants to uncover international fraud schemes. In many cases the fraudsters have access to compromised credit card numbers from US banks. Furthermore, in some cases (identity fraud) the fraudsters have access to complete billing records that allow them to pass an AVS validation. Furthermore, international fraud schemes often use delivery points in the US and later forward the goods to a foreign country, or, like in the auction scheme, they do not even need to receive the goods. Therefore, in many cases these fraudulent orders appear just like legitimate domestic orders.

In both of the case studies reported, the only data element that would have uncovered the scheme was the geographic location of the IP addresses from which the orders were originating. For the online auction fraud case, the scammer apparently connected from various IP located addresses outside the US, including Mexico and Eastern Europe. For the Nigerian scheme, all orders originated from IP addresses that were located in Nigeria.

GeoLocator, the addition to the ClearCommerce fraud prevention product, helps merchants uncover international fraud by leveraging IP-level data. Leveraging the most extensive and accurate IP database available, GeoLocator can resolve, in real-time, the IP address of the browser connection to the country to which the address is assigned. This powerful mechanism enables merchants to set up FraudShield rules specifically designed to uncover international fraud attempts.

The following are just a few examples of rules that a merchant can create using the GeoLocator functionality:

```
IF IP-Country IN "High-RiskCountries-List" THEN Review

IF IP-Country NOT "USA" AND BillingCountry = "USA" THEN Review

IF IP-Country NOT EQUAL TO BillingCountry THEN Review

IF IP-Country NOT "USA" AND ShippingCountry = "USA" THEN Review

IF IP-Country NOT "USA" AND AVS NOT EQUAL TO "G" THEN Review
```

The first rule in the list, screens orders based on the originating country of the connection (IP country), comparing it with a specific list of high-risk countries. The second and the third rule, detect mismatches between the reported billing country and the IP country, while the fourth rule looks for orders shipped to a US location but originating outside the US. Finally, the last rule leverages the AVS response, which enables merchants to distinguish between US and non-US issued credit cards.

The Aftermath of Fraud: Reporting and Prosecuting Fraudsters

Unfortunately, even with the most finely tuned risk management process, an online merchant will not be able to eliminate all the fraud. A question frequently asked by merchants is: “Whom should I call and what steps should I take to report the fraud and when should I report fraud?” Reporting and prosecuting fraudsters is a challenging task, but steps can be taken to facilitate this process.

- **Create an internal policy:** a merchant should establish in-house guidelines about fraud and how to handle fraud cases. Typically, if the fraud is less than an arbitrary threshold, the merchant may set up a policy to simply document the fraud case internally and add the information to a negative file. However, if the amount of the fraud incident is greater than that threshold it may be worthwhile to invest time to determine if a prosecution case can be built.
- **Provide detailed data to support the investigation:** when building a case the merchant should summarize all the information available in an easily readable format, and then report the case to a law enforcement agency. It is important to make sure that the most important details are well documented and easy to read. Two very important details required to investigate an online fraud crime are the IP address and shipping address.
- **Have realistic expectations:** when submitting a case to a law enforcement agent, the merchant should understand that detectives are measured by successful cases solved, so the more details that are provided, in an easy to read and understandable format, the more likely the agency will take the case. The size of the case is also a factor. In a large metropolitan area it is unlikely that the police department will go after small cases of fraud; whereas in a smaller sized town law enforcement might be more aggressive against smaller sized crimes.
- **Report all fraud attempts:** both attempted fraud and perpetrated fraud cases should be reported to the authorities when creating the case for the police.
- **Be helpful and respond in a timely manner:** when working with a law enforcement agent on a case, be helpful and answer all questions in a timely manner.
- **Put the fraudster on a payment plan:** internal investigations may actually identify some of the fraudsters even before creating a case. The merchant can decide to put the fraudster on some sort of payment plan to pay back the damage, or follow the legal route. Reportedly, a direct negotiation is typically more likely to achieve restitution.

- **Restitution:** the likelihood of recovering the goods in a fraud case is, unfortunately, very small. It is estimated that only 10% of the lost merchandise are usually recovered or paid back, and is very likely that an investigation will not have a positive financial payoff. However, in the long term strict prosecution policies will have a deterrence effect on Internet fraud, therefore, whenever possible, merchants should consider investigation and prosecution for fraud cases.

More Information

About ClearCommerce

Austin, TX based ClearCommerce is a leading provider of e-commerce risk management and payment- processing software, serving more than 40,000 businesses worldwide. ClearCommerce Payment software is the most scalable, cost-effective, and efficient solution available to automate processing of all major online payment methods and currencies. ClearCommerce Risk Management combine the industry's most powerful arsenal of online risk management technologies and provides total business user control of the screening process. All ClearCommerce software is designed to provide enterprise class reliability, availability, scalability, and performance and to easily integrate with front-end and back-end business applications.

ClearCommerce Corporation
11500 Metric Blvd., Suite 300
Austin, TX 78758

Phone – 888-725-8612

Email – info@clearcommerce.com

Web Site – www.clearcommerce.com

References

Contact ClearCommerce to request any of the following additional white papers:

- *ClearCommerce Engine Features and Benefits*
- *ClearCommerce FraudAnalyzer*

Document History

Document	<i>ClearCommerce Fraud Prevention Guide</i>
Product Version	ClearCommerce Engine 4.3
Author	ClearCommerce Product Management
Revision Date	January 2002
Previous Versions	<i>ClearCommerce Fraud Prevention Guide 4.1</i>

Notices

ClearCommerce® is a registered trademark and FraudShield™ is a trademark of ClearCommerce Corporation.

Other trademarks are property of their respective owners.