

WHITEPAPER

TrustDefender™ ID

Risk-based Digital Identity Verification

ThreatMetrix™ Cybercrime Defender Platform



Contents

Overview	3
Device Identification	3
1st and 2nd Generation Device ID Strategies	3
Cookies	4
IP Addresses	4
Flash Cookies and Browser Cookies	4
SmartID: 3rd Generation Device Identification	4
SmartID™	5
Hidden Proxy Detection, True IP, and True Geo-location	6
TrustDefender ID Rules Engine	6
Policies and Rules	6
Use Cases and Customizations	7
TrustDefender ID API and Protection of PII	7
Policy Score, Response and Reason Codes	7
Leveraging the ThreatMetrix Global Network	8
Global Velocity Rules	8
Global Reputation	8
Indirect Reputation: Review Status	9
Identity Verification: Step-up Screening	9
Reviewing TrustDefender ID Transactions	10
Reports and Dashboard	10
Searching and Link Analysis	10
Whitelists and Blacklists	10
Queues and Alerts	10
Conclusion	11

Overview

Every organization that transacts with an online visitor needs fraud screening. The costs are simply too high not to: direct loss, chargeback costs, data loss associated with account taker over, and potential brand damage. But a poorly executed fraud defense can have high costs as well. Inaccurate or intrusive fraud screening may alienate customers and users, cause lost sales, and incur expensive manual review costs.

TrustDefender™ ID from ThreatMetrix™ helps organizations meet this challenge. It provides a cost-effective first perimeter of defense against online fraud without changing how visitors interact with your site. It identifies suspicious transactions in real-time, and categorizes the transaction as: accept, reject, or review, allowing organizations to optimize their manual review process. For high-risk transactions, TrustDefender ID can perform additional identity verification.

This paper provides an introduction to TrustDefender ID and explains various aspects of how it is effective against fraud, including:

- TrustDefender ID device identification and SmartID
- Hidden proxy detection and True IP
- Real time fraud screening using the rules engine
- Leveraging the ThreatMetrix Global Network
- Identity Verification: step-up screening
- Reviewing transactions

Device Identification

Accurate, subversion-resistant device identification is a valuable tool for combating fraud. It enables organizations to recognize returning visitors and then further classify them as returning good customers or suspicious visitors.

1st and 2nd Generation Device ID Strategies

To understand how TrustDefender ID reliably identifies devices, it is illustrative to first look at device identification techniques that are no longer effective.

Cookies

In the early days of online commerce, browser cookies were used as a device identifier. By setting a cookie in a visitor's browser, it was easy to recognize that browser whenever it came back. But it was not long before even unsophisticated Internet users learned how to remove cookies. Today, all widely used browsers support "private browsing" modes, so reliance on browser cookies to detect return visitors is not effective at all.

IP Addresses

Organizations next turned to tracking visitors by their IP address, which is readily available to any software processing web requests. But many consumers use dynamic IP addresses, which change frequently, and IP addresses can be spoofed. Today, there are many proxy services available all over the world that will hide a web visitor's true IP address. ThreatMetrix has empirical data that the majority of fraudulent transactions occur over Internet proxies. Relying on an IP address to detect return visitors wishing to hide is not effective.

Browser cookies and IP address represent the 1st generation of device identification.

Flash Cookies and Browser Cookies

Next, device identification techniques relied on the general public's and unsophisticated fraudster's ignorance of Flash cookies (Local Storage Objects) that are not deleted when regular browser cookies are cleared, and are invisible unless you know where to find them.

Today, however, browser private browsing modes temporarily suppress cookies and Flash objects and hence evade re-identification. Additional data collected from the browser through JavaScript and hashed into an identifier were next used to try to recognize returning visitors. Unfortunately, these attributes are easily subvertable and change frequently, rendering this 2nd generation of device identification only partially effective.

SmartID: 3rd Generation Device Identification

TrustDefender ID has developed a much more robust and reliable device identification technique. It combines various persistent cookies, browser and plugin attribute collection, and TCP/IP packet inspection with proprietary analysis and matching algorithms to generate 3rd generation device identifiers in real-time.

Customers simply add simple HTML tags on web pages that will be loaded by their visitors. These tags have no visible impact to the visitor's experience, but they enable both passive (IP/TCP/HTTP Profiling) and active (JavaScript, ActionScript) inspection of a visitor's device.

The requests made by the profiling tags are handled by profiling servers in the ThreatMetrix secure data center. There is no need for TrustDefender ID customers to either configure or maintain the server environment to support it. The TrustDefender ID profiling servers perform a comprehensive examination of individual fields in the TCP/IP packet headers instead of relying on IP address information taken from HTTP header found in web server logs or JavaScript.

During profiling, the data collected from the packets is compared with subsequent packets received to establish true, reliable attribution to the operating system and connection. These attributes are critical contributors to deriving the TrustDefender ID device identifier. Stateful packet inspection delivers more complete, reliable device data not attainable using typical profiling techniques.

SmartID™

The end result of this data collection and analysis is that TrustDefender ID assigns a device identifier to each transaction. If TrustDefender ID recognizes the device, it returns its previous assigned identifier. Otherwise, it is a new device, and TrustDefender ID generates new, unique identifiers. TrustDefender ID actually generates two separate identifiers for a device: an ExactID™ and a SmartID™:

- **ExactID:** an identifier derived from a variety of persistent objects (browser, plugin, HTML5 storage). It is useful for recognizing returning good customers – visitors that are not specifically attempting to cloak themselves. Because it is based on persistent objects, it is 100% accurate in identifying return visitors.
- **SmartID:** an identifier derived from the collection and analysis of multiple browser, plugin, and TCP/IP connection attributes. SmartID is useful for recognizing visitors attempting to cloak themselves. It works even when cookies have been wiped or visitors use private browsing modes.

Both ExactID and SmartID can be used within the rules engines (discussed next). A common pattern is to first attempt to identify a visitor by the ExactID. If TrustDefender ID has seen the ExactID before, then it is a confirmed return visitor.

Otherwise, TrustDefender ID attempts to match by the SmartID. If there is a match, then it is a likely case of a return visitor that has wiped all cookies or is using private browsing. TrustDefender ID returns a confidence score.

If there is no match for SmartID or ExactID, then it is likely a new device, and new values are generated for both identifiers.

Hidden Proxy Detection, True IP, and True Geo-location

The attributes TrustDefender ID collects during profiling and packet inspection also enable detection of visitors connecting through hidden proxies. In most cases, TrustDefender ID can pinpoint the originating IP address of the visitor (not just the IP address of the proxy).

This is critical to determine true geographical location of the visitor. TrustDefender ID uses an IP/Geo database to return the city, region, country, and longitude and latitude coordinates of the visitor. When a proxy is detected, this same information is returned for the proxy location as well. The TrustDefender ID rules engine can use this data to identify suspicious anomalies.

TrustDefender ID Rules Engine

As described above, TrustDefender ID delivers SmartID, ExactID, proxy detection, and true IP geographical information for a web visitor in real time. By themselves, these features are extremely effective for fraud detection. But to catch more fraud, one needs to consider more than a single transaction in isolation. When the past transactional history of a device and other attributes associated with a given transaction are analyzed, patterns emerge. TrustDefender ID includes a rules engine that provides such analysis, in real time.

Policies and Rules

The TrustDefender ID rules engine is a high-throughput rule and pattern recognition engine that draws on device ID and transaction intelligence to assess fraud risk in real-time. Each transaction is evaluated against a policy – a collection of individual rules that identify anomalies. Examples of anomalies that TrustDefender ID rules can identify include:

- Is a hidden proxy being used?
- Are the visitor and the proxy in different countries?
- Has the device made a high velocity of a given type of transaction in a short period of time?
- In the last 24 hours, has the device used a suspicious number of IP addresses?
- In the past month, has the device been used with a suspicious number of names | emails | payment accounts?
- Has the IP address recently been associated with suspicious activity at other sites?
- Has the device been associated with fraud from global reputation reports?

Use Cases and Customizations

TrustDefender ID includes default policies for common use cases. These use cases, and descriptions of threats they protect against, are as follows:

- **New Account Opening:** synthetic or stolen identities
- **Payments:** online credit card, auction, alternative payments and money transfer fraud
- **Login:** account takeover from phished or stolen credentials

Customers can customize and extend these default policies with a graphical, web-based editor that requires no coding. Using the rules editor, customers can customize rule parameters and weights and enable or disable rules. Also, the rules editor can create new policies (from existing ones or from scratch) to cover new use cases or to maintain different policies for specific web properties.

Also, policy changes may be immediately deployed, with no downtime - an important capability for keeping up with changing threat models and new attacks.

TrustDefender ID API and Protection of PII

To initiate fraud screening, ThreatMetrix customers make an API call to TrustDefender ID when processing a web visitor's request. The API call may optionally pass transaction attributes to aid in fraud screening, such as: name, email, transaction amount, or one-way hash of payment account number. These attributes can be used in TrustDefender ID rules to correlate transactions across the global ThreatMetrix network.

Some of these data elements are sensitive, as they represent personally identifiable information (PII). ThreatMetrix has gone to great length to protect this information. The API is made over a secure (SSL/TLS) network connection. The data elements are stored and utilized with a combination of strong one-way hashing and encryption. This ensures that PII data used by TrustDefender ID for fraud is never exposed in clear text within the ThreatMetrix data center.

Policy Score, Response and Reason Codes

When an API call is made to TrustDefender ID, it evaluates the transaction against the specified policy, and then returns the resulting policy score, ExactID, SmartID, IP address and geo information and reason codes (indicators for the matching rules). This provides not only the score, but also the underlying data from which the score was derived. Optionally, the raw device data can also be returned.

The API response data is valuable for a number of reasons. First, it provides a score (and review status flag) for the customer's transaction processing logic to make a decision (accept, reject, review). Second, the entire response data set can be fed into the customer's own rules engine for further analysis and correlation with customer history data to enable additional screening. Third, the raw data can be stored with the customer's transaction system for future forensics and review.

Leveraging the ThreatMetrix Global Network

TrustDefender ID customers benefit from the ThreatMetrix global network in three ways: through global velocity rules, global reputation, and indirect global scores.

Global Velocity Rules

The TrustDefender ID rules engine evaluates a transaction against past transactions in order to identify anomalous behavior. A TrustDefender ID rule can be configured as Local or Global. Local rules only consider past transactions made to your site, while global rules will consider all transactions made within the ThreatMetrix network. Global rules cast a wider net to look for patterns of suspicious behavior.

As an example, consider a rule that identifies devices that have used more than two IP addresses in the last week. When configured as a local rule, the first time a device is seen at your site, it will, by definition only be associated with one IP address (the one associated with the current transaction). However, that same rule, configured as a global rule, might reveal that same device to have been associated with eight other IP addresses recently – all from transactions made at other customers' sites. For privacy, the values of the IP addresses used and the sites the device visited are not disclosed, but the count over a given time period is.

Global rules help illuminate a trail that fraudsters leave behind, even when the trail had not previously been to your site.

Global Reputation

TrustDefender ID includes an API, called assertions, for reporting fraud. Customers can optionally use this API to report suspected or confirmed fraud against a device or various attributes (such as a credit card hash, or IP address). When a transaction is screened, all reputation data associated with the transaction is available for your rules to weight: you can configure how much negative weight to associate with a negative reputation.

Indirect Reputation: Review Status

A final way that TrustDefender ID allows you to leverage its global network is to use a Global Review Status rule. This rule will match if the current transaction is related to other transactions that have

recently been scored as “reject” or “review” by other customers in the network. This is distinctly different from an assertion, because it is not confirmed fraud – it only means that a related transaction scored below a threshold.

Such a rule should be used carefully and weighted moderately (or weighted at 0), as it assumes that other customers in the network have sensible rules and weights in place. Still, it can give some insight to first time visitors to your site that you would not normally have

Identity Verification: Step-up Screening

Some transactions warrant additional fraud screening beyond using contextual device information. Given the additional risk associated with new account creation and or high-value payments, it is especially important to verify the identity information (name, phone, billing address, shipping address) supplied by the visitor for these types of transaction.

In fact, any type of transaction that involves identity information can benefit from automated screening. Identity checks are often the first manual check performed when a transaction is sent to an analyst review queue; if this can be automated, then time and money are saved. However, since this involves an API call to a third-party identity service, there are additional costs.

For these situations, TrustDefender ID has integrated a third-party Identity Verification Service as an optional rule, making precise identity verification an efficient one-step process. This service can be used conditionally, based on the transaction type, the transaction value, or any other risk factor that TrustDefender ID has previously identified, such as originating country, use of proxy or suspicious velocity.

The Verify Identity rule verifies visitor-supplied name, address and phone against an up-to-date third party database. A match, or a partial match, is a positive indicator that the supplied data is valid. No match would be neutral or negative. TrustDefender ID uses the match results to adjust the transaction risk score.

Depending on the results of the identity check, TrustDefender ID can more accurately classify the transaction as: accept, reject, or review. This capability:

- Stops more fraud in real-time
- Reduces manual review costs
- Only incurs additional fees for identity checks for high-risk transactions

Reviewing TrustDefender ID Transactions

As described earlier, the TrustDefender ID API returns a real-time response that includes all the device id, scores, reason codes and device data. In addition, this data persists in the TrustDefender ID data warehouse, and is available for online review for six months. Fraud analysts can use the data warehouse to review suspicious transactions and to analyze transactions in order to improve TrustDefender ID policies.

Reports and Dashboard

TrustDefender ID includes dashboard and reports to monitor transaction volumes, policy score, rule match counts, and geographical origin.

Searching and Link Analysis

Analysts can search for historical transactions by attribute value and date range. Common searches may be saved. Once a transaction is selected, a built-in “related transaction” tab helps identify past transactions that linked. Search results can be exported to CSV files.

Whitelists and Blacklists

An analyst, when reviewing a transaction, may easily add its attributes to whitelists and blacklists. Rules can then use this list information to score future transactions appropriately when matching attributes are encountered.

Queues and Alerts

TrustDefender ID supports queues, and the ability to route a transaction to a queue by configured criteria (such as score, or transaction value amount, or geographic region, et). Queues can then be assigned to an analyst. This helps prioritize the transactions that a given analyst should review.

Alerts provide real-time notification of high-risk transactions (based on configurable criteria). An alert sends a custom email notification from the data warehouse when a matching transaction is loaded.

Conclusion

TrustDefender ID provides online businesses with a crucial first perimeter of defense to protect online transactions. It is an effective and easy-to-implement fraud screening solution that reduces fraud and expensive manual screening

TrustDefender ID's third generation device identification technology, SmartID, provides a digital identity to your web visitors without interfering with their user experience. Its integrated, real-time rules engine analyzes a device's transactional history to look for suspicious anomalies and patterns that indicate risk, utilizing intelligence and analysis across the platform and throughout a global network of sites sharing fraud information.

Risk factors and weights can be configured and changed dynamically to respond to evolving threats. Visitor-supplied identity information can be conditionally verified against a third party identity service. The data and scores returned by TrustDefender ID can be used by itself or in conjunction with other risk management platforms.

TrustDefender ID is part of the ThreatMetrix Cybercrime Defender Platform, the first industry solution that integrates sophisticated malware detection and advanced device identification technologies in a single, unified platform.

Contact Us

USA Corporate Headquarters:

ThreatMetrix Inc.
160 West Santa Clara Street
Suite 1400
San Jose, CA, 95113
Telephone: +1.408.200.5755
Fax: +1.408.200.5799

EMEA Headquarters:

ThreatMetrix B.V.
Laan van Vredenoord 33-39
2289 DA Rijswijk
The Netherlands
Telephone: +31 (0)70 8200 508

www.threatmetrix.com

www.threatmetrix.com/fraudsandends