

WHITEPAPER

Fraud Network Whitepaper

- Verify New Account Originations
- Authorize Payments and Transactions
- Authenticate User Logins

Alisdair Faulkner Chief Products Officer

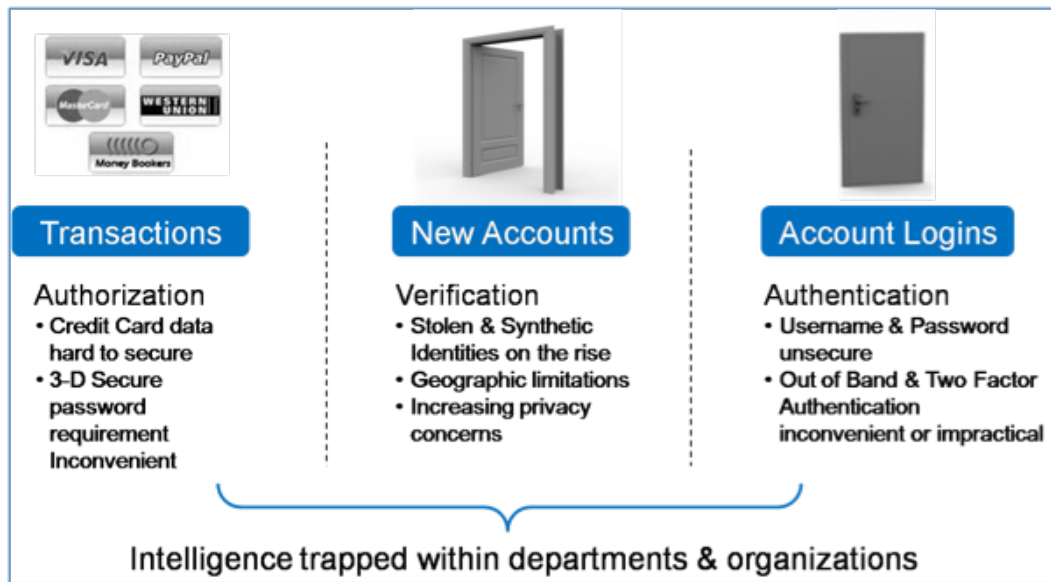
Table of Contents

The Need for A Better Way to Transact Online	3
The Tipping Point	4
High Level Requirements	5
ThreatMetrix Fraud Network	6
Profiling Servers	8
Active and Passive Inspection	8
Device Attribute Collection Methods	9
Attribute Cache	10
Device Identification Engine	11
Transaction Intelligence Engine	13
Risk Engine	14
Transaction and Device Anomalies	14
Behavioral Analytics	14
Reputation Scoring	15
API Servers	15
Conclusion	16

The Need for A Better Way to Transact Online

Every organization that transacts with a customer online must answer two critical questions: “can I trust them,” and “are they who they say they are.”

Today, every company transacting online faces this challenge of authenticating a customer, transaction or visitor in isolation while they fight the collective threat of well-equipped, networked, and technology-sophisticated organized criminals.



Online fraud is an asymmetric problem because fraudsters can quickly adapt their methods to exploit weaknesses in online fraud protection systems. Cybercriminals take advantage of automation, tools, shared knowledge, and expertise to commit crimes while their targets—virtually any organization doing business online—are at a disadvantage due to IT scheduling, shortage of trained people, fixed processes, applications and verification services that do not leverage the investment of others.

The cost of this inefficiency is billions of dollars spent on direct fraud costs, manual review, and fraud applications and services. Additionally there are billions of dollars of lost revenue due to abandonment caused by brand tarnish and protective barriers placed between the consumer’s wallet and the company’s revenue including usernames and passwords, privacy invading questions, tokens and long strings of numbers and credit card details that are hard to remember and at risk of being rejected based on a company’s fraud “policy of the day.”

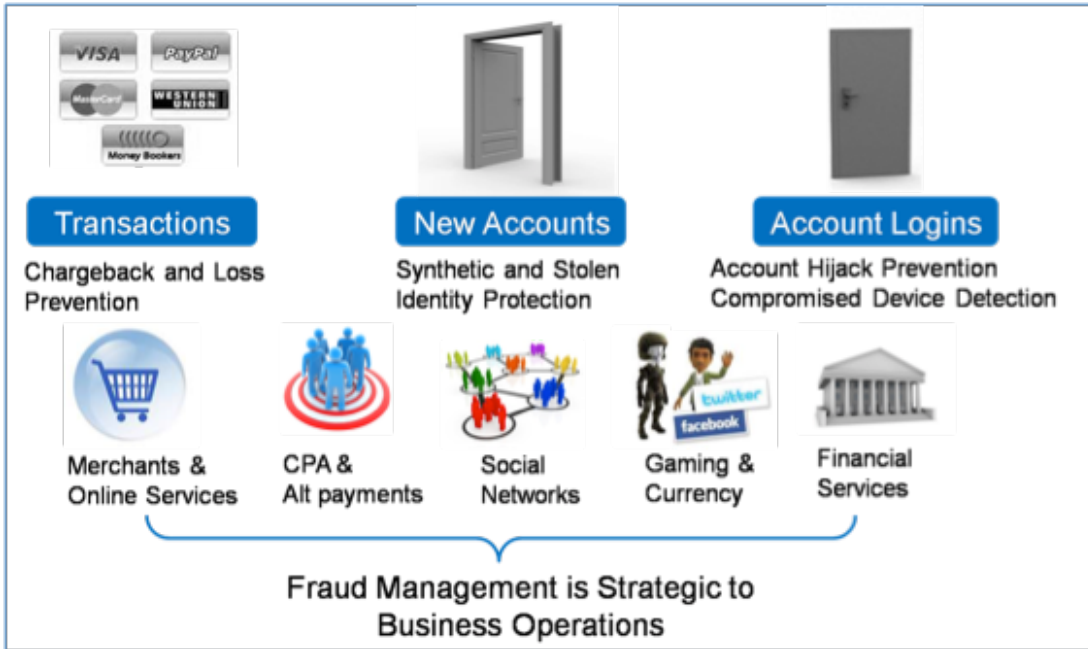
Fraud prevention is a double-edge sword. For every additional field required or verification step a customer is forced to enter, whether it's a lead generation form or shopping cart page there is always a corresponding drop in conversion due to abandonment. This revenue challenge is also compounded by false rejections once a transaction is completed. For every accepted fraudulent order, three to four times the number of domestic credit card transactions is rejected, rising to an average of 10% of all international orders rejected due to suspicion of fraud. Assuming 25% as a conservative estimate of good customers that are falsely rejected, for example travelling or using a spouse's credit card, this means that the cost of not being able to recognize a good customer is easily 1% of revenues and just as costly as direct fraud.

The Tipping Point

Today, depending on the industry and application, fraud attempts are on average 1- 6% of all online transactions. What would it mean to online commerce if fraud attempts reached epidemic proportions such as those in the email/spam realm? In that scenario a staggering 50-80% of all transactions would be fraudulent. This is not only feasible, but several trends suggest that it may be inevitable:

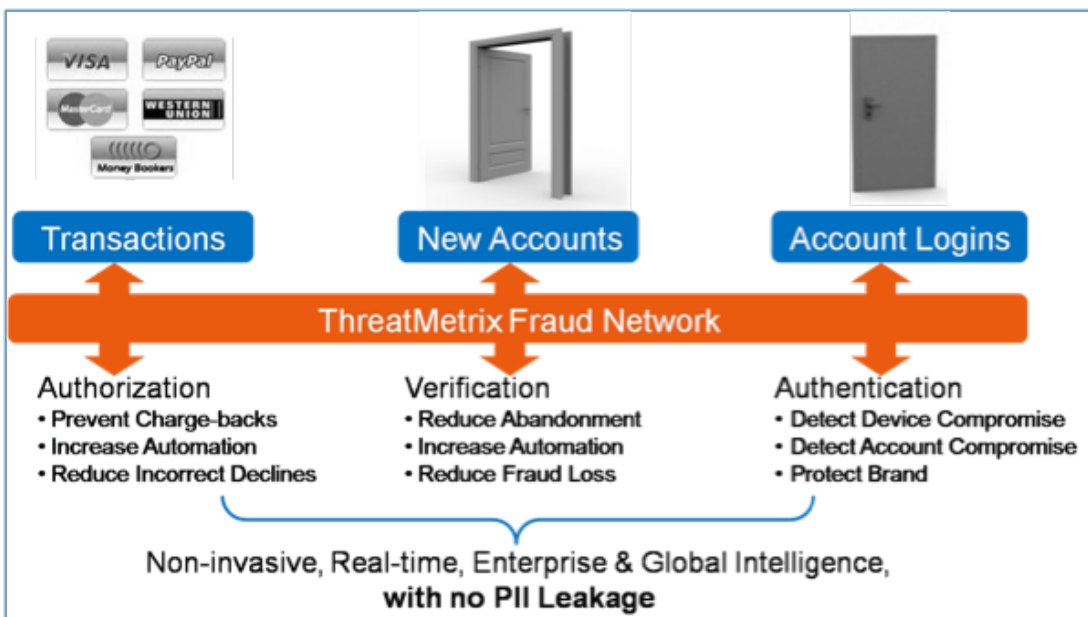
- **Open Access:** Just as an email address and inbox are openly addressable, so too are company website logins, transaction pages and account registrations.
- **Zero Marginal Cost:** Today a stolen credit card number costs as little as 40 cents on the black market, a decade ago this was as high as \$20 - 30 per card. Identity profile information such as social security information, date of birth, mother's maiden name which today retails at \$0.5 - \$3 will soon be essentially free based on the mining and compromise of social networks.
- **Increased Automation:** The advent of botnets and fraud tool kits provide an evolving and ever ready fraud threat platform, allowing fraud attempts to be spoofed from nearly every geography and computer network.

1. Javelin Research suggests New Account Fraud costs \$18 Billion, and existing account fraud 31 billion
2. Anecdotal evidence suggests for every additional field a customer needs to fill out, conversions drop by 9%
3. CyberSource 2010 Online Fraud Report
4. A ThreatMetrix beta customer in the online voice top up space had a fraud attempt rate of 80% of all transactions through its web (PC) channel
5. <http://blogs.zdnet.com/security/?p=2084> Oct 2008 CardCops: Stolen Credit Cards Getting Cheaper
6. <http://www.cbc.ca/technology/story/2009/04/14/credit-fraud.html>
7. ThreatMetrix tracks approximately 10-15m compromised computers at a given time
8. Rock Phish Kit, Metasploit, BeEF (Browser Exploitation Toolkit)



Even without a pandemic doomsday scenario, it's clear that the threat is real and very present, and the burden should not rest solely on the shoulders of fraud prevention and loss prevention staff because fraud management is not just about loss prevention, it is strategic to the operations of most major online businesses today.

High Level Requirements



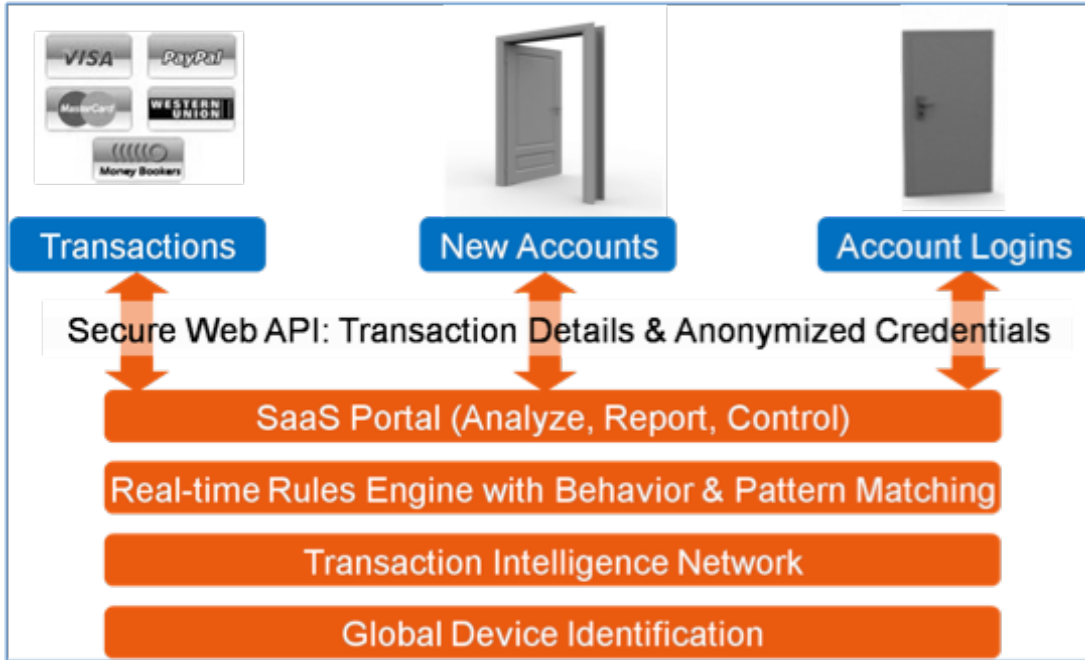
A shared intelligence network that helps to address the “fraud gap” must overcome the following critical challenges:

- **Easy to Deploy:** The network must be transparent to the consumer and easy to integrate with an organization’s legacy applications and processes.
- **No PII Required:** Because organizations face risks and legal obligations regarding data breaches and uses of customer information, a shared intelligence network must be effective without requiring personal customer data to be shared.
- **Real-time, All the Time:** The fraud network must be able to scale in real-time and be available 24 x 7. Solutions that do not operate in real-time may work for physical goods purchased online, but intangible goods that require instant authorization require real-time operation.
- **Enterprise and Global Intelligence:** Enterprises will demand shared fraud intelligence across all online touch-points including purchasing, lead generation, account creation and logins.
- **Fast Detection/Prevention Cycle:** The network must effortlessly provide upgrades as new fraud techniques are detected, without having to wait for an enterprise technology refresh cycle.

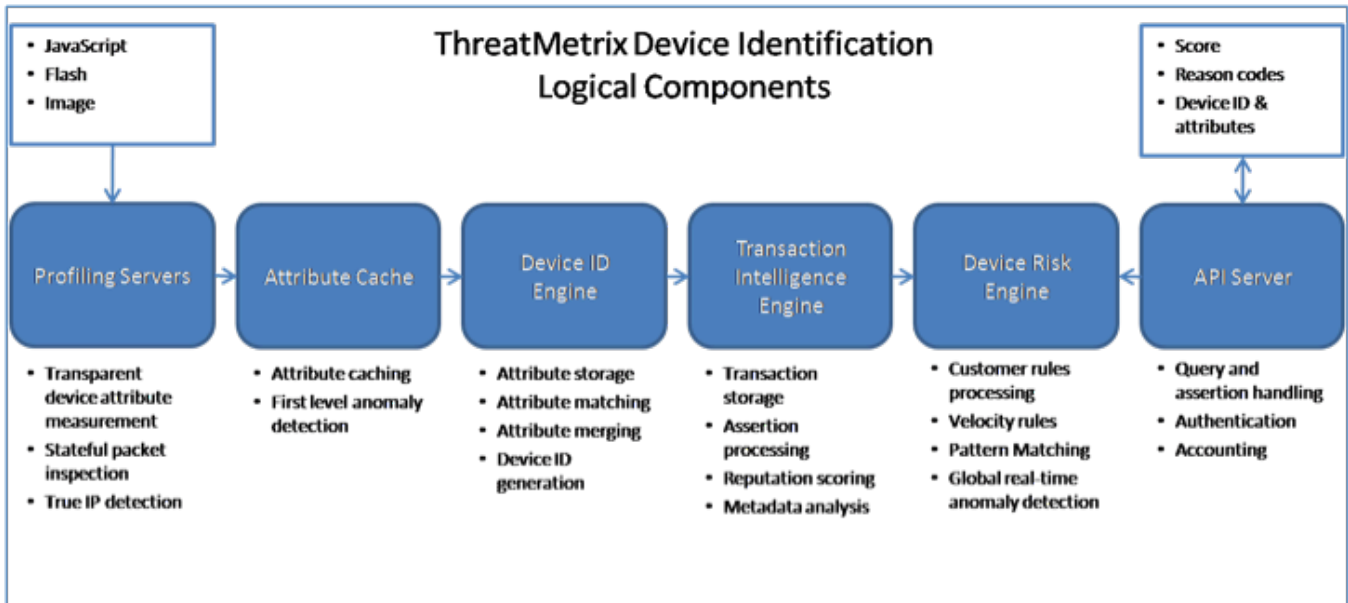
ThreatMetrix Fraud Network

The ThreatMetrix Fraud Network offers companies a better way to verify new account originations, authorize payments and transactions, and authenticate user logins in real-time. The remainder of this document provides an overview of key components of the ThreatMetrix Fraud Network, the design methodologies used and an overview of ThreatMetrix core technologies.

ThreatMetrix is unique in its ability to detect fraud from new and returning computers based on anonymous data mined from a visitor’s computer and TCP/IP packets, then matched in real-time against a global collection of previously profiled computers and their transactions. ThreatMetrix extracts and utilizes more information from more sources than typical first generation device identification technologies that rely on a browser as the primary and in some cases only source of information. A more complete set of information derived from a wider variety of sources makes ThreatMetrix highly resistant to subversion and highly reliable.



The following diagram illustrates the high level components of the ThreatMetrix Fraud Network and their core functions. It does not include ThreatMetrix Data Warehousing and Portal components.



The ThreatMetrix Fraud Network is built upon a distributed architecture. The design provides for real-time data processing and delivery, Internet scalability, shared intelligence across components, redundancy and speed. The key components are:

- **Profiling Server:** Performs both passive (IP/TCP/HTTP profiling) and active (JavaScript, ActionScript) inspection of devices when a user loads a web page that includes ThreatMetrix profiling tags.
- **Attribute Cache Server:** Collects and assembles a complete view of a device's browser, operating system and network characteristics using a combination of Flash, JavaScript and a non-visible image.
- **Device Identity Manager:** Manages logic and processes related to device identities including attribute retrieval, creating unique device identities and matching.
- **Transaction Intelligence Engine:** Processes shared device, transaction, behavioral and reputation history.
- **Real-time Risk Engine:** High-velocity rules and pattern recognition engine detects device risk in real-time based on per-customer and global device transaction histories.
- **API Server:** Customer interface to the ThreatMetrixFraud Network for in-house or third-party risk-based authentication and authorization applications.

Profiling Servers

Active and Passive Inspection

ThreatMetrix Profiling Servers perform both passive (IP/TCP/HTTP Profiling) and active (JavaScript, ActionScript) inspection of a device when a user loads a page that has ThreatMetrix profiling tags inserted. It is a modified web server that performs transparent stateful packet inspections using an image inserted in a page in conjunction with browser and operating system profiling performed using JavaScript and Flash.

ThreatMetrix stateful packet inspection performs 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. 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 ThreatMetrix device identifier. Stateful packet inspection delivers more complete, reliable device data not attainable using typical profiling techniques that

rely on JavaScript to collect and push device attributes through a hidden form field. The key advantages to ThreatMetrix inspection methods are:

- **Reliable:** ThreatMetrix can successfully profile a device relying on passive profiling when active profiling is unavailable—as when JavaScript or ActionScript are disabled.
- **Lower Cost:** ThreatMetrix hosts the remote profiling server safely and securely, saving customers the time and expense of provisioning and managing hardware and software.
- **No Disruption to Customer:** Profiling is accomplished without access to transaction or user data, and without installing applications or applets on the client that would cause a browser warning to the user. Profiling is performed as a background process and immediately stops when the customer leaves a page that has profiling tags inserted. This allows the customer to fill out form data uninterrupted without introducing delays into the transaction flow.
- **Protects Privacy:** ThreatMetrix stateful packet inspection uses anonymous packet header information unlike deep packet inspection which inspects the content (payload) of the packets.
- **Efficient:** Profiled device attributes are cached and deleted if a transaction is not completed.
- **Accurate:** Profiled data is augmented with additional meta-data, such as IP geo information during the device matching and risk analysis routines.

Device Attribute Collection Methods

ThreatMetrix employs a best-effort approach to collect device attributes in pursuit of the most complete and reliable set of device attribution that can be obtained by collecting and examining data from multiple sources. This best-effort approach uses a combination of Flash, JavaScript and a non-visible image to gather a complete view of a device's browser, operating system and network characteristics.

The advantage of a best-effort approach is its ability to work toward obtaining highly accurate and reliable device identification by evaluating and prioritizing all of the data available. The profiling server performs this complex task in seconds—a key requirement for real-time device identification.

The device attribute collection process utilizes redundancy and subversion detection to analyze and rank attributes. Each attribute is measured and then cross-referenced using multiple techniques to identify any discrepancies that help determine attribute reliability. When a device is found to have disabled JavaScript, Flash or images (one or more), or if a visitor leaves a page before profiling is completed, there are sufficient data available to perform the matching process based on the best information available. For example, ThreatMetrix can determine whether a browser string has been tampered with based on inconsistencies in JavaScript, Flash, HTTP and TCP/IP information. Data col-

lection methods that rely solely on JavaScript or web server log files are far less reliable because they are easily manipulated resulting in inconsistent device identification or an inability to derive its identity at all. JavaScript can only report on information captured in the browser, which can be easily manipulated. Web server log files only provide information that is communicated in the HTTP header by the browser—such as the browser agent string and browser language—which are routinely filtered and manipulated by intermediate proxies.

The ability to compare device attributes from different sources is critical using a best-effort approach. For example, the ability to detect the use of an intermediate proxy based on identifying mismatches between browser and protocol attributes is one approach.

In order to accomplish real-time device identification, the data collection process streams attributes from a visitor's PC directly to a separate profiling server. This eliminates the need to collect, store and push device attributes contained in binary large objects (BLOBs) in through a hidden form field. BLOBs—typically several thousand kilobytes—can be very problematic for back-end processes and legacy APIs such as text based payment or order management APIs.

Attribute Cache

The attribute cache servers collect device attribute information from the profiling servers, making a first-pass analysis inspecting for anomalies. The process is managed in temporary in-memory storage until a query for the corresponding device ID arrives. The attribute cache serves as a buffer to make the most efficient and effective use of device attributes including:

- Collecting all the device attributes available regardless of how quickly they arrive from the profiling server (e.g. duration that the visitor is on the profiling web page, connection speed, etc.)
- First-pass attribute comparisons occur in memory where operations are performed at peak speed
- Attributes are collated within a session so comparisons of the same attribute over a block of time can be made; this supports identifying session anomalies indicated by attributes found to have a different state within a single session
- Only devices which are engaged in a genuine transaction are matched. For example, in a new account registration there might be ten devices connecting to a page for every one that actually completes a transaction offering a 10x performance advantage in this example.

Typical examples of anomalies detected on an attribute cache server are the discovery of hidden proxies, bots, browser tampering, session hijacking and cloaking. A closer look at hidden proxy detection illustrates the depth and breadth of inspections performed in the attribute cache:

- Employing proxy bypass methods to cause the device being profiled to directly connect back to the profiling server in order to expose the true IP address
- Detection of a mismatch between the operating system information reported by the browser compared with operating system information reported by the TCP/IP operating system fingerprint
- Examining HTTP protocol fields such as client IP and inconsistencies in HTTP/browser field order
- Detection of removed or modified content in the webpage
- Detection in mismatch in other browser elements including time-zone, language and geo-location
- Filtering out legitimate corporate and ISP proxies
- DNS geo-location mismatches

Device Identification Engine

The device identification engine performs all of the requirements that relate to storing and comparing device attributes, creating a device ID, device matching and exception handling. Here are the core activities that the device ID engine performs:

- Using a device ID, query the attribute cache to retrieve device attributes
- Compare the profiled device attributes against a collection of previously stored device attributes
- Determine whether a device matches a previously known device by executing multiple parallel rule matching strategies
- If a device match is made, updates the state of stored device attributes to reflect a match
- If no match is found, generates a unique device ID
- Generates a warning if there is insufficient or inconsistent data to perform a match or generate a new device ID

The device ID engine outperforms other device fingerprinting technologies that predominantly (in some cases solely) rely on cookies as the source for device attributes, and use hashes of device attributes or cookies for matching. Browser/cookie information is easily and routinely manipulated by fraudsters to hide their device identity—for example, changing the browser agent string in Firefox to falsely report as Windows IE. Because the ThreatMetrix device ID engine utilizes attributes correlated and drawn from multiple sources to match a device, matches are more successful and more reliable. ThreatMetrix uses native device attributes for matching—not hashes created from attributes that allow fraudsters to spawn a new device ID by changing or manipulating the device. Cookies are unreliable as a single source of device identification because they can easily be deleted or copied and distributed—offering fraudsters a convenient way to replicate device identities and use them to their advantage.

ThreatMetrix device ID engine has these unique capabilities and benefits:

- Performs multiple matching strategies in parallel across the entire device attribute collection to yield a device ID in real-time instead of minutes or hours
- ThreatMetrix device ID is globally unique and persistent; every company will see the same device ID for the same computer
- ThreatMetrix device ID matching employs multiple measurement methods instead of one so that a failure in one won't prevent a successful outcome (e.g. if JavaScript is disabled)
- Detect when a device is attempting to cloak its true location and identity
- Device ID is independent of any of the underlying device attributes so when device attributes change over time the device ID remains persistent—for example when a user changes from Firefox to Internet Explorer

A key to delivering real-time device matching and anomaly detection is ThreatMetrix's underlying distributed architecture. Each time a device is profiled and a new device ID created individual attributes are stored and indexed against associated device IDs across multiple distributed file-based storage servers. Partitioning of storage servers is optimized to reduce frequency of collisions on updates as new device attributes are added or updated.

When a new device is profiled ThreatMetrix retrieves cached device attributes from the attribute cache servers and then performs the device matching process using distributed matching rule servers. Each rule server is optimized to match on a specific attribute and logic set, allowing multiple rule matching strategies to be executed in parallel in milliseconds. Running multiple matching rules, each with its own configurable confidence threshold enables ThreatMetrix to match on devices even when underlying attributes are missing (JavaScript disabled), deleted (cookies) or changed (browser updated).

Transaction Intelligence Engine

The transaction intelligence engine is based on a distributed file-based storage architecture used to process shared device, transaction and authentication history.

ThreatMetrix differentiates between shared intelligence that is derived based on behavior patterns including velocity, synthetic identities and statistical anomalies—and reputation which is an explicit assertion of suspected or confirmed fraud.

ThreatMetrix distributed storage architecture is optimized to efficiently process pattern recognition routines across massive data sets that would take traditional SQL databases minutes or hours to process. This enables ThreatMetrix to query and process global device transaction data in real time to enable a global view of a device's behavioral risk at the moment a transaction occurs.

ThreatMetrix maintains a separate data warehouse which is optimized for longer term retention of data for reporting and analytics requirements. The following is a brief summary of inputs to the transaction intelligence platform:

API Servers: Receives real-time queries (event data) and reputation assertions via HTTPS API

- Third party fraud detection software
- Payment processors and gateways
- Login authentication systems
- ERP software (order processing, billing and case management)
- Intrusion detection systems
- Affiliate networks

Submission Servers: Receives batch submission files via SMTP/SFTP e.g.

- Firewall logs
- Spam traps
- Intrusion detection logs
- Firewall logs
- Darknet sensor logs
- Honey pot logs
- Apache logs

The following additional processing is performed on event inputs:

- **Filtering:** Prioritizes and checks the validity of submission formats
- **Lookup/Verification:** Meta data collection and analysis including IP geo databases, Whois data, and BIN number lookups
- **Assertion Processing:** Analyzes submission sources and types and creates assertions

Risk Engine

ThreatMetrix risk engine is a high-throughput rules and pattern recognition engine that draws on device ID and transaction intelligence to assess fraud risk in real-time. Customers can configure the device risk engine to fit their fraud risk profile—an important capability for keeping up with changing threat models. Risk can be assessed based on per-customer and global device transaction and authentication histories.

ThreatMetrix risk engine provides a score based on configurable customer rules, reason codes and the transaction attributes responsible for triggering rules. Customers can decide which rules are active and tailor rules by assigning a weight to their contribution toward an aggregate risk score. Factors include transaction and device anomalies, global and local (customer specific) behavioral analytics and reputation history. Flexible rule sets allow customers to manage risk specifically according to use-case including:

- Online credit card fraud prevention
- Alternative payments and money transfer fraud prevention
- Account opening protection from synthetic or stolen identities
- Account takeover protection from phished or stolen identities

Transaction and Device Anomalies

The ability to detect transaction and device anomalies is essential for assessing device risk and preventing fraud when a device is first encountered by an organization or for the entire network. Examples of configurable transaction and device anomalies include:

- Detection of hidden proxies, browser manipulation, and use of bots and automation
- Mismatch between computer time-zone and true IP geo information
- Mismatch between IP geo and the credit card BIN number

Behavioral Analytics

Examples of configurable pattern recognition rules include:

- Detection of the use of multiple credit cards (hashed) by the same device across multiple sites in a short space of time, both on a per-site and global network basis
- Detection of the same device accessing un-related accounts on a single site
- Multiple proxy IP addresses detected for the same device within a short span of time across the global device identification network
- Login account accessed from multiple time zones in a short period of time

Reputation Scoring

ThreatMetrix reputation is a specific type of shared intelligence defined as an explicit assertion of behavior by an entity such as a merchant or a fraudster's device. Reputation in a fraud context is primarily useful as a warning signal used to increase the risk weighting of a transaction or to flag further investigation by manual review. Customers should not rely on shared reputation information alone to deny access to services due to these inherent limitations of systems that assess risk solely by device reputation:

- Reputation is Transient: Fraudsters can clean or wipe devices, or compromise the device of a previously good customer
- Reputation is Contextual: It can cause false positives where a device may be a "bad actor" in one ecosystem such as gaming, but a good customer in another like travel
- Reputation value is dependent on the timeliness and accuracy of reputation assertions; for example, in an online credit card processing scenario a charge-back can take up to three months before it is discovered
- Reputation systems are at risk of being gamed or exploited, for example if an unscrupulous merchant were to blacklist a customer over a chargeback dispute

ThreatMetrix reputation calculations account for important additional factors to arrive at a more accurate and complete picture of reputation:

- The trustworthiness of the submission source
- Correlation between submission sources
- Submission frequency/volume
- Confidence in the submission (e.g. suspected of fraud versus confirmed fraud)
- Severity of submission
- Submission age

ThreatMetrix reputation is not a static black list—a device's reputation will drift towards a neutral score over time if there are no further reputation submissions seen. Customers have the option of choosing whether to enable reputation assertions, and reputation scores can be weighted or ignored based on customer requirements.

API Servers

ThreatMetrix API servers provide the underlying interface for third-party risk-based authentication and authorization engines to the device identification network. Typical response times to access ThreatMetrix data are less than a second. During a device identity query, customers have the option of providing additional non-personal identifying information for inclusion in the real-time scoring process.

Other than a temporary session identifier allocated during the profiling stage, all other inputs are optional. By eliminating any requirement for the customer to collect data from web server log files or pass a BLOB of attributes from the client machine, customers can integrate ThreatMetrix device identification in days instead of weeks and months.

Conclusion

Online organizations that must verify new account originations, authorize payments and transactions, and authenticate user logins in real-time should consider solutions that are cost-efficient, transparent in web transactions, and respect online privacy.

ThreatMetrix offers a next generation online fraud prevention network that provides a common platform across the entire customer acquisition lifecycle in real-time, without relying on personally identifiable information (PII). Unlike yesterday's siloed web fraud tools for identity verification and authentication, ThreatMetrix uses shared intelligence based on anonymous transaction and authentication history of the device and credentials used across a worldwide network.

For more information, please visit us at:
www.threatmetrix.com