ThreatMetrix®
BUILDING TRUST ON THE INTERNET™

WHITEPAPER

Real Time Trust Analytics

Next Generation Cybercrime Protection

WHITEPAPER

ThreatMetrix®
BUILDING TRUST ON THE INTERNET™

Table of Contents

Assessing Trust in a Zero-Trust World **3**

Identity: the new perimeter of defense 3

Bigger Data or Better Intelligence? 3

Cybercrime Attack Vectors 4

Context Is The New Key 4

Introducing ThreatMetrix Real-Time Trust Analytics **5**

Global Trust Intelligence Network ("The Network") 6

Device Analytics 7

Cross-Platform Device Identification 8

Man-in-the-Middle and IP Spoofing Detection 8

Automated Bots and Malware Detection 9

Identity Analytics 9

Persona Identification and Behavior 10

Link Analysis and Associations 10

Activity Context 11

Behavioral Analytics 11

Velocities and Frequencies 12

Patterns and Anomalies 12

Real-time Business Policy Feedback 13

Trust Evaluations and Analyst Feedback 13

Conclusion **14**

Assessing Trust in a Zero-Trust World

Identity: the new perimeter of defense

As technology has evolved from personal computers to always-connected mobile Internet devices, so too has the perimeter of defense. The first phase of IT security was about securing endpoints, networks, and data residing behind firewalls. The new era of cybersecurity is all about allowing unfettered, yet protected, access to consumer services, business revenues, and productivity assets in a world where these traditional security protections are assumed to be compromised.

Traditional IT Security, developed and installed over the past twenty years, is failing to stop organized cybercriminals. Today's cybercriminals are penetrating perimeter security with stolen or spoofed credentials, and impersonating customer and employee identities. Valid security credentials, whether used by the intended owner or not, easily circumvent traditional IT security solutions, providing cybercriminals easy access to account data, stored credit cards, and valuable IP. Locked-down devices behind firewalls have given way to tablets and smartphones; now, behind the anonymity of an Internet connection, enterprise users and consumers are indistinguishable from criminals.

Every online security analyst and vendor can quote at length the number of new daily data breaches and malware strains targeting personal information. Surprisingly, however, very little data is published on how these stolen identities perpetuate the *cycle* of identity theft—leading to financial and digital asset losses, as well as damaged corporate brands.

Bigger Data or Better Intelligence?

Enthusiastically espoused by big-database vendors, big data security will not fulfill its promise if that data cannot be turned into actionable, real-time insight. Base lining enterprise network behavior can be an important forensic tool, but more essential to stopping cybercriminals is the ability to differentiate between a trusted user and a cyber-threat—in real-time. While traditional big-data solutions can successfully generate possible fraud alerts, they are not able to be actioned in real-time. The resulting delay gives cybercriminals the window of opportunity they need to commit financial fraud, and to access protected personal data as well as valuable intellectual property.

The key to stopping these threats is to execute an automated decision at the point of control, whether at login or within an application, based on a threshold of risk associated with the transaction or access. The goal for this process is to ensure that users, whether valued customers or your workforce, are guaranteed a frictionless online experience while cybercriminals are stopped in their tracks.

Cybercrime Attack Vectors

Know thine enemy. Criminals and nation-states perpetrate cybercrime either by channeling stolen identities and passwords through an anonymized network, or by compromising trusted devices to bypass step-up authentication. Moving from big data to actionable intelligence requires direct, real-time detection of cybercrime attacks based on facts, not heuristics.

COMPROMISED IDENTITY

- **Geo Spoofing:** Criminals use intermediate computers called bots or proxy servers to spoof their true IP address, so that they appear to come from a location that matches the stolen identity or password.
- **IP Spoofing:** Spoofed IP Addresses are continually rotated, so that traditional IP address velocity filters are easily fooled and bypassed.
- **Device Spoofing:** Hackers use automated scripts or manually change device settings, in order to disguise their device fingerprints.
- **Bots:** Criminals use automated scripts to guess passwords, or to register for services en masse.

COMPROMISED DEVICE OR SESSION

- **Man-In-The-Middle:** Hackers use special techniques to intercept or hijack a session—typically, when authentication has already taken place.
- **Man-In-The-Browser:** Criminals use this malware to inject content into the end user's browser. It may steal a one-time SMS password entered into a webpage, or trick a user into revealing private data that can be used to take over their identity or account downstream.
- **Spyware:** This malware may not physically alter the present transaction, but hackers deploy it for future data theft and cross-channel fraud.

Context Is The New Key

Bad things happen to good people, and sometimes good people go bad. In fact, it's safe to assume that any credentials or identity have been, or eventually will be, compromised. In light of this assumption, traditional authentication and reputation systems employing static views of trust no longer stack up. Everyday examples abound. For example, an account compromised by a phishing attack may still require access by the legitimate user. A stolen identity may be used to apply for an illegal credit card, but its legitimate owner may also want to use it. Likewise, a disgruntled employee can gain access based on valid credentials, but is also in a position to abuse the employer's trust—unless proper systems exist to detect suspicious activity.

What's required to address both of these ongoing threats—compromised identities, and compromised devices or sessions—is a new generation of cybersecurity solutions that focuses on *screening the relationship between a user's devices, digital personas, and contextual behavior over time.* Trust, continuously evaluated in the context of each and every interaction, is the only reliable means to establish authentication and access—without compromising convenience for valid users and customers.

Introducing ThreatMetrix Real-Time Trust Analytics

ThreatMetrix Real-time Trust Analytics is the next phase in the evolution of context-aware security. It provides advanced fraud prevention, frictionless authentication, and brand and customer protection, by combining device, identity, and behavioral analytics with collaborative feedback from millions of users across tens of thousands of sites.

ThreatMetrix Real-Time Trust Analytics is proven to achieve:

Advanced Fraud Prevention

- Immediately reduce fraud losses by 90% and cut manual reviews by 70%.
- Prevent account takeover attacks, payments fraud, and new account identity fraud with one integrated platform.
- Directly detect bots, proxies, malware, and stolen identities—based on facts, not heuristics

Frictionless Authentication

- Reduce step-up authentication by 50%.
- Provides 100% coverage with passive two-factor authentication, leveraging device, location, identity, and behavior over time.
- Immediately detect compromised devices and session hijacking attempts.

Brand and Customer Protection

- ThreatMetrix stops customer records from being accessed, and provides an invisible layer of defense against third-party breaches.
- According to Ponemon's 2013 Cost of Data Breach Study, each compromised record costs $277 in the US and $214 in Germany. For e-tailers, banks, or social networks with hundreds of thousands or even millions of stored identities, this figure represents a significant liability.

The remainder of this paper discusses the real-time device, identity, and behavioral data recorded during each user's access, then analyzed—in milliseconds—across the terabytes of anonymized activity and identity intelligence contained in the Global Trust Intelligence Network:

Device Analytics

- Cross-platform device identification
- Man-in-the-middle and IP spoofing detection
- Automated bots and malware (man-in-the-browser) detection

Identity Analytics

- Persona facts and behavior
- Link analysis and associations
- Activity context

Behavior Analytics

- Velocities and frequencies
- Patterns and anomalies
- Policy score feedback
- Trust evaluations and analyst feedback

Global Trust Intelligence Network ("The Network")

Attempting to stand alone against organized, global cybercriminals is akin to battling an organized army with a militia of citizen soldiers. In some cases the militia claims a small victory, but in most cases it loses the war. A shared problem needs a shared and collective solution. That is why the ThreatMetrix Real-time Trust Analytics platform is built on the power of a shared Global Trust Intelligence Network. The ThreatMetrix Global Trust Intelligence Network is the world's largest repository of shared anonymized identity and device recognition in the world. The network collects real-time data from tens of billions of transactions and hundreds of millions of user credentials, across more than 10,000 websites. ThreatMetrix's large installed customer base, which spans global financial institutions, the largest ecommerce sites, and social networks, means that over 80% of interactions screened are recognized as a previously identified and profiled persona.

The biggest challenge of a shared intelligence network, other than handling the sheer scale of data, is achieving the right balance between security and privacy, while negotiating a continually changing legal landscape with respect to what is considered personally identifying information (PII). ThreatMetrix's answer to this escalating issue is that we do not need to know your name to know who you are. The actual identification of an individual is not important. After all, the credentials that are being presented may be stolen. What is important is to recognize whether the anonymized

credentials being presented are consistent, and make sense in the context of the online event occurring at that moment. If for example ThreatMetrix recognizes that an email address and credit card number, even though anonymized, are also in use in five other locations around the world, we don't need to know that person's true identity to know that their credentials have been compromised. ThreatMetrix is the only fraud services provider to provide private-key encryption and non-reversible anonymization, so that even its own employees, even if compelled by law, are not able to access customer or identity data.

When analyzing global device, identity, and user behavior through the lens of a customer's business policies, ThreatMetrix provides more than just a reputation score, which is often out of context and can be highly misleading. Rather, ThreatMetrix provides information that takes into account the context of the online event, resulting in historical evidence of a persona's behavior across all data, and ultimately delivering a full set of attributes which informs its customers of the true nature of any transaction and the individual.

The following sections provide a detailed analysis of the three main subcategories of the network that ThreatMetrix analyzes in real-time across our network:

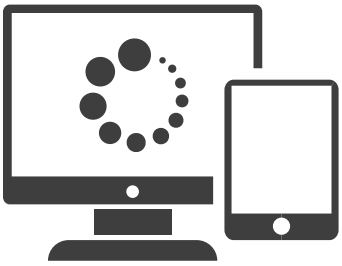

Device Analytics

In any online activity—from applying for credit cards to shopping with them, from logging into a health portal to running a payroll—the user's device substitutes for a physical person. Therefore, the accuracy of any trust decision is directly tied to a provider's ability to uniquely and persistently identify devices, as well as their origins and any anomalies associated with them. Yet nearly every security and online fraud detection system in place today is exclusively reliant on the use of cookies. These cookies are easily compromised by hackers and privacy-conscious users alike. To complicate matters, IP address information is dangerously easy to spoof by using proxies, virtual private networks (VPN) and botnets.

Within one integrated SaaS-based platform, ThreatMetrix Real-time Trust Analytics combines the three pillars of next-generation, context-aware device analytics:

1. Cross-platform device identification
2. Man-in-the-middle and IP spoofing detection
3. Automated bots and malware (man-in-the-browser) detection

Cross-Platform Device Identification

ThreatMetrix provides a holistic platform to recognize and identify visitors across all the platforms they use, from the desktop web to mobile browsers and apps to thick clients. Whether a user is a first-time customer or an employee remotely accessing a network via a VPN, ThreatMetrix's integrated approach accurately identifies users and the devices they prefer.

For mobile and web browser interactions, ThreatMetrix has created the industry's leading Cookieless Device Identification, with patented and non-invasive SmartID™ technology. Using dynamically-updated machine learning models based on passively-collected device and network attributes, ThreatMetrix SmartID dramatically outperforms static device fingerprinting alternatives, and is able to reliably identify returning visitors and known fraudsters on a global basis.

Integrated within the same device identification engine is a mobile application SDK that enables deeper levels of authentication data for IOS and Android smart phones and tablets. For example, ThreatMetrix can provide accurate GPS location information at the time of authentication or purchase. It can also detect if a device has been jail broken—and hence made less secure, prone to infection or hijacking.

For additional protection for known trusted users and remote workers, ThreatMetrix provides a downloadable client for Windows and OS X with mutual authentication capabilities, as well as security posture and compliance enforcement. Minimum password standards, the proper use of encrypted drives, and patch status are all monitored and enforced.

The power of ThreatMetrix's integrated platform is that a single policy and decision engine can be applied consistently across all the ways your users access your site. This consistency dramatically reduces cybercriminals' ability to exploit silos and gaps between technology stacks and departments.

Man-in-the-Middle and IP Spoofing Detection

The ability to recognize returning devices is essential for stopping automated attacks and repeat offenders, as well as for treating valued users with the proverbial kid gloves. Fraud and the majority of account takeover attacks can be prevented the first time, however, by preventing the one thing cybercriminals rely on: the ability to cloak their true IP address, and by extension their location, using proxies, Tor and VPNs. Instead of relying on detecting proxies based on known proxy IP addresses, which are continually changing, ThreatMetrix is the only fraud prevention solution that directly detects the presence of a proxy, VPN, Tor or man-in-the-middle by passively fingerprinting the network packet signature of incoming connections, then using stealth techniques to pierce proxies. This process reveals the true IP address (and thus the true location) of the attacker's device.

Automated Bots and Malware Detection

Cross-platform device identification and IP spoofing detection are table stakes for cybercrime prevention, but what happens if a user connects cleanly with a known and trusted device which now happens to be infected with malware?

In addition to device and packet fingerprinting, ThreatMetrix examines how a device interacts with a page to automatically determine the presence of a bot or man-in-the-browser injection attack. Unlike approaches such as session analytics, which monitor navigation velocity to detect automation, ThreatMetrix is able to directly detect the presence of a scripted attack without false positives, the moment a user requests a page. In the special case of man-in-the-browser attacks, ThreatMetrix uses a unique approach called page fingerprinting, which is based on whitelisting techniques and is able, not only to detect malware, but also to classify the variant of malware with which the end device is infected.

For mobile applications that do not use a web browser, ThreatMetrix provides additional protection through an easy-to-integrate, user-transparent IOS and Android solution that integrates into existing applications. ThreatMetrix Mobile provides comprehensive protection against jailbreak, malware, and malicious applications. Jailbroken, or rooted, devices have some important security features disabled, which mean they are more prone to having malicious applications installed. These applications steal sensitive data and can subvert traditional phone-based out-of-band solutions.

By combining device identification with proxy, bot and malware detection in a single integrated platform, ThreatMetrix Device Analytics has proven to be the most effective solution for stopping cybercrime and recognizing returning users the moment a device initiates its first interaction.

Identity Analytics

Comprehensive endpoint intelligence is necessary, but—unfortunately—not sufficient. In fact, 95% of effective cybercrime protection involves differentiating legitimate users and trusted customers from cybercriminals, and fully understanding that difference.

For example, how many times have you been prevented from accessing an online service just because you are in a different country traveling on business—or, for that matter, for simply using a different device from your own home? Though organizations spend billions to acquire customers online, the expected efficiencies are lost through out-of-date security systems, which are trained to look for bad actors but blind to recognizing valued customers and employees. The only truly effective way to combine convenience and security is by profiling and understanding both sides of the spectrum, squeezing the grey zone to a manageable minimum.

ThreatMetrix Identity Analytics examines each access attempt and transaction in real time across the network, to establish and verify:

- Persona identity and behavior
- Link analysis and associations
- Activity context

Persona Identification and Behavior

ThreatMetrix "Persona" technology provides a real-time identity database, and a sophisticated query and behavior pattern engine, for detecting both trusted customer patterns and behavioral anomalies across the network.

PERSONA FACTS: EXAMPLES

- Employee user names, roles and attributes
- Social profile data collected from social logins
- Product catalogs and prices
- Customer purchase history
- Airline routes
- High risk IPs detected in an external SIEM product
- Lists of third-party breached email names and passwords
- Government restricted trade countries and individuals

PERSONA BEHAVIOR: EXAMPLES

- Classify trusted users based on age and frequency of relationship, using a combination of device, name, account, IP address, email, telephone number, geographic location, amount spent, or any of countless other customizations relevant to your industry.
- Classify high-risk users based on identity anomalies, such as an unusually high number of identities associated with a single device in a short period of time.

Link Analysis and Associations

Who you hang out with says a lot about who you are. So, when screening a login, payment, registration or account update, ThreatMetrix not only examines the device and identity elements being presented; it also interrogates the ThreatMetrix real-time database, to look for connections and associated entities.

For example, a family shares a laptop computer and a tablet. The first time the mother picks up the tablet to log into her favorite online store, ThreatMetrix immediately recognizes the same tablet and identity that has been successfully used at other sites, and that someone else from the same household has successfully logged in using the same device in the last month. Instead of seeing the mother as a complete stranger and potentially putting her through additional verification, the e-tailer can choose to present promotional offers—or to simply streamline the checkout experience based on the recognition of a trusted household.

Conversely, the dad may log into an online banking session from what appears to be his clean and uninfected mobile device. But ThreatMetrix has also seen the dad's credentials used on a home laptop that—without his knowledge—is currently tunneling illegal credit card transactions through a back door, meaning his banking credentials have likely been compromised.

Activity Context

In an online interaction, an identity, whether it be a user name, email, address, credit card number, or password, is really nothing more than a set of labels, attached to a device that is presenting those credentials. In the hands of the authentic user, that identity is the source of new revenue or productive work. In the hands of a criminal, that same identity will result in fraud or a serious data breach—and the subsequent brand damage that results.

Clearly, identity analytics goes beyond just verifying that an identity exists. What's more important is context: how that identity is used in real time, as opposed to individual, per-application, per-site, or global norms. For example, a money transfer would attract a higher level of scrutiny and authentication if it originated from a time zone not previously used for that account and exceeded an aggregate monthly threshold, particularly if the account details were changed in a preceding session. On the other hand, if a credit card is used from the same device and location, and with the same contact information, as a previous card with a long positive history of transactions across the network, it can be automatically approved with a high level of confidence—and, critically, without inconveniencing the user.

Behavioral Analytics

Reputation information is helpful, but actions speak louder than words. ThreatMetrix Real-Time Trust Analytics enables direct insight into global behaviors, without compromising privacy or anonymity. Instead of providing just a global reputation score, or an indicator of "trust" or "no trust," ThreatMetrix enables its customers to create their own risk scores based on anonymized global behavior data—then fine-tune, based on the level of confidence in the signals detected, all in the context of the attempted action. The behaviors that ThreatMetrix customers are able to mine fall into four broad categories:

- Velocity and frequency
- Convergent (co-occurring) and divergent (fluctuating) identity patterns
- Automated policy scores
- Analyst and trust feedback

Since risk is highly contextual, it needs to be treated with the same weight as hearsay. What is more valuable than a globally derived score? For example, a US credit card used online at a European

e-commerce site is orders of magnitude more likely to be fraudulent than the same card used at an online retailer in the US. A publicity-seeking student hacker who defames a website by taking over an administrator's account is still a good potential customer of an airline when buying a ticket to fly home for vacation.

Velocities and Frequencies

Examples of the types of velocity patterns instantly detected by ThreatMetrix on a global, per-site, per-event type, per-device, and per-identity basis include:

- Event velocities: tracking number of events, for example transactions, within a time period for a given entity or combination of entities
- Monetary velocity: tracking aggregate amounts of money moved over a time period relative to an account, device, identity, or location
- Distance Velocities: tracking distance traveled within a time period
- Cookie Velocities: tracking how often cookies are deleted or suppressed for the same recognized device, based on ThreatMetrix SmartID technology
- Identity Velocity: tracking how frequently an identity changes with respect to the same device
- Device Velocity: tracking how frequently an identity is used across multiple devices
- Page Velocities: tracking time spent on a page

Patterns and Anomalies

Trusted user pattern recognition examples include:

- Increase level of trust based on age of persona, device, or any identity element within the network
- Increase level of trust based on age of relationship between identity elements, e.g. device and account association
- Increase level of trust based on number of times a persona, device, and identity elements have been used together within a payment transaction, login or registration
- Increase level of trust based on number of times a persona or device and identity combination has been seen from a given IP address or geolocation

Insider and external threat pattern detection examples include:

- Increase level of risk if device or identity has not been seen before in the network
- Increase level of risk the first time a device has been seen using this account or identity
- Increase level of risk if identity has been used across an excessive number of devices, time-zones, locations, or addresses within a specified period of time
- Increase level of risk if device has been associated with an excessive number of accounts, identities, locations, or addresses within a specified period of time

Real-time Business Policy Feedback

Every interaction observed across the ThreatMetrix network is explicitly scored according to each customer's policies, which in turn are tuned to their specific business requirements—whether it be acquiring a customer, authenticating an employee, or processing a payment. This real-time feedback is then used by ThreatMetrix to implicitly score each identity, thereby building an accurate real-time picture of a persona trustworthiness based on each and every interaction being "voted" on by participants in the network. This scoring and decision feedback is then instantly shared across the network. A transaction at a small online fashion boutique in Australia can be instantly informed by a banking login performed using the same constituent identity elements in New York, just milliseconds prior.

Trust Evaluations and Analyst Feedback

The advantage of access to real-time policy scores is that they are automated, instant, and available for every interaction. In cases where an identity or device has been found to be high-risk or involved in criminal behavior after the original policy score decision, ThreatMetrix closes the feedback loop with truth data received via API or analyst portals. In addition, ThreatMetrix customers have the ability to record authentication results for one or more combinations of entities, using ThreatMetrix Trust Tags capability. For example, if a device and telephone number have passed an out-of-band authentication challenge at an online bank, this entity combination is tagged as 'authenticated' with a context of "SMS OTP" and "banking." That tag can then be shared across the network to enabled frictionless federated trust across organizational boundaries. This anonymized authentication information makes ThreatMetrix the largest source of trusted identities in the world.

Conclusion

Today's sophisticated identity theft and cybercrime attacks exploit the weakest links in your IT and security infrastructure: your customers and employees. It does not matter how high you build your perimeter walls if your border controls cannot distinguish between a trusted user and a cyberthreat. ThreatMetrix Real-time Trust Analytics, delivered through the shared Global Trust Intelligence Network, provides the first industry solution that delivers advanced fraud prevention, frictionless authentication, and customer protection—all within a comprehensive and integrated cybercrime protection platform.

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

V5.10.2014