transmit security

TECHNICAL WHITE PAPER

Detection and Response Services

Technical white paper

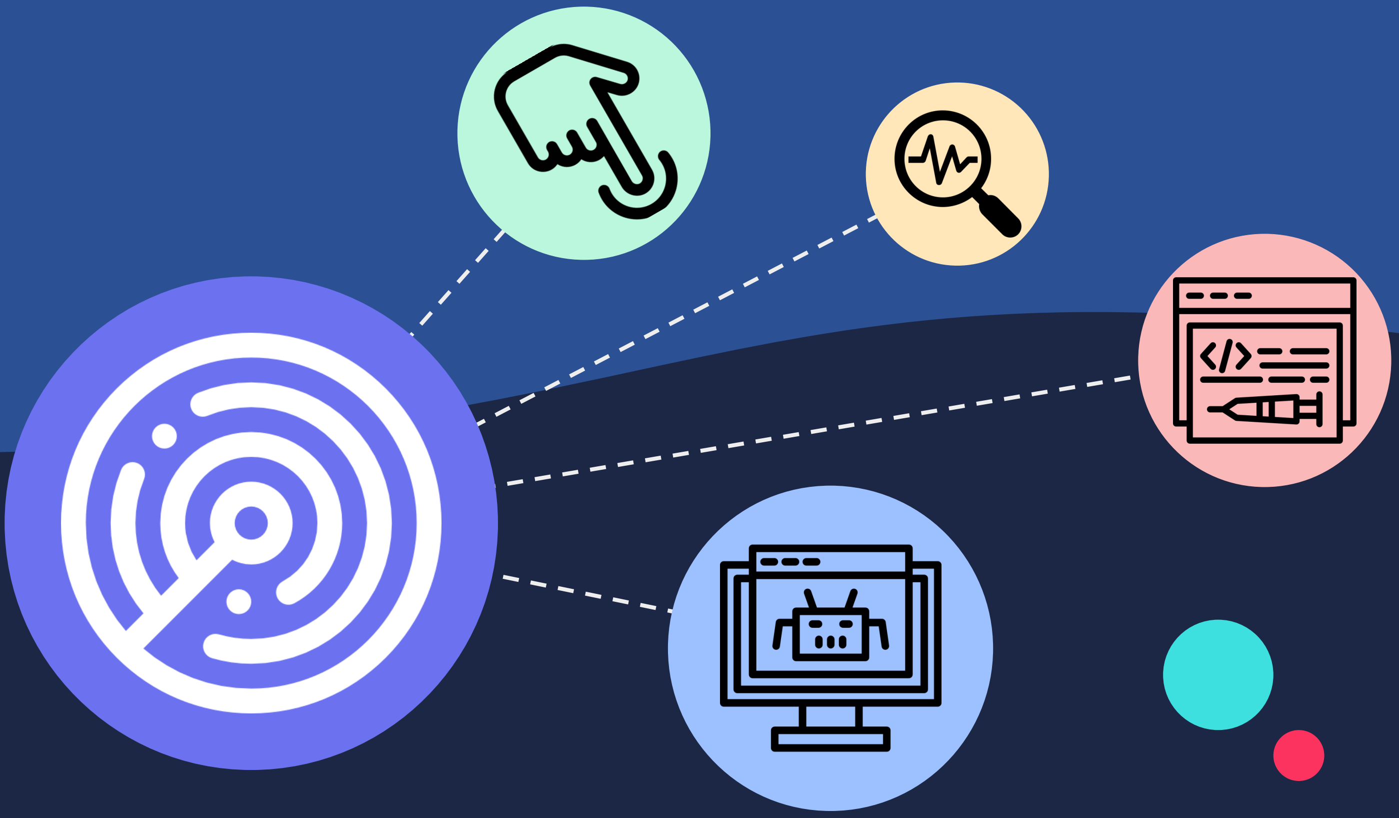

January 2023

Evaluating risk and trust in the new fraud landscape

Changes in the identity and fraud landscape are making it harder than ever to identify suspicious user behavior, and as a result, account takeover (ATO) and new account fraud (NAF) are surging. New regulations and increased demand for consumer data privacy have caused browser providers to strip away many of the mobile and web identifiers used for device identification, while attacks continue to innovate new methods to get around security controls. For example, evasive bots — which use modified web browsers, regularly change IP addresses, imitate human mouse activity and time requests to appear more like legitimate users — were responsible for almost two-thirds of all ATO in 2022, which increased 148% from 2021, according to Imperva's 2022 Bad Bot Report.[1]

In this new landscape, the tools designed to detect fraud are becoming less reliable, forcing businesses to adopt multiple solutions that make data correlation and decisioning an increasingly complex and labor-intensive task. However, man-made decisioning rules that require long tuning cycles are unlikely to defeat machine-assisted fraud. To stay ahead of cybercriminals, businesses must leverage a comprehensive, integrated solution that reduces — rather than amplifies — the need to build and tune custom rules.

[1] Imperva

Table of Contents

Differentiators for detecting risk, trust, fraud, bots and behavior

Transmit Security's Detection and Response Services provide a unified model for fraud prevention that integrates a wide variety of detection methods and analyzes them with machine learning to deliver contextual, real-time recommendations for each user, in each moment of risk, across the entire user journey. In doing so, it improves the experience for the customer, the accuracy of fraud detection tools and the time, cost and effort needed to assess risk and trust.

Three key features differentiate our Detection and Response service from other solutions, enabling improved detection rates and reducing the manual labor required for fraud prevention:

- **Multi-method Detection** leverages hundreds of detection methods applied to a broad range of telemetry to fill the gaps that can result from outdated or missing telemetry streams, weak risk signals and reliance on multiple detection tools that make it difficult to effectively correlate data.
- **Embedded Orchestration** reduces the expertise and knowledge needed to build and orchestrate decisioning rules by providing contextual, real-time recommendations that work out-of-the-box, along with deep insight into the factors that led to those decisions.
- **Dynamic Threat Response** lets you easily respond to new threats without any changes to your application; our adaptive risk engine is continuously updated to accurately detect new attack patterns in real-time, using the Transmit Security Research Lab, where our team of in-house researchers train our machine learning tools on emerging attack cases and evaluate the efficacy of new rules.

Detection and Response is delivered as a cloud-based service that requires no changes to UI or custom code to maintain. A simple integration process, coupled with out-of-the-box recommendations that provide visibility into key risk and trust indicators, enables actionable insights within minutes. As more data is collected, the accuracy and value of these insights rapidly increases.

This paper will provide a deep dive on these capabilities and an overview on the process for deploying and integrating our Detection and Response service with existing providers.

Multi-method detection analyzes a broad range of telemetry

Most fraud detection tools on the market today rely on only a few detection methods — or often, just one. Bot and automation frameworks delineate between bot and human users; behavioral biometric tools evaluate patterns in users' mouse clicks, movements and typing; and device fingerprinting analyzes users' browsers, devices, OS and configurations to establish users' trusted devices and detect suspicious ones. However, even best-in-breed solutions for these frameworks expose gaps in detection due to a narrow focus on specific detection methods and the difficulty of correlating siloed results across different tools from multiple vendors.

Our Detection and Response service's multi-method detection addresses this challenge by combining a wide range of detection mechanisms into a single service powered by hundreds of different telemetry streams. This data is used to build progressive profiles of individual user behaviors that include data points such as mousing patterns, trusted devices, IPs and other telemetry. Customers can receive this raw telemetry via API to use in their internal systems and feed to other solutions to improve their detection capabilities. Alternatively, customers can receive enriched data that is correlated and analyzed to detect anomalies in general usage and specific users' behaviors. This data applied as a basis for real-time, actionable recommendations, as discussed in the next section of this paper.

A sampling of some of the frameworks and telemetry used in the service is shown below.

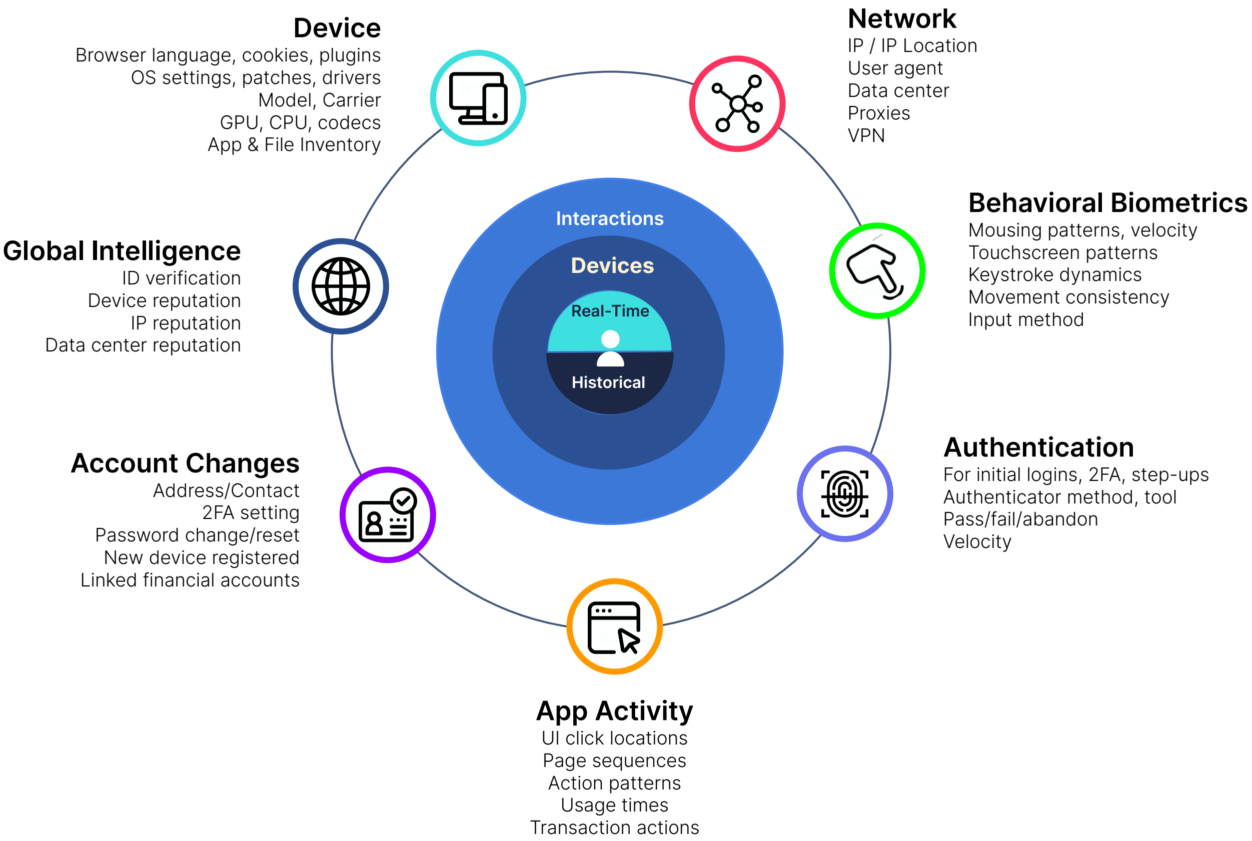

By providing more robust telemetry and consolidating a range of detection methods, Detection and Response enables a comprehensive picture of the trust and risk signals in applications. With it, businesses can reduce their reliance on multiple products or aggregate inputs from other systems to easily correlate and analyze cross-vendor data. This results in more accurate detection and reduces the effort needed to integrate and maintain multiple fraud detection systems.

The next section of this paper will demonstrate how the robust risk and trust signals gathered through multi-method detection are transformed into actionable recommendation via an automated process that delivers out-of-the-box decisioning orchestration.

Embedded Orchestration provides transparent, out-of-the-box decisioning

Creating accurate and robust decisioning logic is a complex task that requires significant expertise, time and effort to build and orchestrate. Fraudsters are constantly adapting their tactics to evade detection, so multiple telemetry signals must be combined — often from multiple sources — in order to create a single reliable indicator of fraud. Determining the most effective way to combine these signals becomes even more difficult when considering that usage patterns vary widely between end users and normal behavior on one page of a website could be an indicator of fraud on another page. And once decisioning logic is built out, lengthy tuning cycles are needed to test and maintain its efficacy — an ongoing process of trial and error that is bound to result in high false positives or false negatives.

Detection and Response circumvents this process with decisioning orchestration that is built and maintained automatically — a process we refer to as Embedded Orchestration. Our machine-learning algorithms analyze signals within the full context of your application and use cases to deliver smart recommendations, which can be accessed via API or within the dashboard of Transmit Security's CIAM Platform to gain a comprehensive picture of risk and trust indicators. Within the dashboard, each user action is given a recommendation of Trust, Allow, Challenge or Deny, along with the key indicators the recommendation is based on.

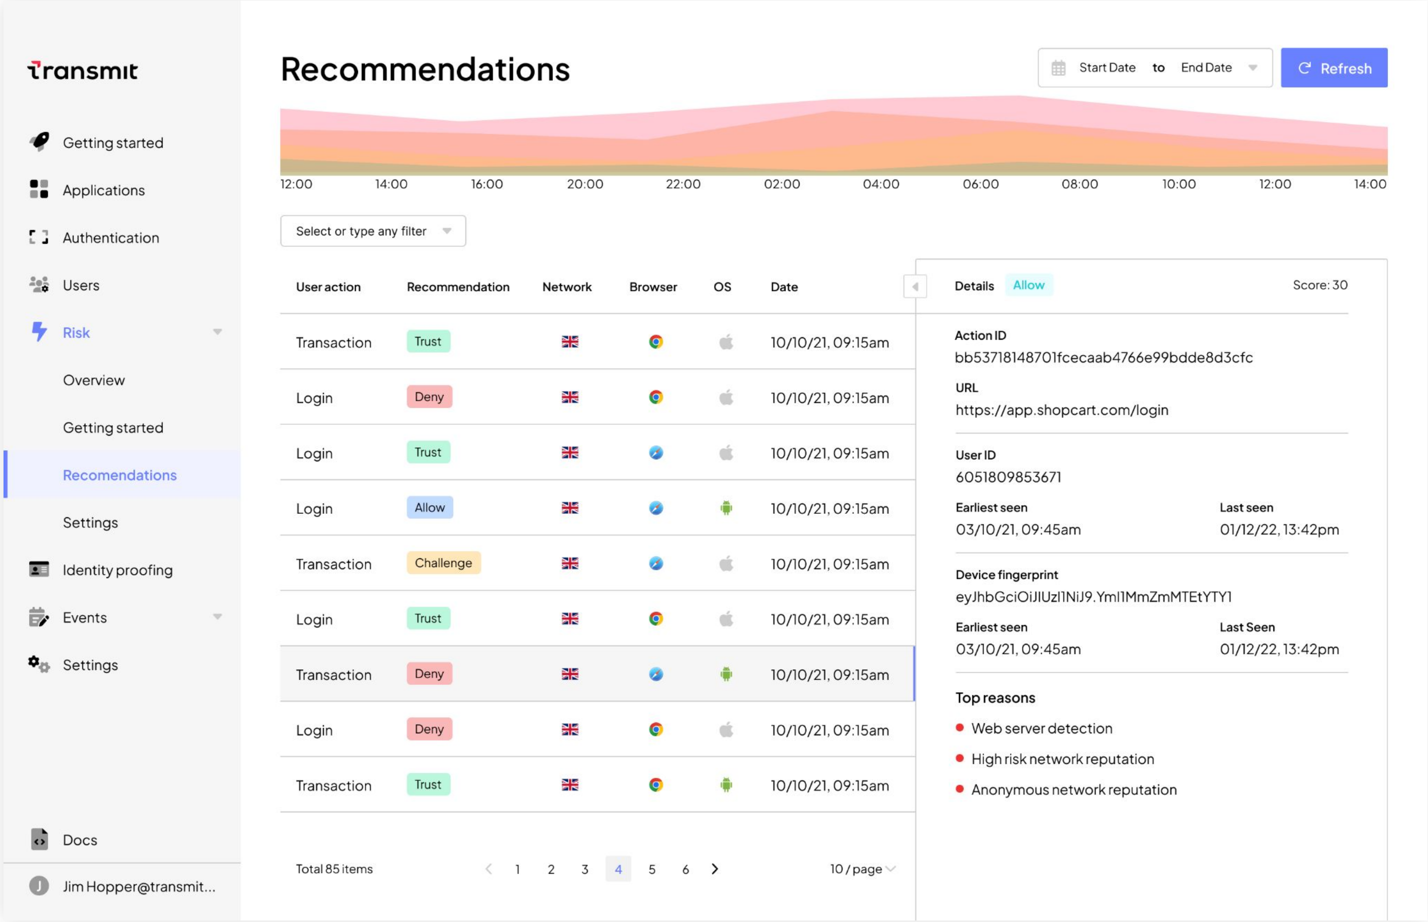

Recommendations

User action	Recommendation	Network	Browser	OS	Date
Transaction	Trust	🇬🇧	◉	🍎	10/10/21, 09:15am
Login	Deny	🇬🇧	◉	🍎	10/10/21, 09:15am
Login	Trust	🇬🇧	🧭	🍎	10/10/21, 09:15am
Login	Allow	🇬🇧	🧭	🤖	10/10/21, 09:15am
Transaction	Challenge	🇬🇧	🧭	🍎	10/10/21, 09:15am
Login	Trust	🇬🇧	◉	🍎	10/10/21, 09:15am
Transaction	Deny	🇬🇧	🧭	🤖	10/10/21, 09:15am
Login	Deny	🇬🇧	◉	🍎	10/10/21, 09:15am
Transaction	Trust	🇬🇧	◉	🤖	10/10/21, 09:15am

Total 85 items ‹ 1 2 3 4 5 6 › 10 / page

Details Allow Score: 30

Action ID
bb53718148701fcecaab4766e99bdde8d3cfc

URL
https://app.shopcart.com/login

User ID
6051809853671

Earliest seen	Last seen
03/10/21, 09:45am	01/12/22, 13:42pm

Device fingerprint
eyJhbGciOiJIUzI1NiJ9.YmI1MmZmMTEtYTY1

Earliest seen	Last Seen
03/10/21, 09:45am	01/12/22, 13:42pm

Top reasons

- Web server detection
- High risk network reputation
- Anonymous network reputation

These recommendations can be used as intelligent action triggers at each risk moment in the user journey, rather than bluntly applying the same controls across an application's entire user base — which can result in high false positives that introduce needless friction for legitimate users and false negatives that open the door to ATO and NAF.

To build the complex decisioning that enables these recommendations, our experts in the Transmit Security Research Lab conduct a machine-learning process called novelty detection, a semi-supervised anomaly detection algorithm that uses clean data sets to define what legitimate usage looks like on individual pages and fields and across application flows, for entire user populations and individual users. Once the parameters for normal behavior are established, the algorithm can assess new data points to detect outliers and leverage labeled data using feature engineering to uncover hidden patterns and determine the most reliable indicators of fraud. Although complex machine-learning models are typically opaque, the use of model explainability methods give us greater transparency into which indicators have the greatest contribution to our machine-learning verdicts. For example, by using SHAP values, we are able to obtain both a global measure of feature importance (which indicates features that are weak vs. strong indicators of overall risk) and a local measure that helps explain the significance of specific features in each individual outcome.

This process enables Detection and Response to not only deliver recommendations, but provide valuable context on which features played the most crucial role in that verdict. With it, businesses can provide fraud analysts with a clear, easy-to-understand view of the signals used to indicate fraud and give developers insight into the rules used in decisioning — as well as the ability to influence those rules. Using the Recommendations API, developers can retrieve a list of all the rules used in decisioning, or send a request with a specific rule ID to retrieve the name and priority level of the rule that triggered a certain recommendation. And although Detection and Response doesn't require developers to write custom rules for decisioning, developers have the ability to create, fine-tune, or override rules, as well as simulate and evaluate the impact of these changes before releasing them to production. More details on this process can be found in the Recommendations API documentation.

In other words, Transmit Security's Embedded Orchestration provides not only out-of-the-box, complex decisioning that can be used to trigger the right action at the right time, but the flexibility and customization capabilities that are normally difficult to obtain in fraud detection systems that leverage complex machine learning algorithms. The next section of this paper will explain how these dual capabilities can accelerate a business's ability to respond to new and evolving attacks.

Dynamic threat response swiftly responds to emerging threats

As new technologies emerge, fraudsters will leverage them to create increasingly sophisticated evasion tactics and zero-day attacks that cannot be detected by static rule sets. As a result, rule sets must be continuously updated to catch the most recent tactics, techniques and procedures leveraged by attackers. However, man-made decisioning rules that rely on custom code require significant time and effort to deploy, update, test and tune, leaving security experts always a step behind cybercriminals and highly vulnerable to new and emerging attacks.

Fraud detection services that leverage machine learning can more quickly adapt to new threats, but these systems are often opaque, resulting in a black box system that cannot be evaluated or tested for efficacy. When training machine learning models, data sets with many varied examples are needed to ensure the model is able to effectively analyze new, unseen data. One single case of fraud provides too small of a data set for the algorithm to learn from it accurately; it would need to be trained on large data sets with variations on the same attack pattern to improve its detection capabilities. This would require a large number of new confirmed attack cases — meaning that businesses would need to wait for an escalating number of customer complaints before obtaining an effective detection mechanism.

Neither scenario is acceptable for responding to new attacks. Quick and accurate detection of emerging threats requires not only the use of machine learning algorithms, but a team of researchers who can test, evaluate and tune those algorithms to assure quality and effectiveness. Our Detection and Response service's dynamic threat response provides this capability through the Transmit Security Research Lab. The team uses threat intelligence and applied data science to install, operate and reverse engineer new attack tools to investigate new fraud techniques. This enables our researchers to continuously train and optimize its algorithms on new attack MOs and uncover patterns used in emerging attacks.

After the accuracy of the model has been confirmed, it can be deployed to detect new attack cases in real time with no changes to the application. And when data is consumed by and fed to other detection services, the algorithm can improve its accuracy as well as the accuracy of those services.

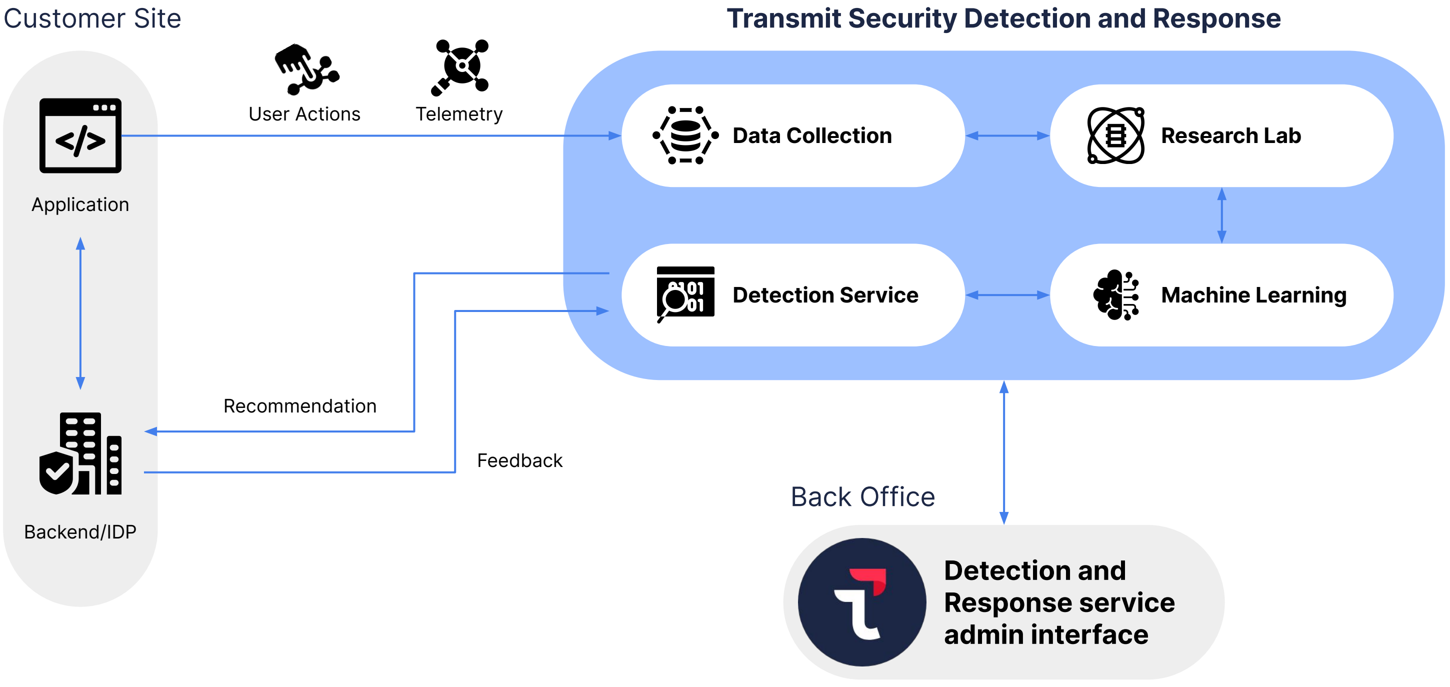

Ease of deployment and integration accelerates time to value

Our Detection and Response service is designed to provide value from the moment it's installed. Using standard APIs, connectors and SDKs, developers can start seeing initial detection results and receive recommended actions in minutes with as little as four lines of code.

To get started monitoring end-user risk levels, developers can load the SDK into their web or native mobile application front end and initialize with the server path and client credentials, which can be obtained from the admin portal. They can then set a user ID as an opaque identifier to track users within the application and add a code snippet to each of the relevant user actions — such as logins, registrations, transactions or account changes — that they wish to monitor. Our Recommendations API can then be used to fetch recommendations, which will appear automatically in the Admin Portal. Full details on this process, along with integration guides for iOS and Android, can be found in our documentation.

Alternately, developers can choose to integrate the service with their business's content delivery network (CDN) or tag management system and get started with no changes to UI and no custom code to maintain. Once integrated with a CDN, the CDN will look for tags that identify relevant actions within user requests and run the code snippets that modify the response after the CDN retrieves it from the backend or cache. The cache is then updated with the modified response, and the response is returned to the end user.

Similarly, with tag management systems (TMS), such as Google Tag Manager, developers can integrate using a single container tag that is loaded in every web page and configured to listen for the user action events that you want to evaluate, which will trigger execution of the code snippet. Once integrated, the tag manager function will query the TMS on the fly to fetch the tag and inject the SDK into the front end once the page is loaded on the client side. Our documentation provides detailed instructions on how to integrate with Google Tag Manager as an example of this.

Following deployment and integration, our Detection and Response service can be tested in monitor mode to assess and tune recommendations, which can be easily done via its Labels API. If false positives or negatives are introduced due to lack of context, developers can use this API to easily provide feedback on these cases via a single call that requires only two inputs: the subject of the label and the type of label ("`KNOWN_MALICIOUS`", "`KNOWN_LEGIT` or "`UNKNOWN`") to send. Additional information about the source of the label, such as manual self-analysis or data from other systems, can be appended to provide additional context, which can be used to further tune the recommendation system. Once businesses are confident with the accuracy of recommendations provided by the platform, they can begin acting on those recommendations using a single API call.

Modular CIAM services strengthen security across the customer lifecycle

Native integration with other services in the Transmit Security CIAM Platform enables orchestration of strong step-ups and challenges such as passkeys and a range of other multifactor authentication (MFA) methods, photo ID comparisons, liveness checks and document analysis.

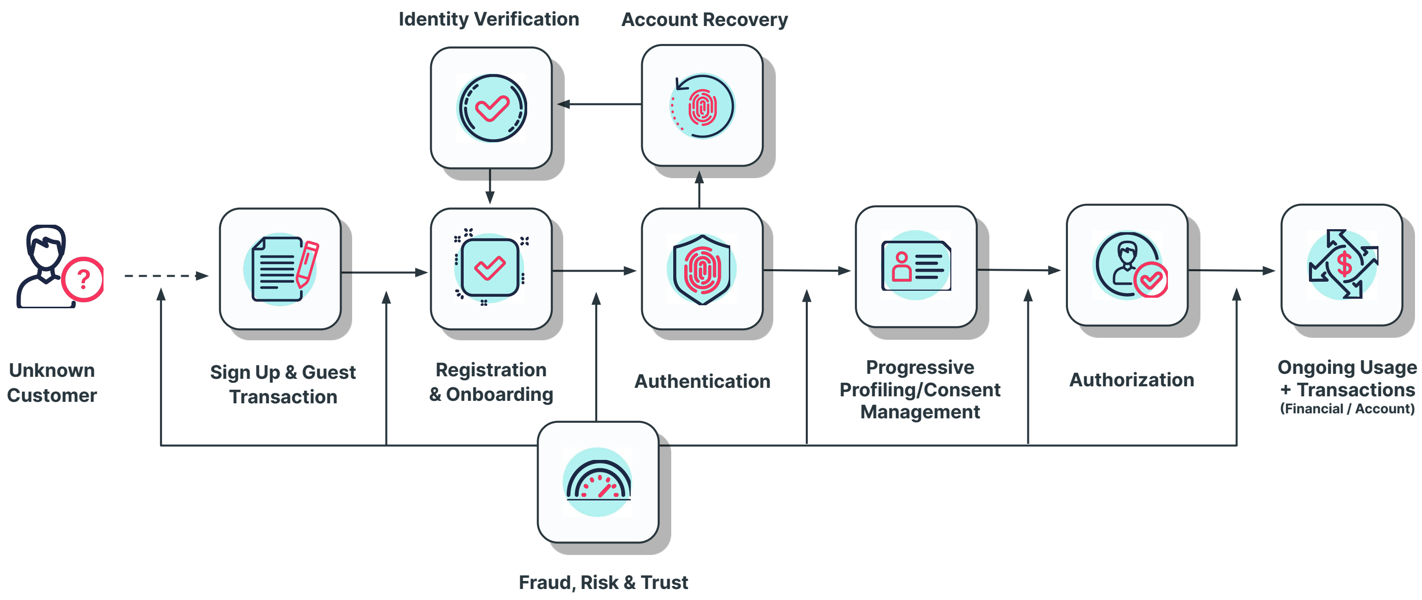
Identity Verification
Account Recovery
Unknown Customer
Sign Up & Guest Transaction
Registration & Onboarding
Authentication
Progressive Profiling/Consent Management
Authorization
Ongoing Usage + Transactions (Financial / Account)
Fraud, Risk & Trust

These controls can be applied to automatically trigger actions at different touch points across the entire user journey, including guest transactions, registration, account recovery, authorization, account changes and transactions. And, as more control points are activated, decisioning rules are tuned, and step-ups are orchestrated, the value of the platform rapidly increases, giving businesses the ability to confidently welcome trusted customers and keep bad actors out.

To learn more about how Transmit Security can help your business stay secure without compromising customer experience through modern authentication, visit transmitsecurity.com.

About Transmit Security

Transmit Security gives businesses the modern tools they need to build secure, trusted and end-to-end digital identity journeys to innovate and grow. CX-focused, cybersecurity-conscious leaders rely on Transmit Security's CIAM platform to provide their customers with smooth experiences protected from fraud across all channels and devices. Transmit Security serves many of the world's largest banks, insurers, retailers, and other leading brands, collectively responsible for more than $1.3 trillion in annual commerce.