trustmi

Security and Privacy

V9 April 2025

Trustmi Platform Overview

Eliminate Socially Engineered Fraud

Trustmi's Behavioral AI defends against attacks that exploit human trust—like phishing and impersonations—by analyzing anomalies across emails, financial systems, and controls. It protects your people and your money, restoring trust where it matters most.

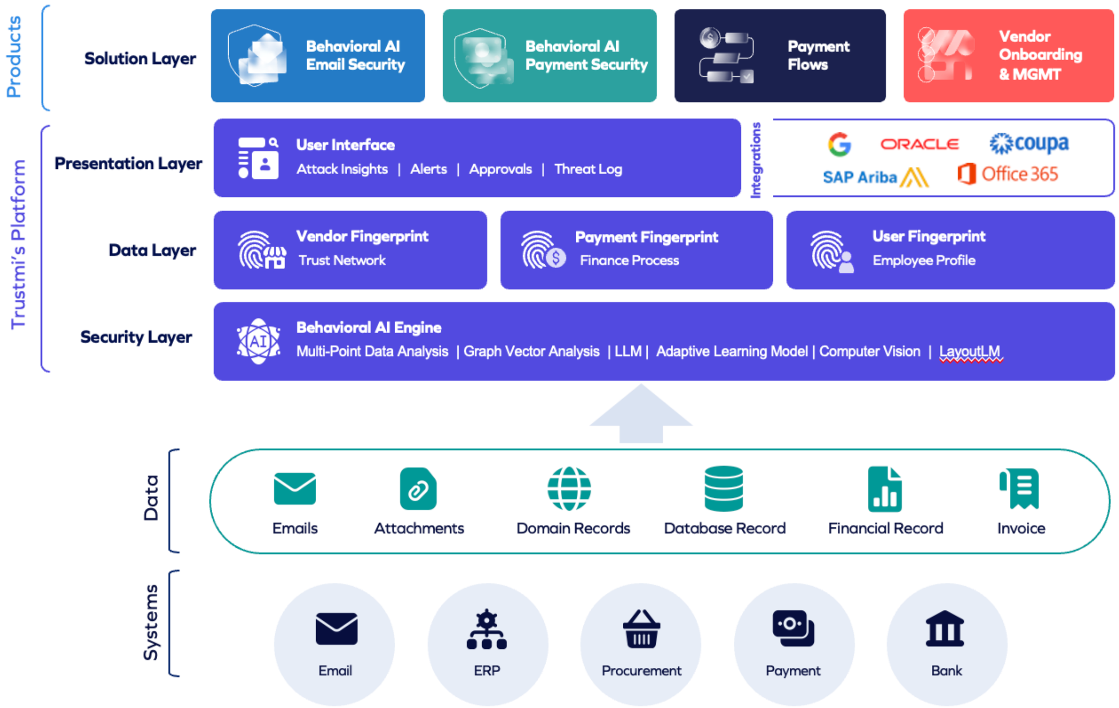

Defense-in-Depth Architecture for Payment Lifecycle Security

Trustmi is built on a **defense-in-depth security model**, applying multiple protective layers across the entire payment lifecycle—from the initial vendor communication and onboarding to payment execution within ERP systems.|

The platform ingests and correlates data from diverse internal and external sources, including ERP records, emails, financial documents, and other system artifacts. This data is further enriched through the **Trustmi Trust Network**, which integrates proprietary intelligence and open-source threat data to enhance visibility and context.
Through this enrichment and analysis, Trustmi generates a **normalized fingerprint** for each entity—such as vendors, payments, invoices, emails, and associated files. Each object is continuously assessed and assigned a **Trust Score** that reflects its composite risk, enabling proactive detection of anomalies and threats throughout the financial process.

Trustmi Security Layers

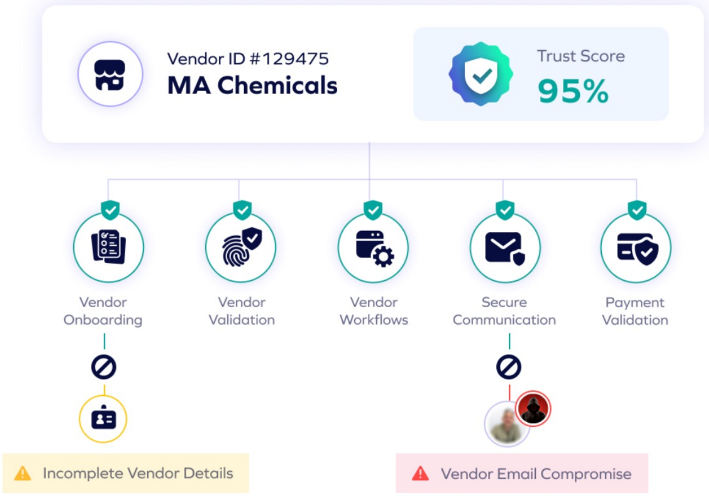

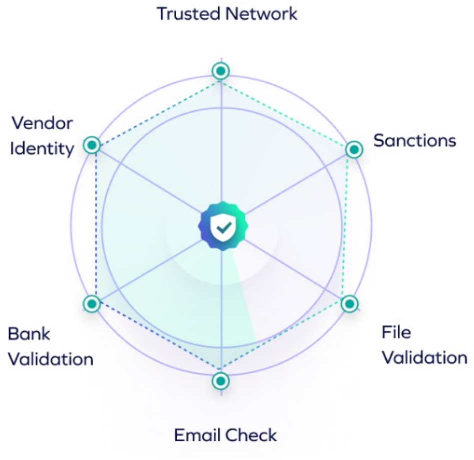

Continuous Risk Assessment and Adaptive Trust Scoring

Trustmi employs **multiple security layers** operating in unison to continuously monitor, analyze, and assess risks across all data points in the vendor and payment ecosystem.
Each interaction—whether a vendor detail, payment instruction, email, or document—is evaluated in real time. Trustmi correlates these data points into connected **object fingerprints**, enabling a **comprehensive, context-aware trust assessment** for every object within the payment lifecycle.

At the core of this process is an **AI-driven, self-learning engine** that adapts based on organizational behavior, historical patterns, and user actions. This feedback loop allows the system to refine trust scoring over time, aligning with each organization's unique operational procedures.

By extending validation beyond static checks to a continuous, adaptive model, Trustmi provides **resilient protection** not only against traditional fraud but also advanced threats such as ERP system attacks, account takeovers, and internal compromise scenarios.

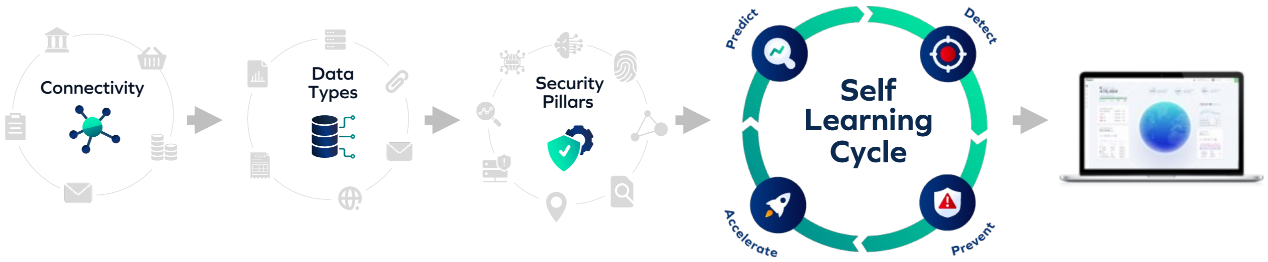

Trustmi Security Architecture

Our solution was built and designed from day one, with enterprise-grade security, implemented with in-depth security layers, including all the standard best practices.
Trustmi is SOC 2 – Type 2 Certified.

Trustmi platform is hosted in AWS Cloud, utilizing the AWS Well-Architected framework and implementing additional security layers around the external interfaces and critical services.
Access is restricted based on the least-privileged role-based access model.

AWS Security Best Practices

- ✓ AWS Security Best Practices: All resources within the VPC are accessed according to the assignment of Users, Roles, and Groups
- ✓ Multi-Factor Authentication (MFA)
- ✓ Access Control Lists (ACLs) and Security Groups
- ✓ Certificate-based authentication
- ✓ Specific access per IP address.

Network Security

- ✓ Access to the VPC is limited by IP address and firewall rules.
- ✓ Access to the VPC is not possible from the public Internet.
- ✓ Access to the VPC is restricted to VPN access only
- ✓ Access to the VPC Resources is via a bastion host
- ✓ The data for each customer is stored in a dedicated DB instance
- ✓ VPC internal network access is hardened based on resource groups
- ✓ VPC environment is properly segmented
- ✓ Access to frontend services via WAF & API Gateway

Trustmi Security layers

Data Protection

- ✓ All data is encrypted at rest
- ✓ Data at transit is encrypted using industry-standard HTTPS
- ✓ Enforced data retention policy

Security practices

- ✓ SOC 2 Type 2
- ✓ Frequent penetration tests & security audits
- ✓ Secure development lifecycle process
- ✓ CI&CD for Production deployment
- ✓ Central auditing for all security events

System availability

- ✓ Trustmi architecture is resilient – No single point of failure
- ✓ Microservice architecture with auto-scaling
- ✓ Backups & DR plans are conducted regularly
- ✓ Data at transit is encrypted using industry-standard HTTPS

AWS Environment diagram

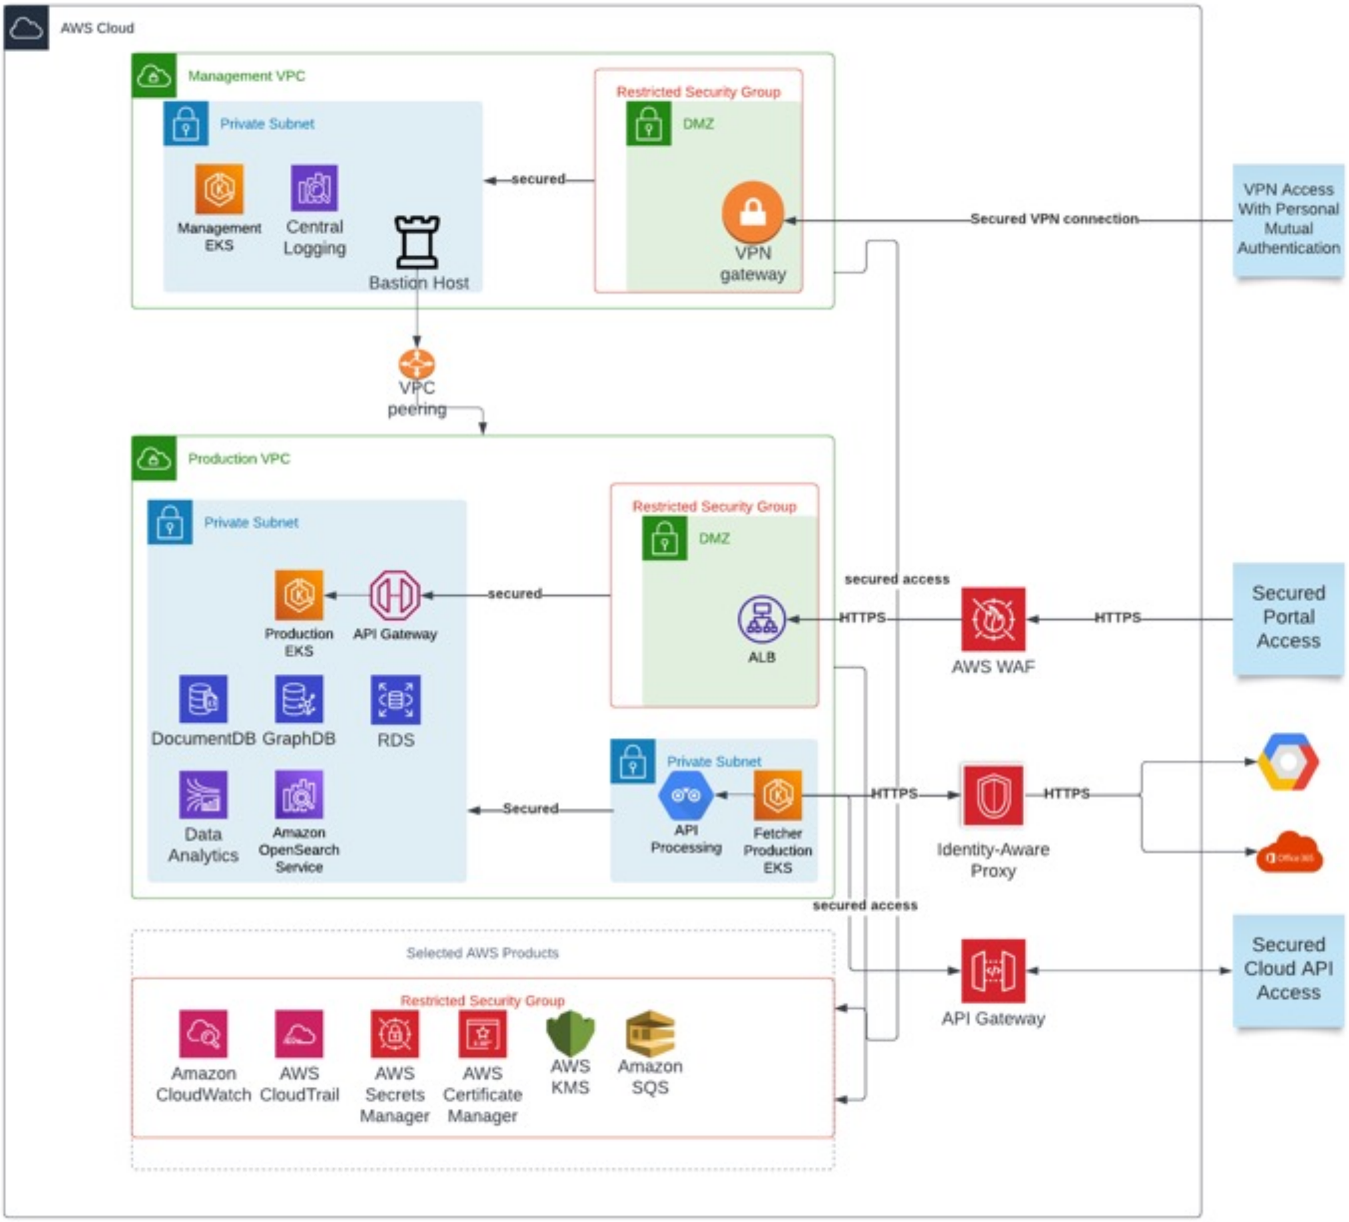

General Data Privacy

Trustmi implements robust measures to ensure that customer data remains secure and private throughout the data processing lifecycle. During analysis, Trustmi evaluates each object (e.g., vendor, payment, invoice) and assigns a **risk verdict** based on its proprietary assessment mechanisms.

To preserve privacy, Trustmi does **not store any raw data**. Instead, a **fingerprint object**—a minimized and privacy-conscious representation of the analyzed item—is generated and stored alongside the associated risk verdict. This fingerprint can be tailored to meet the customer's privacy and compliance requirements. In its most minimal form, the fingerprint consists solely of a unique identifier and the corresponding Trustmi verdict.

Customers have the option to **expand the fingerprint object** to enable additional functionality, such as vendor onboarding workflows or human validation processes.
All data analysis is performed **in-memory using Docker-based isolated instances**, ensuring that no raw personally identifiable information (PII) is retained after a verdict is rendered.

To support fraud prevention efforts, Trustmi provides mechanisms to **retain records of fraudulent verdicts** and any related evidence. These records are maintained in accordance with the customer’s requirements and are accessible for use in fraud investigations and follow-up activities. As part of its service, Trustmi also provides **detailed incident reports** upon request or when applicable.

All Trustmi objects and their verdicts are stored in **customer-isolated environments** and are used solely to deliver protection and services to the respective customer. Trustmi infrastructure is hosted on **Amazon Web Services (AWS)**, with all processing confined to a secure production Virtual Private Cloud (VPC) environment, protected by multiple security layers as outlined in previous sections.

Data Processing

Trustmi data processing occurs in multiple stages. In all locations,
Trustmi servers process all the raw data in memory without persistently storing the raw data. Processing occurs in near real-time and is repeated every few minutes as part of the product pulling cycle from
the Email provider’s API (For example).

The processing occurs as a server service application without human involvement, and once the processing is concluded, the data is purged from the memory and cannot be retrieved.

During the processing cycle, no human-readable output is produced, meaning the process runs as a service and only works for the following processing stage.

The initial Processing focus is on identifying the "interesting" emails that contain payment calls to action and should be considered relevant for further processing.

Examples:

- Payment Requests
- Vendor Registration
- Bank Account Information
- Vendor Change Management
- Invoices

Emails that do not include a payment call to action, such as those mentioned above, are considered irrelevant and ignored.

The second stage of processing focus is identifying the context of the email & parties involved as part of the transaction; this requires the Trustmi security engine to perform a complete analysis of all the data contained in the "Interesting" Email across all its layers, classifying and identifying the transaction various attributes.

The process is built on the same principles as the previous processing Cycle and is performed in a separate microservice in memory; once processing is complete, the raw data is purged.

The Email layers mentioned above include all layers of the email, including attachments, and are crucial for detecting malicious indicators across the entire payment request.

- Email Headers
- Email Body
- Attached Files Metadata
- Attached Files Content

Email-Based Threat Detection and Trust Assessment

The email processing engine decomposes each message into its constituent components—such as sender metadata, message content, attachments, and headers—and performs contextual analysis on the participants, language patterns, and embedded requests. This analysis is benchmarked against the vendor's historical communication profile to identify deviations indicative of malicious activity or fraud.

The engine generates a **Trust Verdict** that assesses the likelihood of the email being part of a vendor impersonation attempt, a business email compromise (BEC), or a spoofing attack originating from a lookalike domain. Anomalies in communication style, frequency, or request patterns trigger alerts and are flagged as potential security risks.

These email verdicts are correlated with broader vendor activity, incorporating additional data sources such as recent changes to payment instructions, account updates, or attempted financial transactions. This integrated approach enables proactive identification of high-risk scenarios—such as fraudulent bank account changes or unauthorized payments—by linking suspicious communications to downstream financial actions.

Email Attributes
Senders, Recipients, Domain, Conversation ID, Email Metadata.

User Attribute
Vendor identity, Signature, Title & Role.

Email Attachments
Financial Documents , Bank letter,w9 , Invoices.

The Trustmi Discovery Algorithm utilizes advanced graph analysis to map inter-organizational email communication flows, enabling deep visibility into the operational relationships and workflows underpinning key financial processes. This includes payment execution chains, vendor onboarding procedures, and approval processes for change requests.

By constructing a dynamic communication graph that represents the interactions between departments and external vendors, the algorithm identifies the implicit structure of a company's financial and procurement operations. This graph is then continuously compared against a known organizational baseline and established internal controls.

Through this comparative analysis, the system can detect anomalous communication patterns that may signify a control failure, unauthorized process deviation, or an account takeover event. These deviations serve as early indicators of social engineering tactics commonly used in fraud schemes, such as business email compromise (BEC) or internal account takeover.

This graph-based intelligence provides an additional, adaptive layer of defense—enhancing organizational resilience against sophisticated manipulation attempts that exploit human and procedural vulnerabilities.

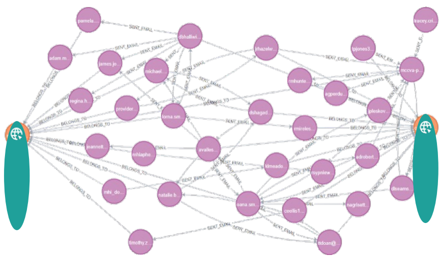

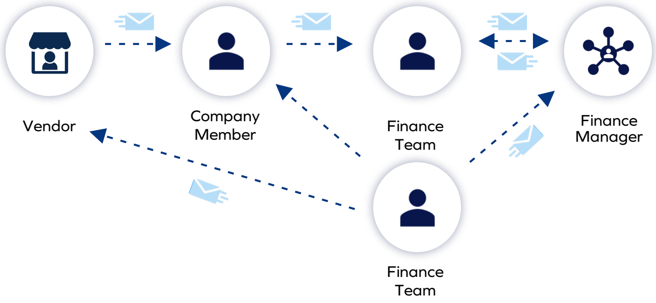

The processed data is broken down and classified into various attributes used to identify the vendor using Trustmi Vendor Repository and perform a security risk assessment on the transaction itself.

The Vendor identification process includes validation & identification of Bank Accounts & Vendor information and is evaluated against the proprietary Trustmi Database stored on a separate layer in the VPC.

Trustmi database contains hashed vendor information and metadata; for example, the Vendor Bank account is used to identify and compare the processed payment request to the vendor Trustmi repository.
The Security engine runs on the outcome of all previous services. Performs an analysis based on the Vendor Data, Customer Baseline, Security rule engine, Domain analysis tools, and many more layers, providing a risk score & suspicious indicator.

Once the processing is complete, only the transaction metadata is stored. All stored data is encrypted using a robust encryption scheme.

Stored Metadata customer record

- Email ID (Unique identifier for the Cloud Provider infrastructure)
- Trustmi Banner (Email Banner text with a status)
- Risk score (Value between 1 – 100)
- Security engine violation (Text containing the security violations)
- Baseline violation (Text containing the fields violating the baseline)

Stored Metadata crowdsourcing

- Vendor ID (Unique identifier for the Vendor)
- IBAN (Hashed using SHA256 with Salt)
- Vendor details (Hashed)
- Timestamp (“Payment timestamp”)

All stored data is encrypted and anonymized and stored within the internal Trustmi AWS environment with strict security policies, system access monitoring, and security prevention layers preventing unauthorized access by implementing the industry best practices.

Example Email Processing

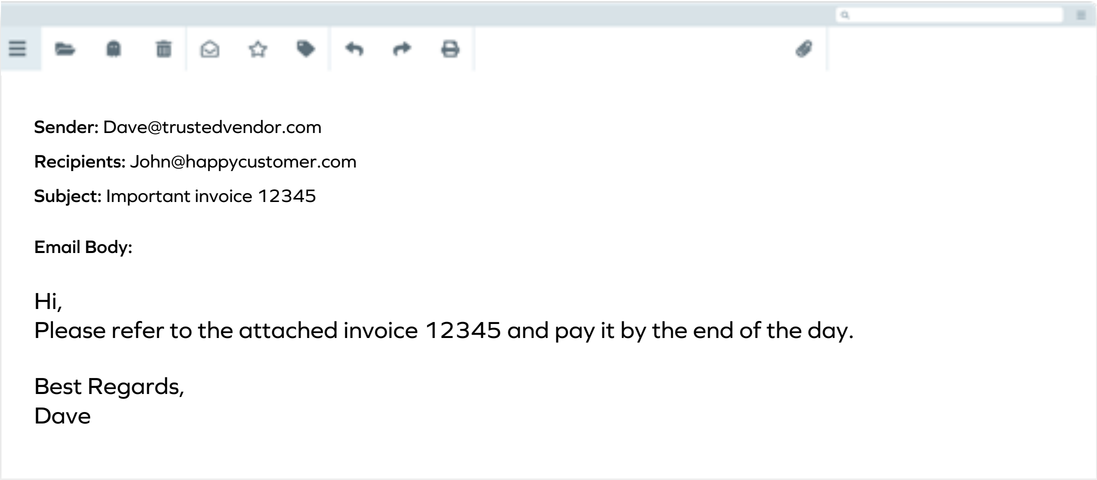

Sender: Dave@trustedvendor.com

Recipients: John@happycustomer.com

Subject: Important invoice 12345

Email Body:

Hi,
Please refer to the attached invoice 12345 and pay it by the end of the day.

Best Regards,
Dave

Attachment Content	Bank details: Bank of America IBAN: BA123412312312354

This email will be processed utilizing the NLP & Keywords engine.
It will be determined as a match for a financial transaction related to payments following a complete analysis process to extract the various labels & entities, and sentiment matching with the vendor fingerprint, behavior baseline, and additional security layers will generate a verdict followed by an incident & verdict creation mapped to the email id.

Data Subject Rights

Trustmi is committed to respecting and protecting the rights of individuals concerning their personal data, as outlined in the SOC 2 Privacy & General Data Protection Regulation (GDPR) and additional applicable laws. These policy provides a framework for managing and responding to data subject rights requests efficiently and lawfully.

Scope

- All data subjects whose personal data is processed by Trustmi, including customers, employees, vendors and third parties.
- All personal data collected, processed, stored, or transferred by the organization, regardless of the medium."

Rights

- **Right to Access:** To obtain confirmation of whether their data is being processed and access to that data.
- **Right to Rectification**: To correct inaccurate or incomplete personal data.
- **Right to Erasure (Right to Be Forgotten)**: To request deletion of their personal data under certain circumstances.
- **Right to Restrict Processing**: To limit how their data is processed in specific situations.
- **Right to Data Portability**: To receive their personal data in a structured, commonly used format and transfer it to another controller.
- **Right to Object:** To object to data processing based on legitimate interests, direct marketing, or automated decision-making.
- **Right to Withdraw Consent**: To withdraw consent at any time, without affecting the lawfulness of processing before withdrawal.
- **Right to Lodge a Complaint**: To file a complaint with a supervisory authority about how their data is handled."

Request Process

- Data subjects can submit requests via: Email: privacy@trustmi.ai
- **Requests must include sufficient information to verify the data subject's identity.**
- Requests will be acknowledged within [30 days] of receipt.
- Trustmi will respond within [60 days] unless an extension is required due to the complexity of the request.
- When applicable Trustmi will notify customers involved in a specific data request followed by a response based on the contract and customer requirements (Direct / Indirect)

* Additional details available on demand

Trust Network

The **Trustmi Trust Network** is a collaborative, AI-driven ecosystem designed to enhance the security of B2B payments by leveraging shared intelligence across a vast network of businesses and vendors. By aggregating and analyzing behavioral data, the Trust Network aims to detect and prevent fraud, particularly socially engineered attacks, before they can cause harm.

Crowd-Sourced Intelligence: The network consolidates data from millions of vendors and businesses, creating a comprehensive repository of trusted behavioral patterns and risk indicators. This collective intelligence enables the system to identify suspicious activities more effectively, even before they are evident in new attacks.

Trustmi ammonizes and tokenizes the Trust verdicts data generated by the system to utilize it to empower its detection capabilities and notify customers about a pending attack before it materializes.

The anonymized data is not associated with any customer and is received as open intelligence red/green indicators, meaning it cannot be traced or identified beyond a public notification similar to other crowdsourcing solutions. Customers can decide to opt out of this service and prevent any sharing of data with the network; however, they will not be able to receive the following benefits as a result.

Benefits

Proactive Fraud Prevention: By leveraging the collective intelligence of the network, businesses can stay ahead of emerging threats and continuously alert customers to upcoming threats

Improved Payment Security: With real-time data correlation and behavioral analysis, with the trust network, you can protect their processes from sophisticated social engineering tactics that are actively being used as part of a “live” fraud campaign

Streamlined Vendor Management: Vendors who would like to share information with The Trust Network, including detailed insights into vendor activity, performance, and risk trends, will not be visible to the customer who opts out of this service

Stored Customer Record

Object ID cf0835f5-bfa9-44b9-90a0-50d735f5-bfa9-44b9-90a0-50d7b95c8b2b

Vendor: Trusted Vendor

Vendor ID: Z271218A

Financial indicators : Bank account Hash

Trust score : Trusted

Security Violation Baseline (‘English score low”, “Urgency Call to action”)

Baseline Violation (“Vendor partial signature”)

Timestamp 01/01/25_18:35:22

Stored Crowdsourcing Record

Vendor ID: Z271218A

Financial Details HASH: 2)_20MgErExOq1tm7EvEM4CGSZknDHcVPbbMds/

Trust score : Trusted

Timestamp 01/01/25_18:35:22