

FNA Papers No. 8

Countering Consumer Fraud and Scams with National Fraud Portals

2024

Florian Loecker
Amanah Ramadiah
Kimmo Soramäki

Countering Consumer Fraud and Scams with National Fraud Portals

FNA Papers No. 8 | 2024

Florian Loecker
FNA

Dr Amanah Ramadiah
FNA

Dr Kimmo Soramäki
FNA

Florian, L., Ramadiah, A., Soramäki K., (2024). Countering Consumer Fraud and Scams with National Fraud Portals. FNA Papers, No. 8-2024

Florian Loecker

FNA | Florian@fna.fi

Florian has 10+ years working in deep tech analytics, most recently as VP at Verisk, responsible for managing software development, product, as well as data science teams for the Arium product. Arium was acquired by Verisk in 2017. He also has deep expertise in creating enterprise-grade analytics and data software products, with deployments to some of the world's largest companies.

[LinkedIn >](#)

Dr Amanah Ramadiah

FNA | Amanah@fna.fi

Amanah is the Head of Analytics and Client Success (Asia) at FNA. She leads a team of data scientists in Asia and leads the delivery and development of various advanced analytics solutions (e.g., supervisory technology, fraud analytics, digital twin simulation). She is also responsible for managing client relationships and providing strategies for successful solution adoption and expansion efforts in Asia. Amanah holds a Ph.D. in Computational Finance from University College London, an MSc in Financial Risk Management from University College London, and a bachelor's degree in Computer Science from Universitas Indonesia.

[LinkedIn >](#)

Dr Kimmo Sormäki

FNA | Kimmo@fna.fi

Kimmo Sormäki is the Founder and CEO of FNA and the author of "Network Theory and Financial Risk". He has over 25 years of experience working with Central Banks and Financial Market Infrastructures. In 1997, Kimmo developed the world's first simulator for interbank payment systems - and since then, has been regularly invited to lead and contribute to simulation and payment system innovation projects with organizations like Bank of England, CLS, Payments Canada and SWIFT. He is a frequent speaker at industry events and has written over 50 articles cited in more than 2,500 academic publications. Kimmo has an MSc in Finance and a DSc in Operations Research, both from Aalto University in Finland.

[LinkedIn >](#)

Contents

- Why a National Fraud Portal..... 5**
- Components of the National Fraud Portal..... 8**
 - 2.1 Pre-settlement Components..... 8
 - 2.1.1 Fraud and Mule Detection with Graph AI..... 8
 - 2.1.2 Complementing FI-level Solutions..... 10
 - 2.2 Post-settlement Components..... 11
 - 2.2.1 Fraud Reporting and Case Management..... 11
 - 2.2.2 Case Validation..... 12
 - 2.2.3 Case Prioritization..... 12
 - 2.2.4 Real-time Money Trails..... 12
 - 2.3 Data Management Components..... 14
 - 2.3.1 NFP Data Hub..... 14
 - 2.3.2 Data Requirements..... 15
- How To Get Started Building a National Fraud Portal..... 17**
 - 3.1 Industry Engagement..... 17
 - 3.2 Pilot Project..... 18
 - 3.3 Building the Data Infrastructure..... 18
 - 3.4 Deployment of the Post-settlement Components..... 19
 - 3.5 Deployment of the Pre-settlement Components..... 19

01

Why a National Fraud Portal

Consumer fraud and scams have arguably become the number one priority for many central banks, consumer protection agencies and other competent authorities within the last year. While the statistics are scattered and lagging, personal experience and high-profile news reports give us an indication that the problem is growing rapidly in almost every country.

At the national level, political and societal pressures to reduce consumer losses are driving political leadership, central banks, and consumer protection authorities to look for solutions. High and growing fraud rates can also lead to the erosion of confidence in digital payments, with damaging consequences for national economic and financial development objectives.

In many countries, drivers for the explosion in fraud can be traced back to one or more of these factors:

- The digitization of payments, new financial services brought by Fintechs, and the introduction of instant or faster payment systems and schemes have allowed criminals to acquire funds faster. Furthermore, they can obscure the source of funds more quickly, moving money intraday through sophisticated money mule networks across banks and payment schemes before taking funds out in e.g. cash or crypto.
- Highly successful and fast-paced financial inclusion initiatives have created great opportunities in many economies while at the same time leaving a large, less financially literate population open to falling victim to fraud and scams.
- Technological developments have increased the sophistication of criminal techniques. For example, fraudsters use Artificial Intelligence (AI) and Large Language Models (LLMs) specifically developed for illicit use cases, such as FraudGPT, to scale up scams and reach consumers in a wider range of languages and geographies.

In today's technological environment, criminals have the advantage. Current efforts to contain fraud and scams are carried out in silos, as cross-bank and cross-payment scheme collaboration is not the norm. As no individual bank has complete visibility of the money muling network with their payment data alone, individual banks are less effective in building predictive models of fraudulent accounts due to the partial nature of their data.

In this paper, we argue that setting up a new Digital Public Infrastructure - a National Fraud Portal (NFP) - is the only way to address fraud and consumer scams efficiently. NFPs provide a technological solution as a shared facility for banks, law enforcement, the Financial Intelligence Unit (FIU), the central bank, the conduct supervisor, and other stakeholders. Further down the line, NFPs can connect to one another as cross-border criminal activity increases (a likely consequence of suppressing fraud domestically).

The National Fraud Portal (NFP) enables:

- The real-time tracing and tracking of fund movements across the banking system that allow banks to recover funds for victims quickly
- The validation and prioritization of cases across the economy using data-driven models
- The faster identification of new mule accounts at a reduced cost
- More accurate methods for fraud detection and risk scoring that employ Graph AI deployed on network data.
- The real-time provision of risk scores and features to banks via APIs, allowing them to improve their fraud models and make faster, more accurate decisions about preventing fraudulent payments before settlement.

In this paper, we detail the **technological components of the National Fraud Portal**. The paper is a result of FNA's work building technology for National Fraud Portals in Southeast Asia, as well as conversations with over 100 institutions across 20 countries that are actively working on, or have an interest in tackling the problem of fraud and scams.

In this paper, we do not cover the equally important topics of:

- **Criminal law** - whether and how criminal proceeds can be restituted to victims
- **Liability** - who should bear responsibility if criminal proceeds cannot be recovered
- **Competence** - who should own and operate the NFP
- **Data privacy law** - what data can be shared between participants - although this can be minimized as it relates to Personally Identifiable Information (PII) with approaches addressed in this note, we see that this can be minimal for the approach discussed in this note
- Other policy-related or legal questions

Instead, we focus on the benefits of using technology to address the problem. At this stage, we emphasize that the proposed technology can be customized to address different requirements along all these topics (e.g. in section 2.3.1, we present alternative approaches to tackle the competence topic). We also note that the NFP complements and enhances existing approaches, which are typically more isolated and siloed at FIs, by adding capability in cross-bank and cross-scheme domains, which are currently exploited by criminals, fraudsters and scammers.

In our view, the operationalization of the NFP must run in parallel with the development of legal and policy frameworks. It should also be noted that all these aspects must be customized at the local level to develop nation-specific features of each National Fraud Portal. We hope that this paper provides helpful input for national and international actors in their own journeys towards combating fraud & scams and making their citizens safer.

02

Components of the National Fraud Portal

The National Fraud Portal (NFP) is a digital public infrastructure for combating fraud and scams at the national level (or even cross-border), built as a shared service with minimal to no requirements for sharing sensitive data. This means all stakeholders, which could include the central bank, payment system operators, law enforcement, and financial institutions, can have access to intelligence outputs of the same system to support their own objectives. This section explains and discusses the underlying functional and technological components of the NFP. We start with pre- and post-settlement components and conclude with data management components.

2.1 Pre-settlement Components

Pre-settlement components focus on preventing a fraudulent payment from being released through the use of risk models and rules, typically at the Financial Institution (FI) level. In so doing, financial loss is prevented before it occurs and, therefore, does not incur costs associated with post-fraud investigation and recovery of funds. It helps FIs maintain positive customer relationships, is a required activity in many jurisdictions, averts the liability of FIs, and, most importantly, keeps participants in the financial system safe from significant financial harm.

The NFP adds significant value to this process, complementing and significantly improving existing processes employed by FIs today. This is due to its better ability to deploy advanced network-level data visibility and models to detect fraud. Using these innovations, banks can make more informed decisions to block the release of fraudulent payments efficiently.

This section is organized into two components: first, we discuss the usage of novel models built on graph AI, and second, we discuss their integration into existing workflows employed by FIs around the world today.

2.1.1 Fraud and Mule Detection with Graph AI

The National Fraud Portal enhances existing capabilities by leveraging Graph AI, Machine Learning Models and feedback loops with post-settlement systems. It is a real-time system providing both summary as well as individual risk scores on a payment-by-payment basis in real-time. These scores are built from several sources:

- A flag or score representing whether the recipient account is or is suspected to be a mule, which is constituted from two sources:
 - Graph AI models score the risk of each account in the economy as a mule.
 - Mule accounts that were flagged as part of the post-settlement investigation workflow (see section 2.2).
- Graph AI models evaluate each new transaction event by analyzing anomalies of the sender and recipient accounts, as well as those in the subcomponents of the network of which they are a part.

Unlike models run at the bank level, network-level payment data provides a unique vantage point for training and running models. FNA's work on real payment data has shown increases in precision and recall of up to 56% and 85%, and a reduction of false positives by 50% by using graph AI deployed on network data, as opposed to traditional methods that use neither. Research conducted by the Bank for International Settlements finds that current siloed data arrangements and rule-based approaches are hindering fraud detection capabilities¹. In another paper² by IBM's Thomas J. Watson Research Center and ABN AMRO Bank, the researchers find that taking into account "graph features" that summarize the position of the sender and receiver in the payment network improves the accuracy of bank-level models by 30%.

Improvements in the performance of models have several advantages. A higher number of true positives means the model improves the detection rate of fraud. Fewer false positives mean lower investigation costs and a lower number of legitimate payments that get blocked. Better models also allow for the faster identification of mules so that fraud losses can be reduced, as well as better information to validate and investigate cases, leading to faster case resolution times. On a national level, these benefits accumulate quickly.

Moreover, as investigation capabilities are improved (see 2.1-2.2), benefits also accrue. With better-labeled data (i.e., data about known mules) used to train fraud detection models, we can further improve the accuracy of the models over time. This virtuous cycle continues as the models become better and better at detecting fraudulent activity. This is particularly advantageous versus more static, rules-based approaches to fraud detection, where fraudsters learn to adapt their behavior to avoid detection. With a machine learning approach, the models continuously adapt to fraudster behavior, limiting criminal actors' ability to avoid detection by changing how they operate.

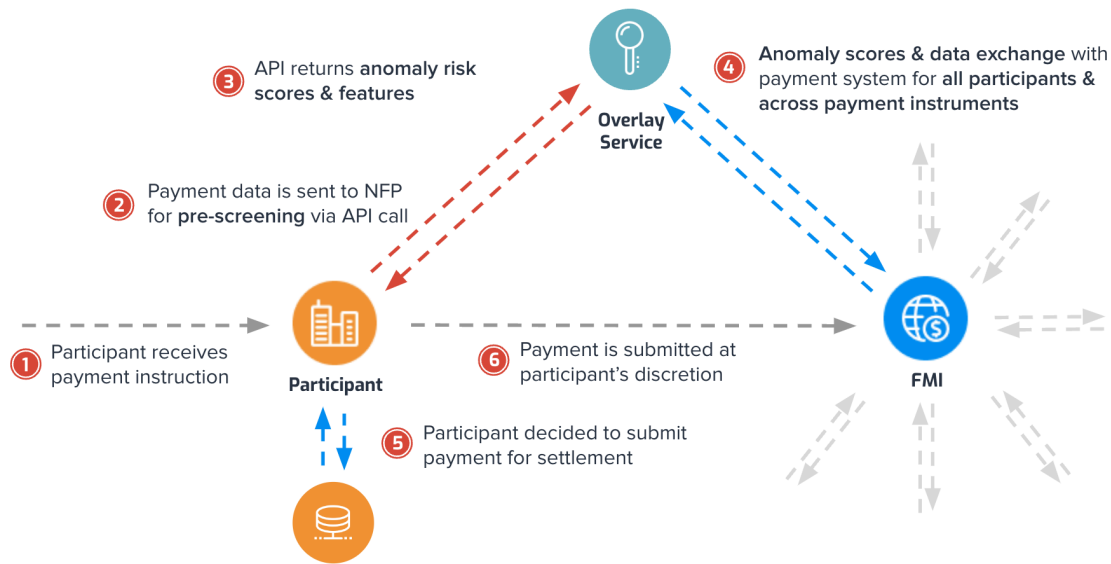
Compared to siloed data and standard statistical modeling approaches, a graph-based approach also reduces the ability of fraudsters to avoid detection. This is because fraudsters are only able to change their own behavior to obscure their activity, but not those of the accounts they are interacting with, such as the victims'. Graph AI-based models capture these neighboring behaviors and interactions within the risk scoring, thus improving detection rates.

¹ BIS (2023)

² Molloy et. al: Graph Analytics for Real-time Scoring of Cross-channel Transactional Fraudin: 20th International Conference on Financial Cryptography and Data Security, FC 2016, Vol. 9603 LNCS, 31 Springer Verlag, 2017, pp. 22-40

2.1.2 Complementing FI-level Solutions

To visualize how NFP Fraud detection complements existing workflows employed by FIs today, it is helpful to look at an example where the NFP is hosted by a Financial Market Infrastructure (FMI). The following chart visualizes the steps taken when scoring a transaction (e.g., by checking whether the recipient of the transaction is a suspected or confirmed mule according to models or investigative discovery):



In the first step, the participant (e.g., a bank or FI) receives a payment instruction from its customer. At this stage (and concurrently with other actions such as payee verification), the participants have already sent data about the payment to the NFP (step 2). It then receives scores for the payment and recipient account via API integration (step 3). Step 3 can also contain other network-level statistics about the payment that are not available to the sending banks, such as network features or features related to the payee account.

In step 4, new payments that arrive from any FI are assumed into the database and are scored. Finally, step 5 takes place inside a FI and contains all other pre-settlement activities, including its internal fraud detection (where present) and merging of results with the scores received through step 3 in line with its internal procedures and processes. Finally, the payment is settled through the usual process (in this case, via the FMI as Step 6).

It should be further noted that NFP Fraud detection can handle thousands of payments per second (using appropriately scaled computational infrastructure), adding minimal latency to payment processing and release, and is therefore able to integrate with and support real-time payment systems.

2.2 Post-settlement Components

Post-settlement solutions refer to workflows that happen after a fraudulent transaction has settled and enters the money laundering and cash-out process. The NFP offers components for improving the capture of initial fraud reports (by banks, police, or other stakeholders) and the validation and prioritization of cases, enhancing the investigation and recovery of funds. Investigation results may be captured and stored in databases, helping pre-settlement activities (e.g., through mule blacklists) and future investigations.



2.2.1 Fraud Reporting and Case Management

Case management refers to the collective technology and processes for creating, assigning, validating, prioritizing, and managing fraud cases. It enables the delegation of investigative tasks to team members, collaboration, and information sharing.

Often, the starting point is a fraud report (forwarded by banks, police, or other national crime agencies) that surfaces in the case management interface and can then be triaged by dedicated staff. This fraud report contains basic information about the fraud, such as its type or scheme, as well as basic information about the account holder (including demographics, as well as previous history of fraud, although that's not required). At this stage, the functionality outlined in 2.2.2 and 2.2.3 is taken into account to gauge the validity of the case and estimate the amount of recoveries available. The combination of the validity of the case and the chances of successful resolution (through recovery) forms the basis for prioritizing a case.

Once a case is prioritized, case officers may be assigned to the actual flagging and recovery workflows. Money Trails is a key component in this stage of the investigation, automating and streamlining this workflow and collecting evidence of money muling.

Once a case is resolved, post-processing is performed to:

- Archive the case for future reference, as well as potentially reopen the case.
- Provide automated reporting (relevant for law enforcement and judicial processes)
- Store validated mule accounts in a database for future reference, e.g., the Fraud detection use case described in 2.4.

Finally, we discuss the technical aspects of case management, such as authentication and authorization, to ensure only specific users can access sensitive personal information and data and carry out privileged actions such as blocking accounts. The details of this depend on the concrete implementation. Looking at shared case management between FIs as an example, only the original FI can see the data in an untokenized format.

2.2.2 Case Validation

Not all cases brought on by consumers are valid fraud or scam cases – instead, they could be related to disputes between the buyer and the seller or reported by criminals themselves to flood the system. Therefore, the quick validation of cases is essential.

Money Trails offers a data-driven method to validate cases. The money trail corresponding to a particular fraud case contains valuable information that helps verify whether the reported incident is genuine (e.g., it directly checks whether muling occurs on the money trail). This information is then summarized for case officers as a case validity score, and policies can be developed to filter cases based on their score.

2.2.3 Case Prioritization

Case prioritization depends on the exact KPIs set for the NFP. These KPIs may change occasionally as the focus shifts to the total value of recovered funds, the number of mule accounts identified, or a combination of these. Case prioritization is a critical capability in achieving KPIs.

A data-driven prioritization system, coupled with Graph AI-generated mule scores, helps calculate a likely estimate for the total amount of recoverable money. This estimate can then influence the prioritization of cases before assigning valuable resources, helping to ensure the highest level of recoveries or achieve other KPIs. This information is then summarized for case officers as a case priority score, and policies can be developed to, e.g. work through valid cases in order of priority.

2.2.4 Real-time Money Trails

Fraudulently acquired funds are often transferred immediately to mule accounts before being moved onward to other accounts acting as mules. Fraudsters may employ several generations of mules to conceal the final destination of the funds before converting the money to crypto or taking it out of the banking system in cash. Given the broad implementation of instant payment systems, there is generally a short window of opportunity to stop the funds before this happens.

To have a chance of success in clawing back ill-gotten funds, investigators need the ability to automatically track and trace money trails emanating from fraud events across the financial system in real-time. The speed of fraudsters’ activities and the slow and manual process of piecing together these money trails in legacy systems results in low recovery rates. In contrast, investigators with complete and immediate visibility of how money moves through the system can react much faster to request that certain accounts or monies on the account to be frozen or blocked to prevent any further movement of these funds, improving case resolution and recovery rates.

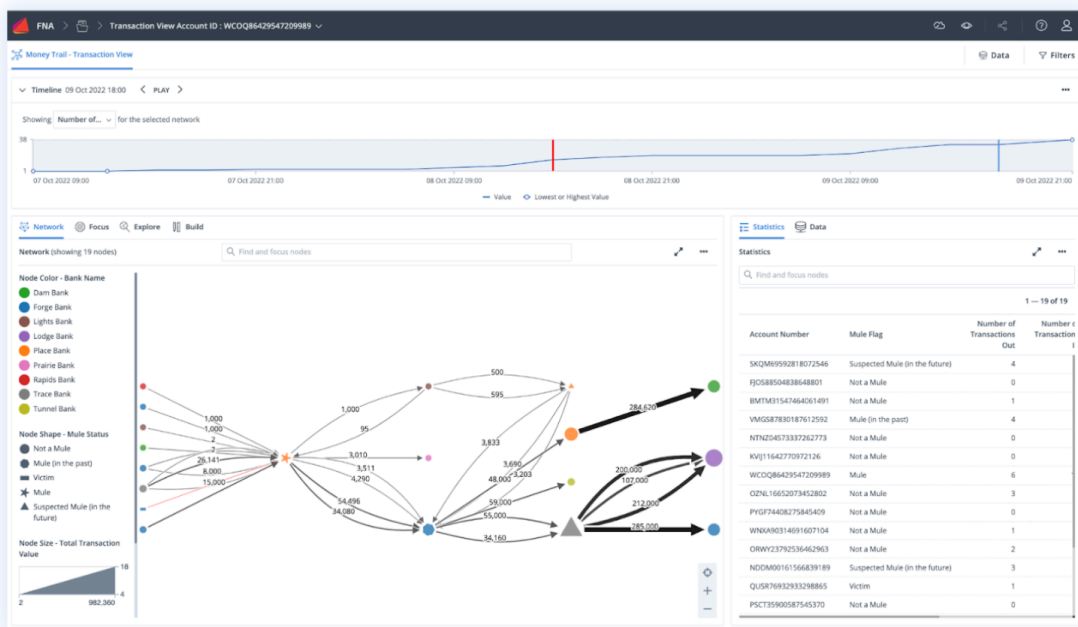


Figure 1: Screenshot of FNA Real-time Money Trails Application

Beyond tracking and tracing suspicious funds, real-time money trails provide additional analytical and data overlays to support investigators. For example, Depending on the legal framework:

- Money Trails may provide a framework for calculating the liability of each mule account holder to the victim (or its subrogates). This could be set as the total amount of transactions that can be traced back to the defrauded account or the set of accounts received by the mule account in question.
- In shared environments such as those with case officers deployed by individual FIs, it is helpful to display the total amount of blocked funds in real-time. This allows investigators to track how much money has been marked as potentially recoverable and thus serves as a progress indicator for the case.

The ability to track and trace activity from reported fraudsters also helps investigators identify undetected mule accounts as well as unreported victims. There are several benefits to this. Firstly,

they may be used to blacklist accounts aiding fraud detection and prevention schemes, as we discuss below. Secondly, the newly identified mule accounts can be used as evidence in other separate or connected cases. For example, the evidence that an individual account that is under investigation is connected to three other newly detected mule accounts may speed up the investigation outcome. Connecting data on separate cases can also uncover wider criminal schemes, as well as identify ‘big fish’ individuals in these schemes. Thirdly, the new information feeds into better-labeled datasets and improved modeling (see Fraud Detection section below).

At the most basic level, case officers manually add transactions to Money Trails. Money Trails can also integrate with automated transaction data feeds and automatically append any new transactions to the existing trail in real-time. Any transaction added by the automatic feed may then be disregarded or flagged as suspicious by investigators.

Once an automated feed is integrated, money trails can become rather large as mule accounts often make payments to other legitimate businesses (e.g., to pay rent or utilities). As a result, a high share of transactions from the mule account may not be relevant to the investigation. Graph AI-generated mule scores assist investigators by filtering and prioritizing accounts, given that Graph AI can reliably detect mules with a low false-positive rate, further accelerating the speed of handling cases.

2.3 Data Management Components

Data Management is key to the public utility and an important enabler for the NFP. After all, the system has to process a large volume of transactions (which are typically in the order of millions or even tens of millions of transactions per day, with pronounced spikes during certain times of the day), meet associated service level agreement (SLAs), and support rigorous requirements on enterprise-grade security and availability.

2.3.1 NFP Data Hub

The NFP Data Hub underpins the NFP, providing a shared service for storing, reading, and managing the data used in the NFP and its component applications. The NFP Data Hub is comprised of:

- a database for storing transactions,
- a database for registering known or suspected mules,
- an incident database, storing e.g., data related to information available in money trails and case management.

The NFP Data Hub can ingest and process data (in particular, transactions) in real-time and at high volumes (up to thousands per second, equivalent to tens of millions per day). It can interface with standard APIs, including HTTP REST, Apache Kafka or JDBC. These APIs are shared between all relevant stakeholders. Section 2.1.2 offers an example of a functional architecture for Fraud Detection.

As a critical infrastructure, the NFP Data Hub offers enterprise-grade security and is highly available. The NFP can be set up as a centralized or peer-to-peer system between FIs. In the former, all services and data hosting are performed in a central hub, whereas in the latter, every participating FI hosts the data, but there is no central hub. A peer-to-peer approach is advantageous when ownership of the NFP is difficult because of data privacy or other reasons.

The hosted data comprises either pre-settlement (payment orders) or post-settlement (settled payments) data, depending on the use case:

- Transactions used for fraud investigation cases (such as Money Trails) are post-settlement.
- Fraud detection analyzes payments before settlement. However, historical data on settled transactions is also highly useful for analytical use cases such as periodically performing mule detection and re-calibrating models.

2.3.2 Data Requirements

A key advantage of the proposed architectures and graph-AI-based models is that there are no hard requirements for any personally identifiable information (PII) or other sensitive data other than in contexts where some PII is already known (such as that the sending and receiving FIs processing a transaction are privy to the account numbers of the payer and payee).

However, it is possible to share more data across FIs if permissible and desired – and indeed, there are real-world examples of such data sharing. In the following sections, we discuss data requirements by various sub-components in the system.

2.3.2.1 Case Management and Money Trails

The table below provides the basic data requirements for the Money Trails (track & trace) solution to be operationalized and also captures whether that data may be obfuscated/tokenized or must be shared in a clear format. Transaction party FIs are the two FIs sending and receiving the transaction data.

Data	Description	Transaction Party FI	Other Stakeholders
Sender	Sender account number and identity	Clear	Tokenized
Recipient	Recipient account number and identity	Clear	Tokenized
Timestamp	Date and time of the transaction settlement	Clear	Clear

Data	Description	Transaction Party FI	Other Stakeholders
Value	Value of the payment	Clear	Clear
Transaction ID	Transaction unique identifier	Clear	Clear

Where data can be obfuscated, this is done by secure one-way hashing using modern cryptographic methods such as SHA256, rendering reconstruction of the original data impossible.

In addition, we have graph-AI-based use cases for account scoring, case validation and prioritization. In these cases, for model validation and possibly training, it is desirable to have labeled fraud data consisting of:

- A database of historical transactions spanning at least six months.
- A database of historical transactions or accounts that are suspicious or (ideally) have confirmed cases of fraudulent transactions or muling activity.

2.3.2.2 Fraud Detection

Fraud detection data requirements are broadly similar to but lighter than those in the previous section for case management and Money trails.

The key difference is that payment orders are scanned pre-settlement, which may occur using obfuscated/tokenized sender and recipient data. In a decentralized framework, this would even occur on-premises at the sending FI.

Any reference data, such as blacklists of mules, are kept in a tokenized format which cannot be reversed.

2.3.2.3 Additional Data

While not mandatory, incorporating additional data elements can enhance both the investigation process and the accuracy of fraud detection. Supplementary fields may include:

- Transactional context (location, device used, ...)
- Security flags or alerts raised during the transaction process
- Cash withdrawal, e-commerce transactions, bitcoin transactions
- Telecommunication data

03

How To Get Started Building a National Fraud Portal

In working to implement National Fraud Portals already, we recognize that this is not purely a technical endeavor, and equally important measures need to be taken to engage stakeholders on a common vision and roadmap. In this section, we discuss our experience in doing so, including potential timelines as well as potential issues of a regulatory or otherwise non-technical nature.

3.1 Industry Engagement

The primary goal of industry engagement is to build a roadmap and shared vision for the NFP and further facilitate knowledge acquisition among key stakeholders regarding the solution's capabilities. These stakeholders may include different departments at central banks, financial institutions or payment system operators, as well as banking/payment system associations, and FIUs. This stage is crucial to give stakeholders information on crafting a comprehensive, national-level anti-fraud portal with operating procedures (e.g., incident reporting, fund recovery, cross-bank data sharing) tailored to their specific needs and that align with the existing laws of their country.

Industry engagement usually entails a 2-3 month period of developing both a prototype of the technology, such as the basic package of the post-settlement solution (Money Trails) and organizing a series of training workshops. Both requirements gathering and development of a roadmap to deliver the NFP can be conducted during this stage.

Given that stakeholder engagement typically needs to occur in the early stages, comprehensive data may not yet be available. For instance, mule/fraud labels are typically not yet compiled, and banks may be hesitant to share data. However, client metadata and relevant aggregate statistics (e.g., total daily volume, value, number of participants, etc.) can be sufficient to create synthetic individual payment data for rapid prototyping and for fostering stakeholder engagement.

Furthermore, a centralized data infrastructure may not yet be established at this stage. Therefore, a turnkey cloud installation that FNA can deploy quickly can be put in place, with subsequent granting of full control of the infrastructure to our customers. Turnkey cloud installations offer a convenient and efficient way to deploy software, enabling customers to leverage our solution without the complexities typically associated with implementation and maintenance.

3.2 Pilot Project

During this phase, typically spanning 4-6 months, the project is initiated to showcase the value of the NFP, including both pre-and post-settlement solutions. Unlike the Industry Engagement phase, empirical data is ideal for this Pilot phase.

While FNA possesses a robust Graph AI fraud detection model that generally performs well, authentic local payment data alongside corresponding mule/fraud labels are needed for validating the model's accuracy and potentially refining it through retraining if necessary. A minimum of 6 months' worth of current payment data, inclusive of related fraud data identifying active fraudulent accounts within the same timeframe, is ideally required. Key work steps in the pilot phase are:

- 1 Initial analysis of the payment and label data
- 2 Graph model training and assessment
- 3 Money Trail visualization configuration
- 4 Detailed Roadmap planning and refinement
- 5 Establish technical integration for stakeholders
- 6 Results presentation and stakeholder discussions

3.3 Building the Data Infrastructure

The subsequent phase involves the development of the NFP Data Hub, a shared service for storing, reading, and managing the data used in the NFP and its component applications. Typically, this phase will span 6-12 months. Depending on stakeholders' preferences, the NFP Data Hub can be deployed either on-premise or in the cloud.

Ideally, this stage should address the data needs for both (1) case management and money trails and (2) fraud detection. However, the approach can also be iterative, contingent upon the NFP Operator's choices regarding the deployment of the National Fraud Portal.

Two roadmap options that the NFP Operator can choose to build the NFP are:

Parallel
Deployment of
NFP

NFP Data Hub that satisfies both data requirements for (1) case management and money trails and (2) fraud detection → in parallel deployment of the post-settlements and pre-settlements solutions

Iterative
Deployment of
NFP

- a. NFP Data Hub that satisfies data requirements for case management and money trails → deployment of the post-settlement solutions
- b. NFP Data Hub that satisfies data requirements for fraud detection → deployment of the pre-settlements solutions.

3.4 Deployment of the Post-settlement Components

After the development of the NFP Data Hub, as outlined in section 3.2, stakeholders have the option to follow a particular path for deploying post-settlement solutions. This begins with the implementation of case management functionality alongside money trails. In the absence of real-time data streaming, stakeholders can initiate data updates in batches, such as at intervals of e.g., 5 minutes.

Considering the diverse range of users from various stakeholders, we advise dividing the deployment into three sub-phases.

1. The first sub-phase involves configuring the case management and money trails according to stakeholders' preferences, typically spanning 1-2 months.
2. The second sub-phase entails initiating a beta release with a select group of users, including the central bank and major banks. This beta phase facilitates user acceptance testing and allows for collection of valuable feedback to enhance the beta version of the NFP, lasting approximately 3-6 months.
3. The third sub-phase involves proceeding with an industry-wide release involving all relevant stakeholders.

Once the NFP post-settlement components have achieved widespread usage, including case management and money trails, stakeholders may consider integrating additional functionalities. One such step involves incorporating a case validation and prioritization solution, leveraging a data-driven approach along with graph-AI generated Mule scores.

3.5 Deployment of the Pre-settlement Components

Pre-settlement solutions focus on monitoring prospective transactions and providing FIs or consumers with flags if the relevant transactions are suspected of being due to fraud or scams. Flags are provided based on two sources, as has been discussed in section 2.1.2: mules confirmed through post-settlement investigation and mules that are flagged by Graph AI.

Confirmed mules are often captured in existing databases, where poor data quality and lack of consolidation at the national level are key issues. These issues may be resolved through a data science workstream during the initial phase of setting up the NFP. This initial set of mules already provides an important foundation for improving pre-settlement solutions via a blacklist. Moreover, as the Money Trails solution is set up and put into production, newly flagged mule accounts are automatically synchronized into this blacklist.

Separate from this, smart systems based on Graph AI typically require several organizational steps. Namely, in addition to the technical integration of scores emitted by the NFP Fraud detection system (see section 2.1.2 for a functional architecture of this), initial and ongoing/periodic validation of resulting risk scores in terms of basic data scientific metrics such as accuracy, precision, recall, rate of false positives, etc. is typically required by regulators and FI-internal risk frameworks.

As a final point, we note that in FNA's work, operation and use of pre-settlement systems typically remain the responsibility and remit of individual FIs, meaning they have discretion in terms of how NFP Fraud detection outputs are used and operationalized in their own institutions.

Countering Consumer Fraud and Scams with National Fraud Portals

Florian Loecker, Dr Amanah Ramadiah & Dr Kimmo Soramäki, 2024

© FNA

FNA | Albert House | 256 - 260 Old Street
London | EC1V 9DD