

Privacy-Preserving Cross-Bank Financial Crime Analytics at Scale

A joint proposal from:
Secretarium and FutureFlow

June 2024

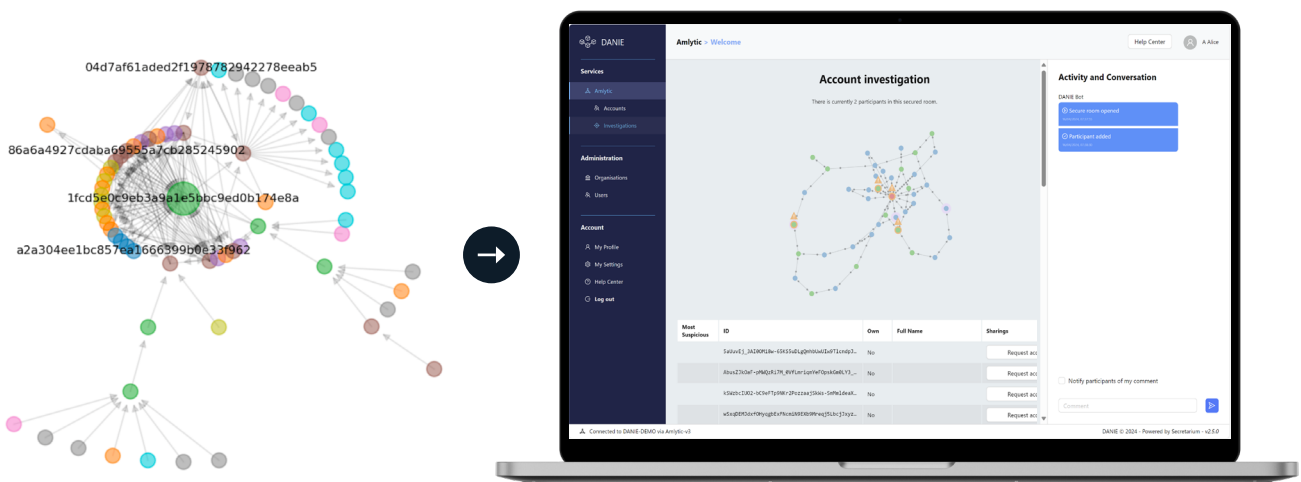
Executive Summary

Collaborative cross-bank Financial Crime Analytics brings a promise of higher detection rates with lower false-positives. However, deploying modern Machine Learning algorithms on bulk data from multiple banks runs into complexities of data sharing and exposes long-known limitations of Machine Learning.

After decades of a KYC-centric in-house AML, financial institutions incorrectly assume that cross-bank AML should necessarily involve sharing of personally identifiable information as well. This assumption raises legitimate data protection concerns, which even Privacy Enhancing Technologies (PETs) struggle to address.

Beyond the data sharing challenges, contemporary Machine Learning-based analytics models require accurate labelled data for training and testing, and produce results that cannot always be adequately explained. Given the historical tendency of financial institutions to 'over-report' suspicion, even the most advanced artificial intelligence trained on financial institutions' isolated or pooled data is bound to produce results that reflect the historically poor track record of spotting and reporting complex financial crime patterns.

In this White Paper, we present Amlytic – an alternative, context-centric approach to bulk-scale cross-institutional Financial Crime analytics at the pre-suspicion level. Amlytic emphasises data minimization on top of anonymization, limiting the scope of shared data sufficiently to leverage PETs in a privacy-respecting manner, in-line with the GDPR and the expectations of society. Designed to work at bank-scale and viable under the existing legal and data protection frameworks, Amlytic is capable of delivering robust and explainable results by using as little of the underlying shared data as possible, and with no requirement for training data.

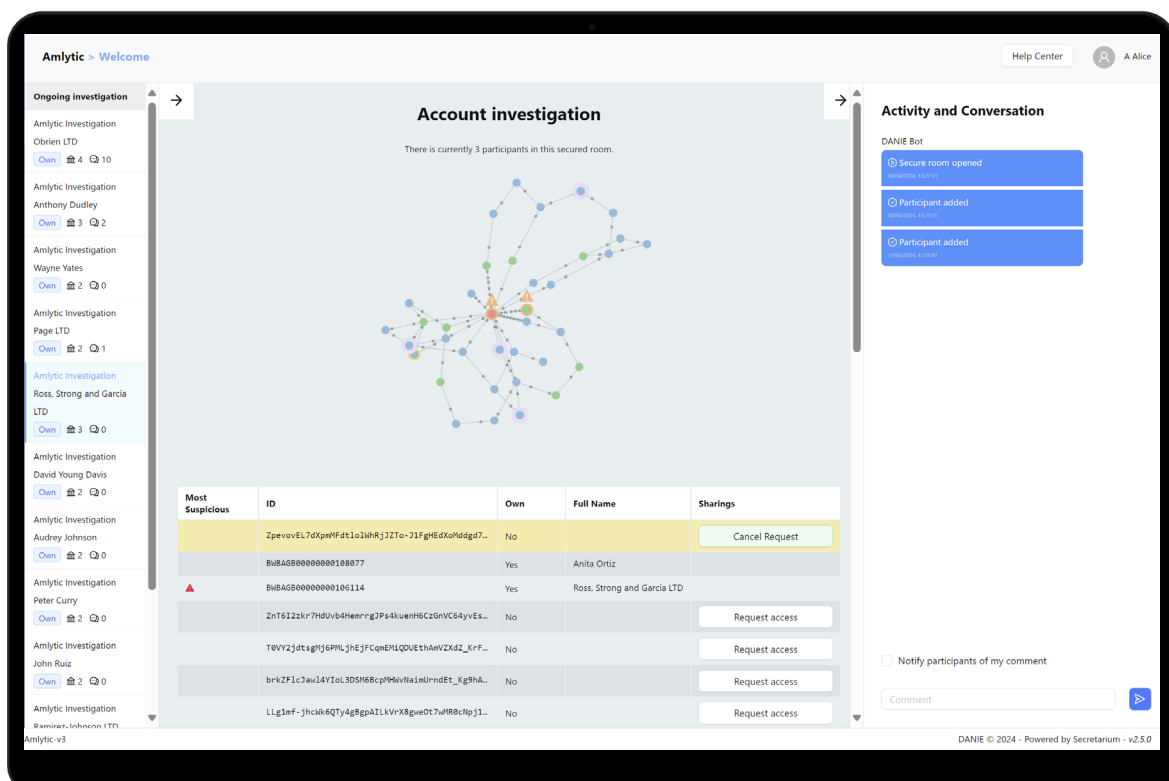


Amytic relies on Privacy-Enhancing Technology provided by Secretarium, as implemented since 2018 as part of the DANIE initiative, a group of banks, data vendors, and other institutions reconciling some of the most sensitive client reference data and securely sharing insights. Secretarium provides the privacy layer that securely manages the obfuscation and integration of each financial institution's minimised data into a linked and de-identified dataset and serves as an interface for post-analysis investigations and collaboration.

Amytic leverages the Transaction Analytics Technology provided by FutureFlow, as implemented for the UK Tri-Bank Initiative in 2019. FutureFlow runs on the linked and de-identified data provided by Secretarium to automatically spot pockets of suspicious activity, as well as to allow the participating financial institutions to explore their existing cases of suspicion in the broader context of the de-identified cross-bank data.

In 2022, Secretarium and FutureFlow showcased Amytic as a comprehensive privacy preserving and analytical platform for the ACPR / Banque de France Confidential Data Pooling TechSprint in Paris, where it was chosen as the winning solution, to be deployed by a group of participating financial institutions. In 2023, we redeployed Amytic for the G-20 TechSprint in Mumbai, where it was chosen as the winning solution for cross-border financial crime prevention.

This paper summarises the factors that set us apart from the competition in these benchmark global industry showcases and demonstrates how a successful country-scale, privacy-forward, and effective AML solution can be implemented today, with production-level technology and under the existing regulatory framework.



Contents

1.0. Industry Challenges	05
1.1. Regulatory Challenges	05
1.2. Technological Challenges	05
2.0. Introducing Amlytic	06
2.1. Broad architecture	06
3.0. Privacy-Preserving Technology	07
3.1. Secure and synchronised data de-identification	07
3.2. SMPP Engine Architecture - On Premise	07
3.2. Data minimisation	08
4.0. Analytics	09
4.1. Overview	09
4.2. Context and depth: proprietary network generation and analysis	09
4.3. Entity suspicion scoring	10
4.4. Group-level investigations	10
4.5. Further investigation: Supervised machine learning and data analytics	11
5.0. Secure Post-Analytic Collaboration	11
5.1. Amlytic dashboard	11
5.2. Secure rooms	12
5.3. Proactive and reactive analysis	13
6.0. Standalone and Multi-Institution Deployment	14
7.0. Case Study	15
7.1. Standalone deployment	15
7.2. Multi-Institution deployment	15
8.0. Privacy and GDPR Compliance	17
9.0. Continuous Learning and Development	17
10.1. Feedback loop	18

Industry Challenges

In this White Paper, we discuss how our joint solution tackles the key regulatory and technological pain-points faced by financial institutions in their fight against financial crime, particularly in jurisdictions subject to strict data protection regulations.

Regulatory challenges

At the Regulatory level, the key obstacle to successful cross-bank financial crime analytics is the historical over-reliance of financial institutions on KYC, CDD, and other types of sensitive personal information for conducting financial crime analysis and investigations. After decades of the KYC-centric in-house approach to AML, financial institutions assume, incorrectly, that the next-generation collaborative approach to AML should necessarily involve comprehensive sensitive data sharing across institutions as well. This assumption clashes with basic principles of data protection, even when considered in conjunction with Privacy Enhancing Technologies. A successful Cross-Bank Financial Crime Platform requires not only data anonymisation, but also meaningful data minimisation, with privacy considerations built into the platform from ground-up, and with analytics capable of delivering robust results by using as little of the underlying shared data as possible.

Technological challenges

At the Technology level, contemporary Machine Learning-based analytics models require accurate labelled data for successful training and testing, which may not always be available, given the weak track record of financial institutions in spotting financial crime and the historic tendency to 'over-report' suspicion for defensive purposes. Poor explainability of traditional ML models further complicates their practical utilisation even in cases where reliable labelled data may be available. In theory, unsupervised analytical models, such as Network Analysis overcome the limitations relating to explainability and lack of accurate labelled data. But in reality, the scale and complexity of the networks underpinning real-life cross-bank transactional environments poses a challenge to even the most robust technologies presently available.

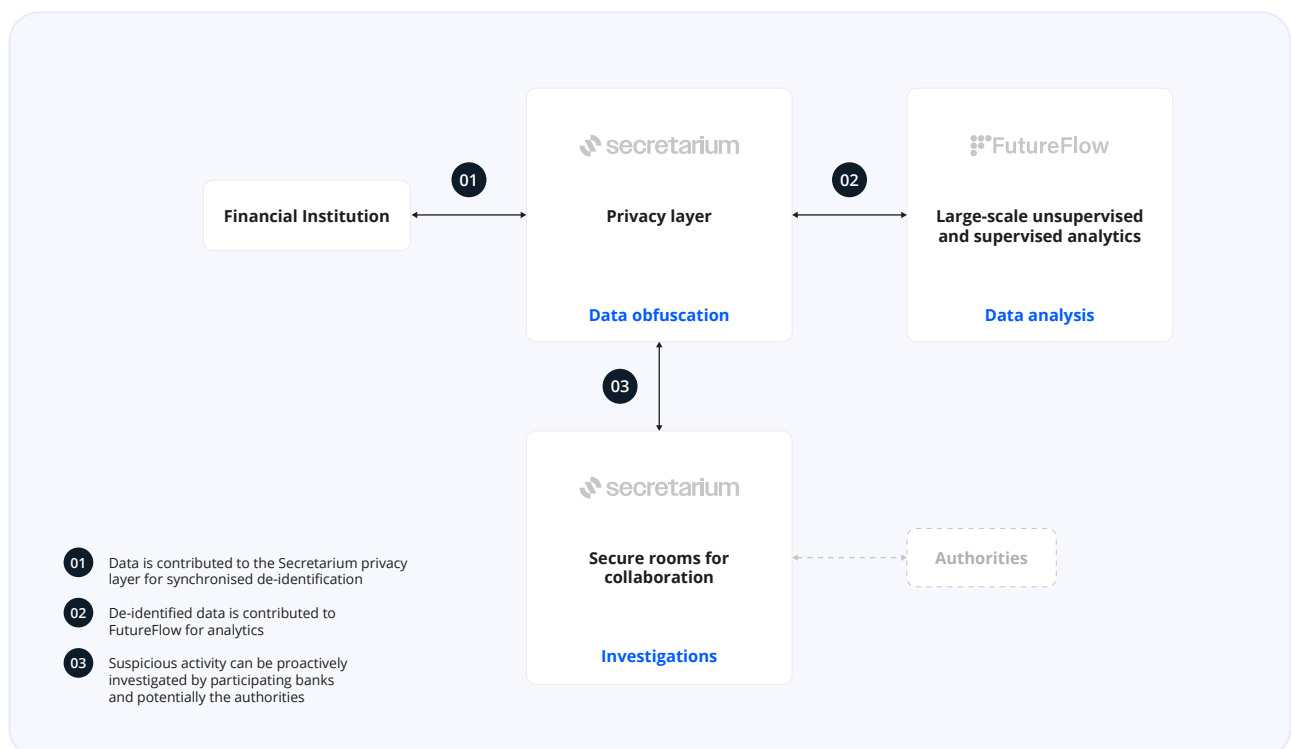
Introducing Amlytic

Our joint work on Amlytic offers an alternative, context-centric approach to Financial Crime Analytics that enables multiple financial institutions to jointly monitor and analyse transactions at scale, at a pre-suspicion level, in order to assist each other in spotting patterns of suspicious behaviour. The solution is viable under the existing legal and regulatory framework in GDPR-centric jurisdictions. The following key principles guiding our work enable the solution to overcome the obstacles and ambiguities discussed above:

- Broad context is more important than individual KYC/sensitive Personal Data on individual Data Subjects
- Unsupervised learning reduces the need for accurate labelled data
- Network size and density reduction greatly enhances analytical outcomes

The platform securely de-identifies, links, and analyses disjoint transactional datasets from multiple financial institutions, safely communicating relevant results to each correspondent institution and the relevant authorities, where applicable.

Broad architecture



The findings of the platform are made available for interactive exploration and analysis by the participating financial institutions via the Secretarium confidential Secure Rooms. These data rooms enable each participating financial institution to securely investigate its own accounts in plain view and, in cases of cross-bank suspicion, to give consent to investigate cross-bank account clusters jointly with other relevant financial institutions and authorities (post-suspicion stage).

Finally, the investigations' feedback is applied for further tuning of the underlying analytical models, enabling continuous improvement of the platform against the evolving and dynamic money laundering landscape.

Privacy-Preserving Technology

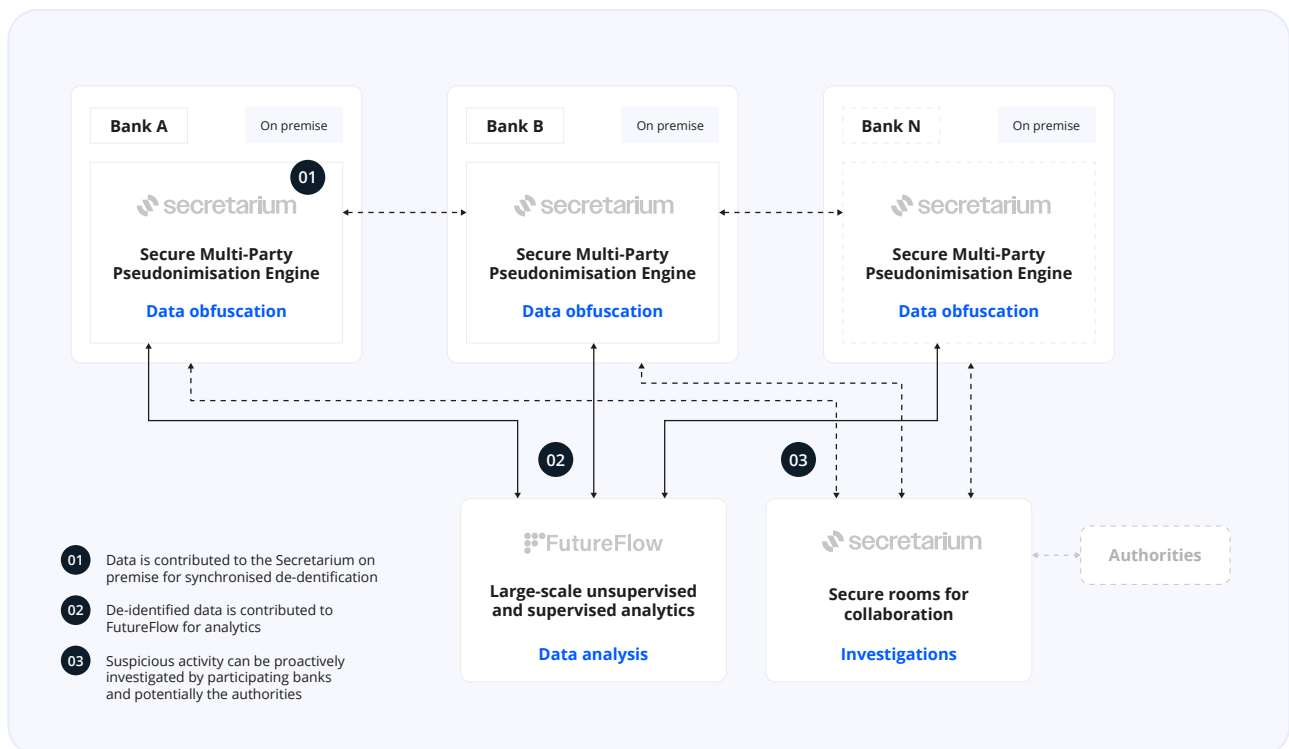
The platform joins transactional data from multiple financial institutions while keeping the data protected. It uses a combination of cryptography and secure hardware, and is designed to provably demonstrate its integrity. The privacy and integrity-driven component provided by Secretarium runs in attestable secure enclaves and ensures security of the data end-to-end (e.g. in transit, at rest and during processing).

Secure and synchronised data de-identification

Cross-institutional data pseudonymisation and synchronisation poses a formidable challenge. Transactions among multiple financial institutions must be pseudonymised deterministically so they can be linked, while preventing any party from being able to re-identify other parties' data. The currently known data obfuscation techniques used in the financial services industry are based on one-way hashing with a secret key, either known to the participating financial institutions or stored securely with a trusted third party¹. This poses non-negligible re-identification risks and operational challenges.

By contrast, Amlytic performs de-identification inside the Secretarium encrypted and tamper-proof layer. Our Secure Multi-Party Pseudonymisation Engine (SMPP Engine) performs de-identification across a set of interconnected enclaves, running on secure hardware. The hardware can be located on-premises with each participating institution or in the Cloud. The SMPP Engine is programmed to run only the designated code, verifiable by all participating institutions.

SMPP Engine Architecture - On Premise



¹ As described by FATF in [Partnering in the Fight Against Financial Crime: Data Protection, Technology and Private Sector Information Sharing](#), published Summer 2022 (see Box 4.1, page 23 and Box 4.3, page 30)

This architecture enables the synchronisation of hardware-generated and hardware-sealed secrets over multiple locations. The secrets are safeguarded, not known to any participant, not even the financial institutions. This approach eliminates the risk of accidental or deliberate re-identification in cases where one financial institution's data is leaked to other financial institutions; it also eliminates the need for a trusted third-party.

When transaction matching is based on a unique account identifier, such as an IBAN or a combination of sort code plus account number, the SMPP Engine can be deployed on premise, guaranteeing that the original data, even encrypted, never leaves each financial institution's data centre, only fully minimised and obfuscated data does.

The SMPP Engine can also be deployed in the Cloud as a managed service, or when more advanced matching options are required, as in fuzzy matching based on multiple data attributes without unique identifiers. For more details on Cloud deployment and advanced matching options, please refer to the Appendix.

Data minimisation

The data minimisation principle is respected at the highest possible level. The de-identified pooled transactions dataset only contains:

- **sender_account_id**: a salted hash of the sender account ID
- **receiver_account_id**: a salted hash of the receiver account ID
- **amount**: the original amount
- **timestamp**: the original timestamp
- **scope_flag**: indicator of an in-system or out-of-system transaction

Optionally, the solution could re-sequence transaction timestamps and smartly round amounts to further diminish re-identification risks; although significant distortions may interfere with subsequent data analysis. More columns could be added to the export, as long as this is considered carefully with data minimisation and privacy by design guidelines (more on this in the data protection section below).

Analytics

Overview

The de-identified transactional data is sent to FutureFlow for analysis. The analytical framework developed by FutureFlow helps the Financial Services industry to overcome three key challenges in large-scale ('bulk') transaction analysis:

- Regulatory barriers to sensitive data sharing across institutions at 'pre-suspicion' level
- Scarcity of reliable labelled data for popular ML-based analytical models
- Large volume, complexity, and density of pooled data

FutureFlow addresses these challenges by switching the attention from Personally Identifying Information (PII) towards the pure flow of funds analysis. By focusing exclusively on flow-of-funds rather than PII, FutureFlow enables a big-picture view of the underlying banking universe, which makes PII and other sensitive data less valuable from the analytical standpoint.

In this respect, FutureFlow relies on Privacy Enhancing Technology exclusively as an enabler of this big-picture view, rather than as a repository of vast amounts of sensitive personal data.

Under such architecture, sensitive PII remains within the walls of each participating financial institution. Only the de-identified data is sent down for subsequent analytical processing. This approach genuinely respects the Privacy-by-Design guidelines set forth as constraints of common Data Protection guidelines, such as the GDPR (more on this in the GDPR section below).

Focusing on the big-picture patterns enables FutureFlow to prioritise unsupervised analytical methods, primarily based on graph analysis. This helps to address the scarcity of accurate and reliable labelled data among the participating financial institutions and improves the level of explainability of the results, which are traceable down to the individual transaction level.

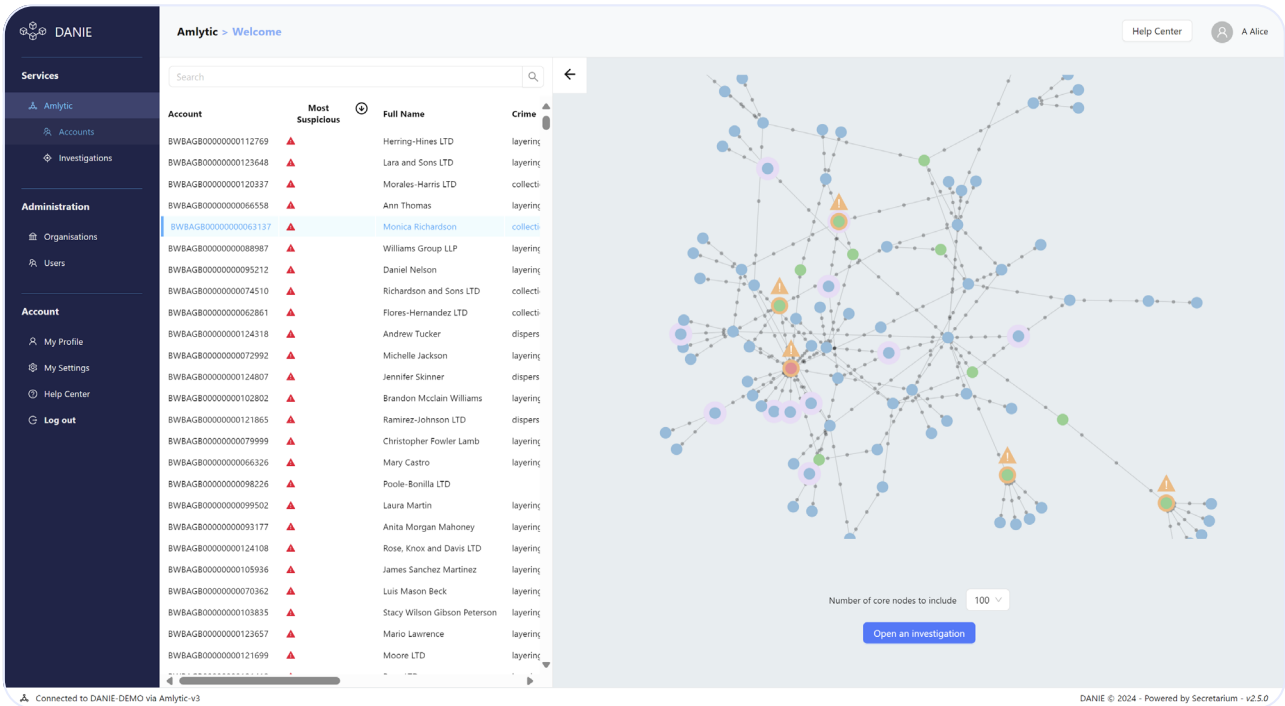
Context and depth: proprietary network generation and analysis

FutureFlow utilises proprietary algorithms to construct a precise picture of the flow of funds in the entire pooled dataset. Based on the knowledge of the flow of funds, FutureFlow constructs a network for each account in the dataset. Breaking down the entire financial network into account-level pieces enables the analytical comparison across the range of accounts and through time.

Our monetary flow algorithms ensure that each network includes only the most relevant connections for each account, greatly reducing the size and density of typical transactional networks. In prior deployments of FutureFlow on both synthetic and real-life data, we consistently witnessed complexity reduction in the order of thousands, allowing us to generate clear analytical insights at bank-scale with industry-standard Network Analysis tools.

Entity suspicion scoring

Each network in the pooled dataset is ranked according to a set of formal network features, which we have empirically identified as being useful in highlighting suspicious activity. In prior deployments on synthetic data with verifiable ground-truth, we achieved near 100% accuracy with negligible levels of false-positives in highlighting malicious accounts in the de-identified pooled data. Our results also helped to accurately invalidate false-positives in the supplied alerts, and to spot suspicious accounts among the banks that do not formally participate in the pooling activity (and on which no intelligence exists by design).



Group-level investigations

The suspicion score enables each participating financial institution to decide what share of the top findings deserve further scrutiny. To facilitate subsequent investigations, the platform automatically segregates its top findings into clusters, enabling group-level, rather than just individual investigations. The clustering approach has also proven to be effective in weeding out false-positives, particularly in the datasets where criminal activity is either deliberately or accidentally mixed with benign activity for some accounts.

While cluster and community detection exercises are common in any network analytics library, our experience once again suggests that this particularly is an area where the size and density of real-life financial networks create significant bottlenecks for their efficacy. FutureFlow executes this cluster detection via proprietary methods, which makes the process viable at bank-scale.

Further investigation: Supervised Machine Learning and data analytics

The scoring exercise produces a proprietary account-level dataset, which is based on a range of formal network features of each account in the underlying pooled dataset.

In prior deployments on synthetic data where ground-truth of the real malicious activity was known, we successfully applied supervised ML models on this dataset to the accuracy levels above 95%.

While FutureFlow does not rely on supervised learning in the core of its processing, these findings could further be helpful downstream to the participating financial institutions for in-house data analytics and ML projects.

Secure Post-Analysis Collaboration

Amlytic Dashboard

The results of the FutureFlow analysis are made available to the participating financial institutions for automated querying or for ad-hoc exploration via the Amlytic Dashboard. The Dashboard is an intuitive web-based interface, which can be enabled on premise or in the Cloud. It effortlessly automates the complexity of the underlying cryptography and secure hardware attestations.

The screenshot shows the Amlytic Dashboard interface. On the left is a dark sidebar with navigation options: Services (Amlytic, Accounts, Investigations), Administration (Organisations, Users), and Account (My Profile, My Settings, Help Center, Log out). The main content area is titled "Account investigation" and includes a network graph with nodes and edges. Below the graph is a table of "Most Suspicious" accounts. On the right, there is an "Activity and Conversation" panel showing a chat log with messages like "Secure room opened" and "Participant added".

Most Suspicious	ID	Own	Full Name	Sharings
	5aUuvEj_3aIT0M18w-65K5SuDlGqhbUwIw9T1cndp7...	No		<button>Request access</button>
	AbusZ3k0aF-pMwQzR17M_0vFLmr1qnyeFopskGm6LY3_...	No		<button>Request access</button>
	kSwzbcIU02-bC9eFTp9Nkr2Pozzaaj5KwS-5mPm1deaX...	No		<button>Request access</button>
	wSxqDEHJdxFOHyagbExFlcm3i9EXb9MneqJ5Lbcj3xyz...	No		<button>Request access</button>

The Amlytic Dashboard serves as an interactive confidential switch-board, connecting multiple institutions into a secure conversation on areas of mutual suspicion in the underlying cross-bank data universe.

In the present space of available Privacy Enhancing Technologies, Amlytic Dashboard represents a unique two-way data obfuscation and de-obfuscation platform through which multiple financial institutions can interactively generate, investigate, and confirm/reject suspicious activity.

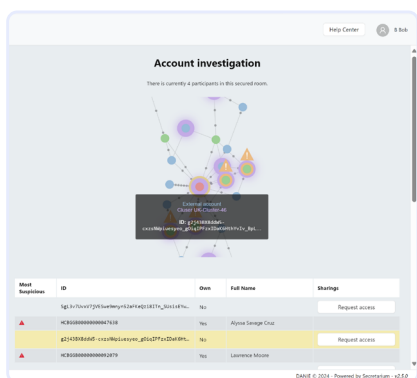
The platform vastly enriches each institution's existing in-house intelligence and makes it beneficial for the entire participating ecosystem, while keeping it private. The data pooling process can run in batch mode as well as in real-time. In the case of a financial institution withdrawing from the reconciliation, all their data is removed and the pooled dataset is re-processed.

The SMPP Engine guarantees that the intelligence returned to each financial institution contains its original data, enriched with de-identified insights from other financial institutions generated by FutureFlow. The Secretarium privacy layer ensures confidentiality and data ownership are respected throughout. The solution does not allow the financial institutions to access the de-identified pooled data as a whole. Still, original data from other financial institutions can be requested in the context of an investigation in a secure room.

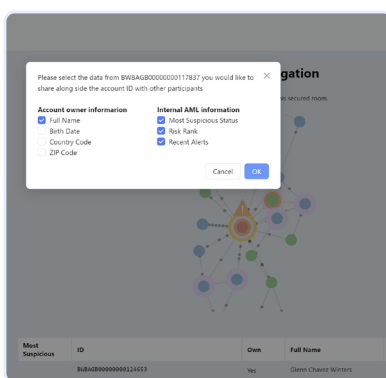
Secure rooms

In cases where collective suspicion is found, more comprehensive data sharing can begin in secure rooms, which are end-to-end encrypted contexts designed for multilateral data collaboration.

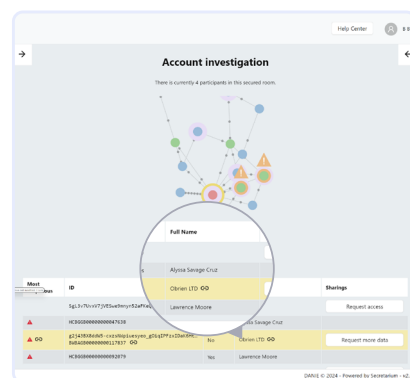
Any financial institution involved in a suspicious flow of funds can open a secure room. Secure rooms are fully anonymous, and will automatically invite peers when data is requested, while protecting all parties' identities. Bilateral data disclosure is consent-based, and parties can choose which attributes they are happy to anonymously disclose.



Bank B requests access



Bank A consents to collaborate



Bank B now sees information

Optionally, a secure room can also be used to invite a local authority to join an investigation and access the financial institutions' original data.

Proactive and reactive analysis

Amlytic offers the banking ecosystem two synergistic approaches to AML investigations with the benefit of the cross-bank view:

Reactive Analysis	Proactive Analysis
<p>Analytics based on a given entry point</p>	<p>Lead generation: no leading intelligence</p>
<p>On-demand exploratory analysis: this is available as a mixture of visual analysis (graphs, clusters, relationship traces, etc.) or numerical analysis (various numerical properties of the network of each account or account group, as well as their relative positioning/rank in the overall universe of accounts, etc.)</p> <p>Resonance building: multiple financial institutions flagging their respective accounts as potentially suspicious, with the platform confidentially cross-referencing the requests from each participating financial institution to indicate areas of resonance and overlap.</p>	<p>The solution proactively generates a set of accounts and account groups as candidates for further investigation and analysis by the participating financial institutions. The solution presents visual and quantitative explanations as to why these accounts were flagged and offers ways to explore them in a broader context.</p> <p>Feature selection and engineering: proprietary account-level dataset for downstream AI/ML model building, exploration, and tuning.</p>

Standalone and Multi-Institution Deployment

Amlytic can be deployed from the start as a multi-bank AML utility platform. Alternatively it can evolve from a standalone in-house deployment by one or more institutions, eventually leading to a more integrated dual or multi-institution setup.

Findings from a range of past deployments on synthetic and real-life data strongly suggest that the industry should expect to benefit from Amlytic even starting with a single-institution deployment, with results improving drastically as the ecosystem grows.

Deployment Type	
Standalone	Multi-Institution
On-premises deployment of FutureFlow at each participating institution	Two or more financial institutions sharing data among themselves using the secure Amlytic platform
SMPP Engine used only for confidential alert sharing among the participating institutions (no transaction data sharing)	SMPP Engine used for transaction de-identification and pooling, and for secure two-way post-analysis communication

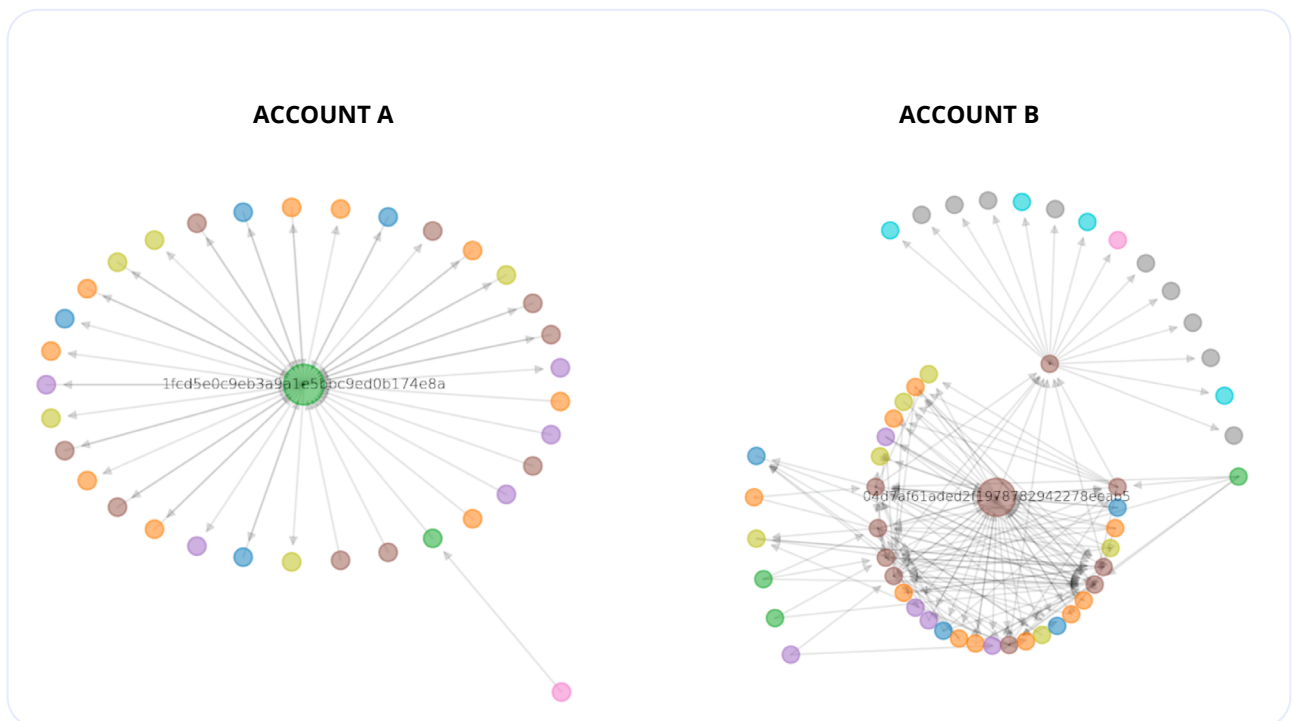
Expected Results	
Standalone	Multi-Institution
Meaningful insights generated on a subset of accounts with significant “on-us” transacting activity.	Superior insight into cross-bank transacting patterns, account relationships, and clusters, even when not all banks participate in the platform.
Better transacting context for accounts associated with alerts and other data shared via the SMPP Engine.	Clear view of each institution’s customers in the cross-bank context. The ability to build confidential ‘resonance’ with the AML intelligence of all other institutions
	Clear and obvious resolution of networks involving accounts from the non-participating third-party banks
	Significant pressure generated on the ‘holdout’ financial institutions to join the established cross-bank utility

Case Study

Below we offer an example of how the analysis of two accounts from two banks evolves across the two proposed deployment models.

Standalone Deployment

The bank deploying Amlytic in a Standalone mode is not generating any meaningful insight on its account A (pictured below): the result is a simple network with mostly one-degree connections and no meaningful relationships and patterns. This is primarily due to the fact that account A transacts almost exclusively with other banks' accounts. In comparison, the peer bank deploying Amlytic in Standalone mode gets a lot more insight on its account B, which exhibits a stronger tendency for "on-us transactions". The resultant network is larger and qualitatively more complex, containing a range of significant inner-relationships and patterns with multiple degrees of separation.




Our experience suggests that even a Standalone deployment will produce a range of robust networks similar to the Account B example, particularly if the anchor hosting bank happens to be a large institution with a significant market share in the underlying banking system. Below we demonstrate how these results will improve when the two institutions join into a Multi-Institution deployment


Multi-Institution Deployment

The two banks start to benefit significantly more from Amlytic when they join into a Multi-Institution deployment, and even more so when other banks join the platform. While the hosting bank got virtually no insight on account A in its Standalone deployment, the same account's network immediately becomes more revealing and informative in the two-bank Multi-Institution deployment (pictured below, left), culminating in a 1,265-edge network in the six-bank Multi-Institution deployment (pictured below, right).

ACCOUNT A: MULTI-INSTITUTION 2 BANKS



ACCOUNT A: MULTI-INSTITUTION 6 BANKS



Furthermore, the connection between the same two accounts A and B, initially invisible in each bank's Standalone deployment, becomes apparent in both the two-bank and six-bank Multi-Institution deployments, highlighted by the underlying obfuscated account IDs on the connecting route in both networks in the pictures above.

The two accounts chosen for this example were both strongly highlighted as malicious and related to each other by our unsupervised algorithms as early as in the Dual-Institution deployment. Despite their relatively small size, both were singled out from over a million of accounts in the underlying dataset, with many of their peer malicious accounts from the same crime syndicate present in the two-bank network. These findings were confirmed in the Multi-Institution deployment, where even further malicious accounts from the same crime syndicate were pulled into the network view, including some accounts from the non-participating banks. All these unsupervised findings were subsequently checked and confirmed against the ground-truth data supplied with the dataset.

Starting from a Standalone deployment, the hosting bank benefits from getting a greater transacting context for its account B and its equivalents. The results improve drastically for each bank when two more more banks join into the Multi-Institution deployment. Our work for the ACPR / Banque de France TechSprint demonstrated that two banks running the Multi-Institution deployment in a multi-bank environment managed to improve their average account network properties by anywhere from 3x to 8x, which speaks drastically in favour of the joint transaction analytics approach proposed in this White Paper.

Privacy and GDPR Compliance

Amlytic is designed to be compatible with existing principles of GDPR. With its ability to generate actionable insights based on a limited schema of obfuscated transactional data, Amlytic simplifies various ambiguities commonly associated with the status of the underlying data, the nature of the applied obfuscation techniques, and the choice of the appropriate Basis for Processing.

Amlytic design and architecture follow the Data Protection principles developed by FutureFlow under the guidance of the UK Information Commissioner's Office (ICO), as part of FutureFlow's participation in the ICO GDPR Regulatory Sandbox in 2019-2020. The ICO's views on the Sandbox work were subsequently published in the Sandbox Final Report².

The Sandbox work centred primarily on defining the status of obfuscated transactional data and choosing the appropriate Basis for Processing that can enable a country-scale cross-bank transaction monitoring system from the data protection point of view. The Sandbox findings are based on a more simplified obfuscation technique involving one-way hashing with a secret key, common for the industry at the time. The subsequent addition of data obfuscation techniques based on Confidential Computing discussed in this White Paper increases further the privacy guarantees appropriate for a successful implementation of a country-level, privacy-centric AML utility platform.

Amlytic follows the conceptual two-tier Data Protection architecture developed in the Sandbox, with the platform acting as Data Processor and the participating financial institutions acting as Data Controllers. Data Controllers must use an appropriate Basis for Processing in order to process sensitive customer data with external parties, unless the data is not related to individuals or is considered fully anonymized.

For stress-testing purposes, we assume that at least some transactional data might relate to individuals, and that it cannot be anonymized sufficiently to be considered non-personal for GDPR purposes, even with the use of Privacy Enhancing Technologies. Instead, to ease the choice of the appropriate Basis for Processing, we emphasise Data Minimisation and adherence to Privacy by Design concepts throughout our technology stack, as described in the Privacy Preserving Technology and the Analytics sections above.

Under our architecture, the basic obfuscated transactional data is analysed in bulk and in obfuscated form at a pre-suspicion level. When specific pockets of suspicion are found and validated, more focused sharing of sensitive data may take place through the Dashboard, with the consent of the corresponding financial institutions.

The bulk processing at the pre-suspicion stage takes place based on Legitimate Interests as Basis for Processing, whereas the more focused sharing and processing post-suspicion is based on Compliance with Law as Basis for Processing. Our strict attention to Data Minimisation and Privacy-by-Design guidelines is meant to help the participating financial institutions to meet the requirements of the Legitimate Interests Assessment Test necessary for this architecture.

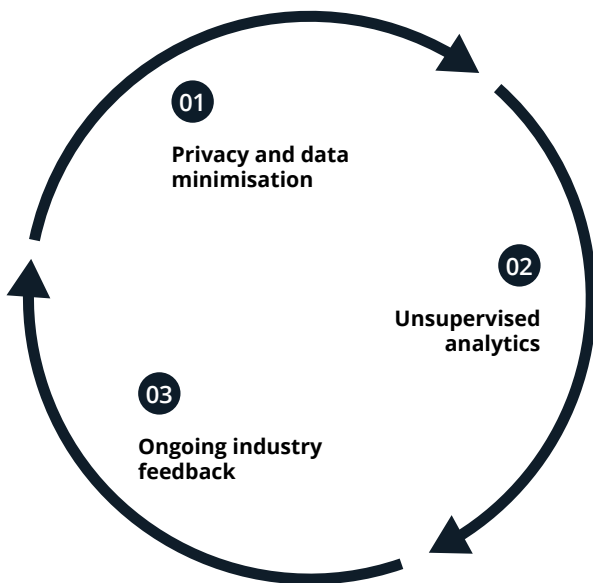
² [Regulatory Sandbox Final Report: FutureFlow](#), published October 2020

Continuous Learning and Development

Amlytic enables the banking ecosystem to work continuously on improving its collective Anti-Financial Crime capabilities through a series of symbiotic steps.

Feedback loop

- 01 The SMPP Engine provided by Secretarium offers a platform for joining datasets from multiple banks without the need for sensitive information disclosure. It is an interface connecting multiple participants in the banking ecosystem for collective exploration and querying of the analytical results.



- 02 Unsupervised analytics provided by FutureFlow enables the banking ecosystem to achieve actionable results with the use of the existing in-house intelligence, as well as via the platform's automatic lead generation capabilities. The data generated by FutureFlow opens further prospects for custom in-house and third-party model training and development.

- 03 The banking ecosystem provides ongoing feedback and validation of the analytical results, enabling further development of new analytical models and improvement of the existing models.

For more information, please contact:

Vadim Sobolevski - CEO at FutureFlow, vadim@futureflow.org

Bertrand Foing - CEO at Secretarium, bertrand@secretarium.org

About Secretarium

Secretarium is a leading data security technology provider. Secretarium is dedicated to transforming data processing through transparent, verifiable, privacy-preserving computing. Secretarium's groundbreaking technology empowers data holders to collaborate and create powerful, meaningful intelligence to solve the globe's most complex challenges.

Secretarium's strategy is firmly rooted in privacy by design principles, harnessing confidential computing and other privacy-enhancing technologies. Secretarium's verifiable transparency assures security and integrity. It greatly facilitates the demonstration to auditors and regulators of the enforcement of security measures and compliance with data protection requirements.

About FutureFlow

FutureFlow brings together multiple banks, regulators, and Financial Intelligence Units into a joint fight against Electronic Financial Crime. Our platform automatically spots anomalous relationship patterns and linkages in pooled transactional data spanning multiple financial institutions.