DARTE SERIES

Berlin 6.0

Initiated by Mariana de la Roche W. (BlackVogel) and Dr. Nina-Luisa Siedler (siedler legal), the DARTE Series is a high-level roundtable format designed to enhance legal clarity around digital assets, focusing on regulation, compliance, data protection, and market integrity.

The Berlin 6.0 DARTE edition took place on June 15th, 2026, in collaboration with the European Commission, Bitvavo, Bundesblock and N26. The roundtable gathered regulators, legal practitioners, and digital asset infrastructure providers to examine the current revision of the FATF Recommendations and their implications for the crypto-asset industry.

The agenda focused on three themes: 1) FATF Recommendations 10, 15, and 20, presented by Nadine Kari (N26); 2) repurposing FATF Recommendation 16, presented by Zornitsa Daskalova (Zodia); and 3) FATF Recommendations 33, 41, and 42, presented by Dr. Agata Ferreira.

A theme that ran through all three discussions was the distinction between measuring how effectively institutions comply with the standards and asking whether the standards themselves achieve what they were built to do. While DARTE sessions typically focus on identifying practical solutions for operating within the current regulatory framework, without directly requesting guidance or legislative change, this edition took a different approach. The discussion was a critical analysis of the FATF Recommendations themselves, examining their effectiveness and proposing concrete adjustments to the framework rather than working around its existing contours.

We extend our sincere thanks to all speakers, participants, and institutional partners for their valuable contributions. The views presented in this report reflect the collective understanding of the participants and do not necessarily represent the official positions of individual attendees or rapporteur

1. FATF Recommendations 10, 15, and 20: Customer Due Diligence, New Technologies, and Suspicious Transaction Reporting

The first topic was presented by Nadine Kari (N26) and addressed three FATF Recommendations that together define the operational core of the AML compliance regime for financial institutions: Recommendation 10 on customer due diligence, Recommendation 15 on new technologies, and Recommendation 20 on suspicious transaction reporting. The framing throughout was that these Recommendations, designed in their current form before the technological and operational realities of today, are now constraining rather than enabling the fight against financial crime, and that meaningful improvement requires both public and private sector actors to coordinate in ways the current framework does not explicitly support nor suggest.

On Recommendation 10, the operational reality of customer due diligence is that the underlying data is siloed across KYC systems, core banking, CRM, transaction systems, and external data providers. This fragmentation weakens machine learning and AI models, creates blind spots in customer risk profiles, and contributes to incorrect identification of suspicious behavior and unfair customer treatment. Given that the banking landscape is more and more online based, these siloes can create a difficult gap in KYC Due Diligence compliance. Supervisors expect clear, risk-based rationales for CDD decisions, including for the application of enhanced due diligence, but many of the ML and AI models in use behave as black boxes whose outputs are difficult to justify in the language regulators require. Recommendation 10 calls for a risk-based approach and ongoing CDD with dynamic risk ratings, but many institutions continue to run static periodic reviews rather than integrating models into life-cycle monitoring, with the consequence that alerts do not automatically update customer risk ratings. A further structural concern is that AI-triggered re-KYC or other alerts may unintentionally signal suspicions to the customer, undermining the regime's tipping-off protections or consumer protection if not data privacy related issues.

On Recommendation 15, financial institutions are deploying machine learning and AI in fraud and AML operations without the formal risk assessment, testing, and inventory that the Recommendation's requirement to assess new technologies implies. Poor training data because of lack of cooperation with peers can cause models to over-flag particular customer segments or geographies, raising conduct and compliance risk under the proportionate risk-based approach the framework otherwise requires. The implementation of the Travel Rule and cross-border data sharing for virtual assets is itself flawed: AI models may operate on incomplete originator or beneficiary information, and divergent local implementations of the Recommendations mean that models trained for one market may not be compliant elsewhere, particularly where expectations on explainability and consent differ.

On Recommendation 20, traditional rules engines generate large volumes of

low-value alerts, and naive AI layers on top can simply reprioritize rather than meaningfully reduce useless cases, leaving regulatory backlogs and delaying suspicious transaction reports. Financial intelligence units provide little to no structured feedback on the usefulness of STRs/SARs filed, which makes it difficult to train models to prioritize high-value cases. Single-institution data also misses patterns that span multiple institutions, limiting the ability of machine learning and AI to capture complex networks or trade-based money laundering schemes.

The solutions presented were focused on data, technology, and structured cooperation. On Recommendation 10, the proposals included building a central customer risk store that merges KYC files, transactional behavior, sanctions and PEP screening, adverse media, and beneficial-ownership registries, so that ML operates on consistent, high-quality features; deploying ML to update customer risk scores in near-real time based on behavior, with automated triggers for EDD, periodic review, or exit; routing high-impact AI decisions such as onboarding denial or off-boarding to experienced compliance officers with structured review checklists and full audit logging to reduce tipping-off risk; and participating in KYC utilities or public-private partnership platforms that share basic CDD data and outcomes, so that models can benefit from broader patterns while respecting privacy and secrecy laws. CDD measures should remain anchored to individual risk assessment of each Financial Institution, and EDD or exit decisions should not be taken on the basis of anything other than individual risk indicators partially depending on the individual institutions' Risk Appetite.

On Recommendation 15, each ML and AI use case in fraud and AML should be treated as a "new technology" in the sense of the Recommendation and not limited to the current version, subject to documented risk assessment, approval, and periodic review. Institutions should maintain a model inventory with defined owners, validation standards, challenge functions, and periodic back-testing, ideally in exchange with peers.

Techniques such as federated learning and secure multiparty computation should be used to allow institutions to jointly train typology models on distributed data without sharing raw personal data, supporting cross-institution detection while respecting secrecy laws. Blockchain analytics should be integrated into transaction monitoring to classify wallet risk, trace funds, and identify mixers or sanctioned addresses, with shared information spaces for high-risk and flagged counterparties. AI-assisted Travel Rule engines should auto-validate counterparty VASP information, flag risky or incomplete originator and beneficiary data, and learn from historical false positives, with shared information among financial institutions and regulators in a secure environment.

On Recommendation 20, the proposals combined scenario-based rules with supervised and unsupervised ML, including anomaly detection, clustering, and network analytics, to focus on truly unusual behavior and materially reduce false positives. ML should be used to score

alerts by likely STR-worthiness and potential ML/TF impact, so analysts focus first on the highest-risk cases. FIUs should be required to share outcome information and provide feedback so that institutions can use it to refine detection models. Cross-institution models of customers, accounts, and counterparties should be built using an encrypted solution to detect complex layering and syndicates while preventing the over-sharing of customer information and reducing the volume of low-quality STRs/SARs.

The discussion that followed was extensive and surfaced significant disagreement about whether the proposed direction is the right one. There was broad recognition that the current regime has failed by its own standards and that improvement is necessary. Anti-money laundering is structurally a post-crime regime, and the resource imbalance between compliance and enforcement is stark: roughly $86 billion is spent globally on compliance against single-digit billions for financial crime law enforcement. The cost-benefit calculus of expanded data collection was repeatedly questioned. Participants observed that AML officers currently collect substantial volumes of information that does not meaningfully support enforcement outcomes, and that further expanding the data perimeter risks creating new privacy exposures and honeypots without proportionate compliance benefit.

Discussions about prevention versus conviction demonstrated that it currently is more a tendency towards conviction and enforcement efforts than a prevention of financial crime and further collaboration could rather lead the development towards prevention.

A central thread was the question of whether the right response to the current dysfunction is more (encrypted) data sharing or different data sharing. Several lines of argument converged on the view that the unit of exchange should not be raw customer data but suspicion or red flags, with proper safeguards against propagating false positives. The discussion led to several examples of legitimate political activists, opposition figures, and individuals with particular national origins being repeatedly de-banked because a flag raised by one institution had effectively propagated through the system. Participants raised concerns that any expansion of suspicion-sharing risks could produce a totalitarian outcome where individuals persecuted by authoritarian governments have nowhere to bank, because the FATF regime is global and includes jurisdictions whose definition of "suspicious" cannot be assumed to be benign. Several participants emphasized that the current framework already disadvantages individuals on the basis of nationality and that adding new propagation mechanisms without robust legal protections would magnify this harm. Participants stressed that any suspicion-sharing architecture must preserve the rule of law, including the data subject's ability to be informed of the flag, to challenge it, and to have it removed across the network where successfully challenged.

A significant strand of the discussion questioned whether the underlying problem is one of data sharing at all. The case was made that the technological

architecture should move away from sharing data and toward sharing proofs of properties of data, using zero-knowledge proofs, verifiable credentials, and federated identity infrastructure of the kind that has been operational in Estonia since 2001. On this view, a financial institution does not need to know everything about a customer; it needs to be able to verify that certain properties hold (that the customer has been KYC'd at another institution, that they are not on a sanctions list, that they have not accumulated suspicion across the network) without the underlying data being transmitted or stored centrally. The technical maturity of these approaches was contrasted with the political and commercial obstacles to adoption, including the unwillingness of banks to rely on KYC performed by other banks due to liability concerns. The European Digital Identity Wallet was identified as relevant infrastructure, although significant questions remain about timelines and implementation.

The discussion also surfaced a foundational concern about the conflation of distinct regulatory objectives within a single regime. Anti-money laundering is typically concerned with large amounts and follows the proceeds of crime, while counter-terrorism financing typically concerns small amounts and operates under different transmission logic.

Cash remains the dominant transmission mechanism for illicit finance, dwarfing crypto by a significant margin, yet much of the regulatory expansion in recent years has been directed at digital assets rather than at the actual centers of risk. This followed up on the previous discussion on prevention and conviction that operate under different logics. The point was made that banks are and should remain providers of financial services rather than law enforcement agencies, and that asking banks and crypto exchanges to act as agents of the state in this way may be a category error from which much of the current dysfunction flows. The five-year data retention period in Germany was raised as a specific example of disproportionate retention given the volatility of the threat landscape and the risk of accumulated honeypots.

A practical example was offered from the crypto industry: the Beacon Network which allows law enforcement to share wallet-level risk signals across institutions in close to real time, taking advantage of the public nature of on-chain addresses to operate outside the GDPR perimeter that constrains analogous arrangements in traditional finance. The example was cited as a model of what intelligence sharing focused on observable activity rather than personal identity can look like in practice. At the same time, participants cautioned against assuming that propagating wallet-level flags is risk-free: as in the bank-account context, the absence of meaningful challenge mechanisms could produce permanent exclusion based on initial errors.

A further recurring theme was the asymmetry between the obligations placed on financial institutions and the responsiveness of the law enforcement system the obligations are meant to support. STRs/SARs are filed under tight statutory deadlines, but the subsequent enforcement action can take months or years, by which point the information is

often stale. Where institutions are required to report fraud-related events, the absence of timely follow-up means that the reporting itself becomes the regulatory output rather than a step toward an enforcement outcome. Joint FIU analysis across Member States was raised as a possible direction, although participants expressed scepticism that it currently functions in practice.

A final point of consensus, even among those who disagreed sharply on the substance, was that the appropriate response to the failure of the current framework is not to add complexity but to step back and reconsider its foundations.

The FATF Recommendations in their current form date in significant part from 2001 and 2003, designed in a technological environment that no longer exists and not in its entirety allow to adapt its content into the new virtual, remote environment where financial institutions, payment providers and banks try to fight financial crime that in return adapts much more to today's technical possibilities needs an update and response in order to efficiently allow prevention of financial crime for the industry and convictions for the enforcement bodies.

The addition of further frameworks on top of a structure that does not work risks compounding the problem rather than solving it.

Calls to Action regarding FATF Recommendations 10, 15, and 20

The key calls to action from the discussion are:

- **Move from data sharing to proof sharing, anchored in verifiable credentials and zero-knowledge infrastructure:** The architecture of CDD and inter-institutional exchange should shift away from the transfer of raw customer data and toward the sharing of cryptographic proofs of properties of that data, supported by verifiable credentials and the European Digital Identity Wallet. This approach allows institutions to verify what they need to verify, including KYC completion, sanctions clearance, and accumulated network-level suspicion, without expanding the attack surface or creating new honeypots, and gives data subjects meaningful control over their own information.

- **Reframe inter-institutional exchange as suspicion sharing rather than customer data sharing, with rule-of-law safeguards:** Where exchange between institutions is appropriate, it should be limited to red flags and suspicion signals, not customer profile data. Any such mechanism must include the right of the data subject to be informed of a flag, to challenge it, and to have it removed across the network

where the challenge is upheld, to prevent permanent exclusion based on initial errors or the weaponization of the framework by authoritarian jurisdictions.

- **Require structured, timely feedback loops from FIUs and harmonize Member State implementation:** Recommendation 20 should be implemented in a way that requires FIUs to provide structured outcome information on the STRs they receive, enabling institutions to refine detection models and reduce the volume of low-value reports. Member State implementations of the Recommendations should be harmonized to remove the current divergence in CDD, AML, and reporting obligations that allows arbitrage within the single market and produces inconsistent enforcement.

- **Treat each AI and ML use case in AML as a new technology under Recommendation 15:** Each AI and ML use case in fraud and AML compliance should be subject to documented risk assessment, model inventory, validation, and periodic back-testing. Federated learning and secure multiparty computation should be supported as appropriate means of cross-institutional typology development that respects secrecy laws, and supervisory expectations on explainability should be harmonized across jurisdictions to allow models to operate consistently across markets.

- **Step back from regulatory expansion and reconsider the foundations of the AML/CFT framework:** The structural problems identified, including disproportionate compliance cost relative to enforcement outcomes, the conflation of AML and CTF logics, and the displacement of crime prevention by post-crime reporting, cannot be addressed by adding further layers to the existing framework. The current revision of the FATF Recommendations is an opportunity for a more fundamental reassessment, including the appropriate role of financial institutions, the proportionality of data retention, and the resource balance between compliance and enforcement.

2. Repurposing Recommendation 16: From Bulk Data Transmission to Intelligence Sharing

The second topic was presented by Zornitsa Daskalova (Zodia) and proposed a fundamental rethinking of FATF Recommendation 16, the Travel Rule, arguing that the obligation should be repurposed from the bulk transmission of personal identity data to participation in a qualifying, vendor-neutral intelligence-sharing network in which financial institutions, virtual asset service providers, financial intelligence units, and law enforcement agencies are full participants.

The framing of the problem rested on a straightforward cost-benefit analysis. Global financial-crime compliance reached approximately USD 206 billion in 2023, yet

Europol assesses that only around 1% of criminal proceeds in the EU are confiscated, and law enforcement accessed less than 3% of the currency transaction reports filed between 2014 and 2023.

The regime prioritises procedural output over operational outcome: it consumes nine-figure annual sums while leaving the overwhelming majority of criminal proceeds untouched. Layered onto this is a specific privacy and fraud cost. Recommendation 16 requires the bulk transmission of originator and beneficiary personal data, including names, account numbers, addresses, and national identification numbers, retained for a minimum of five years, for every qualifying transfer, irrespective of any indicator of suspicion, and often to institutions in jurisdictions that offer no equivalent data protection.

This concentrates sensitive identity data into repositories that themselves become honeypots, turning information collected to deter crime into a standing target for breach and identity-enabled fraud, a privacy cost without a proportionate security benefit and one that offers no purchase at all on unhosted wallets or decentralized transfers. The solution presented was that the obligation be repurposed from transmitting identity data into participating in a qualifying intelligence-sharing network in which financial institutions and VASPs exchange risk information only where defined red-flag thresholds are met, and to which FIUs and law enforcement agencies are full participants.

The model is proven and rights-respecting. Singapore's COSMIC platform permits institutions to share customer information precisely where multiple red-flag indicators cross stipulated thresholds, on an explicit statutory footing. The European Union's Horizon 2020 TRACE project has already co-developed an open-source, ethics-by-design framework whose stated objectives include increasing the efficiency and effectiveness of information sharing among law enforcement agencies, confirming both the policy appetite and the technical foundations for cross-border, rights-respecting sharing.

FATF was identified as uniquely positioned to supply the legal standard such networks require: criteria for a qualifying network, including vendor neutrality, interoperability, participant vetting, and governance, and the safeguards on which lawful sharing depends, including a defined suspicion or higher-risk threshold, pseudonymisation, a mandatory data-protection impact assessment, joint-controller accountability, independent audit, and a redress mechanism. This architecture is neither speculative nor unprecedented: it mirrors Article 75 of the EU's Anti-Money Laundering Regulation, which permits partnership sharing where strictly necessary and subject to precisely those safeguards, and rests on the established basis that processing strictly necessary to prevent fraud is a legitimate interest of the controller under GDPR.

The preventative core of the framework would be preserved: targeted financial sanctions derive from Recommendations 6 and 7, not Recommendation 16, and would remain fully in force.

The discussion that followed engaged closely with both the diagnosis and the proposal, and surfaced sharp disagreement on the architecture but broad agreement that the current model is not delivering. A foundational concern raised was the global nature of the FATF regime: of the nearly 200 jurisdictions represented in the United Nations, few are robustly compliant with the rule of law in the way an intelligence-sharing architecture presupposes.

A red flag raised by an institution in one jurisdiction may mean something very different from a red flag raised by an institution in another, and any sharing architecture must therefore distinguish carefully between participants and between the categories of information shared, to avoid creating a framework that authoritarian governments can weaponize against legitimate dissent or against individuals on the basis of nationality or political activity. A second major thread concerned the architecture of the proposed exchange itself.

A significant strand of opinion held that exchange of compliance information between obliged entities, from one private institution to another, should never have been the design choice, because there is no central certifying authority for obliged entities and a malicious actor pretending to be a CASP could harvest substantial volumes of personal data, leading to further crimes. On this view, any information sharing should occur at the state level, between law enforcement, supervisors, and FIUs, with private sector institutions acting as enablers rather than peer participants in the exchange.

The point was raised that the framework's outsourcing of public authority tasks to the private sector is a structural problem rather than an implementation difficulty: public authorities themselves have acknowledged that they do not currently have the capacity to take the function back in-house, but that this resource gap does not justify maintaining a structure that places enforcement responsibilities on private actors who are not equipped, mandated, or accountable to perform them.

A third thread engaged with the technological alternatives. Several lines of discussion converged on the view that the architecture should move away from sharing data altogether and toward sharing proofs of properties of data, supported by verifiable credentials, zero-knowledge proofs, and the European Digital Identity Wallet infrastructure. On this view, a financial institution does not need to receive the personal data of a counterparty; it needs to verify that certain properties hold, that the wallet has been associated with a verified identity, that the holder is not sanctioned, that the counterparty has not accumulated network-level suspicion, without the underlying data being transmitted or stored centrally. The Estonian X-Road infrastructure was cited as a working example of one-to-many data exchange architecture in operation for over two decades, and the maturity of zero-knowledge proof technology for these use cases was the subject of disagreement, with some participants holding that the technology is ready for production and others reporting that financial institutions consistently feed

back that it is not yet mass-adoption ready in a plug-and-play form.

A further line of discussion engaged with the meaning of "bad actor" and the question of who is entitled to define it. Concern was raised that any architecture in which private actors are required to identify bad actors will produce inconsistent and over-conservative classifications, because financial institutions are reputationally incentivized to err toward exclusion, particularly in jurisdictions with significant penalties for compliance failure. The structural problem was identified as the framing of financial institutions as responsible for preventing crime, when their proper role is the provision of financial services and the support of law enforcement through transparency. The point was made again that anti-money laundering is structurally a post-crime regime, and that prevention is the proper domain of law enforcement, not of obligated entities. The current allocation of responsibility produces the perverse outcome that significant compliance budgets are spent on activities that do not meaningfully reduce financial crime, while law enforcement remains under-resourced for the investigative and prosecutorial work that actually disrupts illicit flows.

A specific concern was raised about the propagation of suspicion across institutions in the absence of meaningful challenge mechanisms. Where a flag raised by one institution is shared into a network, other institutions are likely to act on it conservatively rather than independently re-evaluate, because the cost of re-evaluation exceeds the cost of exclusion. This produces a structural risk of permanent debanking based on initial errors, and case examples were given of political activists having to fight for years and at significant personal cost to recover from incorrect flags propagated through international intelligence channels. Any intelligence-sharing architecture must therefore include the data subject's right to be informed of a flag where this is consistent with the integrity of an active investigation, the right to challenge it, and the right to have it removed across the network where the challenge is upheld.

The interaction between this requirement and the tipping-off rules of existing AML frameworks was identified as an area requiring careful legal design.

The discussion also engaged with adjacent regulatory developments. Article 75 of the EU AML Regulation was repeatedly cited as the relevant legal foundation for partnership sharing within the EU, with participants noting that the framework already permits the model under discussion subject to appropriate safeguards. The Global Coalition to Fight Financial Crime's new digital asset task force was identified as actively working on wallet-level labelling and intelligence sharing, and the broader trend toward intelligence-led, network-analytic approaches was noted as a positive direction even where current data reporting under pilot projects is more statistical than operational.

The point was made that fundamental rights to financial services should be recognized as a baseline, both ethically and practically, because the current ease of opening short-lived accounts has shifted the criminal use of the banking

system toward disposable, single-purpose accounts that do not respond to traditional account-level controls.

A final structural observation closed the discussion in the same way that the previous topic: the current architecture is the product of a 2001 policy decision, taken under a particular geopolitical pressure, that the international community has been unable to revisit despite mounting evidence that the model does not work. The current revision cycle of the FATF Recommendations represents a rare opening for substantive reconsideration, and the case for repurposing Recommendation 16 should be advanced through that channel.

Calls to Action regarding Recommendation 16

The key calls to action from the discussion are:

- **Repurpose Recommendation 16 from bulk data transmission to qualifying intelligence-sharing networks:** FATF should supply the international legal standard for qualifying networks, including vendor neutrality, interoperability, participant vetting, and governance, in which financial institutions, VASPs, FIUs, and law enforcement agencies share risk information only where defined red-flag thresholds are met, building on existing models such as Singapore's COSMIC platform, the EU's Horizon 2020 TRACE project, and Article 75 of the EU Anti-Money Laundering Regulation.

- **Anchor any sharing architecture in robust rule-of-law safeguards:** Any qualifying network must operate on a defined suspicion or higher-risk threshold and incorporate pseudonymisation, mandatory data-protection impact assessment, joint-controller accountability, independent audit, and a redress mechanism that gives the data subject the right to be informed of a flag where consistent with active investigations, to challenge it, and to have it removed across the network where the challenge is upheld.

- **Move from data transmission to proof sharing through verifiable credentials and zero-knowledge infrastructure:** The longer-term architecture should rely on cryptographic proofs of properties of data rather than the transmission of personal data itself, supported by verifiable credentials, zero-knowledge proofs, and

federated identity infrastructure such as the European Digital Identity Wallet, allowing institutions to verify what they need to verify without expanding the attack surface or creating new honeypots.

- **Reframe the role of financial institutions as enablers of law enforcement, not as private agents of crime prevention:** The current allocation of preventive responsibility to obliged entities should be reconsidered, with public funding for law enforcement and FIU capability strengthened through redirected compliance spend or membership contributions from regulated entities, and with the role of financial institutions clarified as the provision of financial services and the support of public authority, not the policing of customers.
- **Preserve targeted financial sanctions and the preventative core of the framework:** The repurposing of Recommendation 16 should be undertaken on the explicit basis that targeted financial sanctions under Recommendations 6 and 7 remain fully in force, ensuring that the preventative core of the international AML/CFT framework is preserved while the data transmission obligation is replaced with a more proportionate and operationally effective intelligence-sharing model.

3. FATF Recommendations 33, 41, and 42: Statistics, Evidence Standards, and Fundamental Rights Safeguards

The third topic was presented by Dr. Agata Ferreira (Institute of Free Technology and Status) and addressed a structural blind spot at the heart of the FATF framework: the absence of any mechanism to evaluate whether the regime achieves what it sets out to achieve. The proposal advanced a shift from a compliance-driven model to an evidence-based and rights-respecting one, anchored in three coordinated reforms: a revision of Recommendation 33 on statistics, the introduction of a new Recommendation 41 on evidence standards, and the introduction of a new Recommendation 42 on fundamental rights safeguards.

On Recommendation 33, the framing of the problem was one of asymmetric information. The current Recommendation captures only "wins" while ignoring costs, collateral damage, and inefficiencies. The regime operates without an empirical basis for decision-making: policy outcomes are largely unknown and it is difficult to substantiate that money laundering is actually being prevented. The Recommendation measures activity rather than effectiveness, inputs rather than outcomes, and does not measure costs at all. This produces a "success theater" in which countries report rising suspicious activity report volumes and asset seizures as evidence of progress, without any visibility into whether the regime is generating serious consequences, real benefits, or negative externalities such as discrimination. Banks file reports

defensively to avoid penalties, reports are not independently audited, and institutions never learn whether their submissions are operationally useful.

On the proposed Recommendation 41, the framing was one of self-assessment and compliance theater. Policy is developed and maintained without evidence of efficacy, on the basis of assumptions rather than data. There is no cost-benefit analysis, and current evaluation relies on countries grading their own implementation against criteria that are formalistic rather than substantive. No systematic, independent review examines whether AML measures actually reduce financial crime, increase the difficulty or cost of laundering, or justify the resources they consume, and no mechanism exists for measuring discriminatory impacts. Mutual evaluations focus on process compliance, whether laws exist and whether reports are filed, rather than on effectiveness, whether crime is reduced or the system is proportionate.

On the proposed Recommendation 42, the framing was one of impunity for arbitrary exclusion and fundamental rights violations. The regime prioritizes compliance over fundamental rights protection and treats non-discrimination, privacy, and financial inclusion as acceptable sacrifices. Where rights violations occur, including discrimination, privacy invasions through excessive data demands, and financial exclusion, there is no remedy, no independent review, and no explanation. Banks bear no cost for wrongfully excluding legitimate customers, and there are no procedural safeguards for those affected. Financial institutions act as gatekeepers to economic participation without accountability for erroneous or prejudicial decisions.

The solutions presented were structured to map onto the three functions a working regulatory framework requires: measurement, evaluation, and rights protection.

For Recommendation 33, the proposal was a move toward radical transparency and evidence-based accountability. Policy should be evaluated against its true costs and benefits to society, with publication and analysis of costs, negative externalities, and effectiveness metrics alongside enforcement activity. Mandatory annual disclosures should include de-banking and refusal statistics disaggregated by customer category to expose discrimination patterns; estimated private and public compliance costs to enable cost-benefit analysis; outcome ratios linking SAR volumes to downstream investigations, prosecutions, asset recovery, and convictions, to verify whether reporting generates actionable intelligence or merely defensive noise; and customer complaints and redress outcomes, to create accountability for arbitrary exclusion. Reports should be independently verified rather than self-reported, the underlying data should be public, and feedback loops should reward effectiveness rather than the volume of reports filed.

For the proposed Recommendation 41, the solution was the establishment of regular, independent evaluation of national frameworks conducted by bodies with no operational responsibility for implementation. Reviews should assess outcomes against defined, measurable

goals, such as changes in asset recovery rates, estimated volumes of laundered funds, and the cost or difficulty of laundering for perpetrators, rather than inputs and outputs. This should include full cost-benefit analysis of compliance burdens and assessment of discriminatory impacts. Results should be published and incorporated into evaluation methodology, with self-assessed compliance metrics replaced by externally validated effectiveness data.

For the proposed Recommendation 42, the solution was a fundamental rights-based implementation framework. National implementations should ensure no "black box" decisions and include procedural safeguards: termination or denial decisions should be subject to due process standards including written reasons, the right to be heard, and independent review. Compliance cannot be invoked to violate fundamental rights, and compatibility with human rights law should be guaranteed. Protection should be reinforced at multiple levels, including effective remedies through independent review and rights to require remediation (compensation or reinstatement), ongoing non-discrimination monitoring and assessment, the introduction of basic banking service and minimum access rights, a reversal of the burden requiring banks to justify exclusion (a presumption of inclusion), and time limits on exclusion with a right of review after a defined period.

The discussion that followed engaged closely with the diagnosis and with the political and methodological challenges of moving the framework in the direction proposed. There was broad agreement that the absence of effectiveness measurement is a structural failure of the current regime, and concrete examples were given of how the existing reporting cascade produces volume without value.

One example cited was a single jurisdiction's authorities reporting that 15,000 SARs from crypto companies in the prior year had produced only five criminal cases, with 2,000 of those SARs concerning identity theft, which is not the function the AML regime is designed to address. The point was made that AML has expanded over time to absorb categories of conduct, particularly fraud, that do not fit its original purpose, with the consequence that the burden grows while the regime drifts further from its stated objectives. A parallel observation drew on training experience with national authorities, where the operational data presented by FIUs and supervisors was characterized as irrelevant statistics that the authorities themselves were frustrated by.

A structural concern was raised about the strategic positioning of any reform. A blunt call to abolish particular Recommendations is unlikely to make progress at the FATF level, whereas a call for evidence-based effectiveness reporting is more difficult to oppose openly, because no public actor can credibly resist the principle that policy should be evaluated against outcomes. The proposal to require jurisdictions to disclose effectiveness data was therefore seen as the politically viable wedge for the broader reform agenda, particularly if the existing incentive architecture, which rewards countries for implementing the standards and removes them from blacklists for procedural

compliance, can be realigned to reward demonstrated effectiveness instead.

A significant strand of discussion engaged with what effectiveness measurement would actually require. Mutual evaluation reports already collect substantial data on convictions, asset recoveries, and procedural compliance, but the assessment of effectiveness has been subordinated to the question of how well a jurisdiction implements the standards, rather than how well the standards themselves work.

The proposal was to require, at the FATF level, an additional metric on the effectiveness of the standards as applied, encompassing not only enforcement output but also cost of compliance, unintended consequences, and the benefit-to-burden ratio of the regime in practice. The point was made that effectiveness measurement cannot be limited to counting convictions but must engage with the full social cost of the framework, including financial exclusion, the depressing effect on legitimate investment, and the disproportionate impact on individuals and communities flagged by reputation-driven institutional risk aversion.

A methodological question was discussed at length: whether effectiveness evaluation requires access to raw, individual-level data, or whether anonymized and aggregated data is sufficient. One position held that genuine research into the effectiveness of the regime requires the linkage of individual-level data across multiple sources, including denial decisions, criminal convictions, labor market outcomes, and subsequent banking access, with the analysis conducted by qualified researchers under strict confidentiality controls through a trusted third party. The argument was made that central banks already handle highly sensitive credit registry data under tightly controlled conditions, with credentialed access, no internet connectivity, and personal device restrictions, and that comparable arrangements could enable rigorous evaluation while preserving privacy.

A countervailing position cautioned that this approach risks reproducing the same problem the reform is meant to solve: justifying further data collection and aggregation as a precondition for evaluating whether existing data collection is justified. The concern was that proposing additional centralized data infrastructure, even for research purposes, opens the door to further expansion of the surveillance perimeter, and that the framing should instead place the burden on the regime to prove its effectiveness with the data already available, rather than on the public to enable further collection in order to make that proof possible.

The interaction between metadata, pseudonymized data, and the practical re-identifiability of individuals through metadata analysis was identified as a significant practical risk requiring careful design.

A further thread engaged with the accountability and redress dimension. The point was made that financial services in most jurisdictions are provided by private institutions, which cannot ordinarily be compelled to do business with a particular

customer. However, in jurisdictions such as Germany, public savings banks have a statutory obligation to provide accounts, and in Canada an access-to-banking law requires private banks to provide accounts to individuals, with the institution permitted to manage the relationship through monitoring rather than denial. These models were identified as practical templates for ensuring that legitimate customers are not excluded from the financial system as a side-effect of compliance pressure, while preserving the institutions' ability to manage risk through ongoing monitoring rather than initial exclusion.

A final structural observation closed the discussion. The case for empowering citizens to use accountability mechanisms was raised, with the practical point that the cost of legal challenge currently puts redress out of reach of most individuals affected by erroneous or disproportionate exclusion decisions. A simpler, accessible, lower-cost accountability mechanism was identified as essential if the fundamental rights dimension of the proposed Recommendation 42 is to function in practice rather than only on paper.

The broader observation was that the original problem the FATF framework was created to solve was the lack of collaboration between nation-state law enforcement agencies, not the construction of a private compliance industry, and that the displacement of public responsibility onto private institutions has produced both the inefficiency and the rights violations the proposed reforms are designed to address.

Calls to Action regarding FATF Recommendations 33, 41, and 42

The key calls to action from the discussion are:

- **Reform Recommendation 33 to require effectiveness, cost, and exclusion statistics:** Mandatory annual disclosures at the FATF level should include de-banking and refusal statistics disaggregated by customer category, estimated private and public compliance costs, outcome ratios linking SAR volumes to downstream investigations, prosecutions, asset recoveries, and convictions, and customer complaints and redress outcomes. Reporting should be independently verified rather than self-reported, with the underlying data published to enable public and academic scrutiny.
- **Introduce a new Recommendation 41 on independent evidence-based evaluation:** Establish regular, independent evaluation of national frameworks by bodies with no operational responsibility for implementation, assessing outcomes against defined, measurable goals including asset recovery rates, estimated

volumes of laundered funds, and the cost or difficulty of laundering for perpetrators. Evaluation must include full cost-benefit analysis and assessment of discriminatory impacts, and self-assessed compliance metrics should be replaced with externally validated effectiveness data.

- **Realign incentives at the FATF level to reward effectiveness, not procedural compliance:** The current incentive architecture rewards countries for implementing the standards and removes them from blacklists for formal compliance, regardless of whether the standards produce results. Incentives should be realigned to reward demonstrated effectiveness in reducing financial crime, with national reputations and access to international financial infrastructure tied to outcomes rather than reporting volumes.

- **Introduce a new Recommendation 42 on fundamental rights safeguards:** National implementations must include procedural safeguards for termination and denial decisions, including written reasons, the right to be heard, independent review, and effective remedies including reinstatement and compensation. The framework should establish minimum access rights to basic banking services, reverse the burden so that institutions must justify exclusion (a presumption of inclusion), introduce time limits on exclusion with a right to review, and require ongoing non-discrimination monitoring.

- **Provide accessible, low-cost accountability mechanisms for affected individuals:** Procedural safeguards under Recommendation 42 must be supported by a simpler, accessible, and affordable accountability mechanism so that redress is available in practice to individuals affected by exclusion or rights-violating decisions, rather than only to those who can sustain the cost of legal challenge.

- **Enable independent research into regime effectiveness under strict safeguards**: Where evaluation requires access to data, this should be facilitated through trusted third parties operating under central-bank-style confidentiality controls, with credentialed researcher access and no expansion of the underlying data collection perimeter. The burden of proving the framework's effectiveness rests on the regime, not on the public to consent to further data collection in order to enable evaluation.

We thank all participants of the Berlin 6.0 DARTE event for contributing to the discussion:

Agata Ferreira (Status), Alireza Siadat (Deloitte), Benedikt Faupel (Bitpanda), Chloe White (Riskmastery.xyz), Claire Balva (ADAN), Esen Esener, Gustav Hemmelmayr (Web3 Foundation), Jennifer Hanley-Giersch (Berlin Risk), Jörn Erbguth, Juan Ignacio Ibañez

(MiCA Crypto Alliance), Magnus Jones (Nordic Blockchain Association), Mariana de la Roche (BlackVogel), Marie Rinke (Möhrle Happ Luther), Martijn Keuzenkamp, Max Bruche (Humboldt University), Nadine Kari (N26), Nina Siedler (Siedler Legal), Pavlina Pavlova (ChainComply), Samuel Jacques-Cloutier (Hash Directors), Stefano De Angelis (Nethermind), Zornitsa Daskalova (Zodia).