BISIH Academia Workshop

Report 2025

Bank for International Settlements

September 2025

Contents

Acknowledgements

We would like to thank everyone who participated in the workshop for their valuable contributions and insights. Below is the list of professors and academics who took part.

	Name	Institution
1	Dan Awrey	Cornell Law School
2	Rainer Boehme	University of Innsbruck
3	Srdjan Capkun	ETH Zurich
4	George Danezis	University College London
5	Darrell Duffie	Stanford University
6	Christof Fetzer	TU Dresden
7	Rod Garratt	University of California Santa Barbara
8	Co-Pierre Georg	Frankfurt School of Finance and Management
9	Ruizhe Jia	Standford University
10	Aljosha Judmayer	University of Vienna
11	Markulf Kohlweiss	University of Edinburg
12	Bingle Kruger	University of Cape Town
13	Eysa Lee	Brown University
14	Anna Lysyanskaya	Brown University
15	Michael Maurer	MIT
16	Hart Montgomery	Linux Foundation
17	Fernando Perez-Cruz	BIS and ETH Zurich
18	Bart Preneel	KU Leuven
19	Judith Senn	University of Innsbruck
20	Srisht Fateh Singh	University of Toronto
21	Dawn Song	UC Berkeley
22	Andreas Veneris	University of Toronto
23	Julius Wenzel	TU Dresden

Below is the list of BIS staff members who contributed to the discussions.

	Name	Position
1	Iñaki Aldasoro	Principal Economist, MED
2	Mike Alonso	Adviser, BIS Innovation Hub
3	Mert Askaroglu	Adviser, BIS Innovation Hub
4	Raphael Auer	Centre Head, BIS Innovation Hub
5	Morten Bech	Centre Head, BIS Innovation Hub
6	Codruța Boar	Adviser, BIS Innovation Hub
7	Italo Borssatto	Adviser, BIS Innovation Hub
8	Alonso Carrillo	Junior IT Developer, BIS Innovation Hub
9	Titose Chembezi	Associate, BIS Innovation Hub
10	Joachim Coche	Head of Middle Office, Banking Systems and Operational Services
11	Jakub Demski	Adviser, BIS Innovation Hub
12	Miguel Diaz	Deputy Head BISIH and Head of Strategy
13	Daniel Eidan	Adviser, BIS Innovation Hub
14	Jon Frost	Head of Innovation and Digital Economy, MED
15	Denise Garcia Ocampo	Senior Advisor, FSI, FSI
16	Haukur Gudmundsson	Adviser, BIS Innovation Hub
17	Stephanie Haffner	Adviser, BIS Innovation Hub
18	Jack Ho	Adviser, BIS Innovation Hub
19	Henry Holden	Adviser, BIS Innovation Hub
20	Karmela Holtgreve	Deputy Head of BIS Innovation Hub and Head of Operations
21	Hiren Jani	Head of Data and Analytics, ITS
22	Friedrich Klinger	Adviser, BIS Innovation Hub
23	Luiz Eduardo Laydner Cruz	Adviser, BIS Innovation Hub
24	Juan Jose Lopez	Junior IT Developer, BIS Innovation Hub
25	Darko Micić	Lead Architect, BIS Innovation Hub
26	Raunak Mittal	Adviser, BIS Innovation Hub
27	Marko Nanut Petric	Adviser, BIS Innovation Hub
28	Keerthi Nelaturu	Adviser, BIS Innovation Hub
29	Vasily Pozdyshev	Senior Advisor, FSI

	Name	Institution
30	Esther Rey Losada	Coordination and Operations Lead, BISIH
31	Rafael Schmidt	Head of MED IT, MED
32	Cecilia Skingsley	Head of BIS Innovation Hub
33	Micah Smith	Adviser, BIS Innovation Hub
34	Jakub Sykulski	Operations Manager, BIS Innovation Hub
35	Vanessa Tampoya Espano	Adviser, BIS Innovation Hub
36	Peter Wierts	Senior Economist, CPMI
37	Septine Wulandini	Adviser, BIS Innovation Hub
38	William Zhang	Adviser, BIS Innovation Hub

1. Executive summary

The BIS Innovation Hub's (BISIH) first academia workshop, held from 10 to 13 June 2025, convened distinguished professors specialising in advanced information technologies, financial technology, economics and law, alongside decision-makers and experts from the BIS. The workshop focused on exploring the application of technology in the financial sector and marked the launch of the BISIH's academic engagement programme.

The workshop combined sessions that highlighted perspectives from senior BIS decision-makers with sessions exploring technologies and their potential benefits. Open discussions between BIS leaders and academics set the context for the consecutive sessions, which focused on key technology themes.[1] These discussions were framed around real-world policy needs, such as security, interoperability, inclusivity and regulatory compliance, ensuring that the discussions were both forward-looking and grounded in practical realities.

As financial systems adapt to rapid technological shifts, a compelling narrative is emerging – one in which intelligence, security, interoperability, and identity converge to reshape the foundations of finance. Large language models (LLMs) and generative artificial intelligence (AI) are driving unprecedented automation and insight. This growing impact places pressure on system architecture and governance to evolve in tandem. In response, composable, scalable infrastructures – built for verifiability and trust – are becoming the backbone of next-generation financial systems. Achieving seamless interoperability is a challenge, as value flows across disparate systems with varying definitions of finality and trust. Privacy technologies, especially when embedded by design, must scale with throughput to meet institutional and user expectations, while

[1] Artificial intelligence, interoperability and architecture, scalability, trusted execution environments, privacy-enhancing technologies, digital identity and decentralised finance.

trusted execution environments offer secure processing where latency and confidentiality collide. At the edge of identity innovation, self-sovereign credentials and cryptographic mechanisms promise user-centric, privacy-preserving alternatives, shifting the role of institutions from issuers to validators. In this evolving landscape, even decentralised finance (DeFi) must be re-examined to try to identify its beneficial aspects. Together, these threads form a roadmap not just for modernisation, but for resilience, inclusion and strategic foresight.

Key insights emerged from the discussions:

- LLMs and AI tools are expected to significantly affect the global economy, irrespective of whether they reach human-level capabilities. While their potential benefits are evident, current systems are highly susceptible to AI-assisted attacks. The declining cost of such attacks is likely to result in a sharp increase in their frequency and scale, raising concerns about whether investment levels are sufficient to balance the benefits and risks of AI adoption.
- Scalable distributed systems require composable architectures that integrate seamlessly into broader ecosystems. Effective governance models must evolve alongside technological advancements to maintain trust and operational stability. Key use cases include cross-border payments, digital identity and programmable money, with scalability solutions addressing risks such as resilience trade-offs and regulatory responsiveness.
- In terms of architecture and interoperability, the discussion highlighted that while technological solutions for integrating diverse financial infrastructures are increasingly available, the primary barriers to achieving interoperability lie in policy and political differences. Modern advancements have made technical integration feasible, whether through centralised or decentralised systems, with or without tokenisation. However, the legal and governance frameworks that regulate cross-system value transfers remain critical to ensuring trust and operational continuity. The integration of traditional and non-traditional systems, such as through financial market infrastructures designed to act as settlement integrators, was seen as a promising path forward.
- Privacy is an important requirement for financial sector users, both individuals and financial institutions. The discussion emphasised that increased demand for privacy is being addressed through a combination of advanced technologies and design principles. “Privacy by design” was identified as the optimal approach, where privacy-enhancing technologies (PETs) are integrated into systems from the outset rather than being retrofitted into legacy infrastructures. Scalability emerged as a key challenge for PETs, particularly in high-throughput environments. To address this, advancements in specialised hardware, such as application-specific integrated circuits (ASICs) and field-programmable gate arrays (FPGAs) and optimised cryptographic techniques, were highlighted as potential solutions.

- Trusted execution environments (TEEs) are emerging as a valuable tool for creating secure and trusted operations in financial systems. Already widely used in non-financial sectors for enhancing security and privacy, TEEs are now being explored for their potential in high-throughput applications, including standard SQL databases, with minimal overhead, making them suitable for real-time payment platforms. TEEs offer the ability to isolate sensitive computations, ensuring confidentiality and integrity, even in less trusted environments. Their scalability and low overhead make them suitable for systems where latency and transaction volume are critical. However, in order for them to be adopted in financial services challenges such as ensuring trust in hardware manufacturers and mitigating security vulnerabilities must be addressed.
- Self-sovereign identity and verifiable credentials offer a pathway to inclusive and privacy-preserving identity systems, giving individuals greater control over their data. However, their implementation requires addressing governance, legal and technical challenges through interdisciplinary collaboration. Central banks were seen as validators and enablers, supporting governance frameworks and infrastructure to safeguard privacy, rather than acting as issuers of digital identities. Concerns about biometric identifiers, including risks of impersonation and the irreversible nature of biometric data, highlighted a preference for decentralised cryptographic mechanisms over reliance on immutable personal traits.
- DeFi presents both opportunities and challenges, with discussions highlighting the need for nuanced regulatory approaches and a clear delineation between stablecoins and tokenised wholesale central bank money.

The workshop fostered interdisciplinary collaboration, generating insights and laying the foundation for continued engagement between academia, policymakers and technologists.

2. Introduction

The BISIH fosters international collaboration on innovative financial technologies within the central banking community. Its mission is to identify critical technological trends, develop public goods to enhance the global financial system, and serve as a hub for innovation among central bank experts. In line with this mission, the BISIH hosted its first academia workshop from 10 to 13 June 2025, marking the beginning of a structured programme of engagement with academia.

It brought together 61 participants, including in computer science, law, economics and finance, as well as other academics. Outgoing and incoming BIS leaders also participated. Open discussions between them and academics set the context for the consecutive sessions based on each key theme. These sessions were individually structured by BIS staff and partner academics. Activities were interactive, diverse and tailored to each theme (to better maintain engagement and collaboration).

Through these sessions, the workshop created a common understanding of central bank priorities and pain points. This exploration of future technologies provided academics a window to directly present their research to the BIS. It also allowed them to tailor some of their work considering the needs of the system – and for the BIS it provided a fresh perspective into frontier technologies and an opportunity to understand how and where they could contribute further. It highlighted where shared research interest exists, supported the creation of personal connections to push the boundaries together, and seeded future collaborations. Indeed, the workshop is part of the BISIH broader academia interaction plan, which will see similar workshops and leverage other models of collaboration going forward.

This inaugural workshop explored different potential technologies that can be used in future financial systems. The key themes were: (i) architecture and interoperability; (ii) scalable distributed systems; (iii) trusted execution environments; (iv) privacy enhancing technologies; (v) verifiable credentials; and (vi) DeFi, with a seventh topic briefly touching on the pervasive influence of AI.

The workshop was not an attempt to design every aspect of the system or prescribe technologies. Instead, it was the first step to generate a long-term dynamic interaction among the participants, with the objective of bringing light to the development of better tools for interested central banks to allow them to fulfil their mandates more efficiently.

This report provides a non-attributable record of the discussions and insights from the three-day workshop. Each section corresponds to one of the key topics explored during the event, offering a structured summary of the presentations, breakout sessions, and collaborative discussions.

Section 3 captures the foundational discussions with BIS leaders, which set the stage for the technical sessions by framing the policy needs and challenges that emerging technologies should address.

The following sub-sections of section 4 explore the seven key topics that shaped the workshop agenda.

The report concludes with reflections on the workshop's outcomes and insights, highlighting the importance of continued collaboration between academia, policymakers, and technologists in shaping a resilient and inclusive financial ecosystem.

3. Defining foundations with BIS Management

The workshop began with a foundational session led by BIS management, focusing on the evolving role of central banks in shaping the future financial system. It addressed key topics such as trust, singleness of money, elasticity of money, and integrity, as well as the challenges posed by technological innovation and the role of central banks in adapting to a rapidly changing financial landscape. They are described in detail in Chapter 3 of the BIS Annual Report 2025.

Participants also explored the distinctions between tokenised wholesale central bank money, stablecoins, and commercial bank money tokens, along with their implications for monetary policy and financial stability. Other topics included data protection, the interplay between private-sector innovation and regulatory oversight, and the formation and preservation of trust in digital money.

The session provided a critical framework for the technical discussions that followed, grounding them in the broader context of governance, public policy objectives, and the evolving role of central banks in the digital age. Several open questions and challenges were raised, which are summarised below.

3.1 Further issues and questions

The discussion explored several critical questions: how trust in digital money can be established, maintained or undermined; the appropriate level of central bank engagement with technological innovation; and whether the financial system effectively addresses illicit activities. It also examined whether stablecoins can meet fundamental requirements of the financial system, such as singleness, elasticity and integrity, and considered whether tokenised wholesale central bank money is necessary to preserve the singleness of money as the use of cash declines.

4. Key topics

4.1. Artificial intelligence

This AI discussion titled Securing the Future: Integrating AI in Financial Systems was structured into two parts namely i) *The path to Artificial General Intelligence (AGI)*; ii) *Towards Building Safe and Secure AI: Lessons and Open Challenges*. The former discussed the current capabilities and limitations of AI, including their potential to automate complex tasks and transform financial operations. It drew attention to the engineering and societal challenges of deploying AI systems, particularly in high-stakes environments. The latter emphasised safeguarding AI systems is not just a technical problem but a sociotechnical challenge requiring collaboration across academia, industry, and policy-making communities.

The path to AGI

As part of this agenda, the module titled "*The Path to AGI*" examines the transformative potential of AI agents and their role in the journey toward AGI. While achieving AGI remains an open question, advancements in LLMs and AI agents have already begun reshaping how tasks are automated and how decisions are supported. The discussion covered the current state of AI agents, their applications, limitations, and the challenges of integrating them into real-world systems, particularly in finance.

Defining AI agents and AGI

AI agents are software systems that can process information and take actions using other systems to pursue goals and complete tasks autonomously or semiautonomously. They are categorised into two types.

- **Narrow agents**: These focus on specific tasks, emulating deliberate, logical thinking, often referred to as "System 2"(slow, deliberate and logical). (Kahneman, 2011) Narrow agents are as effective as their users and are exemplified by organisations like OpenAI, xAI, and DeepMind. They are not AGI but excel in specialised applications.
- **General-purpose agents**: These aim to perform multiple cognitive functions simultaneously, engaging both "System 1" (fast, automatic and intuitive thinking) and "System 2". They emulate broader human-like behaviour, such as controlling a computer as a human would. However, their error rates remain high, and they are still far from achieving true general-purpose capabilities.

AGI, in contrast, refers to a general-purpose AI system capable of performing almost all cognitive tasks that humans can do. While AI agents are seen as a potential gateway to AGI, they remain highly useful even without achieving full general intelligence.

The promise of AI agents

AI agents, particularly multimodal LLMs, have the potential to transform industries by automating complex, human-like tasks.

These applications demonstrate the utility of AI agents in financial systems and regulatory environments, where they can save time, improve efficiency, and support decision-making.

Challenges and limitations

Despite their potential, AI agents face significant challenges. A parallel between the current state of AI and the development of self-driving cars a decade ago can be drawn. While initial progress appeared rapid, the remaining challenges have proven far more complex. Two major problems stand out:

1. **Scale of real-world data**: Low-probability, high-impact events are difficult to model effectively.
2. **Reliability and errors**: AI agents are prone to errors, particularly when performing tasks that require reasoning, consistency, or handling unexpected scenarios.

Broader implications

There are also societal implications of deploying AI agents. LLM-based systems, such as chatbots, generate plausible and grammatically correct responses but may lack the ability to discern truth and morality. This limitation may have significant consequences for how AI is used in decision-making processes and raises ethical questions about its role in society.

Reinforcement learning with human feedback (RLHF) is a method that uses human-generated evaluations to guide and improve the behaviour of AI systems. In this approach, human feedback aids the reinforcement learning process to align AI outputs with desired outcomes or values. This technique has been effective in fine-tuning AI systems, particularly in tasks where human judgment plays a critical role.

While RLHF is a valuable approach for aligning AI systems with human expectations, it may not provide a comprehensive solution. Its limitations become apparent in complex, high-stakes applications where deeper contextual understanding and reasoning are required. The path forward will likely require a combination of techniques, including application-specific engineering, modular system designs, and advancements in interpretability to ensure AI systems are robust, adaptable, and aligned with societal needs.

Key takeaways

For sceptics, the presentation underscores the revolutionary potential of zero-shot capabilities in LLMs, which can perform tasks in seconds that previously required months of effort. For enthusiasts, it serves as a reminder of the limitations and challenges that remain. Simply adding more data may not lead to significant improvements, and engineering solutions should be tailored to specific applications.

While AI agents can automate tasks and improve efficiency, they should be deployed responsibly, with an understanding of their limitations and potential risks.

Towards building safe and secure AI: lessons and open challenges

Spectrum of AI risks

AI systems present a range of risks across several dimensions:

- **Misuse and malicious use**: AI can be exploited for scams, misinformation, non-consensual imagery, cyberattacks, and even bioweapon development. The increasing capabilities of AI make such misuse more accessible and scalable.
- **Malfunction and systemic** failures: Issues such as bias, system malfunctions, and inappropriate deployments can cause significant harm. Broader systemic risks include privacy violations, copyright concerns, labour market disruptions, and environmental impacts.
- **Loss of control**: As AI systems become more integrated into critical infrastructures, the potential for systemic failures due to bugs, adversarial attacks, or vulnerabilities grows significantly.

Typical cyber-attack pattern

Graph 1

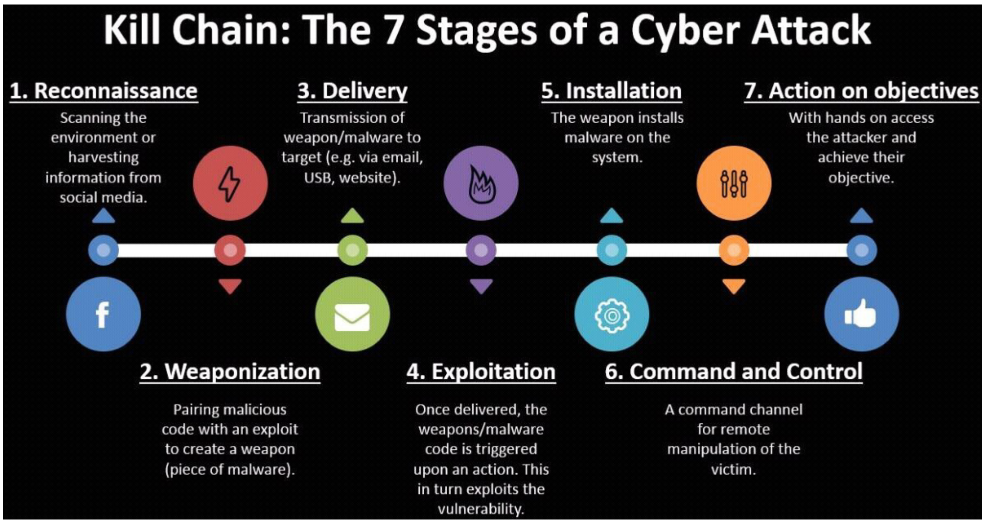

Kill Chain: The 7 Stages of a Cyber Attack

1. Reconnaissance
Scanning the environment or harvesting information from social media.

2. Weaponization
Pairing malicious code with an exploit to create a weapon (piece of malware).

3. Delivery
Transmission of weapon/malware to target (e.g. via email, USB, website).

4. Exploitation
Once delivered, the weapons/malware code is triggered upon an action. This in turn exploits the vulnerability.

5. Installation
The weapon installs malware on the system.

6. Command and Control
A command channel for remote manipulation of the victim.

7. Action on objectives
With hands on access the attacker and achieve their objective.

These risks are exacerbated by the asymmetry between attackers and defenders. Attackers only need to exploit a single vulnerability, while defenders must secure all potential weaknesses, making defence both resource-intensive and time-consuming (Graph 1). History shows that attackers often exploit new technologies quickly, and AI is no exception. As AI systems grow more capable, the incentives for malicious actors increase, along with the severity of potential consequences

Building safe and secure AI systems

To address these challenges, a multi-pronged approach to building safe and secure AI systems is advocated:

- **Proactive defence through secure-by-design systems:** AI systems should be designed with safety and security embedded from the outset. Secure-by-design approaches use formal verification methods to ensure that AI systems meet predefined security properties. This proactive approach contrasts with reactive defences, such as patching vulnerabilities after deployment.

- **Systematic evaluation of AI trustworthiness**: Platforms like DecodingTrust and MMDT are critical for assessing AI systems across multiple dimensions, including robustness, fairness, hallucination, and privacy. These tools provide a comprehensive evaluation framework for large language models (LLMs) and multimodal foundation models, enabling better understanding and mitigation of risks.
- **Adversarial robustness**: AI systems should be resilient against adversarial attacks, such as data poisoning, prompt injections, and jailbreaks. New paradigms are needed to harden systems, including scalable oversight mechanisms, input/output guardrails, and representation control.

AI in cybersecurity

The integration of AI into cybersecurity introduces both opportunities and challenges. Frontier AI systems, which combine symbolic and non-symbolic components, create new vulnerabilities. Misused AI can enhance attackers' capabilities, enabling large-scale phishing campaigns, disinformation efforts, and deepfake generation.

Building defences against cyber attacks

Graph 2

Aspect	Attack	Defense
Cost of failures	● **High tolerance** for failure. ● Can rerun or adjust strategies if an attack fails. ● Exploit probabilistic AI to generate repeated attacks.	● **Low tolerance** for failure due to serious consequences. ● Must ensure accuracy to avoid false positives (disrupt operations) and false negatives (leave threats uncovered). ● Require extensive validation/verification, especially for AI-generated code or patches.
Remediation deployment and required resources	● Target **unpatched and legacy** systems using public vulnerability data. ● Exploit **delays in patch deployment** to launch attacks.	● **Lengthy and resource-intensive** process (e.g., testing, dependency conflict, global deployment). ● Legacy systems take longer to patch, leaving vulnerabilities unpatched.
Different priorities of scalability and reliability	● **Prioritize scalability**, enabling large-scale attacks on huge number of targets. ● Use AI to reduce human effort and automate attacks.	● **Focus on reliability**, making AI adoption challenging due to robustness and transparency limitations. ● High trust in AI is difficult due to unpredictability and errors.

Defensive efforts face significant challenges (Graph 2). Initiatives like BountyBench and CyberGym evaluate AI agents' performance in detecting, exploiting, and patching vulnerabilities. However, the asymmetry between offence and defence remains a major obstacle. Lessons from other fields, such as medical device security, highlight the importance of early risk detection, formal evaluations, and secure system design.

Open challenges and future directions

Several open challenges remain and should be addressed to ensure the safe integration of AI into critical systems.

- **Evaluation and protection in finance**: How can financial institutions ensure the safety and security of AI applications? This includes addressing fraud, compliance, and vulnerabilities in programmable payment systems.
- **Regulatory frameworks**: The proliferation of AI legislation highlights the need for cohesive, science-based policies that balance innovation with safety.
- **Sociotechnical solutions**: Technical solutions alone are insufficient. Collaboration across academia, industry, and civil society is essential to build trust, increase transparency, and enhance societal resilience.

Lessons and predictions

AI will initially benefit attackers more than defenders, as current systems are highly vulnerable to AI-assisted attacks. The cost of attacks is expected to decrease, leading to an unprecedented increase in frequency and scale. However, system defenders can also leverage AI to strengthen security measures, and it remains uncertain which side will ultimately gain the greater advantage.

Drawing parallels with past challenges, such as spam and script-based attacks, urgent action is needed. Organisations must prioritise building secure systems, learning from past experiences, and fostering collaboration to mitigate risks effectively.

4.2. Architecture and interoperability

The session on architecture and interoperability explored how financial systems can evolve to integrate new functionalities and technologies, , while maintaining stability, efficiency, and inclusivity. Participants highlighted the need for a synchronisation layer including a settlement arrangement to enable seamless interaction between new and existing systems without extensive reconfiguration of legacy infrastructures.

Challenges in the current cross-border payment model

The inefficiencies of the current cross-border payment model (Graph 3), which relies on correspondent banking, were a focal point of the discussion. This model involves multiple intermediaries, extensive information exchanges, and frequent updates to ledgers, all of which contribute to delays, higher costs, and operational risks because these are not coordinated and automated. Participants critiqued this system and discussed the potential for unified ledgers and object-based transfers to streamline processes and reduce friction, this improvement can also be achieved through synchronisations and efficient coordination even if the services are provided by many entities.

Governance frameworks and guiding principles

Governance frameworks were identified as essential for managing the integration of new systems while minimising the need for re-contracting and disruption. Participants proposed several guiding principles to shape the evolution of financial system architectures:

1. **Coexistence of systems**: Centralised RTGS, DLT-based ledgers, instant payment hubs, tokenised wholesale central bank money, and other systems will coexist and evolve rather than replace one another entirely.

2. **Evolution without disruption**: New payment infrastructures should be designed to integrate with existing systems without requiring wholesale replacement of legacy systems.

3. **Integrity across systems**: Payments should be designed to ensure atomicity, auditability, and regulatory compliance, avoiding scenarios where transactions fail mid-process or leave funds in limbo.

4. **Interoperability as a priority**: Scalable architecture patterns, interoperability playbooks, and integrity tools are essential for seamless interaction between systems.

Typical cross-border payment involves correspondent banks

Graph 3

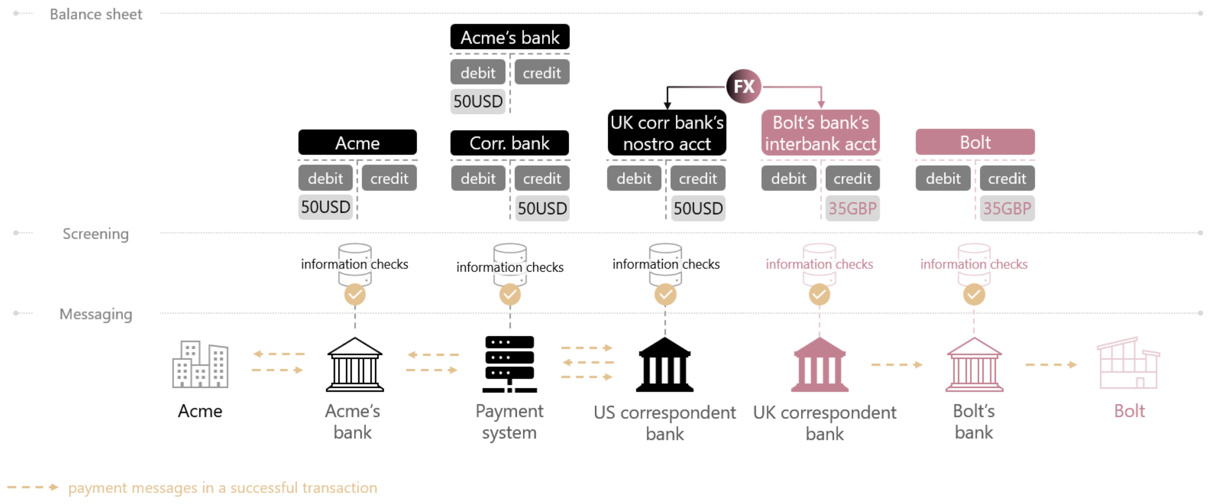

Legal and regulatory dimensions

Another perspective emphasised that interoperability should be treated as a deliberate policy choice to enhance stability and resilience. Legal and regulatory considerations, such as prudential conditions, the legal structure of value (property vs contract), and anti-money laundering (AML) / countering the financing of terrorism (CFT) compliance, were central to that discussion. Participants acknowledged that regulatory alignment across jurisdictions is a significant challenge but crucial for fostering trust and operational continuity.

Approaches to integration

The discussion outlined several technical and operational approaches to achieve integration across diverse financial systems, while maintaining system integrity:

- **Passing value objects between ledgers:** This method ensures that value can move seamlessly across systems without disruption.
- **Burn-and-issue mechanism**: Under this approach, value is destroyed on one ledger and simultaneously issued on another. This ensures that the total value remains consistent across systems, preventing duplication. The Agora model was referenced as an example of this mechanism in practice.
- **Hash time-locked contracts (HTLC)**: These contracts use conditional transactions to ensure atomicity, meaning that a transaction is completed across systems only when specific conditions are met. This prevents incomplete or partial transfers, preserving the integrity of the process.
- **Trusted clearing facilities**: Centralised entities can facilitate clearing and settlement across networks, offering a reliable mechanism for ensuring the finality of transactions.
- **Trusted intermediaries**: These intermediaries can bridge gaps between incompatible systems. These intermediaries help overcome technical or operational barriers to enable value transfers that would otherwise be impossible.

Concerns raised

Participants raised critical concerns about the practical implementation of these approaches. For example, while unified ledgers and object-based transfers offer significant theoretical benefits, questions remain about how these systems would handle complex governance requirements, cross-border regulatory compliance, the scalability needed for high transaction volumes, and the final transfer of value when a user ends up with value in a ledger but there is no liquidity to transfer it to a more desirable one. The role of trusted intermediaries was also debated, highlighting the simplicity of this solution when there are trusted operators, but with some participants questioning whether reliance on centralised entities undermines the reliability of newer technologies.

4.3. Privacy-enhancing technologies

The breakout session on privacy-enhancing technologies (PETs) explored their role in privacy, security, and compliance. These technologies are designed to protect sensitive information while allowing secure, auditable, and regulatory-compliant data processing. They aim to strike a balance between preserving user privacy and meeting requirements for AML/CFT. A key challenge, however, lies in ensuring that PETs enable the detection of illicit activities, such as money laundering, without compromising the privacy of legitimate users.

Several key PETs were discussed, each suited to specific use cases. Zero-knowledge proofs (ZKPs) allow institutions to demonstrate compliance without revealing sensitive data, making them particularly valuable for regulatory checks. Secure multi-party computation (MPC) enables collaborative data analysis while ensuring that individual inputs remain confidential. Homomorphic encryption allows computations on encrypted data without decryption, though its scalability and computational costs limit its practical use. TEEs emerged as a mature solution for isolating sensitive computations, ensuring secure execution even in less trusted environments. Other tools, such as federated analysis and differential privacy, facilitate aggregate data analysis while maintaining user anonymity.

Scalability challenges in PETs

Scalability was identified as a significant barrier, particularly for high-throughput systems like tokenised wholesale central bank money and interbank settlements. While technologies like ZKPs and homomorphic encryption offer theoretical advantages, their computational demands often limit their application in real-time, large-scale systems. Advancements in specialised hardware, such as application-specific integrated circuits (ASICs) and field-programmable gate arrays (FPGAs), as well as optimised cryptographic techniques, were highlighted as potential solutions to these challenges.

TEE use cases, benefits and risks

TEEs were recognised as a practical and ready-to-deploy option, with demonstrated use cases in healthcare and digital payments. By isolating computations from the operating system and hardware, TEEs ensure secure execution even on compromised platforms. However, concerns about their reliance on hardware vendors and vulnerability to physical attacks were noted.

Privacy by design

The workshop emphasised the importance of adopting a "privacy by design" approach, where PETs are integrated into systems from the outset rather than retrofitted into legacy infrastructures. This proactive strategy provides stronger privacy protections and greater scalability. However, achieving this requires significant investment, clear governance, and coordination among stakeholders.

Standardisation and open collaboration

Collaboration among technologists, regulators, and financial institutions was identified as critical for defining requirements and fostering standardisation. Open initiatives, such as hackathons, academic partnerships, and collaboration with standard-setting bodies, were seen as effective ways to accelerate innovation and adoption. Central banks were encouraged to take a leadership role in driving the adoption of PETs by setting standards and fostering collaboration.

Key takeaways

- PETs are mature but should be tailored to specific use cases.

- Scalability is a major challenge for PETs in high-throughput systems. Advancements in hardware and optimised cryptographic primitives are needed to meet the demands of digital cash and interbank markets.
- TEEs are practical but come with trust and governance risks.
- Privacy by design is more effective than retrofitting. Building PETs as platforms from the outset ensures broader interoperability and stronger privacy protections.
- Revocation of privacy should be carefully designed. Systems need mechanisms to revoke privacy under lawful conditions without undermining trust or security.
- Open collaboration accelerates PET development. Partnerships with academic institutions, standardisation bodies, and open-source initiatives can drive innovation and adoption.
- Clear requirements are critical for PET adoption. Policymakers and financial institutions should define precise goals to guide the development of privacy-preserving systems.
- PETs can balance privacy and compliance. Technologies like ZKPs and MPC allow for AML/CFT compliance without compromising user privacy, but trade-offs should be managed.
- Central banks should take a leadership role. By setting standards and fostering collaboration, central banks can drive the adoption of PETs in financial systems.
- Trust in technology complements institutional trust. Transparency in the design and implementation of PETs enhances public trust in financial systems

4.4. Trusted execution environments

The breakout session on TEEs explored their role in enhancing data security, governance, and operational integrity within financial systems. TEEs are specialised, hardware-based security solutions designed to ensure that sensitive data and operations remain confidential and tamper-resistant, even in environments where full trust in data operators is impractical (or undesirable), such as public cloud infrastructures or multi-stakeholder data collaborations.

The architecture of TEEs isolates data, and executable code within an encrypted memory enclave, making them inaccessible to the broader operating system. This isolation protects sensitive information from malicious actors and unauthorised access. An important feature of TEEs is remote code attestation, which allows external parties to verify the identity and integrity of the code running within the environment. This is complemented by safeguards, such as policy enforcement mechanisms, sandboxing techniques to prevent unauthorised communication, and the integration of cryptographic logs or immutable records (e.g., Merkle trees or blockchain-based ledgers) to enhance auditability and accountability.

An example of TEE deployment can be found in Germany's e-health system, where sensitive patient data is hosted on untrusted cloud service providers. TEEs ensure strong privacy and security assurances, preventing unauthorised access or misuse of this data.

Potential use cases in financial systems

In financial systems, TEEs offer several potential applications. One use case is secure and tamper-proof data collaboration. TEEs enable banks to combine data into shared pools for advanced analytics, such as fraud detection, systemic stress testing, and reserve analysis, while enforcing control over their respective datasets. Another application is in digital payments, where TEEs can play roles in the core settlement layer, mobile wallets, payment detail lookup services, KYC and onboarding processes, and offline payment solutions.

Scalability and performance

Scalability and performance were key points of discussion. TEEs can support high-throughput applications, including standard SQL databases, with minimal overhead, making them suitable for real-time systems like payment platforms where latency and transaction volume are critical.

Security

Security concerns, such as side-channel attacks and rollback vulnerabilities, were acknowledged. These risks can often be mitigated through system hygiene, controlled environments, monotonic counters, and distributed consensus protocols, though trust in hardware manufacturers remains essential. Backup and recovery strategies were also considered. Regular security attestation and the use of Hardware Secure Modules for private key storage were suggested to address availability requirements. However, remote key backup was noted as potentially undesirable from a cybersecurity perspective.

Governance

Governance and auditability were identified as important aspects of TEE implementation. For change management, a control policy known as "M of N" was proposed. Under this policy, decisions of approvals require agreement from at least M participants out of a total of N authorised individuals. This ensures that no single person or small group can act unilaterally, boosting accountability and security. Immutable audit logs further support this by providing a clear mechanism to trace and approve data access and modifications.

Legal and policy implications

The workshop also explored the legal and policy implications of using TEEs. A key insight was the relational nature of data, where information about one individual may inadvertently expose information about others. TEEs provide a technical solution to enforce joint control over data usage, enabling stakeholders to define and enforce access policies while retaining ownership. This is particularly relevant in regulatory contexts governed by multiple legal frameworks.

Blockchain technologies

Finally, TEEs were compared with blockchain technologies. While blockchains offer unconditional public verifiability, TEEs provide a more efficient and privacy-preserving alternative for many use cases. The concept of "trusted contracts" executed within TEEs was proposed as a substitute for public smart contracts, particularly when combined with robust attestation and governance frameworks.

Key takeaways

- TEE has transformative potential in building secure, auditable, and privacy-preserving digital infrastructures.
- TEE is a mature and scalable technology and has been deployed in production to process sensitive data from millions of users.
- Further discussions with business and compliance SMEs are necessary to brainstorm more potential financial use cases, including AML and fraud detection. The ideas may be formulated into research papers and proof of concept testing.
- TEEs can operate in a highly isolated (e.g., air-gapped) environment to minimise attack surface.
- TEEs must operate within robust governance structures that realise institutional trust and meet regulatory requirements.
- TEEs can support a scalable and distributed architecture with high throughput, by logically extending across multiple physical nodes.
- Combining TEEs with digital identity systems and other PETs (e.g., ZKPs) enhances privacy and secure data management sharing.
- Establishing standardised certification enhances trust, interoperability and security assurance.

4.5. Self-sovereign identity and verifiable credentials

The breakout session on Self-Sovereign Identity and verifiable credentials centred on four key pillars: stakeholder needs, the dimensions of the digital identity ecosystem, research and development priorities, and interoperability across identity frameworks.

Theme A: Stakeholders and their needs

The digital identity ecosystem involves a wide range of stakeholders, including public authorities (such as central banks and regulatory bodies), private sector entities (notably commercial banks and fintech firms), technical implementers, and end-users.

Central banks were identified as key validators or enablers, rather than issuers of digital identities. Their focus should be on developing governance frameworks and infrastructure that safeguard privacy and promote inclusivity. Centralised control over personal data should be avoided in favour of decentralised governance models.

The private sector, particularly commercial banks, often emphasises the importance of confidentiality in identity systems. Regulatory compliance, including KYC and AML/CFT requirements, significantly shapes their data retention and sharing practices. These institutions typically highlight concerns about sharing sensitive user data in open environments, reinforcing the need for privacy-preserving mechanisms.

For end-users – including citizens, migrants, workers, and patients – the discussions emphasised usability, privacy, and equitable access. Many existing identity systems assume that individuals possess government-issued documentation, which risks excluding underserved populations in rural areas or those without formal records. Mobile-based identity solutions and credentials issued by trusted community entities were suggested as alternatives. Across all user groups, the design of identity systems should prioritise usability and sensitivity to local contexts to ensure meaningful adoption.

Theme B: Dimensions of the digital identity ecosystem

The architecture of digital identity systems encompasses both technical and non-technical dimensions. On the technical side, foundational components include public key infrastructure (PKI), digital signatures, user wallets, and protocols such as Decentralised Identifiers (DIDs) and W3C Verifiable Credentials (VCs). Advanced cryptographic tools, such as ZKPs, were highlighted for enabling selective disclosure and enhancing privacy. Revocation frameworks and reputation-based credentials were also identified as critical for mitigating fraud and maintaining trust.

Non-technical dimensions include governance structures, liability frameworks, and mechanisms for accountability and dispute resolution. Participants stressed the need for identity systems to address broader societal and institutional challenges, such as inclusivity, risk management, and civil resistance. Regulatory frameworks, such as PSD2 and open banking, were cited as useful precedents for designing systems with layered compliance and delegated responsibilities.

The evolution of identity systems was also discussed, moving from traditional document-based models to more fluid, contextual, and multi-attribute constructs. Self-Sovereign Identity architectures were seen as a promising example of this shift, offering users greater control over their identities while supporting privacy and interoperability.

Theme C: Research and development roadmap

A key theme was the gap between interest in self-sovereign identity systems and their practical implementation. While jurisdictions like the European Commission are funding Self-Sovereign Identity experimentation, many institutions remain cautious, limiting their efforts to controlled environments ("Monitoring the landscape, not piloting"). Participants discussed the need for clear indicators to determine when to transition from research to real-world deployment. Suggested criteria included demonstrable user demand, regulatory clarity, the maturity of underlying digital infrastructure and the existence of institutional or market incentives.

Priority research areas identified during the workshop included the design of privacy-preserving architectures using tools like ZKPs, mechanisms for issuing and verifying cross-border credentials, and integration with programmable payment systems. Strategies to mitigate identity duplication and fraud, while maintaining strong privacy protections, were also seen as critical. Central banks were encouraged to lead interdisciplinary research efforts, drawing on expertise from cryptography, legal theory, sociology, and economics to ensure identity systems are both technically robust and socially responsive. An unresolved question was what specific factors or conditions might lead a central bank to transition from exploratory research to real-world deployment.

Theme D: Interoperability across identity systems

Interoperability, both domestic and cross-border, was recognised as a foundational requirement for robust identity systems. Achieving alignment between decentralised and centralised infrastructures is essential for ensuring compliance, operational continuity, and inclusivity. This is particularly important for individuals who migrate across jurisdictions or lack formal documentation.

Proposed solutions included community-based or reputation-derived credentials validated through federated standards, as well as minimal, portable cryptographic proofs to ensure flexibility without imposing rigid frameworks. Participants also advocated for global registries or credential resolvers, potentially leveraging distributed ledger technologies (DLTs), to verify credentials across jurisdictions.

Concerns about biometric identifiers were raised, with participants highlighting risks such as impersonation, surveillance, and the irreversible nature of biometric data. Privacy-centric architectures that avoid reliance on immutable personal traits were strongly preferred.

The session drew attention to existing initiatives, such as the BIS's Project Mandala, which demonstrate how compliance logic can be embedded within programmable financial infrastructures. Participants recommended a modular, interoperable design approach, described metaphorically as a "Lego-style" system. This would allow jurisdictions to tailor identity frameworks to their unique legal, cultural, and technological contexts while ensuring cross-system compatibility.

Key takeaways

- Central banks should focus on acting as validators and enablers of digital identity systems, supporting governance and infrastructure that safeguard privacy and ensure inclusivity, rather than taking on the role of identity issuers.
- User-centric design is critical for ensuring equitable access to identity systems, particularly for individuals without formal documentation. Solutions should be flexible, inclusive, and tailored to local contexts, such as mobile-based identities or credentials issued by trusted community entities.

- Cryptographic tools like ZKPs and robust revocation systems are essential for enabling privacy, selective disclosure and trust in digital identity systems.
- Legal and institutional frameworks should address governance, accountability, and inclusion from the start. They should be treated as core design components, not afterthoughts.
- Interoperability across identity systems is vital, both domestically and internationally. Modular, standards-based approaches, such as "Lego-style" architectures, allow jurisdictions to create systems suited to their specific needs while ensuring compatibility and trust across borders.
- Biometric identifiers face strong resistance due to risks like impersonation, surveillance, and the irreversible nature of biometric data. Decentralised privacy-preserving alternatives are strongly preferred.

The discussions captured in this report highlight the multifaceted nature of building a trustworthy, inclusive, and interoperable digital identity ecosystem. As the global financial system moves toward increasing digitisation, identity is emerging not as a peripheral component but as a foundational layer of digital infrastructure.

While technological solutions such as Self-Sovereign Identity, ZKPs, and verifiable credentials offer promising pathways, the workshop underscored that successful implementation would depend equally on institutional commitment, cross-sector collaboration, and public trust. Moving forward, stakeholders should pursue coordinated efforts that blend innovation with inclusivity, standardisation with local adaptability, and functionality with fundamental rights.

Ultimately, digital identity is not simply a technical or regulatory challenge; it is a societal issue. Its design and governance will shape access, agency, and accountability in the digital economy for decades to come. The insights gathered here provide a foundation for continued dialogue and action toward realizing that vision.

4.6. DeFi

The breakout session on DeFi examined the challenges and opportunities of decentralised finance, focusing on issues of compliance, governance, and the interplay between public and private digital financial infrastructures. The session aimed to foster cross-disciplinary dialogue, raising critical questions rather than providing definitive answers. Two main topics guided the discussion: regulatory compliance in decentralised environments and the dynamics between stablecoins and tokenised central bank money.

Topic 1: Compliance in decentralised environments

Participants explored the feasibility of achieving regulatory compliance in DeFi systems, particularly those without identifiable intermediaries or jurisdictional anchors. A key theme was the need to distinguish between varying levels of decentralisation. Participants recognised that regulatory enforcement faces inherent limitations in decentralised architectures but also discussed emerging ideas for aligning public policy goals with code. Fully decentralised systems may be harder to regulate but could provide essential alternatives in jurisdictions where traditional institutions fail to deliver financial access or trust.

A nuanced, caveat-based regulatory approach was proposed, categorising platforms by their risk profiles and governance structures. Some platforms might adhere to formal regulatory frameworks, while others could operate with clear user warnings. This would require a detailed taxonomy of digital assets and platforms, distinguishing between centralised stablecoins, governance tokens, infrastructure tokens, and partially decentralised services.

The evolving nature of DeFi was also discussed. Many platforms labelled as DeFi incorporate centralised elements, such as admin keys or off-chain governance, making it essential for regulators and researchers to avoid treating DeFi as a monolithic category. A layered understanding of technical architecture, governance models, and user controls is necessary.

The example of Uniswap illustrated the complexity of accountability in decentralised systems. While the protocol is permissionless and open source, its development and interface are maintained by a known team. This raised questions about liability: should the organisation behind a protocol be held accountable for its misuse? While no consensus was reached, participants stressed the importance of fostering community-driven innovation without stifling public goods that enable experimentation and access.

The discussion also touched on adapting DeFi protocols to evolving legal frameworks. Ideas included modular smart contract designs, self-declared compliance parameters, and optional legal wrappers allowing protocols to align with specific jurisdictions. The concept of a "private international law for DeFi," where protocols voluntarily adhere to chosen legal principles, was considered promising but complex.

Ultimately, participants agreed that compliance in DeFi should be addressed across multiple layers, from blockchain and protocol design to asset characteristics and user-facing services.

Topic 2: Stablecoins vs tokenised central bank money (public infrastructure and market dynamics)

The second topic examined the distinctions between stablecoins and tokenised central bank money, focusing on whether these instruments can meet public objectives in similar ways. While both represent digital value, their design, issuance, and institutional backing differ significantly.

Stablecoins are typically issued by private entities with commercial incentives, which can drive innovation but also pose risks related to pricing, governance, and data privacy. Participants noted that stablecoins often charge fees, may limit access, and carry operational or solvency risks. In contrast, well-designed digital cash represents a public good: accessible, fee-free, and backed by sovereign guarantees.

Participants agreed that stablecoins cannot replace tokenised central bank money in delivering the social contract that public money represents. Tokenised central bank money is meant to be a fair, inclusive, and trusted option backed by the state. There was also strong agreement that offline tokenised central bank money should follow the example of physical cash by protecting privacy and allowing for anonymous payments.

It was noted that in many economies, stablecoins have emerged as a practical substitute for tokenised central bank money that are either unavailable or underdeveloped. However, participants expressed caution about assuming stablecoins can permanently fill this gap. Their limitations, particularly in ensuring equitable access and long-term trust, suggest that tokenised central bank money is uniquely positioned to deliver the full benefits of digital public money.

The discussion also explored the role of privacy and data governance. tokenised central bank money was seen to reintroduce cash-like anonymity into the digital age, potentially restoring trust in state-backed systems. However, several participants raised concerns that even tokenised central bank money may struggle to guarantee privacy in practice. Meanwhile, stablecoins, especially those issued by large platforms, might monetise user data, creating incentives that conflict with public interest.

One key observation was the distinction between individual and collective data value. While a single user's data may hold little market value, aggregated data across millions of users represents substantial economic and political influence. This highlighted the importance of designing digital monetary systems that prevent excessive concentration of data and control.

Considering these differences, participants broadly agreed that stablecoins and tokenised central bank money will likely coexist. However, their roles in the ecosystem should be clearly defined, especially in terms of user protections, regulatory obligations, and systemic implications.

Key takeaways

- DeFi raises important questions around how public policy objectives can be met in open, permissionless systems that lack central intermediaries or jurisdictional anchors.
- A caveat-based regulatory approach could allow compliant and non-compliant protocols to coexist but would require a detailed classification of digital assets and system designs.
- The term "DeFi" covers a broad and evolving spectrum, including systems with varying degrees of centralisation; regulatory and technical analysis should consider these layers rather than treat DeFi as a uniform category.

- Fully decentralised protocols may offer critical alternatives in jurisdictions where trust in traditional financial institutions is low, serving as neutral infrastructure for access and resilience.
- Examples like Uniswap illustrate the difficulty of assigning accountability when open-source protocols are replicated or maintained by loosely coordinated groups.
- Participants discussed embedding compliance features in code through modular design, self-declared compliance settings, or legal wrappers, although enforcement remains a challenge.
- Stablecoins and tokenised central bank money are not interchangeable. While stablecoins serve current market needs, tokenised central bank money is designed to deliver a public money system aligned with social policy objectives.
- Participants agreed that offline tokenised central bank money should emulate the privacy and anonymity of cash to maintain trust and accessibility in digital payments.
- There was broad consensus that stablecoins cannot fulfill the role of tokenised central bank money in delivering the social contract between the state and the public.
- The distinction between individual and collective data value highlighted the need for strong privacy safeguards and public oversight in digital monetary systems.
- While stablecoins and tokenised central bank money are likely to coexist, their roles should be clearly defined in terms of access, risk, regulation, and trust.
- The session raised more critical questions than it resolved, reinforcing the need for continued dialogue between academics, technologists, and policymakers.
- Forums like this are essential for building mutual understanding and co-developing credible, adaptable frameworks for decentralised financial infrastructure.
- Regulatory and academic communities should establish regular, structured engagements to turn emerging questions into practical, shared solutions.

The session highlighted the growing complexity of the digital finance landscape, particularly on distinguishing between decentralised infrastructure and centrally governed solutions. Discussions around initiatives like Uniswap, where the solution is highly decentralised but is maintained by a group of developers, raised critical questions about how accountability should be assigned in systems that are open source, widely replicated, and maintained by evolving communities. No definitive resolution was reached, but participants emphasised the importance of not discouraging permissionless innovation, especially where it enables broader access and experimentation. As the BIS and central banks confront rapidly evolving technologies, regular dialogue with academic and open-source communities can help shape regulatory models that are both credible and adaptable.

4.7. Scalable distributed systems technology

The session on scalable distributed systems examined how next-generation financial infrastructures can scale, while maintaining governance, interoperability and user trust. Participants explored both the technical and institutional dimensions of scalability.

Scaling technologies

Participants emphasised that scalability is not just about increasing transaction throughput. It involves addressing multiple factors, including latency, composability and inter-network operability. Several technologies were discussed:

- **Layer 2 solutions**: Technologies like rollups and sidechains can reduce the workload on base chains by handling transactions off-chain. However, concerns were raised about liquidity fragmentation and the erosion of unified trust assumptions.
- **Zero-knowledge proofs**: ZKPs were highlighted as a promising tool for enabling privacy-preserving verification at scalable. While still computationally expensive, they are seen as critical for applications such as digital identity, compliance and settlement systems.
- **Cross-domain messaging**: This was identified as essential for connecting sovereign or application-specific chains, enabling financial systems to operate seamlessly across borders. Participants noted that the future of scalable finance may rely on a network of specialised components interacting securely, rather a single monolithic system.

The concept of “scalability as composability” was introduced, emphasising that systems should not only scale autonomously but also integrate into wider ecosystems without creating friction or systemic risks.

Governance and institutional dynamics

The session also explored how governance and institutional framework intersect with scalability. Several challenges were identified:

- **Governance latency**: Policy and governance frameworks often evolve more slowly than the systems they oversee, creating mismatches between technological innovation and regulatory responsiveness. Participants questioned how systems can remain upgradeable while supporting effective human governance.
- **Institutional scalability**: This refers to the ability of central banks and regulators to adapt their capacity, coordination and technical understanding to engage with increasingly modular and multi-layered architectures.

- **Resilience trade-offs**: Efforts to optimise scalability can sometimes introduce risks to critical system components, particularly in settlement infrastructures. Participants stressed the importance of clearly defining and controlling these failure modes.

Participants highlighted the need for governance models that can evolve alongside technical systems while maintaining trust and operational stability.

Use cases and sector impact

Specific use cases discussed included:

- **Cross-border payments**: Scalable consensus and messaging layers are needed to support low-latency, frictionless settlements across legal jurisdictions.
- **Digital identity and authentication**: Systems handling millions of daily verifications should balance speed, cost, and privacy.
- **Programmable money**: Scalable infrastructure is essential for supporting fine-grained disbursement logic, automated compliance and micropayment flows.

In each case, participants highlighted that scale is not only about accommodating volume, but also about enabling richer functionality and more robust guarantees at speed.

Closing reflections

The session highlighted that scalability is a multifaceted challenge requiring collaboration across system architecture, governance and public policy. Participants noted that scalable infrastructures should be composable, verifiable and designed to minimise systemic risks.

While the session placed important groundwork, participants also identified dimensions that could enrich future discussions. These include opportunities to delve more deeply into interoperability design, explore end-user and market-level considerations, and develop a shared taxonomy to better organise the wide range of technologies and use cases. Expanding regional representation and incorporating illustrative benchmarks may also help anchor conceptual insights in practical realities.

Overall, the session laid important groundwork for further exploration, with participants recognising that addressing scalability effectively will likely require sustained collaboration among academia, policymakers and technologists.

5. Next steps

The workshop provided a platform for collaboration between academia, policymakers and technologists, fostering dialogue on how advanced technologies can address challenges and opportunities in the financial system Its impact is already evident in several actionable outcomes that reflect the workshop's success and its potential for long-term influence.

Many areas were identified where BISIH staff could benefit from academic perspectives. For example, follow-up discussions on stablecoins have provided an opportunity for the leaders of the BISIH's projects to engage with academic specialists, exchanging views and exploring potential avenues for collaboration.

Another outcome is the preparation of a deep dive technical report on TEEs by participating professors. This report will delve into the technical, governance, and policy dimensions of TEEs, reflecting their potential relevance for central bank use cases. The professors will also present their findings at the 2025 Innovation Summit roundtable, further amplifying the workshop's influence on global discussions about financial innovation.

The workshop also inspired the development of a forthcoming paper on joint supervisory data analysis. Multidisciplinary discussions highlighted how encryption technologies could address confidentiality restrictions, enabling collaboration between supervisory authorities without compromising sensitive data. This paper will explore practical solutions to long-standing challenges in data sharing and analysis, offering insights for regulatory cooperation.

Finally, recognising the importance of equipping decision-makers with knowledge of emerging technologies, the workshop has led to plans for additional seminars for central bank governors. These sessions will focus on key topics such as AI, distributed systems and DeFi, helping governors navigate the complexities of the evolving financial landscape.

In conclusion, the workshop facilitated valuable discussions and actionable outcomes, laying the groundwork for continued collaboration. The identified projects and follow-up initiatives will contribute to advancing the understanding and implementation of innovative technologies in the financial system.

6. References

Bank for International Settlements (2023): "Blueprint for the future monetary system: improving the old, enabling the new", Annual Economic Report 2023, Chapter III.

Bank for International Settlements (BIS) (2024): "Artificial intelligence and the economy: implications for central banks", Annual Economic Report, June, Chapter III.

Bank for International Settlements (BIS) (2025): "The next-generation monetary and financial system", Annual Economic Report June, Chapter III.

Bech, M, J Hancock and W Zhang (2020): "Fast retail payment systems", *BIS Quarterly Review*, March, pp 28–9.

Carstens, A, J Frost and H S Shin (2022): "A foundation of trust", *IMF Finance & Development*, September.

Carstens, A and N Nilekani (2024): "Finternet: the financial system for the future", *BIS Working Papers*, no 1179, April

Committee on Payment and Settlement Systems and Technical Committee of the International Organization of Securities Commissions (CPSS-IOSCO) (2012): Principles for financial market infrastructures.

Kahneman, D (2011): "Thinking, Fast and Slow". Farrar, Straus and Giroux.

Garratt, R and H S Shin (2023): "Stablecoins versus tokenised deposits: implications for the singleness of money", *BIS bulletin*, no 73, April.

Perez-Cruz, F and H S Shin (2024): "Testing the cognitive limits of large language models", *BIS Bulletin*, no 83, January.

Rice, T, G von Peter and C Boar (2020): "On the global retreat of correspondent banks", *BIS Quarterly Review*, March, pp 37–52.