emerging tech decoded

LEX CONSULT EDITORIAL

BYTE/WISE

KEY HIGHLIGHTS

MARCH 2026

Artificial Intelligence

Regulatory Updates

- European Commission Releases Draft Code of Practice on Labelling AI-Generated Content
- United States Issues National AI Policy Framework for Federal Legislation
- EU Considers Delays to AI Act Implementation Pending Technical Standards Readiness

Clause in Focus - Model Updates and Change Management

Knowledge Corner - Intended Purpose and Substantial Modification of AI Systems under the EU AI Act

Compliance Snapshot - Internal Access & Use Controls for AI Systems

Privacy

Regulatory Updates

- Supreme Court of India Issues Notice on Challenge to DPDP Act
- Conditional Consent Proposal Gains Attention in EU
- EU CJEU Clarifies Limits on Data Subject Access Requests in *Brillen Rottler* Decision
- China – Cross-Border Data Processing Standards Come into Effect (1 March 2026)
- UK – ICO Fines Reddit £14.47 Million for Children's Data Protection Failures

Clause in Focus - Sub-Processing & Cross-Border Transfer Clauses

Knowledge Corner - Personal Data Breach versus Cybersecurity Incident

Compliance Snapshot - Third-Party Data Sharing Controls

Digital Assets

Regulatory Updates

- Hong Kong Moves Toward Issuance of First Stablecoin Licences
- EU Parliament Committee Examines Digital Assets Challenges
- UAE Approves First USD-Backed Stablecoin Under Central Bank Regulation
- EU ESMA Signals Application of Market Abuse Supervision to Crypto Markets
- UK Advances Cryptoasset Custody and Safeguarding Framework Following Consultation Phase

Clause in Focus - Transaction validation, execution risk & irreversibility

Knowledge Corner - Travel Rule

Dark Patterns Check– Interface Interference

BYTE/WISE

PART ONE

Ai

MARCH 2026

Ai / REGULATORY UPDATES

European Commission Releases Draft Code of Practice on Labelling AI-Generated Content

The European Commission has released a draft Code of Practice on the marking and labelling of AI-generated content, providing a concrete direction on how providers and deployers of general-purpose and generative AI systems should implement disclosure and detectability controls.

Key elements of the draft include:

- Detectability of AI-generated content: Use of technical measures such as watermarking, metadata tagging, or other mechanisms to enable identification of synthetic outputs.
- User-facing disclosure obligations: Clear communication to users where content is generated or materially altered by AI systems.
- Robustness of labelling mechanisms: Controls designed to resist removal, alteration, or circumvention across formats and distribution channels.
- Risk-focused output governance: Identification and mitigation of misuse scenarios, including deepfakes, impersonation, and misleading synthetic media.
- Documentation and traceability: Maintenance of records and internal controls to support accountability and compliance with AI Act requirements.

The draft provides enforceable compliance expectations, requiring organisations to embed transparency, disclosure, and misuse mitigation controls directly into system design, deployment, and monitoring. As discussed in previous Bytewise editions, multiple jurisdictions have already begun taking concrete regulatory steps in this direction, including by China, India and Singapore.

United States Issues National AI Policy Framework for Federal Legislation

The White House has issued a national AI policy framework outlining priorities for the development of a comprehensive federal AI law, signaling a move toward more structured and coordinated AI governance at the federal level.

Key elements of the framework include:

- AI safety and risk mitigation: Focus on harms from advanced AI systems, including misuse of generative AI, child safety risks, and potential downstream harms arising from deployment at scale.
- Accountability across the AI value chain: Emphasis on clearer allocation of responsibility across developers, deployers, and intermediaries, particularly in relation to system design, deployment, and misuse.
- Federal harmonization: Intent to address fragmentation arising from state-level AI regulation and move toward a unified national framework, reducing compliance complexity for organizations operating across jurisdictions.
- Systemic and economic impacts: Recognition of broader risks from large-scale AI deployment, including market concentration, competitive distortions, and wider societal implications.

The framework signals a shift toward more formalized federal oversight of AI but remains at a policy-framing stage. Organizations should view this as an early indicator of the direction of travel for US AI regulation, with more structured statutory requirements likely to follow as federal legislation develops.

BYTE/WISE

EU Considers Delays to AI Act Implementation Pending Technical Standards Readiness

EU policymakers are considering adjustments to the implementation timeline of the EU AI Act, citing delays in the development of harmonized technical standards required to operationalize key compliance obligations. The AI Act relies on these standards to define how requirements, particularly for high risk systems, are to be implemented in practice, including risk management, documentation, and conformity assessment. This development reflects the practical reality that translating legislative frameworks into operational compliance requirements can take time, especially where implementation depends on evolving technical guidance and industry alignment.

Ai / CLAUSE IN FOCUS

MODEL UPDATES AND CHANGE MANAGEMENT

In AI deployments, risk does not arise only at the point of design, but through continuous system evolution. Updates, retraining cycles, and expanded deployment contexts can alter outputs, performance characteristics, and downstream impact, often without any corresponding change to the contractual framework. Without defined controls, such changes may shift risk, expand use beyond scope, or affect compliance obligations without triggering reassessment or accountability.

For this reason, agreements increasingly anchor model updates to two control points: intended purpose and substantial modification. As discussed in the Knowledge Corner, these anchors ensure that system changes are assessed against defined boundaries, preventing updates from introducing unreviewed risks, expanding use beyond scope, or altering compliance obligations without appropriate oversight, reassessment, or control.

Draft Illustrative Clause

"The Provider may implement updates, upgrades, or modifications to the AI System ("Updates") from time to time, provided that such Updates do not materially degrade the agreed functionality of the AI System."

"The AI System shall be used solely for the purposes described in this Agreement or associated documentation ("Intended Purpose"). The Provider shall not implement any change that results in a material deviation from the Intended Purpose without prior notice and, where such change may reasonably impact the Customer's use or compliance obligations, prior written consent."

"A "Substantial Modification" means any change to the AI System that materially affects its intended use, underlying model architecture, training methodology, or data inputs, or that may reasonably be expected to alter outputs or risk profile. In the event of a Substantial Modification, the Provider shall provide advance notice, maintain documentation of such change, and carry out reasonable testing and validation prior to deployment."

make such records available to the Deployer or competent authorities upon reasonable request for the purposes of compliance verification. The Provider shall further ensure that the AI System is designed to support meaningful explainability of its outputs to enable human oversight, in accordance with applicable high-risk AI obligations."

Negotiation Points

Deployer / Customer

- Broaden "Substantial Modification" to capture changes in outputs, deployment context, or user impact
- Require prior approval for high-risk or production-impacting changes
- Include rights to suspend use or trigger reassessment following material changes

Provider / Vendor

- Align "Substantial Modification" with the definition and thresholds under the EU AI Act to limit scope to clearly defined, high-impact changes
- Retain flexibility to implement routine Updates without approval requirements
- Use "commercially reasonable efforts" qualifiers for testing and validation

Balanced drafting should anchor update and change management obligations to regulatory requirements and deployment responsibility, enabling compliance without unintended expansion of liability.

Ai / COMPLIANCE SNAPSHOT

Internal Access & Use Controls for AI Systems

As AI systems are embedded across functions, risks increasingly arise not only from how systems are designed or updated, but from how they are accessed and used within organisations. Internal use, particularly across teams and environments, can extend beyond intended deployment if not appropriately governed.

Global AI governance frameworks and enterprise practices are placing greater emphasis on internal controls around system access, data inputs, and usage boundaries. These controls are critical to ensure that AI systems are used within defined parameters and do not create unintended exposure through day-to-day operations.

What Access and Internal Use Controls Should Cover

- Role-Based Access Restrictions: Access to AI systems should be limited to defined roles, with clear separation across development, testing, and production environments to prevent use beyond intended scope.
- Input and Data Controls: Organizations should establish clear boundaries on the types of data that may be input into AI systems, particularly where personal, confidential, or regulated data is involved.
- Acceptable Use Boundaries: Internal policies should define permissible use cases and restrict deployment in unapproved or high-risk contexts.
- Logging and Monitoring: System access and user interactions should be logged and monitored to enable oversight and timely identification of misuse or anomalous behavior.
- Environment Segregation: Testing, training, and production environments should be clearly segregated to avoid unintended exposure or operational impact.

In addition to the above, organizations should ensure that internal use of AI systems is aligned with documented deployment boundaries. This includes ensuring that employees do not extend system use to unapproved contexts, bypass safeguards, or rely on outputs in ways that were not originally assessed or intended.

Quick Self-Check for Organizations

- Are access to AI systems and permissions structured in line with defined roles, with clear ownership and accountability for system use?
- Can we demonstrate that AI systems are used only within approved purposes and defined deployment boundaries?
- Do we have governance mechanisms to control and review the types of data used as inputs, particularly where it involves sensitive or regulated information?
- Are we able to detect and respond to misuse, unintended use, or deviations from approved use cases in a timely manner?

Organisations should ensure that system use remains aligned with defined deployment boundaries and subject to ongoing oversight. This requires integrating AI-specific controls into broader governance frameworks, including clearly defined ownership, role-based access, data usage boundaries, and continuous monitoring of system use across functions.

Ai

KNOWLEDGE
CORNER

Intended Purpose and Substantial Modification of AI Systems under the EU AI Act

Under the EU AI Act, the compliance framework for AI systems is closely tied to how a system is defined at the point of deployment and how it evolves thereafter. The defined terms "*intended purpose*" and "*substantial modification*" play a critical role in determining the scope and continuity of regulatory obligations.

Intended Purpose

The EU AI Act defines "*intended purpose*" as "the use for which an AI system is intended by the provider, including the specific context and conditions of use, as specified in the information supplied by the provider in the instructions for use, promotional or sales materials and statements, as well as in the technical documentation."

This is a provider-driven, documentation-based concept that anchors regulatory analysis in how the system is presented and positioned for use, rather than in its technical capabilities alone. In practice, the intended purpose determines the context of assessment, the applicable risk classification (including whether the system qualifies as high-risk), and the baseline against which compliance requirements are applied. As a result, the same system may be subject to different regulatory treatment depending on how its intended purpose is defined and communicated.

Substantial Modification

The EU AI Act defines "*substantial modification*" as "*a change to an AI system after its placing on the market or putting into service which is not foreseen or planned in the initial conformity assessment carried out by the provider and as a result of which the compliance of the AI system with the requirements set out in Chapter III, Section 2 is affected or results in a modification to the intended purpose for which the AI system has been assessed.*"

This definition introduces a post-deployment trigger. It focuses on changes that were:

- not originally assessed, and
- that either affect compliance or alter the intended purpose

Routine updates or performance improvements that fall within the originally assessed scope would not typically meet this threshold. The emphasis is on changes that materially affect how the system is used or how risk is evaluated.

These two concepts are structurally linked. The intended purpose defines the system's original regulatory perimeter with conformity assessment, risk evaluation, and associated obligations anchored to this baseline. A substantial modification marks the point at which that baseline is no longer reliable; where a system is altered in a way that changes its intended purpose or affects its compliance with applicable requirements, it may need to be treated, in regulatory terms, as if it were being placed on the market again.

This creates a lifecycle-based approach to compliance:

- deployment within the defined intended purpose → existing compliance framework applies
- modification beyond that scope → reassessment and reallocation of obligations may be required

For organizations, the governance challenge lies in distinguishing between routine system evolution and changes that trigger regulatory consequences. While updates, retraining, and expanded use may be addressed as part of internal development and quality management, regulatory obligations under the EU AI Act arise only where a change qualifies as a substantial modification or alters the system's intended purpose. Systems and teams must therefore be able to identify when a technical or operational change moves beyond normal iteration into a regulatory trigger.

BYTE/WISE

PART TWO

PRIVACY

MARCH 2026

Privacy / REGULATORY UPDATES

Supreme Court of India Issues Notice on Challenge to DPDP Act

The Supreme Court of India has issued notice to the Union Government on a set of Public Interest Litigations challenging key provisions of the Digital Personal Data Protection Act, 2023. The petitions raise foundational concerns regarding the interaction of the DPDP framework with existing legal regimes, particularly the amendment to the Right to Information Act, which is argued to restrict disclosure of personal information even where larger public interest may be involved.

The challenges also question the absence of explicit safeguards for journalistic activity, the breadth of government exemption powers, and the institutional independence of the Data Protection Board.

In addition, concerns have been raised regarding the lack of a compensation mechanism for affected individuals and the scope of what constitutes “personal data” in contexts involving public accountability.

Conditional Consent Proposal Gains Attention in EU

Proposals around “conditional consent” mechanisms are gaining attention within the GDPR ecosystem, reflecting efforts to address structural limitations in existing consent models. The concept seeks to move beyond binary “accept or reject” frameworks by enabling more granular and purpose-specific choices, allowing individuals to selectively consent to certain types of processing, particularly in contexts such as cookie consent, tracking, and personalised services.

These discussions are closely linked to increasing regulatory scrutiny of dark patterns and ineffective consent design, where user interfaces constrain or influence user choice. Conditional consent is being explored as a way to ensure that consent is meaningful, informed, and freely given in practice. While not yet reflected in binding regulatory instruments, this signals a broader shift toward rethinking consent architecture in digital environments.

EU CJEU Clarifies Limits on Data Subject Access Requests in Brillen Rottler Decision

The Court of Justice of the European Union, in its judgment of 13 March 2026 in Brillen Rottler, clarified that data subject access requests (DSARs) under the GDPR may be refused where they are manifestly unfounded or excessive, including in situations involving abuse of rights, even if the request is made for the first time. The case concerned an access request submitted primarily to support a compensation claim rather than to verify the processing of personal data. The Court emphasised that the right of access is intended to enable transparency and verification and cannot be exercised for purposes unrelated to that objective.

The Court also reaffirmed that claims for compensation under the GDPR require a demonstrable causal link between a violation and the harm suffered, and that damages may not be recoverable where the data subject’s own conduct is the primary cause of the alleged harm. The judgment clarifies the application of existing thresholds under the GDPR while maintaining the fundamental nature of the right of access.

Under the Digital Personal Data Protection Act, 2023, comparable situations may be addressed through the statutory duties imposed on data principals, including the obligation not to register false or frivolous grievances or complaints. The Act provides for monetary penalties for breach of such duties, with fines that may extend up to ₹10,000, enabling the Data Protection Board to address misuse of rights mechanisms through an explicit statutory framework.

China – Cross-Border Data Processing Standards Come into Effect (1 March 2026)

China's national standards governing cross-border processing of personal information came into effect on 1 March 2026, providing greater operational clarity on transfer requirements under the Personal Information Protection Law. The standards set out structured expectations across key transfer mechanisms, including security assessments, standard contractual arrangements, and certification routes, and are intended to guide organisations in implementing compliance obligations in practice.

The framework places particular emphasis on accountability of data exporters, requiring organisations to assess the data protection capabilities of overseas recipients, ensure appropriate safeguards, and maintain oversight of personal data after transfer. It also reinforces transparency obligations toward individuals in relation to cross-border transfers. The development signals a continued move toward standardising and operationalising cross-border data governance, with increased focus on control of international data flows and downstream processing.

UK – ICO Fines Reddit £14.47 Million for Children's Data Protection Failures

The UK Information Commissioner's Office (ICO) has fined Reddit £14.47 million for violations relating to the processing of children's personal data, including failure to implement effective age assurance mechanisms and processing children's data without appropriate safeguards. The ICO found that Reddit did not adequately assess or mitigate risks to child users, including exposure to targeted advertising and platform features designed for adults, and failed to ensure that default settings and data use were aligned with the higher protection standards required for children under the UK GDPR and the Age Appropriate Design Code.

The case highlights regulatory expectations around age verification, risk assessment, and age-appropriate design of online services, reinforcing that platforms must actively identify and address risks to vulnerable users rather than rely on general compliance measures. It also reflects a broader trend of intensified enforcement, with regulators increasingly focusing on platform accountability and imposing significant penalties to drive compliance.

Privacy/ CLAUSE IN FOCUS

SUB-PROCESSING & CROSS-BORDER TRANSFER CLAUSES

Across earlier Bytewise editions, vendor-side processing has been addressed through specific control points such as audit and inspection rights, handling of data subject requests, or broader accountability in breach and retention scenarios. Each of these highlights how responsibility for personal data extends beyond the primary organisation to the wider processing ecosystem. This time, we shift the focus to how that ecosystem is structured and governed at the point of data sharing. In particular, sub-processing and cross-border transfer provisions determine who can be brought into the processing chain, under what conditions, and with what safeguards.

These clauses define the baseline controls that apply before data is shared further and ensure that downstream processing remains aligned with the original scope, purpose, and level of protection.

Draft Illustrative Clause

"The Processor shall not engage any sub-processor in connection with the processing of Personal Data without prior written authorisation of the Controller, which may be provided on a specific or general basis. In the case of general authorisation, the Processor shall inform the Controller of any intended changes concerning the addition or replacement of sub-processors, thereby giving the Controller the opportunity to object."

"The Processor shall ensure that any sub-processor is bound by written obligations that are no less protective than those set out in this Agreement, including with respect to confidentiality, security, and data protection."

"The Processor shall not transfer Personal Data outside the agreed jurisdictions without prior written approval of the Controller and shall ensure that any such transfer is subject to appropriate safeguards in accordance with applicable law."

"The Processor shall remain fully liable for the performance of any sub-processor's obligations and shall ensure that Personal Data is processed only for the purposes set out in this Agreement."

Negotiation Points

Data Controller/Data Fiduciary

- Require specific approval rights for sub-processors (not broad general consent)
- Seek visibility into vendor lists and changes over time
- Impose strict flow-down obligations and audit rights
- Restrict or condition cross-border transfers based on risk and regulatory exposure

Data Processor

- Prefer general authorisation for sub-processing to retain operational flexibility
- Limit approval and objection rights to material or security-relevant concerns
- Use standardised safeguards for cross-border transfers
- Avoid extensive or duplicative audit and oversight obligations

These clauses ensure that personal data remains subject to consistent controls across the processing chain, and that transfers to third parties or across jurisdictions do not dilute the level of protection originally agreed.

BYTE/WISE

Privacy/ COMPLIANCE SNAPSHOT

Third-Party Data Sharing Controls

Disclosure of personal data to third parties is a critical control point, requiring organisations to ensure that any sharing remains aligned with the original purpose of collection and subject to defined safeguards. Under both the GDPR and the DPDP Act, disclosures form part of a controlled processing chain rather than standalone events, with responsibility extending beyond the point of disclosure to how personal data is handled by recipients. Under the GDPR, disclosures must be supported by a valid legal basis under Article 6 and remain consistent with the purpose limitation principle under Article 5(1)(b), with controllers responsible for ensuring that recipients process personal data within the defined scope and that cross-border transfers comply with Chapter V safeguards. The DPDP Act adopts a similar approach, requiring processing to remain aligned with notice and consent under Sections 4 to 6, with the data fiduciary retaining responsibility for downstream processing, including disclosures and cross-border transfers.

In practice, the primary risk lies in the incremental expansion of data sharing across vendors, affiliates, and evolving use cases without corresponding review or control. Over time, this can result in disclosures that exceed the originally communicated purpose or introduce unassessed regulatory exposure.

Organisations must therefore implement clear approval processes, contractual restrictions, and ongoing oversight of third-party processing.

Quick Self-Check for Organisations

- Are disclosures of personal data aligned with the purpose communicated at the time of collection?
- Are internal approvals required before sharing data with new or high-risk recipients?
- Are recipients restricted from onward sharing without appropriate authorisation?
- Are data sharing activities documented, including recipients, purpose, and categories of data shared?
- Are cross-border disclosures assessed in advance and aligned with applicable safeguards or restrictions?
- Are responsibilities clearly assigned for approving, recording, and reviewing data sharing arrangements?

Effective management of third-party data sharing requires integrating disclosure controls into broader data governance frameworks. This includes purpose-based restrictions, contractual safeguards, and ongoing oversight of downstream processing, ensuring that data sharing remains controlled as processing ecosystems evolve.

Privacy

KNOWLEDGE
CORNER

Personal Data Breach vs Cybersecurity Incident

Personal Data Breach vs Cybersecurity Incident

The terms "*personal data breach*" and "*cybersecurity incident*" are often used interchangeably in operational discussions. However, they represent distinct scenarios with distinct legal and compliance triggers, and with different reporting obligations and consequences.

Personal Data Breach

Under the Digital Personal Data Protection Act, 2023, "*personal data breach*" is defined as: "*any unauthorised processing of personal data or accidental disclosure, acquisition, sharing, use, alteration, destruction or loss of access to personal data, that compromises the confidentiality, integrity or availability of personal data.*"

This is a data-centric concept. The focus is not on how the incident occurred, but on whether personal data has been compromised in a manner that affects its confidentiality, integrity, or availability.

Once this threshold is met, the consequences are regulatory: organisations are required to notify the Data Protection Board within 72 hours of becoming aware of the breach and to inform affected Data Principals without undue delay, in accordance with the prescribed framework.

Cybersecurity Incident (IT Act / CERT-In Framework)

The concept of a cybersecurity incident emerges from the IT Act ecosystem and CERT-In directions, where the focus is on the security and functioning of systems and networks. A cybersecurity incident typically refers to events such as:

- unauthorised access to systems or accounts
- malware, ransomware, or denial-of-service attacks
- probing, scanning, or attempted intrusions
- compromise or disruption of IT infrastructure

The emphasis here is system-centric i.e. whether the confidentiality, integrity, or availability of systems has been affected or threatened. Crucially, such incidents may occur without any involvement of personal data. These incidents are subject to mandatory reporting to CERT-In, typically within 6 hours of noticing the incident, along with submission of technical details such as logs, system information, and indicators of compromise.

The distinction becomes operational when viewed through both impact and reporting obligations. A cybersecurity incident concerns the security of systems, while a personal data breach concerns the impact on personal data within those systems.

In practice, this leads to three broad scenarios:

- system impact without data impact (for example, a failed intrusion attempt);
- data impact without system compromise (for example, accidental disclosure, misdirected email, or improper sharing); and
- combined impact (for example, ransomware affecting customer databases).

Organisations must therefore assess both dimensions: whether there has been a system-level compromise, and whether that compromise has resulted in impact on personal data. The distinction also affects how and what must be reported. Organisations must be able to distinguish between operational incidents and legally reportable events.

BYTE/WISE

PART THREE

DIGITAL ASSETS

MARCH 2026

BYTE/WISE

Hong Kong Moves Toward Issuance of First Stablecoin Licences

Hong Kong is moving toward issuing its first stablecoin licences, with the Hong Kong Monetary Authority indicating that initial approvals are expected in March 2026 under its new regulatory framework for fiat-referenced stablecoins. The regime requires issuers to maintain full reserve backing, ensure par value redemption, implement segregation of client assets, and comply with prudential, governance, and AML/CFT requirements. Licensed entities will also be subject to ongoing supervision, including obligations relating to disclosure, risk management, and operational resilience.

The framework positions stablecoins as regulated payment instruments rather than unregulated crypto assets, bringing them within a financial regulatory perimeter comparable to traditional payment systems. It also reflects a broader policy objective of enabling controlled innovation while mitigating risks relating to financial stability, consumer protection, and systemic exposure.

EU Parliament Committee Examines Digital Assets Challenges

On 23 February 2026, the European Parliament's Committee on Economic and Monetary Affairs (ECON) published a draft report examining the challenges that digital assets pose for the competitiveness and integrity of the EU's financial system. The draft report includes a legislative resolution exploring how digital assets impact the financial services sector and what this means for the regulatory framework. The analysis focuses particularly on the Markets in Cryptoassets Regulation (MiCA) and the DLT Pilot Regime Regulation, which form the foundation of EU digital asset regulation.

Key Areas of Concern

- Cryptoassets: The report emphasises the need to strengthen data capabilities and supervisory dialogue on significant multi-function groups (MFGs). It calls for alignment of the MiCA policy framework for significant non-bank MFGs and highlights ongoing concerns that cryptoassets are frequently used to evade anti-money laundering, counter-terrorism financing regulations, and sanctions.

- Stablecoins: ECON calls on the European Commission to urgently propose legislation clarifying the legal position on multi-issuance of stablecoins by both EU and non-EU entities, particularly where digital stablecoins issued by different entities are fully fungible and indistinguishable. The report stresses the importance of strong prudential safeguards, robust cooperation arrangements, and enhanced crisis management protocols.

- Tokenisation: The report acknowledges that trial and error remains necessary to understand the full advantages of tokenisation. It emphasises the ongoing need to monitor vulnerabilities as the DLT Pilot Regime continues to develop.

- Interoperability: The draft highlights interoperability as crucial in digital finance. It underscores the importance of assessing legal entity identifier and verifiable legal entity identifier approaches as infrastructure-grade tools.

The draft report signals the EU's commitment to addressing emerging challenges in digital asset regulation while maintaining the competitiveness and integrity of its financial system. The recommendations aim to strengthen the existing regulatory framework under MiCA and related legislation.

UAE Approves First USD-Backed Stablecoin Under Central Bank Regulation

The Central Bank of the United Arab Emirates (CBUAE) has approved the country's first USD-backed stablecoin, marking a significant milestone in digital asset regulation. The approval was granted under the bank's Payment Token Services Regulation (PTSR).

The USDU token will be issued and managed by Universal Digital, a firm regulated by the Financial Services Regulatory Authority (FSRA) of Abu Dhabi Global Market (ADGM). The stablecoin is backed by reserves held on a 1:1 basis in safeguarded onshore accounts at 3 (three) UAE banking partners: Emirates NBD, Mashreq, and Mbank. Digital asset infrastructure firm Aquanow has been appointed as a global distribution partner, enabling institutional access to USDU outside the UAE where legally permitted.

The approval establishes a framework for transparent and efficient digital asset markets in the UAE and potentially serves as a model for other jurisdictions considering similar regulatory approaches to stablecoins.

EU ESMA Signals Application of Market Abuse Supervision to Crypto Markets

The European Securities and Markets Authority has recently signaled that crypto asset markets will be subject to market abuse-style supervision as part of the implementation of the Markets in Crypto-Assets Regulation. Supervisory focus includes detection and prevention of insider dealing, unlawful disclosure of information, and market manipulation, requiring platforms to implement transaction monitoring, surveillance systems, and reporting mechanisms.

This approach aligns crypto trading environments with standards applicable to traditional securities markets, effectively extending market conduct rules into the digital asset space. It introduces stricter expectations on trading behaviour and platform oversight, marking a shift toward treating crypto markets as regulated financial trading ecosystems rather than lightly supervised digital platforms.

UK Advances Cryptoasset Custody and Safeguarding Framework Following Consultation Phase

The UK is progressing its cryptoasset regulatory framework beyond the consultation stage, with a focus on developing rules governing custody and safeguarding of client assets. Building on earlier proposals, the next phase is expected to introduce requirements around segregation of client assets, defined custodial responsibilities, and enhanced accountability for asset protection, aligning crypto intermediaries more closely with traditional financial market standards.

This advancement signals a shift towards implementation, with regulators expected to move toward final rules and supervisory expectations in the coming phase. The focus on custody and safeguarding reflects increasing regulatory priority on asset protection and infrastructure integrity, indicating that crypto market participants should prepare for more formalised and enforceable obligations.

Digital Assets / CLAUSE IN FOCUS

TRANSACTION VALIDATION, EXECUTION RISK & IRREVERSIBILITY

Digital asset transactions differ fundamentally from traditional financial transfers. Once a transaction is broadcast and confirmed on a distributed ledger, it is typically irreversible. Errors in wallet addresses, authentication, or instruction handling cannot be undone through standard contractual or operational remedies. As a result, agreements play a central role in defining how transaction instructions are validated and where responsibility lies in the event of loss.

Given this finality, transaction clauses are structured to allocate risk at the point of execution, particularly between authentication of instructions, validation controls, and user responsibility for accuracy.

Draft Illustrative Clause

"The Service Provider shall execute transactions involving Digital Assets strictly in accordance with Client instructions received through agreed authentication mechanisms."

"Transactions recorded on distributed ledger networks are irreversible and cannot be cancelled or modified once confirmed."

"The Service Provider shall implement commercially reasonable controls to validate transaction instructions. The Service Provider shall not be liable for losses arising from incorrect or unauthorised instructions submitted through valid authentication credentials, except where resulting from its failure to comply with agreed security procedures."

"The Client shall be responsible for verifying the accuracy of transaction details prior to submission."

Negotiation Points

Client / User

- Seek stronger validation controls (e.g., transaction whitelisting, dual approvals)
- Limit liability in cases of credential compromise
- Introduce additional safeguards for high-value transactions

Service Provider

- Anchor liability to valid authentication of instructions
- Emphasise user responsibility for transaction accuracy
- Limit obligations to commercially reasonable validation measures

These clauses ensure that responsibility for transaction execution is clearly defined, particularly in environments where errors cannot be reversed once committed to the network.

Digital Assets

KNOWLEDGE
CORNER

Travel Rule

BYTE/WISE

The "*Travel Rule*" is a regulatory requirement applicable to virtual asset transfers, originating from the Financial Action Task Force (FATF) framework and increasingly implemented across jurisdictions. It imposes obligations on Virtual Asset Service Providers (VASPs) to collect, verify, and transmit information about parties involved in a transaction.

Explanation

The Travel Rule does not appear as a single defined term in most statutes. Instead, it derives from FATF Recommendation 16, which applies to wire transfers and has been extended to virtual assets.

In the context of digital assets, it requires VASPs to:

- obtain originator and beneficiary information, and
- transmit that information to the receiving VASP or financial institution

This information is required to "travel" with the transaction, hence the term Travel Rule.

Scope of Information

In practice, the obligation typically includes:

- Originator information
  - name, account details (such as wallet address), and, in some cases, additional identifiers
- Beneficiary information
  - name and destination wallet or account details

The exact scope varies by jurisdiction, but the objective remains consistent:→ to ensure traceability of transactions and accountability across intermediaries

Operational Character

Unlike many digital asset concepts, the Travel Rule is not merely classificatory; it is operational and transactional.

It applies:

- at the point of transfer, and
- requires coordination between multiple regulated entities (VASPs)

This creates practical requirements around:

- onboarding and KYC alignment,
- secure information-sharing mechanisms, and
- interoperability between platforms

Regulatory Positioning

The Travel Rule sits at the intersection of:

- anti-money laundering (AML) obligations, and
- data-sharing requirements between regulated entities

It effectively extends traditional financial compliance expectations into the digital asset ecosystem, reducing the anonymity typically associated with crypto transactions.

BYTE/WISE

PART FOUR

DARK PATTERNS

MARCH 2026

INTERFACE INTERFERENCE

The Guidelines for Prevention and Regulation of Dark Patterns, 2023 define “Interface Interference” as a design element that manipulates the user interface in ways that a) highlight certain specific information; and b)obscure other relevant information in relation to the highlighted information, so as to misdirect a user from taking an action desired by the user.

Common manifestations include:

- ❌ Designing a light-coloured or less prominent “No” option in response to a purchase or consent pop-up
- ❌ Concealing cancellation symbols or options in tiny font or less visible parts of the interface
- ❌ Changing the expected meaning of key symbols or interface elements
- ❌ Using an “X” icon that opens another advertisement or triggers another action instead of closing the pop-up
- ❌ Highlighting one option while visually downplaying other relevant choices available to the user

Interface interference is problematic because it does not merely influence attention; it manipulates the presentation of choices in a way that can impair user autonomy and distort decision-making.

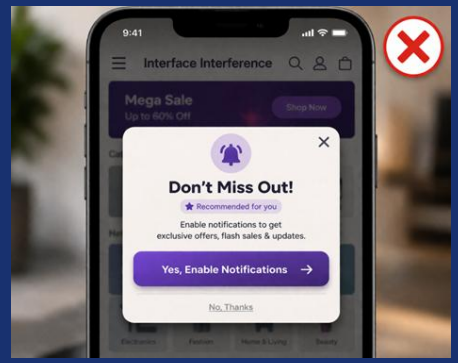

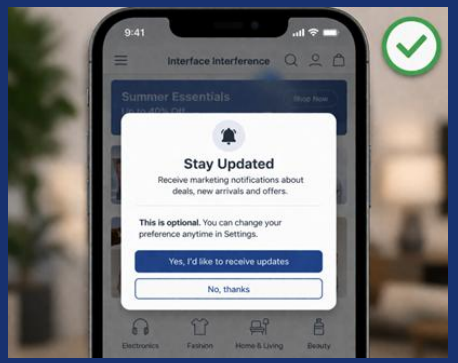

Better practice

- ✅ Ensure that acceptance, rejection, and cancellation options are equally visible and accessible
- ✅ Avoid colour, font, size, or placement choices that make one option disproportionately prominent
- ✅ Maintain consistency in the expected function of icons, buttons, and interface symbols
- ✅ Present all relevant options clearly and without visual manipulation
- ✅ Design consent, purchase, and cancellation flows in a way that supports genuine user choice

Tip for businesses

Review pop-ups, banners, checkout pages, and account settings to identify whether any design element highlights one option while obscuring another. Businesses should ensure that interface design supports clear and neutral user choice, particularly where consent, purchases, or cancellations are involved.

MARCH 2026

COMING SOON

APRIL ISSUE

ISSUE EDITORS

Partner: Naresh Pareek
Consultant: Shravan Kalluri

CONTRIBUTORS

Abhishek Nair
Jash Doshi

PUBLISHING SUPPORT

Vivek Yadav

309-10 Madhava,
C5, E Block, BKC,
Bandra East,
Mumbai 400 051