emerging tech decoded

LEX CONSULT EDITORIAL

BYTE/WISE

KEY HIGHLIGHTS

FEB 2026

Artificial Intelligence

Regulatory Updates

- EU AI Act: Initial Obligations Now in Effect
- Delhi Declaration: AI Summit Reinforces Coordinated Global Governance
- India IT Rules Amendment: AI-Generated Content Takedown Obligations
- South Korea enacts AI Basic Act

Clause in Focus - Auditability, logging & explainability clauses for High-Risk AI

Knowledge Corner - High-Risk AI Systems vs. General-Purpose AI Models with Systemic Risk

Compliance Snapshot - Post-Deployment Monitoring & Model Change Control

Privacy

Regulatory Updates

- California Attorney General Issues Largest CCPA Fine to Date
- EU Digital Omnibus Compromise May Retain Existing Definition of Personal Data
- United States FTC Enforcement Emphasizes Data Retention and Deletion Practices
- Singapore PDPC Updates Advisory on AI and Personal Data Governance

Clause in Focus - Handling Rights Requests & External Communications Clauses

Knowledge Corner - Data Principal Rights and Duties under the DPDP Act

Compliance Snapshot - Data Retention & Deletion Controls

Digital Assets

Regulatory Updates

- UK’s Cryptoasset Regulation: A Structured Step Toward Mainstream Financial Oversight
- Hong Kong Strengthens Its Crypto Rulebook
- Crypto.com Receives Conditional Approval for National Trust Bank Charter
- Stablecoin Use Cases Selected for Regulatory Sandbox Testing in UK

Clause in Focus - Custody Termination & Asset Migration Clauses

Knowledge Corner - Central Bank Digital Currencies

Dark Patterns Check– Subscription Trap

BYTE/WISE

PART ONE

Ai

FEB 2026

BYTE/WISE

Ai / REGULATORY UPDATES

EU AI Act: Initial Obligations Now in Effect

From 2 February 2025, the EU AI Act brought into force the first set of binding obligations, including the prohibition of certain unacceptable-risk AI practices (such as manipulative or exploitative AI uses) and the requirement for organizations developing or deploying AI systems to implement AI literacy measures. Providers and deployers must ensure that personnel involved in the design, deployment, operation, or oversight of AI systems have adequate understanding of system capabilities, risks, and limitations.

These obligations are now enforceable and require organizations to establish internal training, governance documentation, and role-specific awareness measures. They operate alongside the broader EU AI Act framework that must be implemented prior to placement on the market or deployment, including risk-management systems, technical documentation, logging capabilities, human-oversight mechanisms, quality-management systems, and post-market monitoring arrangements.

Delhi Declaration: AI Summit Reinforces Coordinated Global Governance

India recently hosted the February 2026 AI Summit, where participating governments adopted the Delhi Declaration, reaffirming a shared commitment to risk-based AI regulation, safety testing of advanced AI systems, and closer international cooperation on technical standards and incident information-sharing.

The Declaration emphasizes the need for institutional oversight mechanisms, cross-border coordination on AI safety, and support for capacity building to enable responsible deployment across emerging and developing AI ecosystems.

The Summit discussions also highlighted key principles reflected in India’s draft AI governance approach, including lifecycle risk management, accountability across the AI value chain, transparency and content authenticity measures, incident reporting pathways, and proportionate human oversight for high-impact systems. The Declaration was endorsed by more than 20 participating countries and international organisations, reflecting continued alignment among major AI-developing jurisdictions on coordinated oversight of advanced AI systems and the development of interoperable governance frameworks.

India IT Rules Amendment: AI-Generated Content Takedown Obligations

In February 2026, India amended the IT (Intermediary Guidelines and Digital Media Ethics Code) Rules, 2021 to introduce explicit obligations relating to synthetically generated information, including deepfakes and other AI-generated media. The amendment operationalizes disclosure and labelling requirements and strengthens intermediary due-diligence expectations around detection, moderation, and removal of synthetic content.

A key change is the introduction of accelerated takedown timelines, requiring intermediaries to remove flagged unlawful synthetic content within three hours, significantly reducing the earlier response window. The framework also requires platforms to maintain audit logs, deploy detection mechanisms, and prevent re-uploads of previously removed synthetic media.

This development builds on the draft amendment discussed in the *November 2025 ByteWise AI Regulatory Updates*, which first introduced synthetically generated information as a compliance category and signaled enhanced platform responsibility.

South Korea enacts AI Basic Act

In February 2026, South Korea enacted its AI Basic Act, establishing a comprehensive national framework for AI governance. The law introduces a risk-based regulatory structure, including obligations relating to safety, transparency, and oversight for high-impact AI systems, alongside requirements for testing, monitoring, and incident management. It also provides for government supervision of advanced AI development and deployment, including the designation of responsible authorities and the creation of institutional mechanisms for AI evaluation.

Compared with other global approaches, the Act adopts a structure broadly aligned with the EU AI Act's risk-classification model while retaining a stronger policy focus on innovation enablement and national AI competitiveness. Unlike the EU framework, which emphasizes detailed conformity assessment obligations, South Korea's approach places greater emphasis on government coordination, industry support measures, and phased operational guidance. The framework also differs from the more sector-driven approach seen in the United States, positioning the Act as a comprehensive national AI law that combines regulatory oversight with industrial strategy.

Ai / CLAUSE IN FOCUS

AUDITABILITY, LOGGING & EXPLAINABILITY FOR HIGH-RISK AI

Regulatory approaches to AI governance emphasize auditability, logging, and explainability as foundational accountability requirements, particularly for high-risk AI systems. Frameworks such as the EU AI Act position these capabilities as essential to ensuring traceability of system behavior, enabling effective human oversight, and supporting regulatory review.

For high-risk AI, compliance is not limited to outcome-based performance; organizations must be able to reconstruct to a sufficient degree as to how a system operated, what data it relied on, and why a particular output was generated. Contractual clauses addressing auditability and explainability therefore play a critical role in translating regulatory obligations into enforceable operational duties between providers and deployers.

In practice, agreements are increasingly incorporating provisions that:

- require technical and organizational logging mechanisms to record system inputs, outputs, model versions, and material parameters;
- enable regulatory and contractual audits to assess compliance with applicable AI and data protection laws;
- ensure explainability of outputs to the extent necessary to support human oversight, contestability, and regulatory scrutiny; and
- allocate responsibility based on control over model design, training, deployment, and updates.

Draft Illustrative Clause

"The Provider shall implement and maintain appropriate technical and organizational measures to ensure that the AI System is auditable and traceable, including maintaining logs of material inputs, outputs, system configurations, and model versions, for the period required under applicable law. The Provider shall 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

- Seek audit and access rights sufficient to meet regulatory obligations under the EU AI Act and related sectoral laws, especially in case of high-risk AI systems.
- Require clarity on the scope, retention period, and accessibility of logs, particularly where outputs may be challenged or reviewed.
- Ensure explainability commitments are adequate to support human review, user communications, and regulatory inquiries.

Provider

- Limit audit rights to what is necessary for compliance, protecting trade secrets and proprietary model details.
- Define explainability by reference to reasonable and technically feasible measures, avoiding guarantees of full model transparency.
- Allocate responsibility clearly where explainability or logging is impacted by downstream integrations, fine-tuning, or use-case modifications by the Deployer.

Well-drafted clauses should align regulatory expectations with technical feasibility, ensuring transparency and accountability without exposing disproportionate commercial or operational risk.

BYTE/WISE

Ai / COMPLIANCE SNAPSHOT

Post-Deployment Monitoring & Model Change Control

AI compliance does not end at deployment. Both the EU framework and emerging Indian guidance treat AI governance as a continuous lifecycle obligation, requiring monitoring of outputs, incident response, and controlled modification of models after release.

Under the EU regime of the European Union, providers and deployers of high-risk AI systems must operate a post-market monitoring system designed to identify performance issues, emerging risks, and misuse. Monitoring must be proportionate to risk and include logging, corrective action, and regulator notification where serious incidents occur. The European Commission AI oversight framework also expects providers of advanced foundation models to track downstream effects and maintain technical documentation reflecting model updates.

What Monitoring Should Cover

- Continuous performance and accuracy checks in real-world use
- Detection of harmful or anomalous outputs
- Human-oversight escalation procedures
- Incident documentation and corrective measures
- User complaints and regulator-triggered reviews

In addition to the above, AI Model change control is equally critical, as updates to training data, parameters, guardrails, or integrations may alter system behavior. Organizations should therefore maintain a formal change-management process that includes documented approval prior to deployment of updates, regression and safety testing, version control with audit trails, and rollback mechanisms where risks or unintended outcomes emerge.

Quick Self-Check for Organizations

- Do we track real-world outputs and document anomalies?
- Is there a defined trigger for suspending or restricting system use?
- Are model updates risk-assessed before release?
- Can we demonstrate when behavior changed and why?
- Is responsibility assigned between provider, integrator, and deployer?

Post-deployment monitoring and model change control are integral components of AI compliance frameworks. Organizations should ensure that monitoring outputs and change decisions are documented, reviewable, and capable of demonstrating that system behavior remains consistent with regulatory and operational expectations over time.

Ai

KNOWLEDGE
CORNER

High-Risk AI Systems vs. General-Purpose AI Models with Systemic Risk

The EU AI framework distinguishes between model capability risk and deployment risk. A general-purpose AI model is not automatically high-risk; classification depends on how the model is used, while certain models may independently attract systemic-risk obligations due to their scale and high-impact capabilities. As a result, the same model may remain within the GPAI framework in general use but become high-risk when integrated into certain use-cases such as creditworthiness assessment or recruitment screening or fall within systemic-risk obligations at the provider level.

Defined Terms

General-Purpose AI Model (Article 3, EU AI Act) means an AI model, including where such an AI model is trained with a large amount of data using self-supervision at scale, that displays significant generality and is capable of competently performing a wide range of distinct tasks regardless of the way the model is placed on the market and that can be integrated into a variety of downstream systems or applications, except AI models that are used for research, development or prototyping activities before they are placed on the market.

Systemic Risk (Article 3, EU AI Act) means a risk that is specific to the high-impact capabilities of general-purpose AI models, having a significant impact on the Union market due to their reach, or due to actual or reasonably foreseeable negative effects on public health, safety, public security, fundamental rights, or the society as a whole, that can be propagated at scale across the value chain.

GPAI Model with Systemic Risk

A subset of GPAI models whose scale, compute power, or widespread deployment creates significant societal, economic, or security risk. These models attract enhanced obligations such as evaluation testing, incident reporting, and risk-mitigation controls at the model-provider level.

High-Risk AI System (Article 6, EU AI Act) - An AI system shall be considered high-risk where it is:

(a) intended to be used as a safety component of a product, or is itself a product, covered by Union harmonization legislation listed in Annex I, and is required to undergo third-party conformity assessment; or

(b) listed in Annex III, where the AI system is used in sensitive areas such as biometric identification, employment, creditworthiness assessment, education, access to essential services, migration, law enforcement, or administration of justice.

Explanation

Legal obligations are triggered differently under the EU AI Act depending on whether risk arises from the capabilities of a model or from the context in which an AI system is deployed.

A general-purpose AI model (GPAI) is regulated primarily at the model level, with obligations focused on preparing and maintaining technical documentation, ensuring transparency regarding training data and model capabilities, and providing information necessary for downstream providers and deployers to integrate the model lawfully.

Where the scale, compute capacity, or reach of a model creates the potential for significant impact across downstream uses, the model may qualify as a GPAI model with systemic risk, triggering enhanced provider obligations including model evaluation and testing, implementation of risk-mitigation measures, incident reporting, cybersecurity safeguards, and ongoing monitoring of downstream impact across the value chain.

A high-risk AI system is determined by deployment context rather than inherent model capability. When an AI system is used as a safety component of regulated products or in sensitive decision-making contexts focus on system-level governance, including risk-management systems, data governance controls, logging, human oversight, conformity assessment, post-market monitoring, and operational documentation.

Example

A large foundation model capable of language generation across multiple domains may qualify as a general-purpose AI model and, if its scale and reach create broad societal impact, may be classified as a GPAI model with systemic risk, triggering provider obligations such as evaluation testing, incident reporting, and risk-mitigation measures.

If the same model is integrated into a recruitment screening platform or a creditworthiness assessment tool, the resulting application becomes a high-risk AI system because the classification is driven by the use-case. At that stage, deployers must implement conformity assessment, human oversight, monitoring, and documentation requirements, regardless of whether the underlying model itself is systemic risk classified.

Most jurisdictions, including India, do not yet apply this dual classification and instead regulate AI primarily through specific use-based triggers such as harmful outcomes, platform responsibility, or sectoral impact. This approach overlaps with the EU's high-risk concept to the extent that regulators tend to focus on sensitive applications (for example, finance, employment, healthcare, public services), but it differs in structure: outside the EU, obligations are still often imposed through sectoral rules, regulator guidance, or targeted platform/content measures rather than a single statutory taxonomy that separately governs powerful models and high-risk deployments.

BYTE/WISE

PART TWO

PRIVACY

FEB 2026

Privacy / REGULATORY UPDATES

California Attorney General Issues Largest CCPA Fine to Date

The California Attorney General announced the largest settlement to date under the California Consumer Privacy Act (CCPA), reinforcing statutory requirements relating to the right to opt out of the sale or sharing of personal information and the obligation to operationalize consumer preference signals.

The enforcement action focused on whether businesses effectively implemented opt-out requests across account environments and services, rather than limiting compliance to interface-level controls.

Key legal takeaways

- Businesses must ensure opt-out requests are honored across devices, services, and identity contexts where reasonably linkable to a consumer.
- Automated opt-out signals must be treated as valid consumer requests and applied consistently.
- Organizations must ensure that opt-out preferences are communicated to service providers, contractors, and third parties.
- Regulators expect demonstrable evidence that preference orchestration is implemented at the system level, including monitoring and testing.

The settlement is significant because it clarifies that regulators will assess whether user rights function effectively across an organization's systems, rather than accepting interface-level controls as sufficient compliance. As the DPDP framework becomes operational, similar issues should not be treated as low-risk implementation details.

EU Digital Omnibus Compromise May Retain Existing Definition of Personal Data

Following the European Commission's November 2025 Digital Omnibus proposals, which sought to clarify the scope of "personal data" through a more contextual, entity-specific assessment of identifiability, a leaked Member State compromise text indicates that the proposed revision to the GDPR definition of personal data may be removed.

The original proposal aimed to recognize that whether data qualifies as personal data could depend on an organization's realistic ability to re-identify individuals, including access to keys, auxiliary datasets, technical capabilities, and organizational constraints, with particular implications for pseudonymized data and AI development contexts. The reported compromise suggests Member States are cautious about altering the core definition, reflecting concerns about fragmentation, legal uncertainty, and potential impact on data subject rights.

If confirmed, the development signals that while regulators continue exploring contextual approaches to identifiability, the existing GDPR definition of personal data is likely to remain the baseline. For AI, this is significant because organizations cannot rely on narrower interpretations of identifiability to exclude training data or model inputs from GDPR scope, reinforcing the need for robust documentation of re-identification risk and safeguards in AI workflows.

BYTE/WISE

United States FTC Enforcement Emphasizes Data Retention and Deletion Practices

Recent FTC enforcement activity and regulatory commentary in early 2026 highlighted retention and deletion practices as a key area of scrutiny, particularly where organizations retained personal data longer than disclosed or failed to operationalize deletion commitments reflected in privacy notices. The action highlights that retention minimization, effective deletion execution, and demonstrable internal controls are central to assessing whether organizations meet fairness and accountability expectations.

Regulatory focus:

- Increased scrutiny of discrepancies between stated retention policies and operational data practices.
- Expectations that deletion controls extend across backups, vendors, and legacy systems.
- Greater focus on documentation evidencing how retention schedules are defined, implemented, and periodically reviewed.

As DPDP implementation progresses in India, regulatory scrutiny is likely to shift from paper-based compliance to verification of actual operational practice, particularly in relation to retention limitation and deletion obligations. Organizations should therefore ensure that retention schedules, deletion controls, and supporting documentation demonstrate real execution of lifecycle commitments rather than policy-level intent alone.

Singapore PDPC Updates Advisory on AI and Personal Data Governance

Singapore's Personal Data Protection Commission (PDPC) released updated advisory material in February 2026 clarifying governance expectations where organizations deploy AI systems involving personal data, with a stronger emphasis on embedding AI oversight within existing privacy management programmes rather than treating it as a standalone compliance function. The update highlights organizational accountability for automated decision-making, structured documentation of AI use, and integration of risk management across the AI lifecycle.

Key governance expectations highlighted include:

- Documentation of decision-making logic and assignment of organizational responsibility for automated processing outcomes.
- Alignment of AI risk assessments with existing data protection impact assessments and broader governance frameworks.
- Expectations that monitoring, explainability measures, and incident response procedures are incorporated into routine privacy controls.
- Greater emphasis on demonstrating oversight of third-party AI tools and downstream uses involving personal data.

The update reinforces the convergence of AI governance and privacy compliance, signaling that regulators increasingly expect AI controls to operate as part of established accountability, risk management, and lifecycle governance frameworks.

Privacy/ CLAUSE IN FOCUS

HANDLING RIGHTS REQUESTS & EXTERNAL COMMUNICATIONS CLAUSES

Privacy frameworks such as the GDPR and the DPDP Act expect organizations to maintain clear accountability for how requests, complaints, and related external communications are handled across internal teams and service providers.

Contractual provisions therefore should focus on ensuring that requests and communications received by processors or vendors are promptly notified to the controller, that responses are coordinated rather than fragmented, and that assistance is available to support investigation, response, and escalation where legal or regulatory risk arises.

Beyond notification, controllers must retain effective oversight and decision-making authority over the substance and timing of responses. Vendor arrangements should clearly prohibit independent or unilateral communications with data subjects or regulators without prior authorization, except where legally mandated

Draft Illustrative Clause

"The Processor shall promptly notify the Controller of any Data Subject or Data Principal request, complaint, grievance, or communication from a supervisory authority relating to the Services or Personal Data. The Processor shall not respond directly except on documented instructions of the Controller or where required by applicable law and shall provide reasonable assistance to enable the Controller to investigate, respond to, and document such matters, including supporting escalation where regulatory, legal, or operational risk is identified."

Negotiation Points

Controller/Data Fiduciary

- Define strict notification timelines and escalation triggers.
- Ensure coverage of regulator communications and informal complaints.
- Require cooperation across systems, logs, and third parties where feasible.

Service Provider/Data Processor

- Limit assistance to information available within the processor's control.
- Clarify responsibility for identity verification and legal assessment.
- Ensure response obligations do not require disclosure of proprietary information beyond what is necessary for compliance.

Clauses that capture rights requests, external communications, and escalation pathways enable organizations to demonstrate accountability while coordinating response responsibilities across complex vendor ecosystems.

BYTE/WISE

Privacy/ COMPLIANCE SNAPSHOT

Data Retention and Deletion Controls

Under the DPDP framework, retention and deletion operate as lifecycle controls requiring organizations to retain personal data only while the specified purpose continues or where retention is required under law. Deletion decisions are commonly triggered by consent withdrawal, account closure, or inactivity, but must be balanced with regulatory expectations, including minimum one-year retention of logs and processing records and scenarios where organizations may retain certain records for longer risk, audit, or dispute periods (for example up to three years).

Retention and Deletion Controls

- Purpose-linked retention schedules: Organizations should map categories of personal data to specified purposes and define retention periods that reflect legal requirements, contractual obligations, and operational need, ensuring retention is intentional rather than default.

- Trigger-based deletion workflows: Events such as consent withdrawal, completion of a transaction, inactivity, or termination of a service should initiate defined workflows for erasure, anonymization, or restricted storage decisions.

- End-to-end deletion execution: Deletion must extend beyond primary systems to processors and vendors, with controls to prevent re-ingestion of deleted data and to maintain consistency across environments.

- Auditability and minimum log retention: Organizations should maintain audit trails documenting when deletion occurred and why data continues to be retained, recognizing minimum one-year retention expectations for logs and processing records to evidence compliance.

- Backup and extended retention management: Where immediate deletion is not technically feasible (for example backups or archived environments), organizations should implement eventual erasure controls and document justified multi-year retention windows (such as three-year retention for dispute or regulatory review).

Quick Self-Check for Organisations

- Do we have defined retention periods linked to purpose?

- Are one-year log retention and extended retention scenarios documented?

- Is deletion triggered when consent is withdrawn or purpose ends?

- Can we demonstrate deletion across processors and vendors?

- Do we document why data is retained beyond operational need?

Retention and deletion controls demonstrate whether organizations operationalize purpose limitation in practice rather than in policy. Regulators increasingly assess whether retention decisions are documented, justified, and consistently executed across systems and vendors.

Privacy

KNOWLEDGE
CORNER

Data Principal Rights and Duties under the DPDP Act

BYTE/WISE

Data Principal Rights and Duties under the DPDP Act

The Digital Personal Data Protection Act, 2023 establishes a structured interaction model through which individuals (Data Principals) may understand, control, and challenge how their personal data is processed. These rights operate through a request-based framework requiring organizations to provide accessible channels, defined workflows, and documented responses for handling requests and grievances within prescribed timelines.

Alongside rights, the Act introduces duties on Data Principals to ensure responsible exercise of these mechanisms. Individuals are expected to provide authentic information, avoid impersonation, and not file false or frivolous complaints. The Act provides for monetary penalties of up to ₹10,000 for breach of these duties, reinforcing that rights are intended to be exercised in good faith.

Rights of Data Principals

The Act recognizes core rights relating to transparency, control, and escalation, each carrying operational implications for organizations.

- Access (Section 11): A Data Principal may request a summary of personal data processed and details of sharing. Organizations must maintain data discovery capability, coordinated response workflows, and records of disclosures.

- Correction and erasure (Section 12): Individuals may request correction, updating, completion, or erasure of personal data, subject to legal retention requirements. Organizations must validate requests, assess purpose limitation, and execute deletion or restriction across systems and processors where feasible.

- Withdrawal of consent (Section 6(4)): Consent may be withdrawn at any time with comparable ease. This requires consent tracking, suppression controls, and mechanisms to ensure downstream processors cease processing within a reasonable time.

- Grievance redressal (Section 13): Data Principals must be provided readily available grievance mechanisms. Under the DPDP Rules, organizations are expected to operate grievance systems capable of responding within a defined period not exceeding 90 days, supported by escalation pathways and documented reasoning.

- Nomination (Section 14): Individuals may nominate another person to exercise rights in case of death or incapacity, requiring organizations to implement verification procedures and lifecycle controls for nominee requests.

In practice, rights handling depends on intake channels, identity verification, request classification, internal review, and record-keeping sufficient to demonstrate consistent and defensible decision-making.

Example: A user of a digital wallet requests a copy of the personal data held and asks for deletion of an inactive account. The organization verifies identity, retrieves account data from multiple systems, checks whether financial record retention laws apply, deletes data where permissible, and documents the response. If the user disagrees with the outcome, the matter may move into the grievance process.

BYTE/WISE

Duties of Data Principals

The duties of Data Principals (Section 15) shape how organizations operationalize rights workflows. Because individuals must provide authentic information and avoid impersonation, organizations may require identifiers, authentication steps, or supporting information before acting on requests. The duty not to file false or frivolous grievances allows organizations to define thresholds for repetitive or abusive submissions and to document reasons where requests are declined or require clarification.

These duties also support proportionality. For example, where a correction request lacks sufficient detail, organizations may pause processing until adequate information is provided. Similarly, verification measures help ensure that deletion or access requests are not executed on behalf of an unauthorized person. The possibility of monetary penalties up to ₹10,000 reinforces that rights mechanisms are intended for legitimate use and supports organizational controls designed to prevent misuse.

Example: An individual submits a deletion request using an email address that does not match the account on record. The organization requests additional verification before acting. If the individual provides inaccurate information or repeatedly files unsupported complaints after clarification, the organization may classify the request as incomplete or potentially frivolous and record its decision. Where such conduct amounts to impersonation, provision of false information, or submission of frivolous grievances, the Data Principal may be exposed to monetary penalties of up to ₹10,000 under the duties framework.

As is the case with major data protection regimes, the DPDP framework structures rights and duties as mutually reinforcing elements of a balanced governance model. Data Principals are granted enforceable rights to access, correct, and erase their personal data, advancing transparency and individual autonomy in the digital sphere, while the statute presumes that these rights will be exercised honestly and in good faith. Organizations, in turn, must operationalize these rights through clear, time-bound, and documented processes, while retaining the ability to apply proportionate verification measures and safeguards to protect data integrity and prevent misuse. The result is a cooperative architecture in which empowerment and accountability operate in tandem to sustain procedural fairness and trust.

However, the DPDP Act goes further than most major regimes by expressly codifying statutory duties on Data Principals themselves, backed by the possibility of a monetary penalty of up to ₹10,000 for impersonation, furnishing false information, or filing frivolous grievances. While frameworks such as the GDPR permit refusal or reasonable fees for manifestly unfounded or excessive requests, they primarily regulate organizational conduct. By contrast, the DPDP Act embeds reciprocity directly into the statutory design, formalizing individual responsibility within the compliance ecosystem rather than leaving it to procedural discretion

BYTE/WISE

PART THREE

DIGITAL ASSETS

FEB 2026

BYTE/WISE

UK's Cryptoasset Regulation: A Structured Step Toward Mainstream Financial Oversight

In late 2025, the UK Government laid before Parliament the Statutory Instrument under the Financial Services and Markets Act 2000 (Cryptoassets) Regulation 2025 ("**SI**"). This marks a major milestone in the UK's long-awaited crypto regulatory framework and reflects the UK Government's intent to bring crypto assets firmly within the mainstream financial system.

The SI introduces a structured and comprehensive regime. A key feature is the "designated activities regime." Public offers of crypto-assets and admissions to trading on crypto platforms are now classified as designated activities. This means firms must comply with rules to be framed by the Financial Conduct Authority (FCA), even if they are not otherwise authorized financial institutions. This significantly expands the regulatory perimeter.

Disclosure is another central pillar. Issuers and other specified persons will be required to publish a qualifying crypto asset disclosure document outlining the asset's features, governance, risks and potential conflicts of interest. Importantly, liability provisions have been introduced wherein investors can claim compensation for losses arising from misleading statements or omissions in these documents. This aligns crypto markets more closely with traditional securities regulation.

The SI also introduces a market abuse regime covering insider dealing, unlawful disclosure of inside information and market manipulation in relation to qualifying crypto assets. Firms will need systems and controls to detect and prevent abuse. Stablecoins are specifically recognized as a regulated activity, requiring FCA authorisation and enhanced disclosures, particularly around value stabilisation mechanisms.

Hong Kong Strengthens Its Crypto Rulebook

On 24 December 2025, Hong Kong's Financial Services and the Treasury Bureau (FSTB) and the Securities and Futures Commission (SFC) released their consultation conclusions on new licensing regimes for virtual asset (VA) dealers and custodians. The message is clear: Hong Kong is moving decisively to build a comprehensive and credible regulatory framework for crypto businesses.

Under the proposed structure, both VA dealers (platforms and intermediaries facilitating trades) and VA custodians (entities safeguarding digital assets) will require dedicated licences. These frameworks are aimed at strengthening investor protection, ensuring proper segregation and safekeeping of assets, and raising operational and governance standards across the sector. For global investors, this signals greater accountability and transparency.

Regulators have also launched a further consultation on licensing for VA advisory and VA asset management service providers. Following the principle of "same business, same risks, same rules," Hong Kong intends to align crypto advisory and asset management oversight with the standards applicable to traditional securities markets. This approach reduces regulatory arbitrage and ensures that firms offering similar risk profiles are supervised consistently.

These steps complete the core pillars of Hong Kong's virtual asset regime. By regulating trading, custody, advisory, and asset management under clear licensing frameworks, Hong Kong is positioning itself as a trusted international hub for responsible digital asset innovation.

BYTE/WISE

Crypto.com Receives Conditional Approval for National Trust Bank Charter

In February 2026, Crypto.com received conditional approval from the U.S. Office of the Comptroller of the Currency (OCC) for a national trust bank charter, enabling the firm to pursue digital asset custody services under federal supervision, subject to supervisory conditions. National trust bank charters permit institutions to operate as regulated fiduciary and custody providers without engaging in deposit-taking or lending, placing crypto custody activities within an established supervisory framework.

The approval follows a broader trend of digital asset firms seeking trust bank structures to support institutional custody, tokenisation, and settlement services while aligning with regulatory expectations relating to governance, operational resilience, safeguarding of client assets, and third-party risk management. The development also reflects increased use of existing banking licences as a pathway for crypto firms to expand regulated offerings in the United States.

Licensing pathways that bring crypto custody within bank-style supervision are accelerating globally, signalling that institutional-grade governance, asset segregation, and supervisory readiness are becoming baseline requirements for digital asset service providers.

Stablecoin Use Cases Selected for Regulatory Sandbox Testing in UK

In February 2026, UK authorities selected firms, including Revolut, to test stablecoin payment use cases within the regulatory sandbox as part of the government's broader digital asset and payments reform agenda. The sandbox enables firms to pilot issuance, payment, and settlement models under supervisory oversight while regulators evaluate redemption mechanisms, safeguarding arrangements, operational resilience, and interactions with existing payments and e-money frameworks.

The initiative builds on legislative proposals to bring fiat-backed stablecoins within the UK financial regulatory perimeter and reflects ongoing coordination between HM Treasury, the Financial Conduct Authority (FCA), and the Bank of England on the design of the stablecoin regime. Sandbox testing is expected to inform future requirements relating to reserve backing, governance, consumer protection, and the role of stablecoins within payment system infrastructure.

Sandbox-based implementation of stablecoin frameworks is accelerating globally, signalling that testing of issuance models, safeguarding structures, and payment integration is becoming a core step in developing digital asset regulation.

Digital Assets / CLAUSE IN FOCUS

CUSTODY TERMINATION & ASSET MIGRATION CLAUSES

Digital asset custody agreements ("Custody Agreement") are contracts between a client (which can be an exchange, broker, or a private individual) and a custodian under which the custodian safeguards the client's cryptoassets, typically by securely holding and managing the private keys required to access and transfer them. The Custody Agreement sets out the process by which a client may request a transfer or withdrawal of cryptoassets.

Set out below is a sample clause from a Custody Agreement that explains how a client can withdraw cryptoassets, including how instructions must be submitted and verified. It also outlines when a custodian may refuse or delay a transaction, for example, due to sanctions, regulatory non-compliance, or risk concerns. These provisions typically cover safeguards such as whitelisted wallet addresses, transaction limits, and how responsibility is allocated in cases involving compromised credentials or erroneous instructions

Draft Illustrative Clause

1. INSTRUCTIONS FOR WITHDRAWAL OF CRYPTOASSETS

1.1 Authority to act on instructions

The Client authorises the Custodian to act on any instructions regarding the transfer or withdrawal of cryptoassets from any wallet or address maintained by the Client under this Custody Agreement ("Instructions") that the Custodian receives in accordance with the terms of this Custody Agreement.

1.2 Prescribed channels and irrevocability

All Instructions shall be submitted using the channels, processes and security requirements prescribed by the Custodian from time to time (each a "Channel"). Instructions given through a Channel in accordance with this Custody Agreement will be deemed received when actually received by the Custodian and will be irrevocable from that point, subject to the Custodian's rights under this Custody Agreement to reject, suspend or delay execution.

1.3 Right to refuse or delay instructions

The Custodian may refuse to carry out, or may suspend execution of, any Instruction where it reasonably considers that (i) the Instruction is incomplete, unclear, not genuine, or not submitted via a prescribed Channel; (ii) acting on the Instruction would breach, or might reasonably be expected to breach, any applicable law, regulation, sanctions requirement, order or request of a government authority; or (iii) the Instruction would contravene any notified whitelisting requirement, limit or other risk parameter, or any term of this Custody Agreement. The Custodian shall, to the extent permitted by applicable law, notify the Client of such refusal or suspension and the general reason for it.

1.4 Material adverse event and withdrawal right

Without limitation to the foregoing, if the Custodian reasonably determines that a material adverse event affecting either party or the provision of the services has occurred under this Custody Agreement and, on that basis, declines to comply with a withdrawal Instruction, it shall, to the extent permitted by applicable law, promptly notify the Client and allow a reasonable period for the parties to discuss in good faith possible steps to address that event. If the event is not resolved within such period, then, subject to applicable law and any binding order of a government authority, the Custodian shall take reasonable steps within a further reasonable period to enable the Client to withdraw or transfer all cryptoassets held for the Client under this Custody Agreement to one or more addresses or custodians specified by the Client that meet any applicable whitelisting or screening requirements. For the avoidance of doubt, and subject to applicable law, the occurrence or non resolution of a material adverse event shall not, of itself, be used as grounds to permanently withhold or confiscate the Client's assets.

1.5 Reliance on verified instructions

Any Instruction that the Custodian verifies in accordance with this Custody Agreement and the applicable access conditions will be conclusively treated, as between the Custodian and the Client, as having been validly given by the Client or an authorised person appointed by the Client under this Custody Agreement, and the Custodian shall be entitled to rely on and act upon it without further enquiry as to authority or identity, except where the loss arises from the Custodian's fraud, gross negligence or wilful default as determined in accordance with this Custody Agreement.

Key Negotiation Points

Clients:

- Limit the Custodian's discretion to refuse or delay Instructions to clearly defined and reasonable grounds.
- Require prompt written notice and reasons for any refusal or suspension (subject to applicable law).
- Narrow and clearly define what constitutes a material adverse event under the Custody Agreement.
- Ensure suspensions are temporary and include a defined pathway to withdraw the cryptoassets.
- Ensure the Custodian remains liable for losses caused by security failures or failure to follow agreed controls.

Custodians:

- Retain flexibility to refuse Instructions for sanctions, regulatory, AML, fraud, or cybersecurity concerns.
- Ensure Instructions are final once verified through prescribed Channels.
- Mandate exclusive use of secure systems and multi-factor authentication.
- Confine liability to fraud, gross negligence, or wilful misconduct.
- Maintain express rights to freeze or restrict withdrawals where required by applicable law or regulatory direction

Digital Assets

KNOWLEDGE
CORNER

Central Bank Digital Currencies

BYTE/WISE

A Central Bank Digital Currency or CBDC is a digital form of a country's sovereign currency issued and regulated by its central bank. In this article, Central Bank Digital Currency is referred to as CBDC and the Central Bank Digital Currency issued for public use is referred to as Retail CBDC. Unlike private cryptocurrencies such as Bitcoin or Ethereum, a CBDC is fully backed by the monetary authority and carries the same legal status as physical cash.

As economies become increasingly digital, central banks are examining CBDCs as a way to modernize payment systems while retaining regulatory oversight and monetary control.

What is a CBDC

A CBDC is essentially digital cash. It represents a direct claim on the central bank and is denominated in the national currency. This means it maintains price stability and does not carry the volatility typically associated with cryptocurrencies.

CBDCs are generally categorized into two types. Retail CBDC is intended for use by individuals and businesses for everyday transactions. Wholesale CBDC is designed for use by banks and financial institutions for interbank settlements and large value transactions. In this article, Wholesale CBDC refers to such institutional use cases.

Global Developments

Several jurisdictions are actively testing or implementing CBDCs. India, through the Reserve Bank of India, has launched pilot programs for the Digital Rupee for both retail and wholesale segments. China has also conducted large scale trials of its digital yuan.

These initiatives aim to enhance payment efficiency, reduce settlement time, lower costs associated with physical cash management, and improve financial inclusion. For emerging markets in particular, CBDCs may provide broader access to formal financial systems.

Business Implications

For businesses, CBDCs could streamline payment processes. Transactions may settle faster, with reduced reliance on intermediaries. This can improve cash flow management and operational efficiency.

CBDCs may also increase transparency in transactions, which could support compliance and reporting requirements. Governments may benefit from better traceability, potentially reducing tax leakage and informal cash based transactions.

However, there are important considerations. Data privacy and cybersecurity remain key concerns. Central banks must ensure that digital infrastructure is secure and resilient. Additionally, if a significant portion of deposits shifts from commercial banks to central bank wallets, it may impact liquidity within the banking system.

The Road Ahead

CBDCs are not expected to replace physical cash in the near term. Instead, they are likely to coexist with existing payment systems, including bank transfers and digital wallets. The primary objective is to provide a secure, efficient, and regulated digital alternative to cash.

As pilot programs evolve, regulatory clarity and technological robustness will determine the pace of adoption. For businesses and financial institutions, monitoring CBDC developments will be important, as they represent a significant step in the evolution of sovereign money in a digital economy.

BYTE/WISE

PART FOUR

DARK PATTERNS

FEB 2026

BYTE/WISE

SUBSCRIPTION TRAP

The CCPA Guidelines for Prevention and Regulation of Dark Patterns, 2023 define “Subscription Trap” as practices that make it easy for users to subscribe to a service but difficult to cancel, pause, or opt out. This includes design choices that delay, obscure, or complicate cancellation, renewal control, or withdrawal of consent for recurring payments.

Subscription traps are recognized as a dark pattern and an unfair trade practice under the Consumer Protection Act, 2019, particularly where users are automatically renewed without clear notice, cancellation paths are hidden or multi-step, or retention nudges are designed to discourage exit.

Common manifestations include:

- ❌ Hiding cancellation options within account settings or using unclear labels for termination actions.
- ❌ Introducing retention prompts, discounts, or warnings designed to delay or discourage cancellation.
- ❌ Auto-renewing subscriptions without prominent reminders before billing cycles.
- ❌ Requiring users to complete surveys or navigate repeated confirmation screens before cancellation.

Subscription traps are problematic not because recurring billing is impermissible, but because design asymmetry between sign-up and cancellation impairs user control, dilutes meaningful choice over ongoing payments, creates unintended financial commitments, and exploits inertia and friction costs.

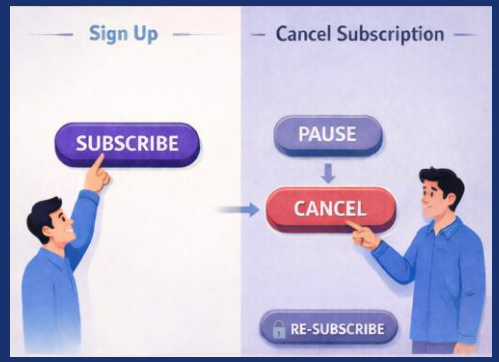

Better practice

- ✅ Provide cancellation pathways that are as simple and accessible as sign-up flows
- ✅ Offer clear pre-renewal reminders with timing that allows users to act before billing
- ✅ Present cancellation, pause, and plan-change options prominently within account interfaces
- ✅ Avoid retention nudges that obscure or delay user exit from recurring services

Tip for businesses

Review subscription journeys to ensure cancellation, pause, and renewal controls are easily discoverable and executable within the same interface used for sign-up. Document renewal notifications, cancellation steps, and retention flows to demonstrate that recurring billing is transparent, user-controlled, and free from interface friction.

COMING SOON

MARCH 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