Data Protection by Design, Cryptography by Default

Building the Quantum Resilient, Compliance-Ready Enterprise

2026

Table of Contents

Executive Summary

Executive Summary

India's digital economy has crossed a structural turning point. The Digital Personal Data Protection Act, 2023, established the substantive rights and obligations for handling digital personal data. The Digital Personal Data Protection Rules, 2025, notified by MeitY (dated 13 November 2025), now operationalise the Act with concrete expectations for notice, consent, security safeguards, breach reporting and cross-border data flows.
For large enterprises and digital platforms, this is not a marginal regulatory adjustment. It redefines how business models, technology architectures and operating processes must be designed:

- Data Fiduciaries are required to implement "reasonable security safeguards", document and enforce retention and deletion practices, and provide verifiable mechanisms for consent, withdrawal and data principal rights.
- Significant Data Fiduciaries (SDFs) carry additional expectations regarding audits, data protection impact assessments, and the appointment of key roles, such as Data Protection Officers.
- Enforcement timelines, breach-notification rules, and sectoral expectations are tightening, with public discourse already highlighting perceived gaps and calling for further clarity in sectors such as telecom and media.

At the same time, the cryptographic foundations that underpin digital trust are under unprecedented pressure:

- The U.S. National Institute of Standards and Technology (NIST) has finalised three post-quantum cryptography (PQC) standards: ML-KEM, ML-DSA and SLH-DSA. Several countries have released resiliency guidelines, urging organisations to start planning migrations.
- These guidelines reflect a consensus that public-key schemes such as RSA and ECC will not withstand large-scale quantum attacks. That long-lived sensitive data is already at risk from "harvest-now, decrypt-later" adversaries.

This paper outlines a strategic blueprint for Indian and global enterprises seeking to move beyond mere compliance and establish cryptography-first, quantum-resilient data protection. The core propositions are:

1. Cryptography is the enforcement layer of DPDP compliance. It is one of the core controls that travels with the data across infrastructure boundaries, can be objectively measured, and can provide audit-grade evidence that reasonable safeguards exist.
2. "Data Protection by Design" must be interpreted as "Cryptography by Default". Without encryption, tokenisation and robust key management at design time, "by design" remains rhetoric.
3. Compliance readiness is an architectural property, not a legal deliverable. The ability to prove compliance depends on how well cryptographic controls, logs and key lifecycles are engineered and governed.
4. Quantum resiliency is a now-task, not a later-task. The time lag between standardisation and full replacement of legacy cryptography is measured in years; organisations that do not start building crypto agility now are effectively accepting future technical debt and regulatory exposure.

The paper is written deliberately from an enterprise leader's perspective, focusing on what boards, CEOs, CIOs, CISOs, and Chief Data Officers should expect from their teams, and how they can steer the transition from fragmented security controls to cryptographic confidence as a strategic capability.

Designing a Compliance Ready Enterprise

- From Control Checklist to Enterprise Capability
  Structural Realignment Across Business, Technology and Risk
  The Compliance-Ready Operating Model
- Extending the Model to the Third-Party Ecosystem
- What a Compliance-Ready Enterprise Actually Looks Like

From Control Checklist to Enterprise Capability

The operationalisation of the DPDP Rules, 2025 compels organisations to treat data protection not as a loose collection of security controls, but as a core enterprise capability. The Rules are explicit in their demand for demonstrable safeguards, defined response timelines, named accountability roles and consistent treatment of personal data across all processing environments. That includes internal applications, shared platforms, outsourced processors, software-as-a-service environments and global delivery centres.

This is a significant step up from earlier practice. Under DPDP, it is no longer sufficient to show that a privacy policy exists, that security tools are deployed, or that occasional audits have been performed. Data Fiduciaries, under DPDP Rules 2025, must be able to demonstrate, with evidence, how personal data is collected, on what basis, for which purposes, under which conditions it is shared, how long it is retained, and what technical and organisational safeguards governed that lifecycle. In other words, data protection becomes an operating property of the enterprise, not a compliance declaration in a policy manual.

Seen this way, DPDP is less a "privacy law" and more a digital operating standard. It forces the organisation to make decisions about data intentional, traceable and defendable. That transition starts with recognising that data protection lies at the intersection of business design, technology architecture, process engineering and governance and not in any one of these domains alone.

DPDP Rules, 2025 elevate data protection into a core enterprise capability, demanding evidence-driven safeguards, accountable roles, and uniform controls across internal systems, SaaS, outsourced processors, and global delivery centres.

Compliance is now proof-based, not policy-based. Organisations must demonstrate how personal data is collected, processed, shared, stored, and secured across its lifecycle—making data protection an operational property, not a check-box.

DPDP becomes a digital operating standard, ensuring data decisions are intentional, traceable, and defensible—spanning business design, technology architecture, process engineering, and governance.

Structural Realignment Across Business, Technology and Risk

For most enterprises, meeting the expectations of DPDP requires structural realignment, not incremental fixes. The law's obligations cut across existing organisational boundaries.

- On the technology side, teams must reassess how personal and sensitive data flows through their enterprise, systems and processes. Legacy systems, microservices, data lakes, AI models, customer-facing apps and partner APIs all need to be examined in terms of what data they ingest, how they transform it, and which downstream systems consume it. Many current integrations were built when data was treated as an internal asset; DPDP forces a re-examination of every such integration through the lens of legality, necessity and risk.
- On the legal and compliance side, teams must move beyond generic privacy clauses toward operational guidance. High-level obligations in the Act and Rules have to be translated into working interpretations: what constitutes a lawful purpose in a particular business context, what retention horizons are justified, how consent should be framed for different channels, and what "reasonable safeguards" mean for specific systems. These interpretations must be stable enough for engineers to implement, yet flexible enough to evolve as case law and regulatory guidance emerge.
- For risk and audit functions, the task is to define what "good" looks like in measurable terms. Merely checking for the existence of policies is not enough; they must establish control expectations, test designs and operating effectiveness, and ensure that key risks such as excessive data aggregation, uncontrolled copies, weak oversight or opaque AI use are visible and managed.
- For product and business teams, the adjustment is equally fundamental. DPDP pushes them away from default "collect everything, use later" behaviour and towards purpose-bound, minimised, explainable data use. That means redesigning customer journeys and internal processes to collect fewer attributes, avoid "just in case" data fields, and structure flows so that personal identifiers are stripped out as early as practicable. In many organisations, this will be the first time that growth ambitions and data protection constraints are handled in a structured, intentional way.

Ultimately, what DPDP is pushing enterprises towards is a more mature form of digital governance: one where business value, regulatory duty and operational feasibility are negotiated explicitly rather than implicitly, and where responsibility for data is shared across the organisation rather than pushed onto a single function.

The Compliance-Ready Operating Model

Sitting above the architecture is the operating model - The way the organisation actually makes decisions, allocates responsibilities and runs its processes. DPDP places responsibility squarely on the Data Fiduciary, even when processing is carried out by vendors or group entities. That demands a level of internal clarity that many organisations have not yet needed to develop.

A mature compliance-ready operating model has a few defining characteristics.

- First, it has well-defined governance roles. The Data Protection Officer, where required, is a function with access to decision-makers and the ability to influence product roadmaps and technology choices. DPO need to work with internal teams with a mandate to understand how datasets containing personal information are used and must champion data minimisation & responsible disposal.
- Second, the operating model has clear decision rights. The organisation knows who approves new data-processing activities, who can authorise new integrations or transfers, who can sign off on retention exceptions, and who is accountable for managing processor relationships. This prevents decisions about data from being made informally or inconsistently under time pressure.
- Third, it has integrated review processes. Privacy, security and architectural reviews are built into product development lifecycles and change-management processes, rather than appended at the end. That integration ensures that DPDP considerations are addressed early, when changes are still inexpensive to make, instead of surfacing late as blockers.
- Fourth, it has repeatable engineering patterns. Rather than leaving every team to interpret DPDP controls on their own, central architecture and security functions provide standard patterns such as reference implementations for consent capture, data minimisation, logging, encryption, tokenisation, access control, etc., that teams can adopt with confidence. These patterns reduce variability and make it easier to prove compliance.
- Finally, it maintains a reporting model tied to real system behaviour, not to policy aspirations. Dashboards and periodic reviews show how data is actually being used, where controls are bypassed or failing, how often exceptions are raised, and how quickly incidents are detected and handled. In this way, the operating model becomes self-correcting: it surfaces gaps and allows leadership to intervene before they turn into regulatory issues.

Extending the Model to the Third-Party Ecosystem

This operating model cannot stop at the enterprise boundary. Indian organisations typically rely on a dense network of third parties such as: SaaS platforms, outsourced operations centres, fintech or Insurtech partners, cloud infrastructure providers, analytics vendors, and global service delivery teams. DPDP explicitly treats these relationships as extensions of the Data Fiduciary's responsibility: outsourcing processing does not outsource accountability.

As a result, contract design, due diligence, onboarding and continuous monitoring become central pillars of compliance-readiness. Contracts must clearly define the purposes for which data is shared, the categories of data involved, retention and deletion expectations, security obligations, the ability to sub-process, and the nature and frequency of evidence the processor must supply. Due diligence must move beyond checkbox questionnaires to include meaningful assessment of how data is stored, segmented, accessed and monitored in the processor's environment.

Onboarding processes need to ensure that data flows to a new processor are consistent with declared purposes and recorded in the organisation's data inventory. Ongoing oversight must include periodic review of logs or attestations, testing of incident-response arrangements, and clarity on how data will be returned or deleted at the end of the relationship. Without this extended governance, enterprises risk being held accountable for practices over which they have little visibility.

DPDP extends the enterprise boundary. Third parties: SaaS providers, outsourced centres, fintech/insurtech partners, cloud vendors, analytics firms and global delivery teams are treated as direct extensions of the Data Fiduciary's responsibility. Outsourcing the processing does not outsource accountability.

Contracts, due diligence and continuous monitoring become compliance pillars. Agreements must define purpose limitation, data categories, retention rules, deletion protocols, security obligations, sub-processing terms and the evidence processors must supply.

End-to-end governance across the vendor lifecycle is mandatory. Onboarding must align data flows with declared purposes; oversight must include log reviews, attestations, incident-response testing and clear off-boarding deletion/return protocols ensuring enterprises remain accountable beyond direct visibility.

What a Compliance-Ready Enterprise Actually Looks Like

When the ideas in this section are implemented properly, the organisation starts to look and feel different in day-to-day work. Data questions stop being escalations that nobody owns and become routine decisions taken within clear boundaries. Product teams know what they are allowed to collect and why. Engineering teams know which patterns to use for which categories of data. Legal and compliance are involved early enough to shape solutions instead of blocking launches at the end. Procurement can tell the difference between a processor that is genuinely in control of its environment and one that is not. Risk and audit work with live signals instead of static checklists.

In practical terms, a compliance-ready enterprise has a small set of recognisable traits:

- Data is mapped and managed, not guessed at. Key datasets, flows and processors are known, documented and periodically updated, rather than reconstructed when something goes wrong.
- Decisions about data are taken in the open. New uses, new integrations and new exceptions go through defined channels, with clear records of who decided what and on what basis.
- Controls are predictable. Similar systems handling similar categories of data are protected in similar ways, instead of each team improvising its own approach.
- Third-party exposure is intentional. The organisation can explain why a particular processor has the access it does, what safeguards apply, and how that relationship is being monitored.
- Evidence is available without a scramble. When asked to show how a particular category of personal data is handled, the organisation can produce logs, configurations and governance records without weeks of manual effort.

Reaching this state does not mean every control is perfect or every risk is eliminated. It means the enterprise has moved from reactive, person-dependent practices to a stable, repeatable way of handling personal data that can withstand questioning from regulators, customers, partners or the board. It is this combination of clarity, discipline and demonstrability that turns DPDP from an external obligation into an internal standard for how the organisation wants to run its digital business.

Technologies Enabling the DPDP Rules

Technologies Enabling the DPDP Rules

The DPDP Act and the DPDP Rules, 2025 may not prescribe a particular technology stack, but they do assume a certain technical maturity in organisations that process personal data at scale. DPDP Rules expects that Data Fiduciaries to implement at least a baseline of "reasonable security safeguards" – encryption, masking or tokenisation of personal data, access control, log monitoring and review, data backups for continued processing, retention of logs and data for investigation, and contractual safeguards with processors to ensure these measures are actually implemented.

These expectations land in an environment where the cost and complexity of data breaches are rising. Recent breach research places the average cost of a data breach in India at around ₹19.5 crore, a 39% increase since 2020, with 35% of breaches involving data in unmanaged sources – "shadow data" – and those shadow-data breaches costing on average 16% more than others. DPDP does not change these dynamics; it changes the standard of justification. When something goes wrong, organisations will be expected to show, with evidence, how personal data was discovered, minimised, safeguarded, monitored and ultimately retired.

This section looks at the technology stack that makes that possible. These set of capabilities, taken together, turn DPDP from a compliance expectation into an engineered property for differentiation of the enterprise:

- Data visibility and lifecycle control
- Protection and control across the stack
- Cryptography and privacy-enhancing computation
- Observability, evidence and automation

The emphasis is on how these capabilities work together, not on any single component.

DPDP Rules to Technology Stack Mapping

If we strip the legal language down to its operational core, the DPDP framework assumes that an organisation can:

- Know where personal and sensitive personal data resides, how it flows and which processes depend on it.
- Shape what is collected, how long it is kept and how widely it is shared.
- Protect that data with safeguards proportionate to its sensitivity and use.
- Observe how systems and users actually interact with it, in near real time.
- Explain and, where necessary, reconstruct events when something goes wrong.

Translating that into a technology agenda yields four families of capability:

- **Data intelligence and lifecycle:** discovery, classification, lineage, minimisation and retention enforcement.
- **Control surfaces:** identity, access, segmentation, application security and configuration management.
- **Cryptography and PETs:** encryption, tokenisation, vaults, and privacy-preserving computation.

Data Intelligence and Lifecycle

In most large organisations, the limiting factor is not the absence of security tools; it is the absence of clean, current visibility into personal data. Over a decade of digital growth, personal and sensitive personal data tends to fragment across:

- Transactional systems and their read replicas,
- Data lakes and warehouses,
- Log stores and observability pipelines,
- Test and pre-production environments,
- SaaS platforms and internal collaboration tools,
- Processes and partner environments.

In such estates, "where is this category of personal data?" is not a trivial question. DPDP's requirement that logs and data relevant to investigations be retained for at least a year only has meaning if the organisation can answer that question reliably.

The foundational technologies here are:

Automated discovery and classification

Tools that can scan structured and unstructured repositories for patterns that indicate personal and sensitive personal data – identity numbers, contact details, behavioural identifiers, location traces – and tag them consistently. The most mature platforms combine pattern-matching with semantic understanding (for example, distinguishing a random string from a customer identifier based on context) and operate across on-premises and multi-cloud environments.

Lineage and flow mapping

Discovery answers "where is it now?". Lineage answers "how did it get there, and what else depends on it?". It traces how a field moves from capture in a front-end, through event streams and microservices, into analytical stores and reports. This is what allows organisations to see that a single identifier might be present in a payment system, a risk model, a fraud rules engine, an analytics mart and a customer-support log – and to treat that pattern as a single risk object rather than six unrelated systems.

Data Intelligence and Lifecycle

Protection and Control Across the Stack

Once personal data is visible and its lifecycle is anchored in reality, the next challenge is to constrain how it is accessed and moved. This is the domain of the protection and control stack: identity, access, segmentation and application-level safeguards.

Key building blocks include:

Identity and access governance

Identity systems provide the vocabulary of people, services, roles and attributes. Used well, they allow the organisation to express DPDP principles as access rules:

- Some roles may see full records; others only pseudonymised or aggregated views.
- Some systems may read from high-risk datasets; others are limited to derived tables.
- Service-to-service interactions are scoped to specific APIs and datasets, not all-access keys.

This is where purpose limitation and minimisation start to become enforceable constraints, not statements of intent.

Segmentation and zoning

Network, environment and account-level segmentation is not new. What DPDP adds is the need to understand where raw personal data, transformed data and non-personal data sit relative to each other. Systems that can export or copy full identifiers deserve a different protective posture than systems that handle only tokens or aggregates. Segmentation becomes a way of localising breach impact and demonstrating that high-risk data was not unnecessarily exposed to lower-assurance environments.

Configuration and posture management

Misconfigurations such as: Open storage buckets, exposed management interfaces, logging debug data in production are a recurring cause of data breaches globally. Tools that continuously inspect infrastructure, platforms and applications for configuration drift therefore become part of the DPDP story: they prevent back-sliding from an engineered, compliant state into an accidental, risky one.

Taken together, these controls do not remove the need for cryptography; they create the context in which cryptography can be applied with precision, because the organisation understands who legitimately needs what and where.

Data Intelligence and Lifecycle

Cryptography and Privacy-Enhancing Computation

Cryptography is the enforcement layer of DPDP compliance. But in practice, the applications of cryptography extends far beyond encryption. In the context of DPDP, Cryptography spans into three concentric layers.

Baseline encryption and key management

The first layer is the most widely use case for Cryptography in the context of an enterprise. The use cases include:

- Encryption at rest: For databases, file systems and object stores containing personal or sensitive personal data.
- Encryption in transit for all channels carrying data.
- Disciplined key management: Key Generation, storage, rotation, revocation and monitoring of key usage.

The DPDP Rules explicitly call out encryption as a minimum safeguard, alongside masking and tokenisation. What distinguishes a mature organisation is not whether these features exist, but how consistently and transparently they are applied, and whether key lifecycles are treated as governance objects rather than operational afterthoughts.

Data vaults, tokenisation and pseudonymisation

The second layer of applied cryptography concerns where raw identifiers are allowed to exist.

A vault-and-token architecture, for example, would ensure that:

- High-risk identifiers (PII, for example) are immediately replaced with tokens or pseudonyms at the first point of capture.
- Only dedicated vault systems ever store the mapping between tokens and real identifiers
- Re-identification is mediated by a narrow set of services, subject to strict policy and logging.
- A majority of systems (analytics, monitoring, reporting, many operational processes) work entirely with tokens or pseudonyms.

This approach aligns strongly with DPDP's emphasis on minimisation, storage limitation and harm reduction. If a breach occurs in a system that never held raw identifiers, the organisation can be rest assured that it can bounce back from the risk of exposure, potentially leading up to reputational and penal damages.

Data Intelligence and Lifecycle

Privacy-enhancing technologies

Beyond these base controls, a class of privacy-enhancing technologies (PETs) is emerging as the bridge between DPDP-grade protection and ambitious data-use agendas. The global PET market is estimated at roughly USD 5 billion in 2025, projected to grow to over USD 12 billion by 2030, at a compound annual growth rate close to 20% – an indication that these technologies are moving from research into mainstream planning. India is one of the fastest evolving mature markets in the world in the development and deployment of PETs.

Key PET families include:

- Differential privacy: Allows organisations to publish and consume statistical outputs with mathematically bounded re-identification risk by adding carefully calibrated noise. This is especially valuable for internal dashboards and external reports that rely on sensitive datasets but do not require individual-level precision.

- Federated and distributed learning: Where models are trained across multiple data silos (for example, across business units or institutions) without centralising raw personal data. Only model updates are shared, and secure aggregation prevents those updates being reversed into individual records.

- Secure multi-party computation and homomorphic encryption: Enables joint analytics when two or more parties want to compute on their combined data but cannot afford to reveal their datasets to each other. Early production deployments in financial and health ecosystems indicate that these techniques are now viable for carefully chosen, high-value use cases.

- Zero-knowledge proofs and privacy-preserving credentials: Allows proof of specific claims – age thresholds, residency, prior verification, enrolment status – without sharing underlying documents or full records every time.

These techniques are still specialised. They will not replace traditional controls. But they give enterprise leaders degrees of freedom: ways to support analytics, AI and ecosystem collaboration without reverting to "copy all the raw data into a new system" as the default pattern.

Cryptography is no longer just the padlock on the door. It is the design material from which these new patterns of computation are made.

Data Intelligence and Lifecycle

Observability, Evidence and Automation

The final technology layer is concerned with whether the organisation can stand behind its own story when its DPDP posture is tested.

Recent breach research shows that nearly half of breaches globally involve customer personal data, and around a third involve shadow data in unmanaged stores; breaches with shadow data cost more and take longer to resolve. In this environment, regulators are less interested in tool inventories and more interested in how quickly an organisation can reconstruct what happened and what safeguards were in place at the time. The following three capabilities define observability:

High-fidelity logging and telemetry

Logging is the essential component that forms the bedrock of Observability. At the bare minimum, Logs and traces must capture: Who or what accessed which datasets and fields:

- Via which applications and APIs,
- From which locations or networks,
- Under which policy decisions (granted, denied, elevated),
- Any configuration changes that affected safeguards (encryption, tokenisation, access rules, retention settings).

Retention expectations in the DPDP Rules, at least one year for logs and relevant data to support detection, investigation and continued p**rocessing, make sense only if these logs are rich enough to answer the questions that will be asked in practice.**

Analytics and continuous controls monitoring

Once telemetry (from logs) is available, analytics and monitoring engines can:

- Highlight unusual movement of high-risk data classes,
- Detect policy violations (for example, clear-text exports where only tokens should exist),
- Monitor for configuration drift (encryption disabled, excessive privileges granted, retention rules bypassed),
- Provide leading indicators of control weakness, not just lagging indicators of incidents.

Data Intelligence and Lifecycle

The Technology Posture of a DPDP-Enabled Enterprise

When these technologies are assembled thoughtfully, the organisation's relationship with personal data changes in recognisable ways:

- Personal and sensitive personal data is mapped and bounded, not reconstructed in a crisis or as and when need arises.
- High-risk identifiers (PII) exist in far fewer systems, protected through field-level encryption, tokenisation and vaulting, rather than being copied indiscriminately.
- Analytics, AI and ecosystem collaborations proceed on transformed, pseudonymised or differentially private data, rather than full raw records.
- Access decisions are made against a shared understanding of data sensitivity and purpose, not just application roles.
- Control logic is captured as code and evaluated continuously, reducing dependence on individual heroics or annual compliance cycles.

Crucially, this posture is not the result of adopting any single "DPDP solution". It emerges when enterprise leaders treat data intelligence, cryptography, privacy-enhancing computation, access control, observability and automation as parts of a single design problem: how to make data protection by design, and cryptography by default, a normal property of how the organisation builds and runs systems.

Quantum Resiliency as an Emerging Regulatory Expectation

Two Families of Quantum-Resilient Controls
Post-Quantum Cryptography: The Default Path for Most Organisations
Quantum Key Distribution and Quantum Networking
Putting It Together: A Coherent Quantum-Resilient Architecture

Quantum Resiliency as an Emerging Regulatory Expectation

Across jurisdictions, the quantum threat has moved from academic discussion into regulatory planning. NIST has standardised three post-quantum algorithms (ML-KEM, ML-DSA, SLH-DSA) and the world now focusing on migration guidance rather than on whether building Quantum Resiliency will be required. The EU, many other countries have published a coordinated roadmap that expects public-sector infrastructures to transition to PQC by the early-to-mid 2030s, framing it as a "timely, comprehensive and coordinated" response to quantum risk. India's National Quantum Mission is expected to launch a national roadmap guideline document for Quantum Resiliency in early 2026. These guidelines advise all enterprises to act from now and avoid a scramble when cryptographically relevant quantum computers appear.

This is matched by sectoral regulators. The UK's NCSC is advising organisations to complete their PQC transition by around 2035 and urges large operators to have discovery and migration plans in place by 2028. Singapore's MAS has issued an advisory requiring financial institutions to inventory cryptographic assets, build crypto-agility and prepare for PQC adoption as part of standard cyber-risk management, explicitly citing quantum computing as a foreseeable risk within the next decade. Although most national timelines for Quantum resiliency and migration of high priority systems converge around 2029-2030, this migration is expected to take many years. So, it's imperative that the early preparations for Migration be done at the earliest.

In practical terms, this means "quantum-resilient" is on its way to becoming a baseline expectation for any system that has long-lived confidentiality requirements including but not limited to: payments, health records, telecom, critical infrastructure, government services, and large consumer platforms. For enterprises already modernising architectures for DPDP compliance, the quantum transition is not a separate project; it is the next horizon of the same design problem: can you still trust your cryptographic guarantees when the underlying assumptions change?

Against this backdrop, several families of technologies are now framed as "quantum-resilient": post-quantum cryptography, quantum key distribution, quantum random number generators, quantum-safe VPNs and tunnels. It is to be noted that these do not play equal roles. Understanding where each of these technologies fit, and where they does not, is essential to avoid expensive investments and vendor lock-in, leading to dead-ends.

Two Families of Quantum-Resilient Controls

At a fundamental level, quantum-resilient technologies fall into two broad families:

1. **Algorithmic defences that run on classical infrastructure:** Cryptographic algorithms, primarily post-quantum public-key schemes, designed to resist attacks from large-scale quantum computers. PQC is deployed in software and hardware much like today's RSA or ECC: inside TLS, VPNs, application protocols, storage systems and PKI.

2. **Physical-layer and hardware-centric approaches – Quantum Key Distribution:** Quantum Key Distribution and quantum networking (using photons and quantum states to distribute keys), QRNGs that use quantum processes to generate entropy, and specialised "quantum-secure tunnels" or QVPNs that blend these ideas with classical networking.

Quantum Resiliency as an Emerging Regulatory Expectation

Across jurisdictions, the quantum threat has moved from academic discussion into regulatory planning. NIST has standardised three post-quantum algorithms (ML-KEM, ML-DSA, SLH-DSA) and the world now focusing on migration guidance rather than on whether building Quantum Resiliency will be required. The EU, many other countries have published a coordinated roadmap that expects public-sector infrastructures to transition to PQC by the early-to-mid 2030s, framing it as a "timely, comprehensive and coordinated" response to quantum risk. India's National Quantum Mission is expected to launch a national roadmap guideline document for Quantum Resiliency in early 2026. These guidelines advise all enterprises to act from now and avoid a scramble when cryptographically relevant quantum computers appear.

This is matched by sectoral regulators. The UK's NCSC is advising organisations to complete their PQC transition by around 2035 and urges large operators to have discovery and migration plans in place by 2028. Singapore's MAS has issued an advisory requiring financial institutions to inventory cryptographic assets, build crypto-agility and prepare for PQC adoption as part of standard cyber-risk management, explicitly citing quantum computing as a foreseeable risk within the next decade. Although most national timelines for Quantum resiliency and migration of high priority systems converge around 2029-2030, this migration is expected to take many years. So, it's imperative that the early preparations for Migration be done at the earliest.

In practical terms, this means "quantum-resilient" is on its way to becoming a baseline expectation for any system that has long-lived confidentiality requirements including but not limited to: payments, health records, telecom, critical infrastructure, government services, and large consumer platforms. For enterprises already modernising architectures for DPDP compliance, the quantum transition is not a separate project; it is the next horizon of the same design problem: can you still trust your cryptographic guarantees when the underlying assumptions change?

Against this backdrop, several families of technologies are now framed as "quantum-resilient": post-quantum cryptography, quantum key distribution, quantum random number generators, quantum-safe VPNs and tunnels. It is to be noted that these do not play equal roles. Understanding where each of these technologies fit, and where they does not, is essential to avoid expensive investments and vendor lock-in, leading to dead-ends.

Two Families of Quantum-Resilient Controls

At a fundamental level, quantum-resilient technologies fall into two broad families:

1. **Algorithmic defences that run on classical infrastructure:** Cryptographic algorithms, primarily post-quantum public-key schemes, designed to resist attacks from large-scale quantum computers. PQC is deployed in software and hardware much like today's RSA or ECC: inside TLS, VPNs, application protocols, storage systems and PKI.

2. **Physical-layer and hardware-centric approaches – Quantum Key Distribution:** Quantum Key Distribution and quantum networking (using photons and quantum states to distribute keys), QRNGs that use quantum processes to generate entropy, and specialised "quantum-secure tunnels" or QVPNs that blend these ideas with classical networking.

Quantum Resiliency as an Emerging Regulatory Expectation

For most enterprises, the first approach (Algorithmic approach) is the most pragmatic approach to achieve quantum resiliency. The second family may add additional protection, but it cannot substitute for a systemic migration of cryptography. That is why national and regional roadmaps from NIST to the EU to TEC converge on PQC as the primary control, with quantum communications treated as complementary or experimental. It is to be noted that China has invested into Quantum Key Distribution and has demonstrated 10,000 Kilometre plus communications between continents. With this divergent perspectives, India needs to balance out both approaches.

Post-Quantum Cryptography: The Default Path for Most Organisations

Public-key cryptography today relies on number-theoretic problems (factoring, discrete logarithms, elliptic-curve discrete logs). Large-scale quantum computers running Shor's algorithm would, in principle, be able to solve these efficiently, breaking TLS, VPNs, PKI and many application protocols that rely on RSA and ECC.

Post-quantum cryptography replaces those building blocks with problems that are believed to be hard even for quantum adversaries. These are principally lattice-based problems, hash-based constructions and a small set of alternative assumptions. NIST's current standards, ML-KEM for key establishment, ML-DSA and SLH-DSA for digital signatures, are designed to be drop-in replacements at the protocol level, running entirely on classical hardware.

From an enterprise-architecture standpoint, this matters for three reasons:

Coverage across use cases: PQC can be embedded wherever public-key cryptography is used today: TLS handshakes, application-level mutual authentication, VPNs, messaging protocols, firmware signing, certificate infrastructures and, over time, encrypted storage formats. It addresses both data-in-transit and many aspects of identity and integrity.

Deployment on existing infrastructure: Because PQC runs on conventional CPUs, organisations can upgrade software stacks, HSM firmware, libraries and devices rather than building a new physical network. Cloudflare and several large operators have already demonstrated PQC-augmented TLS and VPNs at Internet scale without specialised quantum hardware.

Alignment with regulatory roadmaps: Whether in the EU's PQC transition roadmap, TEC's migration guidance, NCSC's 2035 target or MAS's advisory, the central expectation is that organisations will implement crypto-agility and phase in PQC, not re-engineer their networks around quantum channels.

Quantum Resiliency as an Emerging Regulatory Expectation

How PQC fits into "cryptography by default"

From the perspective of "Data Protection by Design, Cryptography by Default", PQC is not an exotic add-on. It is the natural evolution of the control framework enterprises already rely on:

- At the transport layer, PQC cipher suites become part of the standard stack for service-to-service, app-to-API and cross-border connections, including those involving processors and group entities. Hybrid handshakes that combine classical and PQC algorithms are emerging as pragmatic stepping stones.

- In identity and lifecycle management: PQC signatures will, over time, underpin digital certificates, device identities, update signing and transaction authorisation. Many forward-looking OEMs are already signing with PQC today.

- Within application design, PQC becomes one element of a broader crypto-agile architecture alongside symmetric encryption, tokenisation, pseudonymisation and privacy-enhancing technologies. The enterprise goal is not to "do PQC everywhere overnight", but to ensure that systems being designed or significantly re-engineered now can transition without structural rewrites.

The strategic question for leadership, therefore, is not whether PQC is relevant, but how deliberately it is incorporated into roadmaps, vendor contracts and architectural standards.

Quantum Key Distribution and Quantum Networking

Quantum Key Distribution uses quantum states (typically single photons) to establish a shared symmetric key between two endpoints and to detect any eavesdropping on that channel. In theory, this provides information-theoretic security for the key exchange: any interception alters the quantum states and is observable.

Recent milestones are impressive. The EuroQCI programme is building a pan-European quantum communication infrastructure combining terrestrial fibre links and satellites, with QKD as a core technology for protecting critical government and infrastructure communications. In 2025, researchers from South Africa and China established a 12,900 km quantum satellite QKD link, the longest intercontinental quantum link to date, demonstrating one-time-pad-secured image transmission between hemispheres.

These are significant research and sovereignty achievements. But they do not make QKD a general replacement for public-key cryptography in enterprise environments.

Why QKD is not the default enterprise answer

Several independent analyses - from academic reviews to guidance by agencies such as NCSC and NSA - highlight structural limitations of QKD and quantum networking for broad adoption. The core issues are architectural rather than ideological

Quantum Resiliency as an Emerging Regulatory Expectation

Scope of the problem it solves.

As of writing this report, terrestrial QKD solves one problem: how two endpoints with a dedicated quantum channel can establish a shared symmetric key with eavesdropping detection. It does not provide digital signatures, does not address data at rest, does not authenticate endpoints by itself, and does not replace the need for robust symmetric cryptography and key-management. Free-space QKD, while experimental at the moment may offer better flexibility, but will take some time to mature indigenously. Note that China has already demonstrated 10,000 Km+ satellite-based QKD systems.

Dependence on classical or post-quantum cryptography: In practice, QKD requires an authenticated classical channel to prevent man-in-the-middle attacks. That authentication relies on classical or ultimately PQC-based signature schemes or pre-shared secrets. QKD therefore sits inside, not above, a broader cryptographic stack that will still have to migrate to PQC. Relying on PQC for signature presents a dichotomy for QKD since the same PQC can be used to establish a shared secret through Key Encapsulation.

Network-scale limitations: QKD is fundamentally point-to-point. Extending it beyond metropolitan scale requires trusted repeaters (trusted nodes), satellite relays or complex entanglement-based schemes. Key rates over long distances are constrained; building QKD for "last-mile" device-level connectivity is not realistic. As UK's NCSC notes, this means QKD cannot operate at Internet scale and is unsuitable as a general replacement for public-key cryptography in heterogeneous networks.

Operational complexity and cost: Deploying QKD involves specialised optics, quantum-grade components, strict environmental controls and integration with network equipment. For most enterprises, especially in multi-cloud, multi-carrier environments, the cost and operational burden are disproportionate to the incremental benefit over well-implemented PQC-based key establishment.

Vendor Lock-in & Side Channels: Beyond architectural limitations, QKD also carries practical risks that enterprises must weigh carefully: vendor lock-in and susceptibility to implementation side channels. While QKD vendors may emphasise the theoretical security of quantum key exchange and the ability to detect eavesdropping, the security of real-world deployments, however, is constrained not by physics but by engineering. Independent cryptanalysis has repeatedly shown that commercial QKD systems can leak information through side channels that cannot be eliminated by theoretical guarantees alone.

Moreover, QKD offerings are typically proprietary end-to-end stacks: specialised optics, management consoles, link controllers, and vendor-tuned protocols that are not interoperable across suppliers. For large enterprises and public-sector buyers, this creates a long-term dependency cycle, with high switching costs, limited multi-vendor design flexibility, and upgrade paths entirely dictated by the manufacturer's roadmap. In contrast, PQC benefits from open standards, algorithmic diversity, independent validation and broad ecosystem support making it the more sustainable and strategically aligned foundation for enterprise cryptography.

Quantum Resiliency as an Emerging Regulatory Expectation

This is why many current policy documents converge on a pragmatic stance: PQC for broad, scalable quantum resilience; However, as a country, Indian needs to invest in QKD, as QKD is expected to be the backbone for building Quantum Internet.

The "quantum-secure tunnel" and QVPN narrative

Building on QKD, some vendors now promote "quantum VPNs" or "quantum tunnels"—network overlays marketed as quantum-secure. Technically, these fall into one of three patterns:

1. Classical VPNs upgraded with PQC key-exchange (for example, ML-KEM integrated into IKE or TLS).

2. Classical VPNs using keys derived from QKD links between data centres or backbone routers.

3. Hybrid overlays that combine PQC and QKD-derived keys for specific segments.

Only the first approach is broadly scalable; it is simply PQC deployed at the network layer and is a logical part of any crypto-agility roadmap. The second and third are meaningful primarily where QKD infrastructure already exists such as in national backbones, space-based links, or a small number of inter-data-centre connections. For most organisations, "quantum-safe VPN" will and should mean "VPN with PQC-based key establishment and modern symmetric ciphers," not a quantum-hardware-dependent network fabric.

Putting It Together: A Coherent Quantum-Resilient Architecture

When viewed through the lens of enterprise architecture and DPDP obligations, a coherent picture emerges:

- PQC becomes the default upgrade path for public-key cryptography across the stack: TLS, VPNs, application protocols, PKI and critical back-office systems. This is where regulatory roadmaps, national technical guidance and standards bodies are placing their weight.

- Crypto-agility is the enabling design feature, allowing systems to support both current and post-quantum algorithms, rotate quickly as standards mature, and avoid single-vendor or single-algorithm lock-in.

- QKD and quantum networking are strategic, not ubiquitous: QKD's current maturity, cost structure and architectural constraints make them impractical for widespread enterprise deployment. The requirement for engineered quantum channels, trusted repeaters, specialised optics and tightly controlled physical conditions restricts their viability to a narrow set of national or scientific environments. While research advances, including satellite-based QKD and integrated quantum networking, may broaden applicability over the next decade, these technologies should be treated as long-horizon strategic bets, not near-term building blocks. Even where they eventually find a role, they will complement rather than displace PQC, which remains the only scalable mechanism for upgrading the cryptographic foundations of heterogeneous, internet-scale enterprise systems.

Quantum Resiliency as an Emerging Regulatory Expectation

For Indian and global enterprises planning for quantum risk, the strategic message is straightforward:

Treat PQC migration as a core part of your "Data Protection by Design" agenda, design for cryptographic agility, and view quantum hardware technologies as specialised tools to be considered case-by-case, not as prerequisites for being quantum-resilient.

In that landscape, cryptography-first solution providers and internal platform teams have a critical role: abstracting the complexity of PQC pilots into governed services and patterns that can be consumed consistently across the organisation. The enterprises that succeed will be those that make quantum resilience an architectural property of their digital estate, rather than a collection of disconnected experiments.

Summary

Summary

India's digital economy is entering a period where two structural transformations are unfolding simultaneously: the operationalisation of the DPDP Rules, 2025, and the global transition to quantum-resilient cryptography. Each of these shifts, by itself, would require organisations to rethink longstanding assumptions about data, architecture and trust. Together, they mark the beginning of a new governance era, one in which digital systems must be provably compliant, cryptographically robust, operationally predictable and resilient to future computational realities.

The findings across this report point to a consistent theme: modern digital enterprises can no longer treat data protection and cryptography as isolated technical functions. DPDP demands evidentiary governance, quantum computing challenges the mathematical basis of that governance. Both require leaders to act with deliberateness - rebuilding trust not as a by-product of tools, but as an engineered property of the organisation.

DPDP transforms data protection into an enterprise-wide operating discipline

The Act and Rules do not merely impose procedural duties; they redefine the organisational understanding of data. Compliance now depends on whether an enterprise can demonstrate, with evidence, how personal data is collected, minimised, transformed, safeguarded and ultimately retired. Visibility, accountability and repeatability - not breadth of tooling - become the central measures of readiness.

This elevates conversations inside organisations: from whether a control exists, to whether its behaviour can be proven; from static policies to engineered safeguards; from siloed compliance work to coordinated architectural and operational decisions.

A compliance-ready enterprise requires coordinated realignment across business, technology and governance

The report highlights that DPDP cannot be implemented by any single function. Technology teams must rationalise data flows and eliminate uncontrolled spread of identifiers; legal and compliance must translate Rules into operational guidance; risk and audit must define measurable control baselines; and business and product teams must redesign interactions around minimisation and purpose limitation.

When this works, enterprises begin to operate differently: decisions become evidence-backed, exceptions are traceable, and personal data stops being an unmanaged design material.

Key Findings

A DPDP-aligned technology centres on clarity, containment, cryptography and control

Through data-intelligence platforms, vault-and-token architectures, PETs, identity governance, segmentation, continuous controls monitoring and lifecycle enforcement, organisations build a digital environment in which high-risk data is bounded, transformations are intentional and logs explain system behaviour rather than obscure it.

Importantly, cryptography emerges not as a checklist item but as the backbone of enforcement. Encryption, key governance, tokenisation, signing, pseudonymisation and privacy-preserving computation form the substrate on which DPDP becomes demonstrable. This is the operational meaning of "Data Protection by Design, Cryptography by Default."

Quantum computing introduces a long-horizon but structurally certain disruption

Quantum computing does not create a new class of attacker; it erodes the assumptions underlying the current global trust fabric. Because regulated data and many mission-critical systems must remain secure for decades, the threat is inherently long-tail: data exfiltrated today may be decrypted years later.

Global regulators have therefore shifted their stance: quantum resiliency is not a speculative R&D concern but a predictable engineering obligation. Roadmaps across NIST, NCSC, BIS, EU and TEC all converge on the expectation that organisations must begin preparing for a full transition to post-quantum cryptography within this decade.

PQC is the only scalable, ecosystem-aligned path to quantum resiliency

Across all credible analyses, PQC stands out because:

- PQC replaces exactly the primitives that quantum enabled attacks will break.
- PQC runs on classical hardware across cloud, APIs, edge and enterprise environments.
- PQC scales to internet-wide, heterogeneous, many-to-many architectures.
- PQC aligns with regulatory timelines and vendor roadmaps.

Key Findings

- PQC supports hybrid migration that avoids disruption.

QKD, quantum networking will remain important research and national-infrastructure domains, but they are structurally mismatched to the topology and economics of enterprise systems. Their eventual role, where justified, may be complementary, not foundational.

Taken together, these findings define the mandate for organisations building durable digital trust: engineered compliance for today; cryptographic agility for tomorrow.

Digital enterprises are entering an era in which trust must be engineered with the same intent and discipline as any other core capability. The systems that succeed will be those built on clarity of purpose, governed by transparent principles, and reinforced through architectural choices that endure beyond individual technologies or trends. As organisations expand, collaborate and integrate more deeply across ecosystems, the expectation is no longer mere operational soundness but demonstrable reliability, measured in how consistently they protect data, uphold commitments and adapt to emerging constraints. The challenge for leadership is to convert these expectations into practice: to build environments where responsibility is shared, safeguards are embedded, and confidence is earned through behaviour, not declared through policy. The foundations laid today will define not only the resilience of tomorrow's digital platforms, but the credibility of the institutions behind them.

Data Security Council of India (DSCI), is a not-for-profit, industry body on data protection in India, setup by nasscom, committed to making the cyberspace safe, secure and trusted by establishing best practices, standards and initiatives in cyber security and privacy. To further its objectives, DSCI engages with governments and their agencies, regulators, industry sectors, industry associations and think tanks for policy advocacy, thought leadership, capacity building and outreach activities.

JISA Softech is a deep-tech, cybersecurity and cryptography-focused technology company committed to strengthening digital trust by delivering advanced PKI, encryption, and data-protection solutions. Established with the vision of making high-assurance security accessible and scalable, JISA develops, manufactures, and supports a comprehensive suite of Hardware Security Modules (HSMs), key management systems, and authentication solutions under its flagship brand, CryptoBind.

4th Floor, NASSCOM Campus, Plot No. 7-10, Sector 126, Noida, UP -201303