ARTICLE

From Boardroom to System Architecture: The Company Secretary's Techno-Legal Role under the DPDP Act, 2023 and DPDP Rules, 2025

With the notification of the Digital Personal Data Protection (DPDP) Rules, 2025, the Indian corporate landscape has shifted from "interpreting the law" to "Techno-Legal Compliance." For the Company Secretary (CS), this transition presents a unique challenge and opportunity: the mandates of the Act—consent, erasure, and grievance redressal—are legal in principle but entirely technical in execution. A Board Resolution to "Ensure DPDP Act and Rules Compliance" is merely the starting point; the finish line lies in database architecture, API integrations, and user interface (UI) workflows. This article provides a "Techno-Legal" framework for Company Secretaries to guide their organizations through the implementation phase, translating statutory obligations into system specifications that Chief Information Officers (CIOs) and IT teams can execute.

CS Raju S. Surapuraju, ACS
Practising Company Secretary
Chennai, Tamil Nadu
rsurapuraju@gmail.com

INTRODUCTION

A Board Resolution to *"Ensure DPDP Act and Rules Compliance"* is merely the starting point; the finish line lies in database architecture, API integrations, and user interface (UI) workflows. This article provides a **"Techno-Legal framework"** for Company Secretaries to guide their organizations through the implementation phase, translating statutory obligations into system specifications that Chief Information Officers (CIOs) and IT teams can execute.

The Company Secretary's expertise has long been rooted in: governance frameworks, board minutes, and statutory compliance. However, the **Digital Personal Data Protection Act, 2023 (DPDP Act)** and the notified **Digital Personal Data Protection Rules, 2025** have fundamentally altered this landscape.

Figure 1: Techno-Legal Governance Stack

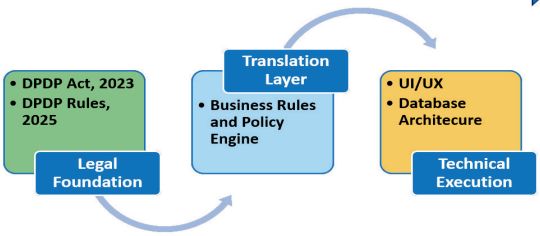

As illustrated in **Figure 1**, effective governance is no longer just a legal mandate — it is a **Techno-Legal Governance Stack**.

While the Act and Rules (**the Legal Foundation**) provide the 'What' and the 'How'— specifically regarding **Data Erasure Timelines (Rule 8)**, **Breach Reporting Formats (Rule 7)**, and **Consent Manager Standards (Schedule I)** — these mandates must be translated into **Business Rules** and ultimately executed in **Database Architecture**.

As we stand in early 2026, the organizations are now in the critical implementation phase. A primary friction points often observed in organizations at the Translation Layer between the Legal function and the IT function:

- **The Legal Team says:** "We must obtain verifiable consent."
- **The IT Team asks:** "Is a simple Boolean flag (is_consented = True) in the user table sufficient? Or is a separate Immutable Log Table required — complete with timestamps and SHA-256 Hash Values — to ensure the record is tamper-proof for legal audits?"

The ability to answer the above is critical for ensuring compliance. The gap between a legal requirement and a software feature is where non-compliance risks breed. This article aims to empower the CS to act as the 'Techno-Legal Governance Architect' for privacy, breaking down the DPDP Act and DPDP Rules into actionable steps.

THE ARCHITECTURE OF "CONSENT" (SECTION 6 & RULE 3)

Under the DPDP Act, consent is the bedrock of processing personal data. The Act mandates that consent must be free, specific, informed, unconditional, and unambiguous. While these are familiar legal terms to comply, their implementation necessitates precise technical specifications.

1. **The "Notice" (Section 5 & Rule 3): Multi-Lingual UI Architecture**

   While **Rule 3** mandates that the Notice must be "itemised" and "understandable independently" of other terms, **Section 5(3)** of the Act explicitly requires

that this option be available in English and all 22 languages specified in the Eighth Schedule of the Constitution.

a. **The Techno-Legal Challenge:** It is no longer sufficient to have a static privacy_policy.html in English.

b. **The System Solution:** The CS must verify that the Content Management System (CMS) has a **"Localization Engine."** The system must detect the user's device language preference and dynamically render the Notice in English or other language specified in the Eighth Schedule of the Constitution before consent is sought.

c. **Audit Requirement:** If a user selects "Hindi" as their interface language, does the **Audit Log (Consent Artifact)** record that they were shown Notice_Version_2.1_Hindi? If the log shows they agreed to the English version, we may be in violation of **Section 5(3)**.

2. **The "Consent Artifact": Moving Beyond the Checkbox**

A simple "tick mark" in a database (e.g., Consent = True) is no longer sufficient. To meet the burden of proof required by the Act, we must treat consent as a transactional record, often called a **"Consent Artifact."**

a. **Data Structure Recommendation:** The CS should guide the IT team to store a **JSON (JavaScript Object Notation)** log for every consent provided, containing:

Who: User ID

What: Specific purpose ID (e.g., "Payroll", "Marketing," "Billing", "Other Services")

When: UTC Timestamp

How: Version ID of the Notice displayed at that time

Validity: Expiry date (if applicable)

The "Digital Evidence" Standard (JSON Schema)

This JSON structure represents the "Consent Artifact" that the IT system should generate every time a user interacts with a Privacy Notice. It is the ultimate piece of evidence for a Section 63 BSA certificate (The Bharatiya Sakshya Adhiniyam, 2023).

ANATOMY OF A "CONSENT ARTIFACT" - AUDIT LOG RECORD

JSON

```
{
 "consent_artifact": {
  "artifact_id": "CONSENT-2026-X892",
  "timestamp": "2026-01-26T10:00:00Z",
  "techno_legal_metadata": {
   "statute": "DPDP Act, 2023",
   "relevant_sections": ["Sec 5(3)", "Sec 6"],
   "relevant_rules": ["Rule 3", "Rule 4"]
  },
  "provenance": {
   "data_principal_id": "user_7741",
   "notice_version": "v2.1_Hindi",
   "hash_of_notice_text": "sha256:e3b0c442..."
  },
  "user_action": {
   "status": "GRANTED",
   "timestamp_ms": 1739527200000
  },
  "evidence_integrity": {
   "hashing_algorithm": "SHA-256",
   "record_hash_value": "f123456789abcdef...",
   "bsa_compliance": "Section 63(4) Ready"
  }
 }
}
```

The inclusion of **Legal Metadata** (Sections and Rules) directly within the JSON record transforms a simple IT log into a "Compliance Artifact." Furthermore, while many IT teams prefer hashing entire log files for efficiency, the Techno-Legal CS should insist on individual record hashing. This ensures that each consent event is a standalone **'electronic record'** under the **Bharatiya Sakshya Adhiniyam (BSA), 2023**, making it significantly easier to produce a valid **Section 63 Certificate** for a single user's dispute without exposing the entire database log to the Data Protection Board.

1. **Techno-Legal Tip:** The "BSA" Standard Under the *Bharatiya Sakshya Adhiniyam (BSA), 2023*, electronic records are treated as primary evidence, provided their integrity is provable. The Company Secretary must ensure that "Consent Artifact" logs are secured using hashing algorithms (such as SHA-256).

In the event of a legal challenge regarding consent, the organization must produce a certificate under Section 63(4) of the BSA. This section mandates the identification of the electronic record and a certification of the proper operation of the computer system. In the digital domain, the Hash Value serves as the definitive technical fingerprint to satisfy this requirement, proving the record has not been tampered with since the timestamp of consent.

THE RIGHT TO ERASURE: OPERATIONALIZING THE DATA PURGE PROTOCOL (SECTION 12 & RULE 8)

While Section 12(1) grants the right to correction, which is a standard database update operation, **Section 12(3): The Right to Erasure** presents a unique architectural challenge.

The Act stipulates that upon a Data Principal's request; the Data Fiduciary must erase personal data unless retention is necessary for a specified purpose or strictly for legal compliance.

For the Governance Professional, the objective is to define a "Clean Deletion" protocol that satisfies the statutory request of the Data Principal without compromising the organization's audit trail or compliance evidence.

1. **The Use Case: Execution of a Termination Request**

   Consider a scenario where a Data Principal initiates a formal account closure and exercises the Right to Erasure pursuant to Section 12(3).

   a. **The Legal Mandate:** The Data Fiduciary is statutorily obligated to remove the Data Principal's Personally Identifiable Information (PII) — such as Name, Email, Mobile, and Residential Address — from the active production environment.

   b. **The Technical Risk:** Executing a standard SQL command (e.g., DELETE FROM Customer_Master WHERE User_ID = 'X') creates a critical compliance vulnerability. Such a "Hard Delete" operation inadvertently destroys the **Consent Artifacts** and historical logs required to prove that the prior processing was lawful, leaving the entity defenceless during a regulatory audit.

2. **The Solution: "Soft Deletion" Architecture & Log Retention (Rule 8)**

   Rule 8(3) resolves this conflict by establishing a dual-state requirement: while the primary personal data must be erased, the "associated traffic data and other logs of the processing" must be retained for a minimum period of **one year**.

   a. **The Instruction to IT:** The system architecture must reject immediate "Hard Deletion." Instead, it must trigger a **"Soft Delete"** workflow consisting of three stages:

   - **Identity Masking:** Pursuant to **Rule 6**, the system must overwrite PII fields in the live database with irreversible, generic placeholders (e.g., converting **"User_Name"** to Deleted_User_UUID_8942) or NULL values, effectively severing the link between the data and the individual.
   - **Log Preservation:** The system must strictly retain timestamped metadata (e.g., "User_UUID logged in at [Timestamp]") and the original Consent Artifacts.
   - **Archival Migration:** These anonymized logs should be automatically migrated to a "Secure Compliance Archive "— ideally configured as Write-Once-Read-Many (WORM) storage — physically separate from the active marketing database.

> Under the DPDP Act, consent is the bedrock of processing personal data. The Act mandates that consent must be free, specific, informed, unconditional, and unambiguous. While these are familiar legal terms to comply, their implementation necessitates precise technical specifications.

This architectural segregation ensures that if the Data Protection Board inquires, **'Under what authority was this user's data processed prior to erasure?'**, the Data Fiduciary can produce the original Consent Artifact (potentially dated years prior) as primary evidence under the *Bharatiya Sakshya Adhiniyam*. This serves as the immutable proof of lawful processing, even though the Data Principal's active profile no longer exists.

3. **The "48-Hour" Pre-Erasure Protocol (Rule 8(2))**

   For scenarios involving automated erasure (e.g., policy-driven deletion of inactive accounts), **Rule 8(2)** mandates a preventative notification mechanism.

   a. **The Requirement:** The system must trigger a notification to the Data Principal at least **48 hours prior** to the permanent erasure of data.

   b. **The System Logic:** The automated scheduler (Cron Job) must execute a conditional check:

   Logic: IF (Account_Status == 'Marked_For_Erasure') AND (Time_To_Execute < 48 Hours) -> TRIGGER_NOTIFICATION ("Urgent: Data scheduled for permanent erasure in 48 hours pursuant to Rule 8(2).")

4. **The Deletion Cascade: Orchestrating Downstream Processor Compliance**

Finally, Section 12(3) mandates that the obligation to erase extends beyond the Data Fiduciary to all downstream Data Processors.

a. **The Action:** Upon processing a valid erasure request, the primary system must initiate an API-driven cascade signal to all third-party vendors (e.g., Cloud Storage Providers, Email Marketing Agencies), instructing them to purge the specific PII from their sub-systems.

b. **The Compliance Check:** Governance Professionals must ensure that Data Processing Agreements (DPAs) with vendors include specific Service Level Agreements (SLAs) requiring the instantaneous execution of these automated deletion signals.

Figure 2: Automated Data Erasure Cascade – Right To Erasure

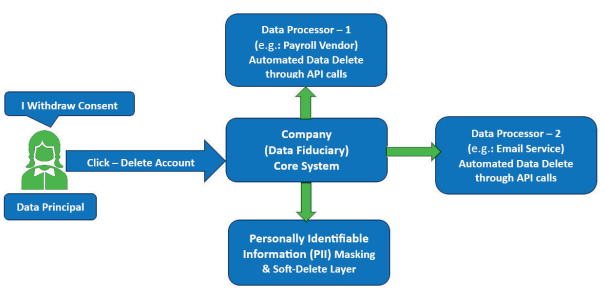

SECURITY SAFEGUARDS & THE BREACH NOTIFICATION ENGINE (SECTION 8 & RULE 7)

The Act mandates a dual layer of protection: **"appropriate technical and organisational measures"** under **Section 8(4)** for general compliance, and specific **"reasonable security safeguards"** under **Section 8(5)** to prevent data breaches. As a Company Secretary, we must audit the **architecture of security**.

1. **Access Control Lists (ACLs) and RBAC**

   "Need to know" is a legal concept; **"Role-Based Access Control" (RBAC)** is the technical implementation.

   a. **The Question for IT:** "Can a junior developer access the live production database containing real customer names?"

   b. **The Requirement:** Operational data should be **masked** for internal staff. Only the automated system should have clear-text access. The CS must inspect the **"Privileged Access Management" (PAM)** logs as part of the internal audit.

2. **Encryption: At Rest and In Transit**

   a. **In Transit:** Data moving from the User's access point — be it a mobile app, browser, or third-party portal — to the Organization's server must be encrypted (e.g., HTTPS/TLS 1.3).

   b. **At Rest:** Data sitting on the hard drive must be encrypted.

   c. **The Techno-Legal Prerequisite:** Encryption is rendered ineffective if the encryption keys are stored next to the data.

3. **The Breach Notification Engine (Rule 7)**

   While **Section 8(6)** of the Act mandates the duty to report, **Rule 7** defines the "Techno-Legal" execution. It creates a dual-notification protocol that must be triggered the moment an organization becomes "aware" of a breach.

   a. **The Data Principal Intimation (Rule 7(1)):** The notice to users must be in "plain manner/ language" and include specific safety measures the user can take. This requires the **Grievance Portal** to have pre-drafted, multi-lingual templates that can be dynamically populated with breach details.

   b. **The Data Protection Board (DPB) Staged Reporting (Rule 7(2)):**

      Stage 1: Immediate Alert [Rule 7(2)(a)]: Requires "nature, extent, timing, and location of occurrence and the likely impact."

      Stage 2: 72-Hour of becoming aware of the breach (or within such longer period as the Data Protection Board of India (DPBI) may allow on a request made in writing in this behalf) **Comprehensive Report [Rule 7(2)(b)]:** This is where the technical burden peaks. The report must include the "root cause" (events and reasons), findings on the perpetrator, and remedial measures taken.

   c. **The CS Role:** We cannot manually gather this metadata during a crisis. The **"Incident Response Platform" (IRP)** must be pre-configured to auto-populate the specific fields required by Rule 7.

   d. **The Technical Hurdle:** If the IT dashboard merely shows "Database Unauthorized Access," but Rule 7(2)(b) requires a "report regarding intimations given to affected Data Principals," the system must be able to cross-reference the database logs with the communication logs (Email/SMS/App/Others) to prove that every affected user was notified "without delay." Without this **automated cross-referencing**, the CS cannot sign off on a compliant 72-hour report.

THE SIGNIFICANT DATA FIDUCIARY (SDF) & THE SYSTEM AUDIT

For organizations classified as **Significant Data Fiduciaries (SDF)**, the bar is higher. We are required to appoint an independent Data Auditor and conduct Data Protection Impact Assessments (DPIAs).

1. **The DPIA as a "Product Spec"**

   A Data Protection Impact Assessment (DPIA) is not a generic risk document; it is a technical stress test of the product architecture.

   a. **The Scenario:** The Company wants to launch a new AI-driven recommendation engine personalized on user behaviour.

   b. **The DPIA Stress Test:** The assessment must analyse the **"Training Data Pipeline"**

      The Question: "If a Data Principal withdraws consent under Section 6(4) or demands Erasure under Section 12, can we surgically remove their data from the AI model?"

      The Technical Reality: Current AI architectures (specifically Neural Networks) store information in "weights," not database rows. True "Machine Unlearning"— removing a specific user's **mathematical trace** without retraining the entire model — is currently a scientifically unsolved problem for most complex models.

2. **CS Guidance (The "Privacy by Design" Fix)**

   a. **The Risk:** If the AI model "memorizes" PII, and we cannot erase it upon request, the Company is in violation of Section 12.

   b. **The Solution:** The CS must advise the Board to enforce **"Anonymization at Source."** The Engineering Team must scrub/mask all PII *before* it enters the training set. If the AI never learns "Who" the user is, only "What" they did, the need for unlearning is mitigated.

THE AUTOMATED GRIEVANCE REDRESSAL (SECTION 13 & RULE 14)

While "Consent" and "Erasure" focus on the happy path of user interaction, **Section 13** addresses the friction: the Grievance Redressal Mechanism. Under the Act, a Data Principal has the right to approach the Data Protection Board (DPB) only *after* exhausting the internal grievance options.

Rule 14 of the DPDP Rules, 2025, transforms this from a soft skill into a hard technical specification. It mandates that the Data Fiduciary must respond within a "reasonable period" (not exceeding 90 days, though industry best practice is **15 to 30 days**) and requires the issuance of a unique reference number.

1. **The "Identifier" Validation Protocol (Rule 14(1), (Rule 14(3) & Rule 14(5))**

   Rule 14(1)(b) mandates that the Data Fiduciary must publish the specific "particulars" required to identify the user. **Rule 14(5)** defines these "Identifiers" as technical sequences — ranging from a **Customer Identification File (CIF) Number** to an **Application Reference Number** or **Enrolment ID**.

   a. **The Compliance Gap:** A generic "Contact Us" form that only asks for "Name" and "Message" is technically non-compliant. It fails to capture the unique Identifier required by Rule 14(5) to map the user to their data accurately.

   b. **The Techno-Legal Fix:** The CS must instruct the UI/UX team to update the Grievance Portal. The form must include a **validated field** for the specific Identifier (e.g., "Enter 12-digit Customer ID").

      Backend Logic: The system must instantly validate this ID against the Master Database before allowing the ticket to be created.

2. **The "Unstructured Repository" Risk (Rule 14(3))**

   Rule 14(3) explicitly demands that the company - within a reasonable period not exceeding ninety days under its grievance redressal system for responding to the grievances of Data Principals and shall, for ensuring the effectiveness of the system in responding within such period, implement appropriate technical and organisational measures.

   a. **The "Shared Inbox" Failure:** Utilizing a basic email alias (e.g., *privacy-grievance@company.com*) is a potential violation of Rule 14(3). An email inbox is an "unstructured repository" — it lacks the built-in technical measures required to automate aging alerts, enforce escalation workflows, or mandate identity verification.

   b. **The Risk - The Communication Auditability Gap:** Relying on manual email creates a dual-layered compliance risk. First, an **Inbound Failure** occurs when a user's grievance is **obscured** by spam filters or lost within fragmented threads, leading to a statutory deadline breach. Second, an **Evidentiary Weakness** arises when the organization attempts to prove compliance; while a manual email screenshot is a **fragile** exhibit that can be challenged for authenticity, a system-generated log provides immutable, timestamped proof of the entire communication lifecycle.

3. **The "Ticketing System" Architecture**

   To satisfy Rule 14(3), the CS must mandate the implementation of a **Privacy CRM or Ticketing System** (similar to Customer Support, but protected for privacy).

   a. **Step 1: Auto-Acknowledgement (The Statutory Receipt):** When a user submits a complaint, the system must instantly generate a unique **"Ticket ID"** (Reference Number) and email it to the user.

      Legal Value: This serves as the definitive evidence that the organization acknowledged the complaint "without delay.

   b. **Step 2: The SLA Countdown Clock:** The system must have a hardcoded **Service Level Agreement (SLA)** timer.

      Techno-Legal Logic: While Rule 14(3) allows up to 90 days, relying on the maximum limit is high-risk. The system timer should be configured to an internal target of 15 days (T-minus-15). This ensures that the grievance is resolved well before a Data Principal can escalate the matter to the Data Protection Board.

   c. **Step 3: Escalation Matrix:** If the ticket is still "Open" as the deadline approaches, the system should automatically trigger an escalation alert to the Company Secretary's dashboard. This prevents **"human error"** from turning into a **"statutory violation."**

4. **The "Closure Report" Artifact**

   Finally, the resolution cannot be conveyed via informal channels. The system must automatically generate a PDF **"Closure Report"** summarizing the investigation findings and the final decision.

   a. **The Strategic Importance:** In the event of an appeal to the Data Protection Board (DPB), this system-generated Closure Report serves as the **Data Fiduciary's primary defence exhibit**. It provides verifiable proof that the organization acted with due diligence, successfully authenticated the user's Identity (Rule 14(5)), and responded strictly within the statutory timeline.

CONCLUSION: THE NEW GOVERNANCE ARCHITECT

The paradigm of managing privacy is **evolving beyond** static policy documents. As demonstrated by the **"Techno-Legal Governance Stack"** (Figure 1) and the **"Automated Data Erasure Cascade"** (Figure 2), true compliance with the DPDP Act, 2023 and Rules, 2025 lies not in the boardroom resolution, but in executing on the server.

For the modern Company Secretary, this represents a **significant strategic shift**. We must **expand** the traditional role of a Compliance Officer who **primarily** interprets the law, to become a 'Techno-Legal Governance Architect' who helps embed compliance into the system. This requires **extending**

our engagement beyond the boardroom to **understand the operational realities** of the technical infrastructure. By asking the right questions — about SHA-256 Hash Values, JSON logs, and Purge Cycles — the CS ensures that the organization's legal intent is faithfully translated into technical reality.

Ultimately, the gap between a legal mandate and a software feature is where liability lives. It is our new professional duty to bridge that gap, ensuring that the digital interactions are evidence-ready under the **BSA, 2023** and aligned with engineering best practices like **IS 17428** and **ISO 27701**. As we move into this new era of digital governance, the most valuable professional will be the one who can stand in a boardroom and explain the legal risks of a "privacy policy," and then walk into the server room and explain the architectural necessity of "immutable audit logs." It is time for us to move from drafting resolutions to designing techno legal frameworks.

ANNEXURE: TECHNO-LEGAL READINESS FOR THE COMPANY SECRETARY

To effectively bridge the gap between compliance mandates and technical execution, the following framework serves as a direct alignment for the CS to use during interface meetings with the CIO/CTO:

Legal Requirement	Legal Reference (Act & Rules)	Technical Question for the CIO/CTO	Desired Answer (The "Pass" Criteria)
Notice (Multilingual)	Section 5(3)	"Does our CMS support dynamic rendering of the privacy notice in 22 regional languages?"	"Yes. We use a headless CMS where we update the legal text centrally, and it reflects across all regional interfaces instantly."
Notice (Itemized)	Section 5(1) & Rule 3	"Is our notice 'independent' of other terms and does it provide an itemized list of data collected?"	Yes. It is architected as a standalone module with specific data-point mapping (not embedded within general T&Cs).
Consent Logging	Section 6(1) & Rule 4	"Do we log the specific version of the terms the user agreed to?"	"Yes. The database stores the Notice_Version_ID, Language_ Code, and a timestamped Hash Value of the consent artifact."
Data Erasure	Section 12 & Rule 8	"Is our database 'purge cycle' set to run efficiently upon request?"	"Yes. We run daily 'Soft Delete' scripts. PII is masked immediately, but traffic logs are migrated to the 'Immutable Audit Vault' for the mandatory 1-year retention."
Breach Reporting	Section 8(6) & Rule 7	"Can our dashboard auto-generate the fields required for the DPB notification?"	"Yes. The Incident Response Platform maps directly to Rule 7 fields and tracks user notifications for the compliance report."
Grievance Redressal	Section 13 & Rule 14	"Do we have a ticketing system with a hardcoded SLA timer for privacy complaints?"	"Yes. The system requires a valid Identifier (Rule 14(5)) and triggers an escalation alert to Legal if a ticket remains open past Day 12 (internal warning)."

REFERENCES:

1. Primary Legal Statutes

i. The Digital Personal Data Protection Act, 2023 (Act No. 22 of 2023).

ii. The Digital Personal Data Protection Rules, 2025 (Ministry of Electronics and Information Technology).

iii. The Bharatiya Sakshya Adhiniyam, 2023 (Act No. 47 of 2023); specifically, Sections 57, 61, and 63 regarding electronic evidence.

2. Technical Standards

iv. ***IS 17428 (Part 1): 2020*** *– Data Privacy Assurance: Engineering and Management Requirements (Bureau of Indian Standards).*

v. ***ISO/IEC 27701:2019 –*** *Privacy Information Management System (PIMS) (International Organization for Standardization).*

3. Professional Standards

vi. ICSI Auditing Standards (CSAS-1 to CSAS-4) – The Institute of Company Secretaries of India.

Appendix: Glossary of Techno-Legal Terms

vii. ***API (Application Programming Interface)*** *A software connector that allows two different applications to communicate and exchange data automatically.*

viii. ***Cron Job*** *An automated background task scheduler that executes specific software scripts at pre-set times (e.g., midnight data purges).*

ix. ***Hash Value*** *A unique digital "fingerprint" of a file used to verify its integrity; any alteration to the file changes the hash, making it crucial for digital evidence (BSA Section 63).*

x. ***Immutable Audit Vault (WORM)*** *"Write-Once-Read-Many" storage where data cannot be modified or deleted once written, ensuring tamper-proof compliance logs.*

xi. ***JSON (JavaScript Object Notation)*** *A standard, lightweight text format used for storing and exchanging data in human-readable "key-value" pairs (e.g., {"Consent": "True"}).*