GRG EDUCATIONAL INSTITUTIONS

GRG Educational Institutions

INSTITUTIONAL AI GOVERNANCE POLICY

Effective from Jan 2026

Issued by the

Chairperson

GRG Educational Institutions

Applies to all institution-approved AI systems — generative AI assistants, source-grounded research tools, and AI-enabled learning, assessment & proctoring platforms — regardless of vendor. See Appendix B for the current tool inventory.

Field	Detail
Document Title	Institutional AI Governance Policy 2026
Document Owner	AGM IT & DS
Approving Authority	Governing Council / Principal
Classification	Internal — Staff & Governing Body
Effective Date	01 January 2026
Next Scheduled Review	01 January 2027 (or upon material change to platforms, law, or risk profile)
Version	1.2
Standards Referenced	ISO/IEC 27001:2022, ISO/IEC 27002:2022, ISO/IEC 27701:2019 (Privacy Information Management), ISO/IEC 27017:2015 and ISO/IEC 27018:2019 (Cloud Security & PII in Cloud), ISO 22301:2019 (Business Continuity), ISO/IEC 42001:2023 (AI Management), ISO/IEC 23894:2023 (AI Risk Management), CERT-In Directions 2022, India DPDP Act 2023 and DPDP Rules 2025, Google Workspace for Education security guidance

Revision History

Version	Date	Author	Summary of Change
0.1	Oct 2025	IT Governance Working Group	Initial draft circulated for stakeholder review
0.2	Nov 2025	AGM IT&DS	Added category-based AI Governance Framework
1.0	Jan 2026	Governing Council	Approved for institution-wide rollout
1.1	Aug 2026	AGM IT&DS	Standards gap-analysis remediation: added CERT-In 6-hour reporting (Section 12), cited DPDP Rules 2025 and 72-hour Board timeline (Cover page, Section 12), and named a Grievance Officer role (Section 3).
1.2	Aug 2026	AGM IT&DS	Standards gap-analysis remediation, remaining findings: added ISO/IEC 27701, 27017/27018, 22301 and 23894 references; PCI DSS data row and vendor check (Sections 6, 10); broadened cross-border privacy language (Section 6); encryption minimums (Section 6); vulnerability-scan/pen-test cadence (Section 11); anonymous reporting channel (Section 12); Risk Tier and Last Reviewed columns (Appendix B).

Table of Contents

1. Purpose and Scope

This policy establishes the institution's requirements for information security and artificial-intelligence (AI) governance for the 2026 academic year and beyond. It applies to all staff, faculty, students, contractors, and third parties who access institutional systems, data, or any institution-approved AI system.

The policy governs, without limitation, the institution's core cloud productivity environment and every AI system operating on top of it — spanning generative AI assistants, source-grounded research tools, and AI-enabled learning, assessment, and proctoring platforms — together with all networks, Endpoints, and cloud services used to deliver teaching, assessment, research, and administration. The policy is written to be vendor-neutral: it governs by category of AI capability and risk, not by product name, so that it remains valid as specific tools are added, replaced, or upgraded. The specific tools currently deployed against each category are maintained in Appendix B and reviewed independently of this policy body.

This Institutional AI Governance Policy is designed to satisfy international and nationally recognized information security management standards, to reflect the cloud provider's published security guidance for education customers, and to incorporate emerging best practices in AI governance — including principles for managing AI-related risk, promoting the responsible use of generative AI in educational settings, and ensuring sound oversight of AI management systems.

2. Policy Statement and Objectives

The institution is committed to protecting the confidentiality, integrity, and availability of its information assets, and to deploying AI systems in a manner that is lawful, transparent, fair, and subject to meaningful human oversight. The objectives of this policy are to:

- Maintain a certifiable Information Security Management System (ISMS) aligned to ISO/IEC 27001:2022.
- Ensure safe, pedagogically sound, and privacy-respecting use of every institution-approved AI system, regardless of vendor.
- Protect the personal data of students and staff, with particular care for minors and Restricted-tier data.
- Preserve academic integrity in an environment where generative AI is available to students and staff.
- Maintain institutional resilience against cyber incidents, and readiness to respond and recover.
- Build a culture of security and AI literacy across the community through ongoing training.

3. Governance Structure and Roles and Responsibilities

Effective governance requires clear ownership. The table below summarises the principal roles under this policy; a full RACI matrix is maintained separately by the AGM IT&DS.

Role	Key Responsibilities
Governing Council / Principal	Ultimate accountability for the ISMS and AI governance programme; approves policy and risk appetite; signs off Statement of Applicability.
AGM IT&DS (ISMS Owner)	Day-to-day ownership of information security, this policy, risk register, incident response, and vendor security assessments.
Grievance Officer	Published point of contact for Data Principals under the Digital Personal Data Protection Act 2023; receives and responds to data-protection grievances and rights requests. Name and contact details to be designated by the Governing Council and published in

Role	Key Responsibilities
	Appendix A. Pending legal confirmation of whether the institution meets Significant Data Fiduciary criteria under DPDP Act 2023 Sections 10–11, which would additionally require an India-based Data Protection Officer, periodic independent data audits, and expanded DPIA obligations.
AI Governance Committee	Cross-functional body (IT, academics, legal, student representative) that reviews, approves, and monitors all institution-approved AI systems against the category framework in Section 8.
System / Application Owners	Configure and maintain each platform's tenant/admin settings — including any generative AI, research-tool, or assessment-platform features — in line with this policy.
Faculty & Staff	Use approved systems appropriately, complete mandatory training, and report incidents or suspected AI misuse.
Students	Comply with acceptable-use and academic-integrity rules, including AI-use disclosure requirements.
Internal / External Auditors	Independently verify control effectiveness and compliance at least annually.

4. Definitions

- ISMS — Information Security Management System, the overall framework of policy, process, and controls.
- Restricted data — the institution's highest data-sensitivity tier (see Section 6).
- Generative AI — AI systems that produce text, summaries, or other content from prompts or source material.
- Source-grounded AI — AI systems that answer questions or summarize strictly from a defined set of user-supplied documents.
- AI-proctoring — automated monitoring of an exam session using webcam, screen, or behavioural signals to flag possible integrity concerns.
- Human-in-the-loop — a mandatory human review step before an automated system's output is allowed to produce an academic, disciplinary, or employment consequence.
- Shadow AI — any AI tool accessed via a personal account or outside institutional management, not covered by an institutional Data Processing Agreement or admin-console visibility (see Section 8.9).
- DPIA — Data Protection Impact Assessment, a structured risk assessment for processing that is likely to result in high risk to individuals.

5. Information Security Management System (ISO/IEC 27001:2022 Alignment)

The institution operates a risk-based ISMS following the Plan-Do-Check-Act cycle described in ISO/IEC 27001:2022 Clauses 4–10. Risk assessment and treatment are performed at least annually, and whenever a new AI system in any category defined in Section 8 is introduced or materially changed. Results are recorded in the institutional risk register and reflected in a maintained Statement of Applicability (SoA). Given the volume of personal data processed under Section 6 (student records, minors' data, health data, and proctoring signals), the ISMS is progressively being extended to align with ISO/IEC 27701:2019 (Privacy Information Management), giving privacy controls the same certifiable footing as the security controls below.

5.1 Annex A Control Domains

Annex A Domain (ISO/IEC 27001:2022)	# Controls	Institutional Implementation Focus
A.5 Organizational Controls	37	Information security policy suite, roles (above), supplier/vendor management for all cloud and AI vendors, threat intelligence, ICT continuity, legal/regulatory register.
A.6 People Controls	8	Screening, terms of employment, security awareness & AI-literacy training, disciplinary process, remote-work rules, confidentiality agreements.
A.7 Physical Controls	14	Secure server/network rooms, campus device custody, clear-desk/clear-screen in exam and admin offices, equipment disposal (data wiping).
A.8 Technological Controls	34	Endpoint protection, MFA & SSO via the institutional identity provider, logging/monitoring (incl. AI usage logs), secure configuration of all approved AI systems, data leakage prevention, cryptography, capacity & change management.

5.2 Cloud and AI Service Controls

For cloud and AI services specifically, the institution applies the following baseline controls, mapped to Annex A. Cloud-service risk assessments are conducted against ISO/IEC 27017:2015 (cloud service security controls) and ISO/IEC 27018:2019 (protection of PII in public clouds) as the independent assurance baseline, rather than relying solely on a vendor's own published security guidance:

- A.5.23 (Cloud service security) — every cloud and AI vendor is subject to a documented cloud-service risk assessment before use and annually thereafter.
- A.5.34 (Privacy and protection of PII) — all AI tools are configured so that Restricted-tier personal data cannot be freely pasted or uploaded without a compensating control (see Section 8).
- A.8.16 (Monitoring activities) — admin console and API logging is enabled for every approved AI system to support incident investigation.
- A.8.23 (Web filtering) and A.8.7 (Malware protection) — applied consistently across campus networks and managed Endpoints.

6. Data Classification and Handling

All institutional data must be classified into one of four levels. Classification determines where data may be stored, which systems it may be processed by (including AI tools), and applicable retention and disposal rules.

Level	Examples	Handling Baseline
Public	Prospectus, published research, marketing content	May be posted openly; no special controls beyond integrity checks.
Internal	Timetables, internal memos, non-sensitive staff communications	Restricted to institutional accounts; not for external sharing without approval.
Confidential	Staff HR records, financial data, unpublished research, contracts	Access on least-privilege basis; encryption at rest/in transit; no upload to non-approved AI tools.

Level	Examples	Handling Baseline
Restricted	Student PII, grades, health/disability records, exam content, proctoring recordings, minors' data	Highest control tier: MFA-gated access, DPIA required, retention limits, never used to train third-party AI models, breach-notification obligations apply. Minimum encryption: AES-256 at rest, TLS 1.2 or higher in transit.
Payment Data	Cardholder data or payment tokens arising from admissions, tuition, or fee transactions, whether processed directly or via a third-party gateway	Governed by PCI DSS v4.0 regardless of whether card numbers are stored on institutional systems; Finance/IT to confirm scope, and any payment processor is subject to the Section 10 vendor checklist plus PCI DSS evidence.

Restricted-tier data — including exam content, proctoring recordings, health information, and records concerning minors — must never be entered into a generative AI assistant or source-grounded research tool outside of the institution's managed tenant configuration, and must only be processed by an AI-enabled learning, assessment, or proctoring platform under the terms of its signed Data Processing Agreement. Applicable regulatory regimes (including India's Digital Personal Data Protection Act 2023 and DPDP Rules 2025; GDPR/FERPA where the institution has EU or US-linked students or partnerships; and Australia's Privacy Act 1988 or New Zealand's Privacy Act 2020 where the institution has Australian or New Zealand students, staff, or partnerships) should be confirmed with legal counsel and mapped to this classification scheme.

7. Acceptable Use of Institutional Systems

Institutional cloud productivity, learning-platform, and AI accounts are provided for teaching, learning, research, and administrative purposes and remain institutional property.

7.1 Permitted Use

- Use of institutional email, document, generative-AI, research-tool, and learning/assessment-platform accounts for coursework, research, communication, and administration.
- Reasonable incidental personal use that does not breach law, this policy, or consume disproportionate resources.

7.2 Prohibited Use

- Sharing account credentials, disabling MFA, or circumventing device-management controls.
- Uploading Restricted-tier data to any AI tool not explicitly approved for that data type.
- Using generative AI to impersonate another person, fabricate research data, or produce content that violates academic-integrity rules without disclosure.
- Attempting to bypass an exam platform's proctoring controls during an assessment.
- Using a personal account on an external, non-approved AI tool (e.g., ChatGPT, Claude.ai, Copilot, or similar) to process Confidential or Restricted institutional data, or as the system of record for grading, admissions, or disciplinary decisions (see Section 8.9).
- Installing unauthorized software or connecting unmanaged third-party apps to institutional accounts.

8. AI Governance Framework

This section sets governance for institution-approved AI systems by category of capability and risk, rather than by product name, so the framework remains valid as specific tools are added, replaced, or upgraded. It is consistent with emerging AI-governance principles: transparency, human oversight, fairness, accountability, and data

minimization. The tools currently deployed against each category are listed for reference in Appendix B; that inventory is reviewed and updated independently of this policy body.

8.1 Guiding Principles

- Transparency — students and staff are told when AI is used in teaching, assessment, or decision-making that affects them.
- Human oversight — no AI output alone may determine a grade, admissions outcome, disciplinary finding, or employment decision (see 8.6).
- Fairness and bias mitigation — AI-assisted assessment tools, particularly automated proctoring flags, are periodically reviewed for disparate impact.
- Data minimization — only the data necessary for the tool's purpose is processed; Restricted data is withheld unless expressly authorized.
- Accountability — the AI Governance Committee maintains a tool inventory mapped to these categories and reviews new AI requests before adoption.

8.2 AI System Categories and Risk Tiers

Every institution-approved AI system is classified into one of the following categories, which determines its baseline governance requirements:

Category	Description	Data Typically Processed	Risk Tier	Governance Owner
Generative AI Assistants	General-purpose tools used for drafting, planning, summarizing, and administrative productivity	Internal & Confidential content; Restricted data prohibited	Medium	AGM IT&DS + AI Governance Committee
Source-Grounded Research & Study Tools	Tools that answer questions or summarize strictly from a defined set of user-supplied documents	Course materials, notes, research documents (no Restricted data)	Medium	AI Governance Committee
AI-Enabled Learning, Assessment & Proctoring Platforms	LMS, coding-lab, placement, and examination systems that use AI to deliver, score, or monitor academic work, including automated exam proctoring	Student academic records, exam responses, proctoring signals such as webcam/screen data (Restricted)	High	AGM IT&DS + Examinations Controller
AI-Assisted Recruitment / HR Tools	Tools used to support shortlisting or screening of candidates or staff	Candidate and staff personal data (Restricted)	High	AGM IT&DS + HR

8.3 Generative AI Assistants

Tools in this category are approved for lesson planning, drafting, research assistance, and administrative productivity within a managed institutional tenant. Administrators must confirm, via the vendor's admin console, that conversation content is not used to train the vendor's general-purpose models and that data-processing terms applicable to education customers are in force. Staff and students must not paste Restricted-tier data (see Section 6) into prompts, and any AI-generated content used in official communications or grading feedback must be reviewed by a human before release.

8.4 Source-Grounded Research and Study Tools

Tools in this category are approved for study aids, literature summarization, and research support drawn strictly from documents the user has the right to upload. Users remain responsible for verifying AI-generated summaries and citations against the original source before relying on them academically. Restricted-tier data (e.g., another student's personal records) must not be uploaded as a source.

8.5 AI-Enabled Learning, Assessment and Proctoring Platforms

Tools in this category underpin learning management, coding labs, placement preparation, and AI-proctored online examinations. Any platform that uses webcam, screen, or behavioral signals to monitor an exam session is treated as High risk and subject to additional safeguards:

- Informed consent — students are notified in advance of what the proctoring system monitors, how long recordings are retained, and how flags are reviewed.
- Human review — every integrity flag raised by an automated proctoring system is reviewed by a designated invigilator or Examinations Controller before any academic-integrity process is initiated; no automated flag alone results in a penalty.
- Retention limits — proctoring recordings and biometric-adjacent signals are retained only as long as needed for grade appeals or integrity investigations, then securely deleted per the institution's retention schedule.
- Accessibility — accommodation pathways exist for students whose environment, disability, or connectivity may trigger false-positive flags.
- Vendor assurance — each platform vendor's security posture, sub-processors, and DPA terms are reassessed at least annually against the checklist in Section 10.

8.6 Human-in-the-Loop Requirements

The following decisions require documented human review before being finalized, regardless of which AI system produced supporting output:

1. Final course grades and progression decisions.
2. Academic-integrity findings arising from AI-proctoring flags or AI-detection tools.
3. Admissions and scholarship decisions.
4. Staff or candidate recruitment shortlisting decisions supported by an AI-assisted recruitment tool.
5. Any disciplinary action against a student or staff member.

8.7 AI Literacy and Academic Integrity

Students must disclose material use of generative AI in submitted coursework in line with each course's stated AI-use policy, ranging from unrestricted use to full prohibition depending on the learning objective. Faculty are supported with guidance and exemplars to set clear, course-level AI-use expectations, and annual AI-literacy training is provided to staff and students alike (see Section 14).

AI should support human thinking, not replace it. Consistent with this principle, and to give the disclosure requirement above practical effect at the level of a single assignment, faculty must classify each significant assessment into one of the following categories, communicated clearly in the assessment instructions:

- **A. AI Prohibited** — AI tools may not be used. The assessment is intended to demonstrate independent knowledge or capability.
- **B. AI Minimal** — Limited use may be permitted, such as basic language or formatting assistance, as specified by the faculty member.
- **C. AI Permitted** — Students may use specified AI tools for defined purposes, with appropriate acknowledgement and critical evaluation.
- **D. AI-Integrated** — AI is intentionally incorporated into the learning task, and students are expected to demonstrate their ability to use, evaluate, and critically engage with AI.

This category-based approach reflects emerging practice of giving students explicit, assessment-level AI expectations rather than relying on a single blanket institutional rule.

8.8 New AI Tool Onboarding

Any request to adopt a new AI tool, or to expand an existing tool's use beyond its currently approved category and scope, must go through the AI Governance Committee using the vendor risk checklist in Section 10 and, where personal data is materially affected, a DPIA led by the AGM IT&DS. The risk-tier assignment (Low/Medium/High) required by Section 8.2 is determined using ISO/IEC 23894:2023 (AI risk management guidance) or the NIST AI Risk Management Framework 1.0 as the documented methodology, and the resulting assessment is recorded per tool in Appendix B rather than left to an undocumented judgment call. On approval, the tool is added to the inventory in Appendix B under the appropriate category from Section 8.2.

8.9 Use of External and Non-Approved AI Tools (e.g., ChatGPT, Claude, Copilot, and similar)

Staff and students may encounter, and in many cases already use, general-purpose AI tools that the institution has not procured, configured, or covered under a Data Processing Agreement — for example, personal accounts on ChatGPT, Claude, Microsoft Copilot, Gemini's free consumer tier, or comparable services. This policy does not attempt to prohibit awareness or personal use of such tools outside institutional business, but it does set firm limits on their use in an institutional context:

- No Confidential or Restricted data — nothing classified Confidential or Restricted under Section 6 (student records, grades, exam content, health data, HR data, unpublished research, contracts, etc.) may be entered into any external AI tool that has not been through the onboarding process in Section 8.8 and added to Appendix B, regardless of the tool's reputation or the individual's intent.
- No institutional decision-making — output from an external, non-approved AI tool must never be the basis for a grade, admissions or disciplinary decision, official institutional communication, or financial/contractual commitment, without independent human verification and, where applicable, routing through an approved system of record.
- No official account or integration — external AI tools must not be connected to institutional email, Drive, LMS, or exam-platform accounts (e.g., via browser extensions, OAuth grants, or API keys), and must not be used to auto-process institutional documents in bulk.
- Personal accounts, personal risk — where permitted incidental use occurs (e.g., general research or personal upskilling), it must use the individual's personal account and personal device or a clearly personal context, not an institutional identity, and the institution provides no support, warranty, or data-protection guarantee for that use.
- Same disclosure rules apply — students using any external AI tool in coursework are bound by the same AI-use disclosure requirements in Section 8.7 as they would be for an approved tool; being off the approved list is not an exemption from academic-integrity disclosure.
- Monitoring — the institution may apply data-loss-prevention (DLP) and network-monitoring controls (Section 5.2, A.8.16) to detect Confidential or Restricted data being submitted to unmanaged AI Endpoints, and will treat confirmed exfiltration as a security incident under Section 12.
- Path to approval — staff or departments who find a genuine need for a specific external tool (including ChatGPT, Claude, or similar) should request it through the AI Governance Committee under Section 8.8 rather than adopting it informally; approved requests are added to Appendix B with an appropriate Section 8.2 category and risk tier, at which point institutional-context use becomes permitted under this policy.

9. Access Control and Identity Management

Access to institutional systems, including every category of AI system defined in Section 8, is provisioned on a least-privilege, role-based basis through a centrally managed institutional identity.

- Single sign-on (SSO) via the institutional identity provider is mandatory for all integrated third-party applications, including AI-enabled learning and assessment platforms, where technically supported.
- Multi-factor authentication (MFA) is mandatory for all staff accounts and strongly enforced for student accounts, particularly those with access to Restricted data or exam systems.
- Joiner-mover-leaver processes ensure accounts and AI-tool entitlements are provisioned, adjusted, and revoked in step with enrolment or employment status.
- Privileged administrative access to any AI system's admin console or institutional dashboard is limited to named individuals and logged.

10. Third-Party and Vendor Risk Management

Every cloud and AI vendor used by the institution (see Appendix B for the current inventory) is subject to the following ongoing assurance checklist, applied at onboarding and reviewed annually:

Assessment Area	Requirement Before Onboarding
Data Processing Agreement	Signed DPA covering purpose limitation, sub-processors, and deletion on termination.
Data Residency & Storage	Confirmed hosting location(s) and whether data leaves the institution's approved jurisdictions.
Model Training Use	Written confirmation that institutional data is not used to train the vendor's general-purpose models.
Security Certification	Evidence of ISO/IEC 27001 (or equivalent) certification or a satisfactory third-party security assessment.
PCI DSS (Payment Vendors)	For any vendor processing admissions, tuition, or fee payments (see Section 6, Payment Data), evidence of PCI DSS v4.0 compliance (e.g., an Attestation of Compliance or Self-Assessment Questionnaire) is required before onboarding.
Human Oversight	Vendor supports a human-review workflow before any automated output triggers an academic or disciplinary consequence.
Accessibility & Fairness	Documented bias-testing approach (e.g., for AI-proctoring flags) and accommodation pathway for students with disabilities.
Exit / Portability	Defined process to export institutional data and confirm secure deletion at contract end.

11. Network and Endpoint Security

- Managed Endpoints (staff and lab devices) run current Endpoint protection and receive security patches within institution-defined SLAs.
- Campus and guest networks are segmented; networks running the institution's online-assessment platform are isolated from general guest traffic during assessment windows.
- BYOD devices accessing institutional accounts must meet minimum security requirements (screen lock, OS patch level, no jailbreak/root) enforced via Endpoint management.
- The network and application estate is scanned for technical vulnerabilities on at least a quarterly basis (ISO/IEC 27002:2022 Control 8.8), with an independent third-party penetration test conducted at least annually and findings tracked to remediation via the risk register.

12. Incident Response and Breach Notification

The institution maintains an incident-response plan covering conventional cyber incidents and AI-specific incidents (e.g., an exposure of proctoring data, or a data-leakage event from a generative AI or research tool).

Phase	Key Actions	Owner
1. Detect & Report	Any staff/student reports suspected incident via IT helpdesk, security@ mailbox, or a hosted anonymous whistle-blower channel for those reluctant to report through a named route; automated alerts from monitoring tools.	All users / SOC tooling
2. Triage & Contain	AGM IT&DS classifies severity, isolates affected accounts/systems, preserves logs (incl. AI-platform audit trails).	AGM IT&DS
3. Assess & Notify	AGM IT&DS assesses personal-data and cyber-incident impact. Where the incident falls within CERT-In's mandated categories, it is reported to CERT-In within 6 hours of detection (IT Act 2000 s.70B; CERT-In Directions, 28 April 2022), with AGM IT&DS (or designated CISO function) as filing owner. Where personal data is affected, an initial notice is issued to the Data Protection Board of India and affected Data Principals without delay, followed by a detailed report to the Board within 72 hours (DPDP Rules 2025, Rule 7), alongside any other applicable law timelines.	AGM IT&DS
4. Eradicate & Recover	Remove root cause, restore from verified backups, re-issue credentials, patch vulnerabilities.	IT Operations
5. Post-Incident Review	Root-cause analysis, lessons-learned report to Governing Council, update risk register and Statement of Applicability.	AGM IT&DS + AI Governance Committee

13. Business Continuity and Backup

- Institutional data hosted with the primary cloud provider benefits from that provider's built-in redundancy; institution-critical configuration and AI-platform academic records are independently backed up on a defined schedule.
- Recovery Time and Recovery Point Objectives (RTO/RPO) are defined per system and tested at least annually, including a tabletop exercise covering loss of access to a core AI platform during an exam period. Business continuity planning, the RTO/RPO framework, and the tabletop exercise programme are aligned to ISO 22301:2019 (Business Continuity Management Systems).

14. Security Awareness and AI-Literacy Training

- All staff complete mandatory annual security-awareness training, including phishing-simulation exercises.
- All staff and students complete AI-literacy onboarding covering the approved AI categories in Section 8, the limits on external/non-approved tools such as ChatGPT or Claude (Section 8.9), prohibited uses, and academic-integrity expectations.
- Role-specific training is provided to Examinations staff on the AI-proctoring review workflow, and to the AI Governance Committee on emerging regulatory developments.

15. Monitoring, Audit, and Continuous Improvement

15.1 Monitoring

The Institution will monitor the use of AI tools to ensure they are used responsibly, securely, and in line with this policy. The IT Department may track system usage, access logs, and security alerts to identify risks and ensure compliance. Content of prompts or AI outputs will not be reviewed unless required for a security investigation, policy violation, or legal/audit purpose.

15.2 Audit

The Institution will conduct regular internal and external audits to check compliance with AI policies and security controls. Audits will review system usage, data protection, licensing, access control, and incident management. Audit findings will be reported, and corrective actions will be taken and tracked until completion.

15.3 Continuous Improvement

The Institution will regularly review and update this policy to keep up with changes in AI technology, regulations, and institutional needs. Feedback, audit results, and incident reports will be used to improve AI practices. The IT Department and AI Governance Committee will support ongoing improvements and awareness activities to ensure safe and effective use of AI tools

16. Compliance, Exceptions, and Enforcement

Breach of this policy may result in suspension of system access, disciplinary action under applicable staff or student conduct procedures, and, where warranted, referral to law enforcement or regulators. Any exception to this policy must be requested in writing to the AGM IT&DS, include a compensating control, and be time-bound and formally approved.

17. Policy Review and Version Control

This policy is reviewed at least annually, and immediately following any material change to an institution-approved AI system, a significant security incident, or relevant new legislation. The AGM IT&DS is responsible for maintaining the version history recorded on the cover page and the tool inventory in Appendix B.

Policy Approval

This policy, including all amendments recorded in the Revision History above, is formally approved for institution-wide adoption by the Approving Authority named on the cover page.

Name: R. NANDINI Designation: Chairperson, Governing Council

Signature: _____ Date: 4/09/26

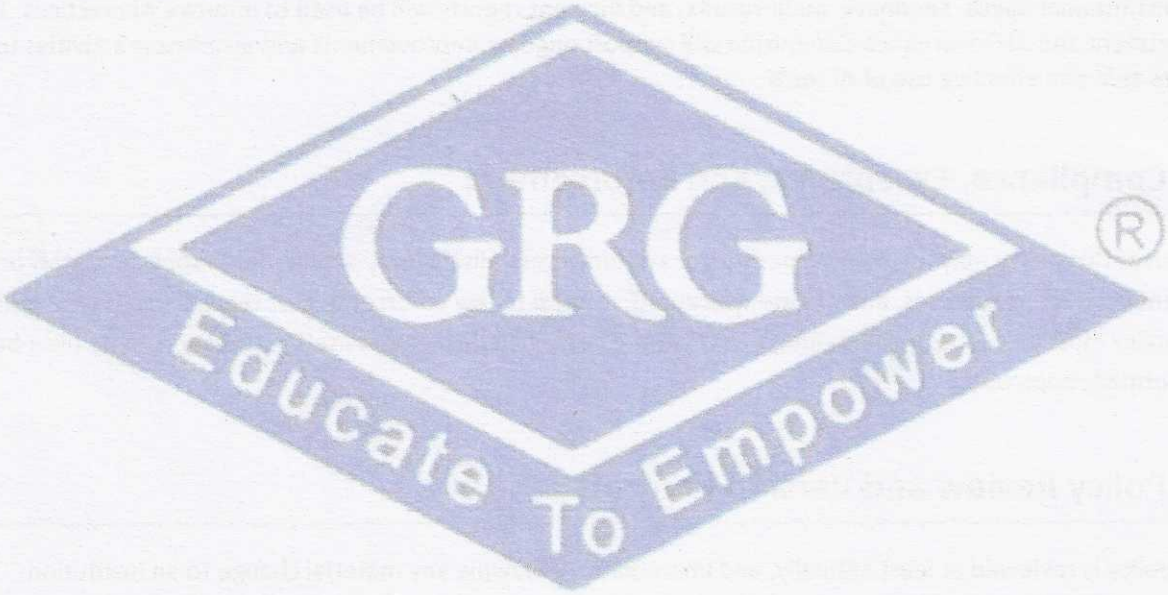

Appendix A: Acceptable Use Acknowledgement

I confirm that I have read and understood the Institutional IT & AI Governance Policy 2026, including the requirements governing use of all institution-approved AI systems described in Section 8, and agree to comply with its terms.

Name: _____ Role: _____

Signature: _____ Date: _____

Appendix B: Current AI Tool Inventory (Reference)

This appendix lists the specific AI tools presently deployed by the institution and maps each to the category framework defined in Section 8.2. Unlike the policy body, this inventory is expected to change over time as tools are added, replaced, or upgraded — it is maintained and reviewed by the AGM IT&DS and AI Governance Committee independently of the policy itself, and does not require a full policy revision to update.

Deployed Tool	AI Category (Section 8.2)	Vendor	Risk Tier (Section 8.2)	Last Reviewed	Review Cycle
Google Gemini Pro (Workspace for Education)	Generative AI Assistant	Google	Medium	Jan 2026	Annual
NotebookLM	Source-Grounded Research & Study Tool	Google	Medium	Jan 2026	Annual
AMYPO Suite — DotLab, DotPrep, DotExam/AmypoShield, DotRecruit, DotCoder	AI-Enabled Learning, Assessment & Proctoring Platform (DotRecruit also maps to AI-Assisted Recruitment/HR)	Amypo Technologies	High	Jan 2026	Annual