

Artificial Intelligence (AI) Usage & Governance Policy

Policy Changes Log

Version	Issued On	Effective Date	Purpose of Change
1	10 th September, 2026	10 th September, 2026	New policy release

Table of Contents

1. Purpose	5
2. Scope	5
3. Definitions	5
3.1 Artificial Intelligence (AI)	5
3.2 Public AI Tools	6
3.3 Enterprise AI Tools	6
3.4 Private AI Services	6
3.5 GIG Information or Data	6
3.6 Consequential AI Decision	7
4. Core Policy Principles	7
4.1 Human Accountability	7
4.2 Data Protection & Confidentiality	8
4.3 Approved Enterprise AI Tools	9
4.4 Accuracy, Transparency, Explainability & Ethical Use	9
5. AI Risk Classification & Approval Framework	9
5.1 AI Risk Classification Matrix	11
5.2 General Requirements Across All Risk Levels	12
6. AI Solution Acquisition and Development	12
6.1 Procurement of AI-Enabled Solutions	12
6.2 Development Approval	14
6.3 Training Data Management	14
6.4 Use of External AI Models and Software Components	14
6.5 Secure Development Practices	14
6.6 Protection of AI Instructions and Configurations	14
6.7 Documentation and Transparency	15
6.8 Change and Version Management	15
6.9 Human Oversight and Compliance	15
6.10 Secure AI Development Standards	15
6.11 Selection and Registration of Foundation Models	15
6.12 AI Features Embedded in Third-Party and SaaS Applications	16
7. AI Deployment and Lifecycle Management	16
7.1 Pilot Evaluations	16
7.2 Deployment Approval	17
7.3 Validation Before Deployment	17
7.4 Monitoring and Review	17
7.5 Incident Management	17
7.6 Logging and Audit	17
7.7 Retirement and Decommissioning	17
8. AI Governance, Oversight and Continuous Improvement	18
8.1 Governance Ownership and Oversight	18
8.2 Review, Monitoring and Effectiveness Checks	18
8.3 Issue Management and Corrective Actions	19
8.4 Continuous Improvement of AI Governance	19
9. Reporting	19
10. Training	20
11. Violations	20
12. Policy Review	20
13. Annexure	20
13.1 AI Strategy and Operating Model	20
13.1.1 Purpose and Intent	20
13.1.2 AI as a Decision Support Capability	20
13.1.3 AI Operating Model	21
13.1.4 Central Enablement	21
13.1.5 Federated Business Ownership	21
13.1.6 Joint Alignment	21
13.1.7 Value Orientation and Scaling	21
13.1.8 Shared Intelligence Across the Group	21
13.1.9 Sensing and Management Intelligence	21
13.1.10 Leadership Responsibility	22
13.1.11 Relationship to Other Policy Sections	22
13.2 Standards and Controls Mapping	22

14. Policy Endorsement 24

1. Purpose

This Artificial Intelligence (AI) Usage & Governance Policy ("Policy") establishes mandatory requirements for the responsible, secure, and compliant use of Artificial Intelligence (AI) technologies across the Godrej Industries Group ("GIG").

The objectives of this Policy are to:

- Maintain human accountability in AI-assisted decision-making
- Enable safe and responsible adoption of AI technologies.
- Protect GIG's confidential, personal, and proprietary information.
- Ensure compliance with applicable legal and regulatory requirements.
- Provide a structured risk-based governance framework.

GIG recognizes that AI offers material operational, commercial, and analytical benefits, and this Policy is designed to enable and best derive that value. GIG also acknowledges that AI systems, including productivity tools, can interact with Personal Data, business-sensitive information, and regulated processes in ways that are not always visible to end users. This Policy aims to ensure that such interactions are understood, managed and accountable.

2. Scope

This Policy:

- applies to all GIG employees, contractors, consultants, interns and service providers
- applies to all AI systems, whether internally developed, fine-tuned on GIG data, embedded within enterprise systems (e.g., ERP, CRM), provided or developed by third-party vendors, or delivered via enterprise-approved generative AI platforms (e.g., Claude Enterprise, ChatGPT Enterprise)
- applies to all GIG Information Assets and enterprise environments

This Policy applies globally across all GIG entities.

3. Definitions

3.1 Artificial Intelligence (AI)

Refers to software systems that use machine learning models, large language models (LLMs) (including public and private/enterprise LLMs), statistical models, knowledge-based systems or algorithmic techniques to generate outputs, make predictions, support decisions, automate processes or perform tasks that would typically require human intelligence.

AI within GIG includes systems designed to:

- **Generate content:** Such as produce text, code, reports, presentations, communications, images, or summaries using generative AI models (e.g., Large Language Models such as ChatGPT, Copilot, or similar tools).
- **Specialized or Fine-Tuned Language Models:** Domain-specific LLMs that are trained, fine-tuned, or configured using GIG data, industry data, or proprietary knowledge bases to support functions such as procurement, HR, manufacturing, legal and compliance documentation, or customer service.
- **Analyse and Predict Outcomes:** Such as to perform forecasting, demand planning, pricing analysis, predictive maintenance, financial modelling, or risk assessment using structured or unstructured data.
- **Automate Decision-Support Processes:** Such as to assist in resume screening, quality deviation detection, compliance flagging, customer complaint triaging, or lead prioritization.
- **Optimize Operational Systems:** Such as to improve manufacturing processes, supply chain efficiency, logistics, or inventory planning through algorithmic optimization.
- **Monitor and Detect Anomalies:** Such as to identify safety risks, environmental deviations, fraud indicators, system irregularities, or abnormal operational patterns.

3.2 Public AI Tools

Public AI Services means AI services that are publicly available and not covered by a GIG-approved enterprise agreement or subject to GIG-defined security, privacy, compliance, and governance controls, regardless of whether such services are free or paid.

Examples include but are not limited to:

- ChatGPT (Consumer)
- ChatGPT Plus (Personal Subscription)
- Claude (Consumer)
- Claude Pro (Personal Subscription)
- Gemini (Consumer)
- DeepSeek
- Perplexity Consumer

3.3 Enterprise AI Tools

Enterprise AI Services means AI services provided under a GIG-approved enterprise agreement that incorporates contractual, technical, and operational controls for security, privacy, compliance, identity management, audit logging, and data protection.

Examples include but are not limited to:

- ChatGPT Enterprise
- Claude Enterprise
- Gemini Enterprise
- Microsoft 365 Copilot Enterprise
- Higgsfield Enterprise

3.4 Private AI Services

Private AI Services means AI systems deployed within a dedicated environment controlled by or exclusively provisioned for GIG, where data processing, retention, storage, access, and model governance are subject to GIG-approved controls.

3.5 GIG Information or Data

Means any data of a member of GIG including (without limitation) personal data, financial data, Intellectual Property, operational data, customer data, strategic data and any other information and know-how of or pertaining to any member of GIG and contained, transmitted or generated using an Information Resource. Below are the key definitions that are relevant to this Policy:

- **Personal Information:** means any data about an individual who is identifiable by or in relation to such data and includes any data defined as "personal data", "personally identifiable information", or any analogous term under applicable laws, including the Digital Personal Data Protection Act, 2023 (DPDPA), and rules framed thereunder. Examples: Email Address, Mobile Number, Customer Information, HR records.
- **Confidential Information:** means any non-public, proprietary, or restricted information of or relating to GIG, its directors, officers, employees, customers, vendors, suppliers, business partners, and other stakeholders, where unauthorized disclosure could result in commercial, operational, legal, financial, or reputational harm.

Confidential Information includes, but is not limited to:

- Internal business plans and strategies
- Financial and commercial information
- Contractual information
- Customer and supplier information
- Personal Data

- Internal reports and analyses
- Policies, procedures, and methodologies
- Information received under confidentiality obligations
- Intellectual property and proprietary know-how not classified as Sensitive Data
- **Sensitive Information:** means a subset of Confidential Information that requires enhanced protection due to its strategic, legal, regulatory, security, competitive, or reputational significance. Unauthorized disclosure, processing, or access to such information may result in significant harm to GIG, its stakeholders, or affected individuals.

Sensitive information includes, but is not limited to:

- Unpublished Price Sensitive Information (UPSI);
- Research and Development (R&D) information;
- Trade secrets, proprietary algorithms, and intellectual property;
- Product formulations, designs, technical specifications, and source code;
- Merger, acquisition, investment, and strategic transaction information;
- Security-sensitive information, including vulnerabilities, security architecture, encryption keys, credentials, and incident investigation data;
- Sensitive Personal Data or personal information subject to enhanced legal or regulatory protection; and
- Any information designated by GIG as Restricted, Highly Confidential, or subject to enhanced access controls.

3.6 Consequential AI Decision

Certain AI-assisted activities may involve use of AI systems to make or influence decisions that have a material impact on the rights, interests, opportunities, or access to services of individuals or business partners. Any use of an AI system that results in a Consequential AI Decision will be categorized as High risk under Section 5 of this Policy below. A Consequential AI Decision includes, without limitation:

- the use of AI to assess, rank, score, or determine the eligibility of any prospective dealer, distributor, vendor, supplier, or other business partner, where the AI-generated output materially influences the acceptance, rejection, suspension, or restriction of such relationship;
- the use of AI to flag, suspend, restrict, or deny access to any customer account, transaction, loyalty programme, product, service, or digital platform based on a risk, fraud, or trustworthiness assessment generated by an AI system;
- the use of AI to determine or influence pricing, discounts, rebates or other commercial terms offered to customers or business partners; and
- the use of AI to make or materially influence decisions relating to recruitment, hiring, performance evaluation, promotion, disciplinary action, or termination of employees or contractors.

4. Core Policy Principles

4.1 Human Accountability

AI systems are assistive tools. Final decision-making authority and accountability remain with designated human employees.

Human oversight shall be exercised in relation to all AI systems used by GIG. Human oversight must be meaningful and effective, and must be capable of identifying, questioning, and overriding erroneous, biased, unlawful, or unintended AI-generated outputs.

The nature, intensity, and frequency of such oversight shall be proportionate to the risk, impact, and intended use of the relevant AI system. In determining the appropriate level of oversight, GIG personnel should have regard to the risk categorization and illustrative use cases set out in Section 5 of this Policy.

Notwithstanding the foregoing, final decision-making authority and accountability for any business decision, action, or outcome assisted by AI shall remain with designated GIG personnel.

4.2 Data Protection & Confidentiality

Users may input or upload the following only into AI tools expressly approved and made accessible by Corporate IT, using the reference below:

Information Type	Consumer AI Tools	Enterprise AI Tools	Private AI Services
Public Information	☑ Allowed	☑ Allowed	☑ Allowed
Internal Information	✗ Not Allowed	☑ Allowed	☑ Allowed
Personal Data	✗ Not Allowed	☑ Allowed, subject to applicable laws and approved use cases	☑ Allowed
Confidential Information	✗ Not Allowed	☑ Allowed, except where classified as Sensitive Data	☑ Allowed
Sensitive Data	✗ Not Allowed	⚠ Restricted. Only where expressly approved by GIG and appropriate contractual, legal, and security controls exist	☑ Allowed
UPSI	✗ Not Allowed	✗ Not Allowed	☑ Allowed
R&D Information	✗ Not Allowed	✗ Not Allowed	☑ Allowed
Trade Secrets	✗ Not Allowed	✗ Not Allowed	☑ Allowed
Product Formulae/Recipes	✗ Not Allowed	✗ Not Allowed	☑ Allowed
Source Code & Proprietary Algorithms	✗ Not Allowed	✗ Not Allowed	☑ Allowed
Security Architecture & Vulnerability Information	✗ Not Allowed	✗ Not Allowed	☑ Allowed
Encryption Keys, Passwords, Secrets & Credentials	✗ Not Allowed	✗ Not Allowed	✗ Not Allowed unless explicitly approved secure AI solution

The use of AI (including the processing of Personal Data in this regard) will be subject to GIG's internal data protection and handling policies, including the ISMS Policy and the Acceptable Usage Policy, in addition to applicable data protection laws. This includes the Indian Information Technology Act, 2000 ("IT Act") and the DPDPA and their respective rules.

Where an Enterprise AI Service or Private AI Service processes, stores, or transmits GIG data outside India, the applicable vendor contract shall document the safeguards adopted to meet GIG's obligations under the DPDPA and other applicable cross-border data transfer requirements. Corporate IT and Legal shall jointly maintain a record of the hosting location(s) and data residency posture of each approved Enterprise AI Service.

Choice of Build Platform — Hosted Frontier Models vs. Private/Self-Hosted Models: The data classification table above shall also govern the choice of build platform when GIG develops AI-enabled applications (see Section 6.11). As a general principle:

- Public, Internal, and Personal Data (with appropriate contractual safeguards) may be built using hosted frontier models available through GIG's approved cloud AI platforms (e.g., Microsoft Azure AI Foundry).
- Sensitive Data, UPSI, R&D information, trade secrets, product formulae, and source code shall default to Private AI Services (including models self-hosted on GIG-owned or GIG-controlled GPU infrastructure) unless a specific exception permitting use of a hosted frontier model is approved in accordance with Section 6.11.
- Use cases requiring fine-tuning on GIG's proprietary or confidential data shall default to Private AI Services, given the elevated risk of data exposure where fine-tuning is performed on a multi-tenant hosted platform without a contractually confirmed no-training-use clause.

4.3 Approved Enterprise AI Tools

Only AI tools formally approved, notified and made available by Corporate IT may be used for business purposes.

Currently approved platforms include:

- Microsoft Copilot Enterprise
- Claude Enterprise
- ChatGPT Enterprise
- Gemini Enterprise
- Higgsfield Enterprise

This list reflects platforms under active enterprise agreement with Corporate IT as of the date of this Policy and supersedes any prior version of this list. Corporate IT shall maintain the authoritative, current list of approved AI tools and shall update it promptly as agreements are added, changed, or discontinued.

Use of unapproved AI tools for business purposes is prohibited.

Requests for new AI tools must be made to Corporate IT and will undergo IT security and risk review prior to approval.

4.4 Accuracy, Transparency, Explainability & Ethical Use

Employees must:

- Review and validate AI-generated outputs before relying on them;
- Ensure accuracy in regulatory, legal, financial, and external communications;
- Avoid generating misleading, discriminatory, offensive, anti-competitive, or unlawful content;
- Not impersonate individuals or fabricate data using AI;
- Ensure that AI-assisted decisions and recommendations are explainable, understandable, and capable of being justified to relevant stakeholders, regulators, customers, or auditors, where applicable;
- Use AI systems responsibly and, where reasonably practicable, select solutions that promote efficient resource utilization, minimize environmental impact, and support the organization's sustainability objectives.
- Treat AI-generated output as non-authoritative without human validation.
- Use AI systems responsibly and, where reasonably practicable, select and operate AI solutions that minimise their environmental footprint, including energy consumption, water use and associated carbon emissions. This may include computational optimisation, use of energy-efficient infrastructure and, where feasible, preference for AI services and data centres supported by renewable or lower-carbon energy sources.

5. AI Risk Classification & Approval Framework

AI risk classification shall be determined based on the intended business use case, decision impact, and potential operational, financial, regulatory, reputational, or safety consequences.

Risk is based on the use case and impact — not on the specific AI tool being used.

The use of enterprise-approved platforms (e.g., Microsoft Copilot Enterprise, ChatGPT Enterprise or Claude Enterprise) reduces infrastructure and data security risks but does not eliminate business, regulatory, or operational risks.

5.1 AI Risk Classification Matrix

Risk Level	Definition (Impact-based)	Representative GIG Examples	Approval	Key Controls (Minimum)
Prohibited	AI makes or executes decisions with legal, financial, safety, or regulatory impact	Safety shutdowns (Chemicals), regulatory submissions, autonomous payments, employee termination; AI generating legal conclusions relied upon without qualified review; AI performing social scoring of employees, contractors, suppliers, or customers; AI inferring sensitive characteristics of individuals (e.g., political opinions, religious beliefs, sexual orientation). AI systems shall not be used or deployed to: • use subliminal, deceptive or purposefully manipulative techniques intended to materially distort an individual's behaviour or impair informed decision-making; • exploit vulnerabilities of individuals or groups arising from factors such as age, disability or socioeconomic circumstances in a manner that may result in material harm; • perform social scoring or evaluate the trustworthiness of individuals based on social behaviour or personal characteristics where this could result in unjustified or disproportionate treatment; or • conduct unauthorised biometric surveillance or real-time remote biometric identification.	Not Permitted – No Approval Pathway Exists	Not Permitted – No Approval Pathway Exists
High	Consequential AI Decision; Input/Upload of Sensitive and Confidential information in AI Systems	Pricing recommendations, recruitment shortlisting, drafting documents for regulatory filings	CISO + BU IT Head + Business Head	Mandatory human decision; Enterprise AI Approved Services; Private AI Services for Sensitive Data; bias/fairness review; explainability validation; documented validation; full audit logs, override & periodic audits
Medium	AI supports business decisions; no direct regulatory or safety impact	Sales forecasting, predictive maintenance alerts, marketing analytics	BU IT Head	Only Enterprise AI Approved services; No Sensitive Data input; Human review; no auto-execution

Low	AI for productivity or assistance only	Email drafting, summaries, translations, internal search	Business manager	Human review; approved Enterprise AI Services; No Sensitive Data and Confidential Information input
-----	--	--	------------------	---

5.2 General Requirements Across All Risk Levels

- All AI use-cases must use IT-approved platforms.
- High & Medium risk AI systems must be formally registered with Corporate IT.
- Where uncertainty exists, the higher risk category shall apply. AI shall not replace an accountable human decision-maker.

Notwithstanding the risk classification assigned to a use case under the Table in Section 5 of this Policy above, enhanced human review, validation, approval, or other controls may be required where the AI system is used in connection with:

- research and development activities, including product formulations, technical specifications, patentable inventions, trade secrets, or other confidential know-how;
- legal, regulatory, compliance, audit, or investigative activities;
- cybersecurity, information security, or incident response activities;
- material financial reporting, treasury, tax, or accounting activities; or
- safety-critical operations, manufacturing processes, or environmental, health, and safety functions.
- Data Protection Impact Assessment (DPIA): Where an AI system is classified as High-Risk and processes Personal Data, a DPIA shall be conducted prior to deployment, documented, and made available to the Data Privacy/DPO function, consistent with GIG's data protection and DPDPA compliance framework.
- Detection of Unapproved ("Shadow") AI Use: Corporate IT Cyber Security Team shall implement technical controls (e.g., AI Security solutions, Data Leakage Prevention, Cloud Access Security Broker, network-level restrictions) to detect and prevent the use of unapproved AI tools for business purposes, in addition to the policy prohibition in Section 4.3.
- Sector-Specific Regulatory Overlay: Business Units operating in regulated sectors shall apply additional sector-specific controls over and above this Policy, including (without limitation) RBI guidance on fair lending and digital lending practices, and SEBI (Prohibition of Insider Trading) Regulations applicable to the handling of UPSI by any GIG entity.
- AI Access Offboarding: Upon an employee's or contractor's exit, or change of role, access to all AI tools and any associated chat history or stored prompts shall be revoked or removed as part of GIG's standard offboarding process, consistent with the Identity and Access Management Procedure.

6. AI Solution Acquisition and Development

This section applies when GIG procures, develops, trains, customizes, or significantly modifies an AI system. It defines the minimum requirements to ensure AI systems are acquired, developed, tested, and managed securely, responsibly, and in compliance with applicable laws and Company policies.

6.1 Procurement of AI-Enabled Solutions

As AI capabilities become increasingly embedded in enterprise software, cloud services, and third-party platforms, procurement activities involving AI-enabled solutions shall comply with the following requirements:

AI Capability Assessment

Before procuring an AI-enabled solution, the Business Owner, BU IT along with Corporate IT Cybersecurity Function shall assess:

- The AI capabilities available within the solution.
- How GIG data will be processed, stored, or used by AI features.
- Whether AI features can be disabled, restricted, or configured to meet business requirements and risk appetite.

- The solution's classification under the AI Risk Classification Framework.

Contractual Requirements

Where applicable, contracts for AI-enabled products and services shall include:

- Restrictions on the use of GIG data to train, improve, or develop the vendor's AI models without GIG's explicit consent.
- Notification requirements for material changes to AI functionality, data processing activities, or hosting arrangements.
- Sufficient transparency and information to enable GIG to meet its legal and regulatory obligations.
- Appropriate protection of GIG intellectual property, data rights, and licensing requirements.
- Data privacy and security controls commensurate with the risk classification of the solution and applicable legal requirements.
- Clear allocation of ownership of intellectual property rights in AI-generated output, and, where GIG intends to use such output externally or commercially, confirmation that such use will not infringe third-party rights.
- Where GIG data is processed or stored outside India, the safeguards required under the DPDPA and other applicable cross-border data transfer requirements (including contractual clauses, where applicable).
- For material AI solutions, environmental considerations may form part of vendor assessment, including, where relevant, the energy efficiency, water use, carbon intensity and renewable energy practices associated with the AI model or supporting data centre infrastructure.

Review of Existing Contracts

- BU IT, Corporate IT Cybersecurity Function and BU Legal shall periodically review material vendor contracts to assess AI-related data processing activities, risks, and contractual obligations arising from new or updated AI capabilities introduced by vendors.

6.2 Development Approval

- All AI development projects shall be registered with BU IT & Corporate IT Cybersecurity Function before development or model training activities commence.

Each project shall have:

- A Business Owner responsible for the business use case.
- A Technical Owner responsible for implementation, security, and ongoing oversight.
- Development activities shall comply with the AI Risk Classification and Approval Framework defined in Section 5.

6.3 Training Data Management

- Data used to train or improve AI models shall originate only from approved and documented sources.

Development teams shall maintain records of:

- Data source.
- Data owner.
- Licensing or usage rights.
- Intended purpose of use.
- Personal Data, Sensitive Data, Confidential Information, or commercially sensitive information belonging to GIG or third parties shall not be used for AI model training unless expressly approved and permitted under applicable laws and Company policies.
- Personal Data processed by GIG as part of its business operations, including employee, contractor, customer, supplier, shareholder, partner, and applicant information, shall not be used for model training without appropriate approvals, legal basis, and safeguards.

6.4 Use of External AI Models and Software Components

- AI models obtained from external providers shall be reviewed for authenticity, licensing compliance, and security before use.
- AI development teams shall use only trusted and approved software libraries, frameworks, tools, and repositories.

6.5 Secure Development Practices

- AI models, training processes, configurations, and related documentation shall be treated as confidential Company assets.
- Access to AI development environments and related resources shall be restricted to authorized personnel.
- Development and testing environments should be separated from production environments where practical.
- Internal implementation details shall not be disclosed externally without appropriate authorization.

6.6 Protection of AI Instructions and Configurations

- System prompts, model settings, and configuration instructions used in AI systems shall be protected as confidential information.
- Confidential Company information should not be unnecessarily embedded within prompts or system instructions.
- AI systems shall be designed to prevent the unintended exposure of confidential information, prompts, model configurations, or other sensitive internal information.

- Users shall exercise caution when interacting with prompts, documents, or content intended to manipulate AI behaviour, extract information, or bypass established controls.

6.7 Documentation and Transparency

Development teams shall maintain documentation covering:

- Intended purpose.
- Training data sources.
- Performance measures.
- Known limitations.
- Approved use cases.
- For High-Risk AI Use Cases, sufficient documentation shall be maintained to support oversight, accountability, auditability, and regulatory review.
- Where appropriate and proportionate to the risk classification of the use case, AI systems should provide sufficient transparency to enable meaningful human review and oversight.

6.8 Change and Version Management

- AI models, datasets, and configurations shall be maintained under version control.
- Significant changes shall be documented to enable traceability, reproducibility, and rollback where necessary.

6.9 Human Oversight and Compliance

- Human oversight shall remain integral to AI development, testing, and evaluation activities.
- AI systems shall not be developed or trained to support activities that violate applicable laws, regulations, or Company policies.
- Personnel involved in AI development shall complete required training relating to responsible AI, AI security, and applicable legal obligations.

6.10 Secure AI Development Standards

- AI development teams shall comply with AI security requirements and development standards approved by Corporate IT.
- Corporate IT shall maintain and periodically update approved AI security standards, frameworks, and implementation guidance to reflect evolving technologies, emerging threats, regulatory requirements, and industry best practices.
- Where applicable, AI-enabled systems should be developed and secured in accordance with recognized industry standards, frameworks, and guidance approved by Corporate IT.

6.11 Selection and Registration of Foundation Models

GIG makes frontier and open-weight foundation models available for building AI applications, including hosted models accessed through GIG's approved cloud AI platform (e.g., Microsoft Azure AI Foundry) and models self-hosted on GIG-owned or GIG-controlled GPU infrastructure (Private AI Services). This Policy governs the selection and use of such models by criteria rather than by naming specific models, as the available models change frequently.

- Any foundation model used to build a GIG application, whether hosted-frontier or self-hosted, shall be registered under the Development Approval process defined in Section 6.2, with a Business Owner and Technical Owner and a risk classification under Section 5.
- Corporate IT Cybersecurity Function shall maintain a living Foundation Model Registry, listing currently approved models by tier (hosted-frontier vs. private/self-hosted), which shall be reviewed and updated at least quarterly. This Policy shall be read together with the current Registry; the Registry, not this Policy, is the authoritative source of which specific models are approved at any point in time.
- A new foundation model may only be added to the Registry following joint sign-off from BU IT and the Corporate IT Cybersecurity Function, based on a documented assessment covering: (a) data residency of the hosting region; (b) the subprocessor chain, where the model is hosted by or accessed through a third

party; (c) applicable licensing terms, particularly for open-weight models; and (d) ownership of ongoing patching and vulnerability management, particularly for self-hosted models.

- Selection between a hosted frontier model and a Private AI Service for a given use case shall follow the data classification and build-platform guidance set out in Section 4.2.

6.12 AI Features Embedded in Third-Party and SaaS Applications

Many third-party and SaaS applications used by GIG increasingly offer embedded AI capabilities, whether included at no additional cost or as a paid add-on. Such features shall be governed as follows:

- **Default-Off:** AI features embedded within an existing or new third-party or SaaS application shall not be enabled until they have been reviewed and approved by Corporate IT. Enabling such a feature directly (e.g., by a local administrator) without such approval is a violation of this Policy.
- **Free-Tier Scrutiny:** The fact that an embedded AI feature is offered free of additional cost shall not be treated as reducing risk; free AI add-ons more frequently reserve the right to use customer data for model training or improvement than paid or enterprise tiers. Before enabling any embedded AI feature (free or paid), Corporate IT shall confirm the vendor's data usage and training terms specific to that feature.
- **Subprocessor Visibility:** Approval to enable an embedded AI feature shall confirm which entity actually processes GIG data in connection with that feature (which may be a model provider other than the SaaS vendor itself), and shall assess that entity under the same third-party and vendor risk requirements set out in Section 6.1.
- **Change Trigger:** A vendor's introduction of a new or materially changed AI feature within an existing third-party or SaaS application shall be treated as a material change requiring review under Section 6.1 (Review of Existing Contracts), and shall not be enabled until that review is complete.

7. AI Deployment and Lifecycle Management

This section defines the requirements for testing, deploying, monitoring, reviewing, and retiring AI systems used within GIG.

7.1 Pilot Evaluations

Business units may wish to test new AI tools or AI use cases before obtaining full approval.

- Any pilot evaluation must be approved by BU IT and Corporate IT Cybersecurity Function before it begins.
- Business Owner along with BU IT and Corporate IT Cybersecurity Function shall define the scope, duration, data usage restrictions, and reporting requirements applicable to the pilot.

Each pilot shall have:

- A clearly defined objective and scope.
- A start and end date.
- A designated Business Owner and Technical Owner.
- Wherever possible, pilot evaluations should avoid the use of Personal Data, Confidential Information, or other sensitive information unless specifically approved.
- Findings, lessons learned, and any identified risks or issues shall be shared with Corporate IT Cybersecurity Function at the conclusion of the pilot.
- Pilot deployments shall not be moved into production until all required approvals have been obtained under the AI Risk Classification and Approval Framework.

7.2 Deployment Approval

- AI solutions shall not be deployed for production or business use without obtaining the approvals required under the AI Risk Classification and Approval Framework.
- Approved AI solutions shall be recorded in the Corporate IT Cybersecurity Function's AI inventory.
- Significant changes to an approved AI solution or use case, including changes to purpose, functionality, data used, automation, or risk classification, may require reassessment and re-approval before implementation.

7.3 Validation Before Deployment

Before an AI solution is deployed:

- The Business Owner and Technical Owner shall validate that the solution performs as intended and that applicable security, privacy, compliance, and risk management requirements have been addressed.
- Corporate IT Cybersecurity Function may require additional review or approval for High-Risk AI Use Cases.
- Where appropriate, AI solutions should provide mechanisms for human review, intervention, or override before important decisions are acted upon.

7.4 Monitoring and Review

Following deployment:

- Business Owners shall continuously monitor AI solutions to ensure they deliver intended business outcomes and operate as expected, with mechanisms in place to detect and remediate model drift or performance degradation over time.
- Corporate IT Cybersecurity Function may monitor AI solutions to support security, compliance, operational oversight, and risk management activities.
- High-Risk AI Use Cases shall undergo periodic review to identify unexpected outcomes, bias, performance issues, regulatory concerns, or other emerging risks.
- Significant concerns identified during monitoring shall be investigated and addressed in a timely manner.

7.5 Incident Management

- Any AI-related incident, including inaccurate outputs, system failures, security concerns, unintended consequences, or unauthorized disclosure of information, shall be reported through GIG's established incident reporting processes.
- Appropriate corrective actions shall be taken following investigation and review.
- Where an incident involves Personal Data, Confidential Information, or cyber security concerns, it shall also be managed in accordance with GIG's Incident Management Procedure, Data Breach Response Procedure, Cyber Incident Response Procedure, and any applicable legal or regulatory reporting requirements.

7.6 Logging and Audit

- Business Owners and Corporate IT Cybersecurity Function may maintain records of relevant AI system activity where necessary to support operational monitoring, security, compliance, investigations, and audit requirements.
- Logging and monitoring activities shall be conducted in accordance with GIG's Information Security, Data Privacy, Records Retention, and Monitoring policies, as applicable.
- Access to logs and monitoring information shall be restricted to authorized personnel.

7.7 Retirement and Decommissioning

When an AI solution is no longer required:

- The solution shall be retired in a controlled manner.
- Relevant records, documentation, configurations, and related assets shall be retained, archived, or securely removed in accordance with GIG's Information Security and Records Retention requirements.
- Access to retired AI systems shall be removed where appropriate.

- Any residual risks arising from the retirement of an AI solution shall be assessed and addressed before decommissioning is completed.

8. AI Governance, Oversight and Continuous Improvement

8.1 Governance Ownership and Oversight

GIIG maintains a defined governance structure to oversee the responsible use of AI across the organization. This structure is modelled on the tiered Information Security governance structure set out in Section 5.1 of the ISMS Policy, with oversight cascading from the GMC through a dedicated AI Governance Council to Corporate IT for day-to-day execution, as set out below:

Level	Composition	Responsibility
GMC	Group Management Committee	Receives periodic inputs from the AI Governance Council on Group-wide AI adoption, risk posture, and material incidents; approves Group-level AI strategy direction.
AI Governance Council	Group CISO; Head - Corporate Audit; Head - Group Legal; all BU IT Heads	Owns and updates this policy; defines the overall AI usage and governance framework for GIIG; provides strategic input to the GMC on Group-wide AI adoption and governance; reviews High-Risk AI use cases and material incidents escalated by Corporate IT.
Corporate IT Cybersecurity Function	Corporate IT Cybersecurity Team along with BU IT Teams	Executes day-to-day AI governance: tool approvals, use-case registration, maintenance of the Foundation Model Registry and AI inventory, monitoring, incident response, and lifecycle controls; escalates High-Risk items and emerging risks to the AI Governance Council.

The AI Governance Council shall meet at least six monthly. Material high-risk incidents or emerging risks shall be escalated to the Council within 5 business days of identification. Legal, Data Privacy/DPO, and other relevant functions shall be consulted by the AI Council, as appropriate, on matters within their respective domains, including legal permissibility, contractual review, and Personal Data processing oversight. Reference may also be made to the provisions of the Annexure under Section 13 when determining roles and responsibilities.

Each AI system or use case must have:

- Business Owner accountable for business outcomes arising from the use of AI systems and its outputs, including accountability for ensuring that AI outputs are used within approved use cases and is subject to appropriate human oversight as enshrined under this Policy, and
- Technical Owner accountable for implementation, security, and operational controls.

Final accountability for decisions supported by AI remains with designated human decision makers, consistent with the Human Accountability principle defined in this Policy.

8.2 Review, Monitoring and Effectiveness Checks

GIIG will periodically review AI usage and governance practices to assess whether:

- AI systems continue to operate within approved risk boundaries.
- defined controls remain effective; and
- AI usage remains aligned with organizational expectations and external obligations.

These reviews may include:

- monitoring of AI system behaviour and outputs;
- review of logged usage and incidents;

- assessment of higher-risk AI systems; and
- identification of improvement opportunities.

Findings from such reviews may be shared with relevant stakeholders and management, as appropriate. Material AI governance matters, including significant AI incidents, emerging risks, high risk use cases, or material vendor changes, may be escalated to senior management through existing governance channels where appropriate.

8.3 Issue Management and Corrective Actions

Where AI-related issues are identified — including misuse, control gaps, unintended outcomes, or policy deviations — GIG will:

- assess the nature and impact of the issue;
- determine appropriate corrective or mitigating actions; and
- track actions to closure, proportionate to risk and impact.

Corrective actions may include changes to controls, approvals, usage restrictions, training, or governance processes.

8.4 Continuous Improvement of AI Governance

GIG recognizes that AI technologies, risks, and regulatory expectations continue to evolve. Accordingly, AI governance practices will be continuously improved through:

- lessons learned from incidents or near-misses;
- feedback from reviews and assessments;
- introduction and changes in laws, regulations, or industry guidance; and
- advancements in AI technology and risk management practices.

This Policy will be reviewed periodically to reflect these changes and to ensure continued relevance and effectiveness. Governance effectiveness may be assessed through indicators proportionate to organisational maturity, including policy deviations, incident trends, training completion, review completion rates, remediation timelines, or recurring control weaknesses.

9. Reporting

Employees must report suspected AI incidents including:

- AI misuse
- Data exposure
- Ethical concerns
- Security incidents
- AI system malfunction or unintended behaviour

Through:

- Corporate IT Service Desk – Corp-servicedesk@godrejinds.com
- ISMS Compliance – corpit.compliance@godrejinds.com
- Data Privacy – data.privacy@godrejinds.com

10. Training

Mandatory AI awareness training shall be provided to:

- All employees.
- New joiners (as part of induction).
- Relevant role-based teams where applicable.

AI awareness training shall include, at a minimum:

- a. An overview of GIG's obligations under applicable law in the context of AI use.
- b. Requirements and obligations under this Policy.
- c. Practical guidance on lawful and responsible use of approved AI tools, including data minimisation and avoidance of sensitive data inputs.
- d. Role-specific training for personnel involved in AI development, deployment, or vendor management.

11. Violations

Non-compliance with this Policy shall constitute a breach of the Code of Conduct, the Acceptable Usage Policy and may result in disciplinary or other action applicable to such a breach under such Code / Policy.

12. Policy Review

This Policy shall be reviewed annually or earlier if required, including due to:

- Regulatory developments;
- Introduction of significant AI technologies; and
- Audit findings or risk events.

13. Annexure

13.1 AI Strategy and Operating Model

13.1.1 Purpose and Intent

This section describes the strategic intent and operating principles for the adoption and use of Artificial Intelligence (AI) across Godrej Industries Group (GIG). This section is intended to:

- Provide context for how AI is expected to create value across the Group
- Clarify operating principles for coordination between Corporate IT and Business Units
- Support consistent and scalable AI adoption without introducing additional mandatory controls beyond those defined elsewhere in this Policy

The requirements governing AI usage, risk management, security, approvals, and compliance remain defined by Sections 4 through 8 of this Policy.

13.1.2 AI as a Decision Support Capability

AI within GIG is intended to function primarily as a decision support capability, enabling improved insight, responsiveness, and management effectiveness. Accordingly:

AI initiatives are expected to focus on supporting recurring and material business decisions, rather than isolated or experimental use alone.

AI usage is intended to enhance human judgment by improving access to information, identifying patterns, highlighting risks, and supporting timely intervention.

- AI outputs are not intended to replace accountable human decision makers, consistent with the Human Accountability principles defined in Section 4.1.

Illustrative examples of decision support use include demand sensing, risk identification, pricing analysis, operational monitoring, and issue prioritization.

13.1.3 AI Operating Model

GIG intends to operate AI capabilities through a hybrid operating model that balances central enablement with federated business ownership.

13.1.4 Central Enablement

Corporate IT, in collaboration with Information Security, Compliance, Legal, Risk, and other relevant functions, provides centralized enablement for AI adoption. This includes:

- Enterprise approved AI platforms and shared infrastructure
- Common governance standards, guardrails, and reference architectures
- Group level oversight for higher risk AI usage
- Coordination and escalation mechanisms for material AI related risks

13.1.5 Federated Business Ownership

Business Units retain ownership of business specific AI initiatives and are expected to:

- Define the business context and intended outcomes for AI usage
- Integrate AI output into business processes under human oversight
- Assign Business Owners and Technical Owners as defined in Section 8.1
- Apply AI outputs in a manner consistent with this Policy and applicable approvals

13.1.6 Joint Alignment

Certain aspects of AI adoption — such as prioritization of high impact initiatives, reuse of capabilities, and resolution of cross cutting issues — are intended to be aligned jointly between Corporate IT and Business Units to ensure consistency with Group level objectives.

13.1.7 Value Orientation and Scaling

AI adoption within GIG is intended to progress beyond isolated experimentation toward repeatable and scalable value creation, proportionate to risk and impact. Accordingly:

- For Medium and High-risk AI use cases, Business Owners are encouraged to articulate the intended business outcomes associated with AI usage.
- AI initiatives may contribute value through cost efficiency, revenue enablement, risk reduction, resilience, or improved management effectiveness.
- AI systems that demonstrate sustained usefulness or impact may be considered for broader adoption or reuse, subject to applicable governance and risk review.

This section establishes directional expectations and does not mandate uniform financial return metrics for all AI usage.

13.1.8 Shared Intelligence Across the Group

GIG recognizes that operating as a diversified group provides opportunities for responsible sharing and reuse of AI learnings and capabilities. Where appropriate and permissible:

- AI approaches, models, insights, or patterns developed within one part of the Group may inform or support adoption elsewhere.
- Any such reuse remains subject to data protection, confidentiality, intellectual property, and risk classification requirements defined in this Policy.
- Corporate IT may facilitate visibility of reusable AI capabilities to support efficiency and consistency.

13.1.9 Sensing and Management Intelligence

AI adoption within GIG is intended to increasingly support continuous sensing and management intelligence, complementing traditional reporting and analysis. Where applicable:

- AI systems may be used to surface emerging trends, anomalies, or risks in operational, financial, customer, or compliance data.

- AI generated insights are intended to support informed management action under human oversight, rather than automated execution.
- Such usage remains subject to monitoring, review, and lifecycle controls defined in Section 7.

13.1.10 Leadership Responsibility

Senior management and business leaders remain accountable for:

- Ensuring AI is used in alignment with GIG's values, risk appetite, and strategic objectives
- Exercising appropriate judgment when relying on AI supported insights
- Maintaining governance discipline in AI assisted decision making

AI is intended to strengthen management effectiveness, not dilute accountability.

13.1.11 Relationship to Other Policy Sections

This section:

- Provides strategic and operating context for AI adoption
- Does not override or replace any mandatory requirements defined in Sections 4 through 8

Should be read in conjunction with:

- AI Risk Classification and Approval Framework
- AI Development and Lifecycle Management
- AI Governance, Oversight, and Continuous Improvement provisions

13.2 Standards and Controls Mapping

ISO 42001 Clause	ISO Intent	NIST AI RMF Function	NIST RMF Intent	GIG Policy Section(s)
4.1	Organizational context for AI	MAP	Establish AI context and use environment	1. Purpose; 2. Scope
4.2	Stakeholder expectations	MAP	Identify impacted stakeholders	1. Purpose; 4. Core Policy Principles; 8. Governance
4.3	Scope of AI management system	GOVERN	Define governance boundary	2. Scope (implicit)
4.4	AI management framework	GOVERN	Establish AI risk governance	5. AI Risk Classification; 8. AI Governance & Oversight
5.1	Leadership commitment	GOVERN	Accountability and oversight	8.1 Governance Ownership & Oversight
5.2	AI policy	GOVERN	Organizational AI risk policy	Entire Policy
5.3	Roles & responsibilities	GOVERN	Clear ownership of AI risks	6.2 Development Authorization; 8.1 Business & Technical Owner
6.1	AI risk identification	MAP	Identify and categorize AI risks	5. AI Risk Classification Framework
6.2	AI objectives & planning	GOVERN	Risk-informed objectives	1. Purpose; 8.4 Continuous Improvement

6.3	Planning changes	MANAGE	Adapt governance to change	12. Policy Review; 8.4 Continuous Improvement
7.1	Resources	GOVERN	Enable AI risk governance	8. Governance & Oversight
7.2	Competence	GOVERN	Skilled workforce for AI risk	6.9 Human Oversight & Compliance; 10. Training
7.3	Awareness	GOVERN	Risk awareness culture	4. Core Principles; 10. Training
7.4	Communication	GOVERN	Risk communication & escalation	9. Reporting; 8. Governance
7.5	Documented information	GOVERN	Traceability & transparency	6.7 Documentation & Transparency; 6.8 Change and Version Management
8.1	Operational control	MANAGE	Enforce risk controls	5. Risk Framework; 7.2 Deployment Approval
8.2	AI lifecycle management	MANAGE	Risk control across lifecycle	6. AI Development; 7. Deployment & Lifecycle
8.3	Data governance	MAP	Understand data risks	4.2 Data Protection; 6.3 Training Data Management
8.4	Transparency & explainability	MEASURE	Assess AI trustworthiness	6.7 Documentation & Transparency
8.5	Human oversight	GOVERN	Human accountability	4.1 Human Accountability; 5 Controls; 6.9 Oversight
8.6	External AI systems	MAP / MANAGE	Third-party AI risk	4.3 Approved AI Tools; 6.4 Use of External AI Models
9.1	Monitoring & evaluation	MEASURE	Monitor AI performance & risks	7.4 Monitoring and Review; 7.6 Logging and Audit
9.2	Internal audit	MEASURE	Evaluate control effectiveness	8.2 Review & Effectiveness Checks (implicit)
9.3	Management review	GOVERN	Senior oversight of AI risks	8.1 Governance Oversight (implicit)
10.1	Incident & corrective action	MANAGE	Respond to AI incidents	7.5 Incident Management; 8.3 Corrective Actions
10.2	Continual improvement	MANAGE	Continuous risk improvement	8.4 Continuous Improvement; 12. Policy Review
Annex A	Governance & risk controls	GOVERN / MANAGE	Risk-based AI governance	5 Risk Matrix; 8 Governance
Annex A	Secure AI development	MANAGE	Reduce technical AI risks	6.5–6.10 Secure AI Development
Annex A	Monitoring & logging	MEASURE	Detect drift & misuse	7.4 Monitoring; 7.6 Logging

Annex A	Decommissioning	MANAGE	Safe AI retirement	7.7 Decommissioning
Annex A	Ethics & fairness	MAP / MEASURE	Identify societal impacts	4.4 Accuracy & Ethical Use

14. Policy Endorsement

The policy is endorsed by the Board Management Committee

Aasif Malbari

Managing Director & CEO
Godrej Consumer Products Limited