WELSPUN WORLD

Group Artificial Intelligence (AI) Policy

Controlled Enablement | Data Protection | Governance | Regulatory Compliance

Document ID	WEL-ISMS-POL-AIP	**Classification**	Internal — Restricted
Document Owner	CISO / Group IT	**Version**	1.0
Effective Date	18-Jun-2026	**Next Review**	Annually or upon regulatory change
Applicable Frameworks	ISO 27001:2022 | ISO 27701 | ISO 42001 | NIST CSF 2.0 | NIST AI RMF 1.0		
Applicable Regulations	DPDP Act 2023 & Rules 2025 | CERT-In 2022 | SEBI CSCRF 2023 | GDPR | UK GDPR | EU AI Act 2024 | IT Act 2000		

Document Version History

Version	Date	Author	Change Description
1.0	18-June-2026	Vinayak Laud	Initial issue — Core ISMS AI Policy

1. Purpose

This Policy establishes the governance, security, and usage requirements for Artificial Intelligence systems used within the Group. It ensures that AI adoption enhances business value while protecting organisational data, intellectual property, OT/ICS infrastructure, employee privacy, and the Group's regulatory obligations under DPDP, CERT-In, SEBI CSCRF, GDPR, and applicable international frameworks.

Welspun is committed to the responsible, ethical, and sustainable use of Artificial Intelligence across its lifecycle. Welspun will minimize the environmental impact of Artificial Intelligence by promoting energy-efficient systems, utilizing data centres with low ecological footprints, renewable energy, and responsible resource management.

2. Scope

This Policy applies without exception to all employees, consultants, and contractors of all Company Group entities. It covers all AI systems like Traditional predictive and analytical machine learning models, Generative AI systems, Agentic AI and autonomous workflow systems etc. included but not limited to (Refer Annexure – 1) — whether internally developed or third-party — and all AI-enabled enterprise applications used for company purposes.

This Policy applies to IT systems only. AI systems deployed in or adjacent to OT/ICS environments are additionally governed by the Critical — OT/Safety risk tier (Section 7) and must comply with IEC 62443 requirements. No AI system may interface with OT systems without completing the Critical — OT/Safety approval pathway.

3. Approved AI Usage Model

- Only AI tools approved by Group IT and the CISO function may be used for company work. Approved AI systems include: (a) the Group-approved enterprise AI platforms — ChatGPT Enterprise and Claude Enterprise — registered in the AI Use-Case Registry; (b) entity exception platforms approved under the Multi-Platform Governance Model (Clause 23); (c) open-source LLM deployments approved under Clause 24.
- Public or personal AI platforms (including free-tier or personal subscriptions) are prohibited for any official purpose.
- No AI tool may be deployed or used without formal onboarding approval from Group IT.
- Employees are personally responsible for ensuring their AI usage complies with this Policy.

4. Acceptable Use Guidelines

Permitted Uses

Employees may use approved AI tools for:

- Drafting non-confidential documents, communications, and reports
- Summarisation of non-sensitive information
- Research assistance on non-sensitive topics
- Coding support for non-proprietary or approved code
- Process automation ideation and workflow design

Prohibited Uses

Employees must NOT use AI platforms to:

- Upload or paste confidential contracts, agreements, or legal documents
- Share unpublished financial data, forecasts, or earnings information (SEBI PIT Regulations apply)
- Upload or process personal data of employees, customers, or third parties without DPO clearance and DPIA completion
- Share proprietary source code, trade secrets, or OT/process data
- Generate misleading, discriminatory, or unethical content
- Impersonate individuals or misrepresent AI-generated output as human-created without disclosure
- Produce AI-generated content for regulatory filings or Board submissions without mandatory human review
- Access or use AI tools from OT/ICS network segments or from devices connected to plant control networks

Employee Responsibility: Employees are accountable for the accuracy and appropriateness of any AI-generated content they use or share professionally. AI output must be reviewed, validated, and owned by the employee before use.

5. Data Classification & AI Usage Rules

AI usage must comply with the Group Data Classification Policy – WEL-ISMS-POL-DCP. Sensitive, confidential, financial, regulated, personal, or OT/process data shall not be shared with AI systems unless explicitly approved through the designated governance process.

6. Security Requirements

All AI systems and usage must comply with the technical controls defined in Welspun_World_Group_AI_Security_Standard_Policy_v1.0. Key requirements: secure access architecture with enterprise identity and MFA; DLP enforcement; encryption in transit and at rest; 180-day log retention and SIEM integration; AI-specific threat mitigation; periodic security testing including AI-specific VAPT (Clause 25); incident response integration including CERT-In 6-hour notification; third-party AI vendor risk management.

7. AI Risk Tiering

Risk Tier	Description	Scope Examples	Approval Requirement
Low	Informational and assistive tasks with no personal data and no consequential decisions	Email drafting, summarisation of internal non-sensitive documents, non-proprietary coding assistance	Standard user registration in AI Inventory
Medium	Process automation with human oversight; analytics on non-sensitive data	Maintenance scheduling, internal report generation, procurement analytics	Manager + IT sign-off; AI use-case form
High	Decision support influencing business outcomes; processing of personal data; customer-facing AI	Quality control AI, HR analytics, customer service AI, predictive demand planning	CISO + Business Owner + DPO approval; DPIA mandatory; annual bias audit

Critical — IT	Autonomous decisions; financial or regulatory impact; AI in regulated submissions	AI in financial reporting, investor communications, regulatory filings, UPSI-adjacent workflows	AI Governance Committee + Risk Committee; quarterly review
Critical — OT/Safety	ANY AI interacting with or influencing OT/ICS systems, safety-critical plant equipment, SCADA, PLCs	Predictive maintenance AI with actuator control, AI-driven process optimisation in plant, robotics AI	AI Governance Committee + OT Security Manager + Plant Safety Officer; IEC 62443 assessment mandatory; no autonomous actuation without physical safety interlock

7.1 AI Impact Assessment

The AIIA shall evaluate, at minimum – for High, Critical — IT, or Critical — OT/Safety tier:

a) The intended purpose of the AI system and the decisions or actions it will inform or automate

b) The individuals, groups, communities, or organizational functions that may be affected by the AI system's outputs or decisions;

c) The potential adverse impacts of the AI system across the following dimensions: fairness and non-discrimination, safety and physical harm, privacy and data protection, transparency and explainability, security and resilience, environmental impact, and economic or social impact;

d) The likelihood and severity of each identified adverse impact, taking into account the AI system's intended and reasonably foreseeable uses;

e) The risk mitigation controls proposed to reduce identified impacts to an acceptable level, mapped to the applicable risk tier controls in Clause 7;

f) The residual risk after mitigation and the Business Owner's acceptance of that residual risk.

For Low and Medium tier AI use cases, a simplified impact screening may be performed.

8. Monitoring & Compliance

- AI activity may be logged and monitored in accordance with applicable privacy and employment laws. Employees in EU/UK jurisdictions will receive a monitoring notice in accordance with GDPR obligations.
- Violations of this Policy may result in disciplinary action up to and including termination. A detailed consequence matrix will be defined within 60 days of Policy approval.
- Periodic audits of AI usage will be conducted by Internal Audit, including pre-assessment for SEBI CCI purposes.
- Shadow AI usage — accessing unapproved AI tools — is a policy violation subject to enforcement.
- Defined performance metrics appropriate to the use case (e.g., accuracy, precision, recall, false positive/negative rates) established at deployment and tracked continuously or at defined intervals; Automated or manual drift detection mechanisms to identify when model performance deviates materially from baseline metrics established at deployment or last review; A documented threshold and escalation procedure: when monitoring detects performance degradation or bias emergence beyond the defined threshold

9. AI Incident Classification & Reporting

Any employee who suspects or detects an AI-related security event must report it immediately to IT Security via the established incident reporting channel. The Incident Management Policy (WEL-ISMS-POL-IMS) governs classification and escalation. Failure to report a known or suspected AI security incident is itself a policy violation and may attract regulatory liability under CERT-In Directions 2022 and DPDP Act Section 8(6).

10. Governance Structure

Role / Body	Responsibility	Regulatory Accountability
CISO	Policy ownership, security oversight, incident escalation, AI Governance Committee chair, CERT-In notification authority	CERT-In reporting; SEBI CSCRF; ISO 27001/42001
Data Protection Officer (DPO)	DPDP and GDPR compliance oversight; DPIAs; consent mechanism review; data principal rights management; cross-border transfer approvals; Art.22 GDPR human review oversight	DPDP Act DPO obligations; GDPR Art.37–39; ISO 27701
AI Governance Committee	Use-case approval, risk tier review, quarterly AI monitoring, policy amendments, EU AI Act conformity oversight, SEBI CSCRF AI audit coordination. Charter to be constituted within 60 days of Policy approval.	SEBI CSCRF Domain 1; ISO 42001; EU AI Act
Business Owners	Compliance within their functions, user training, incident reporting, AI use-case sponsorship	Accountable for AI risks in their domain
Group IT	Technical controls, platform management, access provisioning, OT/IT boundary enforcement	ISO 27001 Annex A technical controls
Legal & Compliance	Regulatory mapping; SEBI CSCRF audit coordination; CERT-In notification support; cross-border transfer legal basis; IP indemnity review	IT Act; SEBI PIT Regulations; GDPR; EU AI Act
OT Security Manager	Governance of AI systems in or adjacent to OT/ICS environments; IEC 62443 compliance; safety-critical AI risk assessment	IEC 62443; Functional Safety (IEC 61508/61511 where applicable)
Internal Audit	Periodic independent compliance assessment; SEBI CSCRF CCI audit support; AI control testing	SEBI CSCRF CCI; ISO 27001 Clause 9.2

11. Third-Party AI Vendors

Third-party AI vendors must complete the Group's security due diligence process and execute appropriate contractual controls including a DPA before being approved for use. For EU/UK personal data: EU SCCs (Module 2) or UK IDTA must be executed. For DPDP purposes: data residency in India or approved jurisdiction must be confirmed. The Group retains full accountability for data shared with or processed by any vendor.

Vendor AI Risk

a) The supplier must disclose, upon request, whether AI systems are used in the delivery of services to the Group and, if so, the nature of the AI processing, the data involved, and the model provenance (vendor, training data sources, and known limitations).

b) Procurement and vendor management functions must include AI usage disclosure as a standard item in vendor onboarding questionnaires and contract renewals.

c) For commercial third-party AI vendors, the Group retains the right to request model provenance information including the identity of the underlying model provider, the training data categories, and any sub-processors involved in AI inference.

12. Training & Awareness

All employees authorised to use AI tools must complete mandatory AI Security Awareness training annually. Training will cover acceptable use, data protection obligations (DPDP and GDPR), prompt hygiene, CERT-In incident reporting obligations, and HR AI risks (Clause 21). Completion will be tracked and non-compliance may result in access revocation.

13. AI Content Accountability

Employees using AI-generated content in official communications, reports, client deliverables, or regulatory submissions are personally accountable for the accuracy, appropriateness, and factual correctness of that content. AI tools may produce errors or hallucinations; employees must apply professional judgement and verify all AI-generated output before use. AI-generated content must not be included in SEBI filings, stock exchange disclosures, Annual Reports, Board submissions, or legal documents without explicit review and sign-off by the relevant function head.

Where AI-generated outputs may incorporate or be derived from third-party copyrighted material, employees must exercise professional judgement and must not use AI outputs that reproduce substantial portions of identifiable third-party works without Legal clearance.

14. Business Continuity

In the event that one or more approved AI platforms are temporarily unavailable, employees should continue work using conventional tools or other Group-approved AI platforms where available. No use of unapproved or personal AI platforms is permitted as a workaround during outages. Group IT will communicate planned maintenance windows. Business Continuity plans (WEL-ISMS-POL-BCP) include AI platform availability as a covered scenario.

15. Review & Updates

This Policy will be reviewed annually or upon significant technological or regulatory changes (including amendments to DPDP Rules, CERT-In Directions, SEBI CSCRF, or GDPR guidance). The CISO is responsible for initiating reviews and presenting amendments to the AI Governance Committee for approval. Material changes require CFO notification.

16. AI Ethics Framework

The Group is committed to the responsible and ethical use of AI. The following principles govern all AI deployments:

- Fairness & Non-Discrimination: AI systems used in decisions affecting employees, customers, or third parties must be assessed for bias. Systems used in HR decisions, customer eligibility, or resource allocation must undergo documented bias testing before deployment and annually thereafter. Protected characteristics must not be inadvertently encoded in AI decision logic. See Clause 21 for HR-specific controls.
- Explainability: AI systems influencing significant business, HR, or customer decisions must produce an explanation of the basis for their output in human-understandable terms. Black-box AI outputs must not be used as the sole basis for material decisions affecting individuals. Explainability documentation is a mandatory component of the AI use-case registration for High and Critical tier use cases.
- Human Oversight: AI systems may not replace human judgement for decisions that are consequential for individuals, legally regulated, or safety-critical. Human review and override capability is a mandatory control for High and Critical tier AI systems. See Clause 21 for HR-specific human oversight requirements.
- Accountability: Every AI system must have a designated owner who is accountable for its governance, performance, compliance, and ethical operation throughout its lifecycle.
- Prohibited AI Uses (absolute): Social scoring of individuals; real-time biometric surveillance of individuals in publicly accessible spaces; AI systems that exploit vulnerabilities of specific groups; subliminal or manipulative AI-driven influence techniques; AI-generated deepfakes used without disclosure. These align with EU AI Act Article 5 prohibited practices.

17. Agentic AI & Autonomous Workflows

Agentic AI refers to AI systems that can take autonomous multi-step actions, call APIs, write and execute code, trigger workflow steps, or operate without per-step human approval. The following additional controls apply:

- Agentic AI may only be deployed after completing the AI use-case registration and achieving the appropriate risk tier approval. Agentic AI with access to financial systems, HR systems, or external APIs is automatically classified as Critical — IT tier.
- Authorisation Scope: agentic AI must operate within a strictly defined, least-privilege authorisation scope.
- Human-in-the-Loop: agentic workflows involving financial transactions, regulatory submissions, vendor communications, or customer data must include a mandatory human review and approval gate before execution.
- Rollback Mechanisms: agentic AI deployments must have a tested rollback or undo mechanism for any workflow step that modifies data, sends external communications, or triggers financial actions.
- Agentic AI must not be deployed in or adjacent to OT/ICS environments without completing the Critical — OT/Safety approval pathway and IEC 62443 assessment.
- Logging and Auditability: Actions taken by agentic AI systems must be recorded in an immutable action log that is separate from the standard AI interaction log.
- Multi-Agent and Orchestration: Where multiple agentic AI systems are chained, orchestrated, or operate in a multi-agent pipeline.

17.1 MCP Servers, Connectors, and Plugin Governance

Where agentic AI systems connect to internal data sources, enterprise applications, APIs, or external services via Model Context Protocol (MCP) servers, connectors, plugins, function-calling interfaces, or similar tool-integration mechanisms

- Registration and Approval
- Least-Privilege Access and Scope Boundaries
- Authentication, Secrets, and Credential Management

- Third-Party Plugin and Connector Vetting

18. AI Decommissioning & Model Retirement

When an AI system or vendor is decommissioned, the following process must be completed:

- The AI system must be formally retired in the AI Use-Case Registry. Access provisioned for the system must be revoked within 24 hours of decommissioning decision.
- All company data retained by the AI vendor must be deleted within 30 days of contract termination. Written confirmation of deletion must be obtained from the vendor.
- AI model artefacts developed internally must be deleted or archived in accordance with the Data Retention & Disposal Policy and DPDP storage limitation requirements.
- A post-decommissioning security review must be conducted to identify any residual data exposure, active API keys, or orphaned integrations.
- The decommissioning must be documented and reported to the AI Governance Committee.

19. Exception Handling — Use of Tools Other Than the Approved AI Platform

The Group operates a Controlled AI Platform Strategy. All AI usage must occur through approved AI systems. This clause governs exceptions for tools not covered by the Group-approved enterprise AI platforms or the Multi-Platform Governance Model (Clause 23).

19.1 Circumstances Where an Exception May Be Considered

- The required capability is demonstrably unavailable in any approved AI system.
- The business need is time-critical or project-specific and cannot wait for the standard procurement cycle.
- The tool under consideration is a commercial, enterprise-grade AI product — not a free-tier, personal, or consumer-facing platform.
- The intended use does not involve Confidential, Financial (Unpublished), Personal Data, Source Code (Proprietary), or OT/Process Data. No exceptions to this data restriction will be granted.

19.2 Exception Request Process

Submit via AI Exception Request Form (WEL-ISMS-FRM-AI) to the CISO Office through the GRC/ITSM portal. CISO screens within 2 business days. Security & privacy assessment within 5 business days. CISO approves (Low/Medium risk) or AI Governance Committee approves (High/Critical risk) within 3 business days of assessment. All exceptions are time-limited to a maximum of 90 days and not automatically renewable.

19.3 Conditions Attached to All Approved Exceptions

- Exception is personal and role-specific — may not be shared, delegated, or extended to other employees without a separate request.
- Employee must complete a targeted security briefing before first use. Completion is logged.
- Usage must be logged and monitored via CASB. Any data classification violation detected during the exception period results in immediate revocation.
- A DPA must be in place with the tool vendor before use.
- Employee remains personally accountable for all outputs per Clause 13.
- The exception must be documented in the AI Use-Case Registry and reported to the AI Governance Committee at the next quarterly review.

19.4 Grounds for Immediate Revocation

- The employee uses the tool with data categories that were explicitly excluded at the time of approval.
- The tool vendor suffers a confirmed or suspected security incident affecting company data.
- CASB or DLP monitoring detects policy violations attributable to use of the excepted tool.
- The employee's role changes such that the original business justification no longer applies.
- ChatGPT Enterprise or Claude Enterprise introduces the capability cited as the justification for the exception.

Zero-exception reminder: No exception process exists for tools used with Confidential, Financial (Unpublished), Personal Data, Source Code, or OT/Process Data. These restrictions are absolute and cannot be overridden by any approval, regardless of seniority.

20. AI Guardrails

The AI Guardrails Framework defines the technical and operational boundaries within which all approved AI systems at Welspun World must be configured, operated, and monitored. Guardrails are not restrictions on productivity — they are the engineering and governance controls that make safe, sustained AI adoption possible at scale. They prevent the AI platform from being used in ways that expose the Group to data leakage, regulatory breach, reputational harm, or safety risk, while preserving the full productive capability of the platform for all permitted uses.

This Framework applies to all approved AI systems across all Group entities — the Group-approved enterprise AI platforms, approved entity exception platforms (Clause 23), and approved open-source LLM deployments (Clause 24). Where a specific guardrail cannot be technically enforced on a particular platform, compensating controls must be documented and approved by the CISO before that platform can continue to be used.

Governing authority: The AI Governance Committee is responsible for approving, reviewing, and amending this Framework. The CISO Office is responsible for operationalising, monitoring, and annually testing the guardrail controls defined here.

21. HR Data — Specific Risks & Safeguards

HR data represents the highest-risk AI exposure intersection within the Group: it combines sensitive personal data (protected under DPDP Act, GDPR, and employment law) with the functions where AI adoption pressure will be greatest — recruitment screening, performance management, payroll analysis, and workforce planning. The controls in this clause are mandatory and apply in addition to the general risk tier requirements in Clause 7.

HR AI sits at the highest risk intersection: sensitive personal data + high AI adoption pressure + direct individual impact. All HR AI use cases are classified as minimum Medium risk tier. Any HR AI system influencing individual employment decisions requires DPO approval and a completed DPIA.

HR Domain	Risk Exposure	Governance Control
Payroll & Compensation	Risk Salary benchmarking prompts may expose individual pay bands, increment decisions, or precise CTC data. Once	Control Classified as CONFIDENTIAL. DLP upload block enforced. Any AI processing of payroll data requires DPO clearance and a completed DPIA.

	entered into a cloud AI tool without an enterprise DPA, this data has left the corporate boundary permanently.	Enterprise DPA with the AI vendor must explicitly exclude compensation data from model training.
Recruitment & Screening	Risk AI-assisted CV screening carries a material discrimination risk: models trained on historical hiring data may perpetuate existing biases around gender, ethnicity, age, or educational background. Classified as High risk tier.	Control DPIA mandatory before any AI recruitment tool deployment. Bias audit conducted before go-live and annually thereafter. Human review gate is a hard requirement on all shortlisting decisions — no candidate may be rejected solely by AI output.
Employee PII & Health Data	Risk Medical records, biometric data, disciplinary files, and union membership information are special-category personal data under GDPR Article 9 and highly sensitive under DPDP. Any AI processing of these categories triggers the most stringent regulatory protections.	Control Prohibited from AI processing unless Critical tier approved by AI Governance Committee. Requires DPO and Legal sign-off, written consent notice to the affected employee, and a documented Article 9(2) GDPR exception condition.
Performance & Exit Data	Risk AI-generated appraisal summaries, attrition predictions, or performance trend analysis that influence individual outcomes are subject to the automated decision-making protections in DPDP Act Section 8(6) and GDPR Article 22.	Control Human override is mandatory — no performance decision may be taken solely on AI output. The AI system must be explainable at the level required by the use case. Every affected employee has the right to contest an AI-influenced decision through a documented process.

The AI License Request Form (WEL-ISMS-FRM-AI) — which is mandatory for all HR AI use cases — captures the specific data categories involved, the business justification, and the KPIs against which benefit will be measured. Licenses are time-limited to 90 days, after which a benefit realisation review determines whether to renew or revoke.

22. Cost Optimisation & Controlled Adoption

The AI governance framework is not designed to restrict AI adoption — it is designed to make adoption sustainable, auditable, and cost-effective. A structured license governance model ensures that every AI expenditure is justified by a quantified business case, reviewed at 90 days, and renewed only where measurable benefit has been demonstrated.

22.1 License Governance Model

Step 1 — Request	Step 2 — Assess	Step 3 — Measure
Employee submits the AI License Request Form with a quantified business justification, defined KPIs, and a learning plan. Vague requests are returned for clarification. The form captures:	CISO office screens for data risk within 2 business days. DPO reviews if personal data is involved. Security and privacy assessment completed within 5 business days. Approval is granted by CISO	Mandatory 30-day onboarding check-in with the AI Centre of Excellence. Mandatory 90-day benefit realisation review against the KPIs declared at the time of request. Mandatory 6-month

tool, data categories, estimated hours saved per month, and a 90-day benefit target.	(Low/Medium risk) or AI Governance Committee (High/Critical). Maximum license term: 90 days.	license renewal gate. Licenses that cannot demonstrate measurable benefit are not renewed — cost without value is eliminated.	

22.2 Adoption Guardrails

Guardrail	How It Works
Approved tools only	Enterprise SSO provisioned for all Group-approved enterprise AI platforms. All public AI platforms — ChatGPT free tier, personal Claude, personal Gemini accounts — blocked at DNS and CASB. Employees on managed devices cannot access unapproved AI tools from the corporate network.
5-tier risk classification	Every AI use case is classified Low / Medium / High / Critical-IT / Critical-OT before deployment. Low tier requires standard registration. Medium requires manager and IT sign-off. High requires CISO, Business Owner, and DPO approval with a mandatory DPIA. Critical-IT requires the AI Governance Committee. Critical-OT additionally requires the OT Security Manager and Plant Safety Officer with IEC 62443 assessment.
Exception process	A structured six-step exception process exists for legitimate business needs that cannot wait for the standard procurement cycle. All exceptions are time-limited (maximum 90 days), non-renewable without fresh assessment, and subject to zero-exception data restrictions — no exception will ever be granted for Confidential, Financial (Unpublished), Personal Data, Proprietary Source Code, or OT/Process Data.
AI Guardrails Framework	System prompt standards, output constraints, topic restrictions, agentic workflow limits, and jailbreak mitigations are being developed by the CISO Office and Group IT for AI Governance Committee approval within 90 days. Once approved, the framework will govern how all Group-approved enterprise AI platforms may be configured and constrained across all use-case categories and user roles.

23. Platform Strategy — Group & Entity-Level Guidance

Welspun World comprises multiple legal entities, each of which may have adopted AI tools independently. The Group platform strategy must accommodate this reality while establishing a clear, governed framework that permits multiple enterprise AI platforms to operate in parallel — each subject to the same security, data protection, and governance standards — with a time-bound path to full enterprise-grade deployment across all entities.

23.1 The Multi-Platform Governance Model

Group IT adopts a Multi-Platform Governance Model. Track A covers the Group-approved enterprise AI platforms — ChatGPT Enterprise and Claude Enterprise — a defined set of enterprise AI tenants procured and governed at Group or entity level, each representing a fully governed deployment target for specific use cases. Track B is the entity exception pathway — a structured, time-limited process that allows existing deployments on non-approved platforms to continue by use case, subject to security controls, while migrating to enterprise-grade contracts.

Track A — Group-Approved Enterprise AI Platforms	Track B — Entity Exception Pathway
What it is:	What it is:

ChatGPT Enterprise and Claude Enterprise — procured and governed at Group or entity level, configured to enterprise security standards, and provisioned to authorised users through enterprise SSO. These platforms collectively represent the governed target state for all AI use cases.

Who it applies to:

All entities for all new AI use cases from the date of CFO endorsement. Entities without existing AI deployments onboard directly to ChatGPT Enterprise or Claude Enterprise as directed by Group IT.

Timeline:

ChatGPT Enterprise and Claude Enterprise tenants operational within 60 days of CFO endorsement. Group IT leads procurement, platform registration, and SSO integration for each approved platform.

A governed exception process for entities already using a platform not yet registered as a Group-approved enterprise AI platform. Existing use cases may continue — by individual use case, not as a blanket approval — but must migrate to the enterprise version of the same platform within 90 days, following which the platform may be submitted for registration under Track A.

Who it applies to:

Entities currently using AI platforms on non-enterprise subscriptions or platforms not yet registered under Track A. Each use case must be registered and approved individually — not the platform as a whole.

Migration requirement:

Enterprise contract (ChatGPT Enterprise or Claude Enterprise) with DPA and no-training confirmation must be in place within 90 days. Personal and team-tier subscriptions must be decommissioned at migration.

23.2 Use-Case-by-Use-Case Exception Registration

The entity exception pathway (Track B) operates at the use-case level, not the platform level. An entity cannot register "OpenAI" or "Claude" as an exception — it must register each specific use case it is running on that platform. This approach forces a review of which workflows genuinely require the current platform, and which can be deployed on an already-approved Group enterprise AI platform.

Step	Stage	What Happens
1	**Use-case inventory**	Entity submits a complete inventory of all active AI use cases on non-Group platforms to the CISO Office within 30 days of CFO endorsement. Each use case must state: platform, data categories, business function, estimated users, and monthly cost.
2	**CISO risk screening**	CISO Office screens each use case against the data classification matrix. Use cases involving Confidential, Financial (Unpublished), Personal Data, or OT/Process Data are referred for immediate remediation — they may not continue on non-enterprise platforms.
3	**Exception approval**	Approved use cases receive a 90-day exception to continue on the current platform, conditional on: (a) upgrading to the enterprise version within 90 days; (b) executing a DPA with the vendor; (c) confirming no model training on company data.
4	**Enterprise migration**	Within 90 days, the entity upgrades to the enterprise version, executes the required DPA, and notifies CISO Office. The use case is then registered in the Group AI Use-Case Registry under the enterprise platform.
5	**Ongoing parallel use**	Entities may continue to run their specific use cases on multiple approved AI platforms indefinitely, provided all active use cases are registered, running on enterprise contracts, and subject to the same six-layer security controls as all Group-approved enterprise AI platforms.

Key principle: the exception is granted to the use case, not to the employee or the entity. An entity running five OpenAI workflows must register five separate exceptions. This creates the visibility needed to manage cost, risk, and migration progress.

23.3 Entity-Level vs Group-Level Governance

Governance Element	Group Level	Entity Level
Platform procurement	ChatGPT Enterprise and Claude Enterprise: registered and governed by Group IT	Entity may run approved exceptions on enterprise-tier platforms
Policy & risk framework	Uniform across all entities — non-negotiable	Entities implement locally; CISO Office audits compliance
AI Use-Case Registry	Group-level registry: all use cases across all entities	Entity submits registrations; Group consolidates
DPA & vendor contracts	Group negotiates master DPA for ChatGPT Enterprise and Claude Enterprise	Entity responsible for enterprise DPA on exception platforms
AI Governance Committee	Group Committee: High/Critical tier approvals	Entity may establish a local AI Working Group for Low/Medium tier
VAPT & Audit	Annual Group-wide CERT-In empanelled audit covers all platforms	Entity-specific AI VAPT for on-premises and exception deployments

24. Open-Source LLMs — Azure-Hosted & On-Premises GPU

Building on open-source Large Language Models (LLMs) is encouraged as an exception pathway for specific use cases — particularly where transaction volumes make cloud API costs disproportionate, where data sensitivity makes cloud egress unacceptable, or where low-latency inference is required for manufacturing or operational workflows. Open-source LLMs are not a shortcut around governance; they are a governed alternative with their own controls and must comply with all six security layers referred in WEL-ISMS-POL-AIS.

24.1 Two Deployment Paths

Path A — Azure AI Foundry / Azure OpenAI Service (Managed)	Path B — On-Premises GPU (Self-Hosted)
What it covers: Open-source models (Llama, Mistral, Phi, Falcon etc.) deployed through Azure AI Foundry or Azure Machine Learning, running within the Group's Azure tenant. Data never leaves the Azure boundary. Why it is encouraged: Leverages existing Azure enterprise agreement and security controls. No new vendor contract needed. Data residency in Azure India region satisfies DPDP requirements. Significantly lower per-call cost at high volume.	What it covers: Open-source models hosted on Group-managed GPU servers in data centres or approved colocation facilities. Suitable for highest-sensitivity use cases or manufacturing inference where sub-100ms latency is required. Why it is encouraged: Zero data egress — all inference stays on-premises. No per-call cost at high volume. Enables high-frequency manufacturing AI (quality inspection,

Key controls: Entra ID authentication. Azure Private Endpoint (no public internet exposure). Azure Monitor integration with Group SIEM. Model accessed via internal API only — no direct internet access.	predictive maintenance) without cloud latency. Full model control. Key controls: Network-segmented GPU server. Encrypted model storage (AES-256). PAM-controlled access. Serving layer (vLLM / Ollama / Triton) patched within 30 days of critical CVEs. Inference logs to Group SIEM. Quarterly energy/sustainability reporting.

24.2 Model Selection & Vetting

Not all open-source models are production-safe. Before any model is deployed — on Azure or on-premises — it must pass a six-criterion vetting checklist:

Vetting Criterion	Requirement
Commercial licence	Must permit commercial use (Apache 2.0, MIT, or Llama community licence with commercial permission). Legal confirms licence compatibility before deployment approval.
Training data transparency	Provider must publish training data sources. Models known to have been trained on unlicensed personal data or proprietary code are prohibited. Legal assesses IP liability.
Security assessment	Model undergoes adversarial testing: prompt injection susceptibility, jailbreak resistance, output safety, and absence of known backdoors. Conducted by CISO or appointed assessor before go-live.
Bias & fairness assessment	Bias test aligned to the intended use case. A model that performs acceptably for document summarisation may exhibit bias in HR or customer-facing contexts. Context-specific testing is mandatory.
Model card or equivalent	Provider must publish architecture details, intended use, known limitations, evaluation results, and misuse guidance. Absence of a model card is disqualifying for High/Critical tier use cases.
Active maintenance	Model must be under active maintenance with a security patch process. Unmaintained models (no updates in 12 months) are prohibited for new deployments and must be scheduled for retirement.

24.3 Approval Process

Open-source LLM deployments follow the same AI Use-Case Registration process as cloud platforms, with two additional steps: model vetting (§24.2) and an infrastructure review before provisioning. The CISO approves Low/Medium risk deployments; the AI Governance Committee approves High/Critical. Quarterly review is mandatory for all on-premises deployments. All open-source LLMs — on Azure or on-premises — must appear in the Group AI Use-Case Registry.

25. AI-Specific Penetration Testing (VAPT)

The annual CERT-In empanelled security audit required under CERT-In Directions 2022 must include AI-specific penetration testing — not just the standard network and application VAPT scope. AI systems introduce attack vectors that conventional VAPT does not test: prompt injection, model extraction, adversarial inputs, and

routing layer manipulation. Refer WEL-ISMS-POL-AIS document defining the minimum AI VAPT scope for the Group.

26. ISO 42001 — AI Governance Standard for Entities

ISO/IEC 42001:2023 is the international standard for Artificial Intelligence Management Systems (AIMS). It provides a structured framework for organisations to demonstrate that their AI systems are governed responsibly — covering risk management, ethics, transparency, and continuous improvement. Adoption of ISO 42001 is strongly recommended for all Welspun World entities deploying AI at scale, and is required for entities deploying High or Critical tier AI systems.

Approval

Role	Name	Signature	Date
Head of Cyber Security	***Vinayak Laud***		
Group CTO	***Anil Nimbargi***		
Textile CIO	***Subramanyam Gupta GV***		
MD	***Dipali Goenka***		

Annexure – 1

List of AI Systems

1. Traditional predictive and analytical machine learning models
2. Generative AI systems, including large language models, image generation, and code generation tools
3. Agentic AI and autonomous workflow systems as defined in Clause 17
4. Retrieval-Augmented Generation (RAG) systems that connect AI models to internal knowledge bases, document repositories, databases, or enterprise data sources to augment inference
5. AI systems extended via plugins, connectors, Model Context Protocol (MCP) servers, or similar integration mechanisms that enable AI tools to access internal data sources, APIs, or enterprise systems
6. AI features embedded within third-party SaaS, ERP, HRMS, procurement, or operational technology
7. Fine-tuned or custom-trained models developed using Group data
8. Open-source AI models deployed under Clause 24
9. AI systems used in supply chain, procurement, or vendor-facing contexts by or on behalf of the Group
10. AI in operational technology (OT), manufacturing, and industrial control systems as governed by the Critical — OT/Safety tier (Clause 7).