

Policy on Governance and Responsible Use of Artificial Intelligence (AI)

DOCUMENT
IDENTIFICATION : PO-006

LATEST
REVISION : 1.1

REVISION DATE : 06/10/2026

Policy on Governance and Responsible Use of AI Systems	Department	: IT Security
	Document No.	: PO-006
	Revision No.	: 1.1
	Rev. Release Date	: 06/10/2026

Document Control Information

Document Title	:	Policy on Governance and Responsible Use of AI Systems
Document No.	:	PO-006
Document Revision	:	1.1
Effective Date	:	06/10/2026
Document Classification	:	JSW Internal

Document Publication History

Document Prepared By	:	IT Security Department
Date of Creation	:	16/02/2026
Document Reviewed By	:	Manish Sehgal
Document Approved By	:	Harish Mehra
Initial Release Revision	:	1.0
Initial Release Date	:	16/02/2026

Document Revision History

Revision No.	Revision Description	Change Author	Change Reviewer	Change Approver	Revision Date
1.0	Initial Release	IT Security Department	Manish Sehgal	Harish Mehra	16/02/2026
1.1	Consolidated logically grouped JSW policy with manufacturing, OT, lifecycle and agentic AI requirements.	IT Security Department	Manish Sehgal	Harish Mehra	06/10/2026

Policy on Governance and Responsible Use of AI Systems	Department	: IT Security
	Document No.	: PO-006
	Revision No.	: 1.1
	Rev. Release Date	: 06/10/2026

Confidentiality Agreement

This document is copyrighted, and all rights are reserved. This document may not, in whole or in part, be copied, photocopied, reproduced, translated, or reduced to any electronic medium or machine-readable form without prior consent, in writing, from an authorized representative of JSW Group. This document is for internal use only and may, in whole or in part, be provided to anyone outside of Company, including customer, clients, or prospects after taking an approval from an authorized representative of JSW Group.

Note: The controlled master copy of this document is on the shared drive / Intranet Portal of JSW Group. Printed copies are not controlled. If you are working from a printed copy, please verify with the published version in the shared drive / Intranet Portal of JSW Group to ensure it is the latest revision.



Authorized Signatory

Name: Harish Mehra, Group CIO

Policy on Governance and Responsible Use of AI Systems	Department	: IT Security
	Document No.	: PO-006
	Revision No.	: 1.1
	Rev. Release Date	: 06/10/2026

Table of Contents

1. Introduction
2. Purpose
3. Scope
4. Monitoring and Review
5. Objective
6. Policy Commitments
 - 6.1 Governance, Accountability and Human Oversight
 - 6.2 Privacy, Data Protection, Intellectual Property and Cybersecurity
 - 6.3 Fairness, Reliability and Transparency
 - 6.4 Responsible Use, Generative AI, Agentic AI and Prohibited Practices
 - 6.5 Environmental Sustainability
 - 6.6 Manufacturing, OT, Safety and Product Quality
 - 6.7 Lifecycle Governance and Third-Party Assurance
 - 6.8 Incidents, Exceptions, Training and Enforcement
7. RACI Matrix
8. Measurements and Monitoring
9. Reference
10. Annexure
 - Annexure A – Terms and Definitions
 - Annexure B – Abbreviations
 - Annexure C – AI Risk Classification and Approval
 - Annexure D – Minimum Approval Evidence

Policy on Governance and Responsible Use of AI Systems	Department	: IT Security
	Document No.	: PO-006
	Revision No.	: 1.1
	Rev. Release Date	: 06/10/2026

1. Introduction

Artificial Intelligence (AI) enables JSW Group to improve productivity, decision-making, product quality and operational performance. Its use can also create risks to people, information, intellectual property, industrial assets, business continuity and the environment. This policy establishes proportionate governance so that AI delivers business value within approved and accountable boundaries.

JSW Group includes JSW Steel Limited, JSW Energy Limited, JSW Infrastructure Limited, JSW Cement Limited, JSW Paints Limited and other Group companies and subsidiaries. Implementation shall follow the applicable entity's governance and delegated authority.

2. Purpose

To establish mandatory requirements for the responsible development, procurement, deployment, use, monitoring and retirement of AI systems, protecting JSW information and operations while enabling innovation and practical business adoption.

3. Scope

This policy applies to JSW business and support functions, plants, offices and information assets hosted by or on behalf of JSW across geographies, and to employees, consultants, associates, contractors and third parties acting on its behalf. Application to joint ventures and other entities shall follow management control and contractual obligations.

It covers predictive analytics, machine learning, computer vision, optimisation, AI-enabled digital twins, generative AI, copilots and agents; internally developed, purchased, cloud-hosted, open-source and embedded AI in ERP, MES, engineering, quality, EHS, HR, finance and cybersecurity systems. Pilots and proofs of concept are in scope when they involve JSW data, systems, people or assets.

4. Monitoring and Review

The IT Security Department shall coordinate review at least annually with Digital/IT, Legal, Data Privacy, Risk and business owners, involving HR, Engineering/OT, EHS and Quality where relevant. Earlier review is required following material legal, technology, business or risk changes, significant incidents or evidence of ineffective controls.

Business and system owners shall maintain evidence of compliance. Findings and deviations shall have an accountable owner, due date and escalation route. Internal Audit shall provide risk-based independent assurance through established audit arrangements.

5. Objective

Enable safe, secure, lawful, reliable and fair AI use with measurable business benefits, meaningful human oversight, clear ownership and controls proportionate to the impact and autonomy of each use case.

6. Policy Commitments

JSW shall apply the following commitments throughout the AI lifecycle. An approved tool does not automatically authorise every data category, integration or use case. Existing security, privacy, engineering, safety, quality and delegated-authority requirements continue to apply.

6.1 Governance, Accountability and Human Oversight

6.1.1 Each AI system shall have a named business owner accountable for its purpose, benefits, risk, approved use and ongoing outcomes, and a technical owner responsible for its engineering and operational lifecycle. AI shall not obscure human accountability.

6.1.2 Digital/IT shall maintain an AI register with business owners. Register systems and use cases before pilot or production, including embedded AI that is enabled. Record purpose, owners, provider/model, data, deployment location, integrations, autonomy, risk tier, approvals, review dates and retirement status. General use of an approved productivity service may be registered centrally; individual prompts need not be

Policy on Governance and Responsible Use of AI Systems	Department	: IT Security
	Document No.	: PO-006
	Revision No.	: 1.1
	Rev. Release Date	: 06/10/2026

separate inventory entries.

6.1.3 Classify each use case before development, procurement or deployment using Annexure C. Classification shall consider safety, people, legal rights, product quality, financial impact, environmental consequences, critical operations, sensitive data, autonomy, scale and reversibility. A proof of concept does not waive required reviews.

6.1.4 Approvals shall follow delegated authority. Low-risk use may proceed within an approved enterprise service and documented usage boundaries. Moderate-risk use requires business-owner approval and applicable Technology, Security and Data Privacy reviews. High-risk pilot and production use require executive approval through the Group CIO/designated governance forum, with relevant specialist concurrence. No new governance council is assumed by this policy.

6.1.5 Document the intended purpose, users, operating environment, limitations, data sources, tools, decision rights, prohibited actions and permitted autonomy. Material changes require reassessment, validation and reapproval at the authority corresponding to the revised risk.

6.1.6 Assign competent human reviewers with enough information, time and authority to challenge, override, stop or escalate AI-supported decisions. High-impact or irreversible decisions shall not rely solely on AI. Define the appropriate human-in-the-loop, human-on-the-loop or human-in-command oversight.

6.1.7 Where workers or individuals are materially affected, provide appropriate notice, accessible challenge and correction channels, and HR/Legal review. AI shall not be the sole basis for hiring, termination, discipline, promotion, compensation or medical-fitness decisions. Plan role-specific training and reskilling where AI materially changes work.

6.2 Privacy, Data Protection, Intellectual Property and Cybersecurity

6.2.1 Use data only for an approved and lawful purpose respecting data privacy in the development, deployment and use of AI. Apply data minimisation, quality checks, ownership, lineage, classification, retention and secure disposal. Obtain Data Privacy/Legal review where personal data is involved, including notice, consent or other applicable legal basis, rights handling and cross-border processing as required.

6.2.2 Confidential, personal, regulated, customer, supplier, engineering, drawing, formulation, source-code and plant/process information shall be processed only in a service and configuration explicitly approved for that data and use case. No JSW information shall be entered into unapproved public or consumer AI services.

6.2.3 Confirm rights to training, fine-tuning, grounding, prompt and output data. Protect trade secrets and assess generated content, code and designs for licensing, infringement and confidentiality before reuse or publication. AI output shall not be assumed to be original, accurate or legally protected.

6.2.4 Provider use of JSW prompts, uploads, interactions or outputs for model training requires explicit approval and contractual controls. Retention, deletion, data location and subprocessors shall be assessed before onboarding. De-identification and synthetic data shall be evaluated for re-identification and hidden bias where relevant.

6.2.5 Apply approved architecture, secure configuration, managed identities, least privilege, role-based access, MFA for supported human access and segregation of duties. Protect models, system instructions, APIs, connectors, secrets, embeddings, vector stores, training data and logs.

6.2.6 Protect the cybersecurity of systems used in the development, deployment and operation of AI. Assess prompt injection, jailbreaks, poisoned data, model theft, sensitive-data leakage, unsafe outputs, insecure plugins, supply-chain compromise and denial of service. Treat retrieved content and model output as untrusted; validate before execution and restrict tools and outbound transfer through approved controls.

6.2.7 Perform security testing, vulnerability management and adversarial testing proportionate to risk. Log material decisions, privileged access and tool actions with access control and tamper protection. Logs shall minimise sensitive content and follow approved retention requirements.

Policy on Governance and Responsible Use of AI Systems	Department	: IT Security
	Document No.	: PO-006
	Revision No.	: 1.1
	Rev. Release Date	: 06/10/2026

6.2.8 Maintain incident response, tested fallback, rollback and recovery appropriate to business criticality. OT assets shall not be exposed directly to public AI services or receive unvalidated AI output through uncontrolled connections.

6.3 Fairness, Reliability and Transparency

6.3.1 Validate data quality, representativeness and provenance for the intended use. Assess and mitigate unjustified bias, in data, models and outputs, particularly where people or stakeholders may be materially affected. Avoid unlawful inference of sensitive characteristics.

6.3.2 Define acceptance criteria before release for relevant accuracy, false accepts/rejects, robustness, latency, uncertainty, safety and operational performance. Test realistic conditions, edge cases, degraded sensors, missing data, adversarial inputs and foreseeable misuse. High-risk systems require validation sufficiently independent of development.

6.3.3 Monitor relevant performance, bias, drift, hallucination, availability and control effectiveness across affected plants, products, operating conditions or groups. Define thresholds, corrective actions and suspension criteria; revalidate periodically and following material changes.

6.3.4 Communicate material AI involvement, intended purpose and limitations where required by law, contract or risk. Provide explanations and provenance appropriate to the decision and audience, subject to privacy, confidentiality, security and intellectual-property constraints.

6.3.5 Users shall verify material facts, calculations, citations, code and engineering parameters before reliance. External publications, regulatory submissions, legal documents, product claims and engineering deliverables require qualified human review. Disclose synthetic media where necessary to prevent deception.

6.4 Responsible Use, Generative AI, Agentic AI and Prohibited Practices

6.4.1 Use enterprise-approved services and accounts for authorised purposes. Complete appropriate acceptable-use training before access. Treat generative output as a draft or recommendation unless automated execution is expressly approved. AI-generated code requires secure development review, testing and approval before release.

6.4.2 Each agent shall have a named owner, registered system identity, defined objective, approved tools and bounded authority. Default to read-only access; write or execute privileges require documented need and approval. Apply action allowlists, transaction limits, approval gates and segregation of duties.

6.4.3 Human approval is required for high-impact safety, financial, legal, employment, external communication, master-data, production-control or code-deployment actions. Any preauthorised routine automation shall have documented limits, appropriate risk approval and a practical means to intervene; it shall not waive high-impact decision review.

6.4.4 Agent records shall capture material instructions, tool calls, data accessed, approvals, outputs and resulting transactions proportionate to risk. Provide tested stop/containment and rollback mechanisms where feasible, rate and resource limits, loop detection and controlled credential lifetimes. Agents shall not expand permissions, delegate outside approved boundaries or alter their guardrails without authorised change control.

6.4.5 Prohibited uses include unlawful manipulation, exploitation of vulnerabilities, social scoring, unlawful discrimination or surveillance, malicious impersonation, fraud, harassment and unlawful infringement of third-party rights.

6.4.6 AI shall not fabricate inspection, maintenance, test, regulatory or approval records; bypass security, safety interlocks, quality gates or segregation of duties; or independently change production logic, access rights or safety limits without authorised change control.

6.4.7 Unregistered systems, unapproved integrations, untested deployment into live production and unapproved autonomous control of safety-critical operations are prohibited. No exception may authorise



Policy on Governance and Responsible Use of AI Systems	Department	: IT Security
	Document No.	: PO-006
	Revision No.	: 1.1
	Rev. Release Date	: 06/10/2026

unlawful activity or unacceptable safety risk.

6.5 Environmental Sustainability

6.5.1 Consider energy, compute, infrastructure and resource consumption when choosing models and providers. Use appropriately sized models and efficient operation aligned with JSW sustainability objectives, without compromising safety, reliability or security.

6.5.2 Validate evidence supporting claimed energy, emissions or sustainability benefits. AI shall not suppress adverse readings, substitute for required statutory monitoring or create unsupported environmental claims.

6.6 Manufacturing, OT, Safety and Product Quality

6.6.1 AI shall follow existing Engineering, OT, EHS, functional/process safety, maintenance, Quality and management-of-change requirements. It shall not replace statutory accountability, competent engineering judgement or required certified safety functions.

6.6.2 AI shall not be the sole protective layer against a hazardous event. Emergency stops, safety instrumented functions and certified interlocks shall remain independent of non-safety AI; any alternative requires competent safety engineering, applicable certification and formal approval.

6.6.3 Bound recommendations and control actions within approved operating envelopes. Validate loss of model, sensor, data, network or cloud service against a defined safe or controlled state. Operators shall have clear authority and practical means to reject, override or stop unsafe actions.

6.6.4 Use simulation, shadow mode or phased trials appropriate to risk before production control. Validate process optimisation and digital-twin assumptions; unverified simulation outputs shall not directly change plant control parameters.

6.6.5 Predictive maintenance shall support authorised work decisions and shall not remove preventive-maintenance obligations without approved change. Safety computer vision shall be tested for lighting, occlusion, PPE and worker variation, respect privacy, and supplement required physical safeguards.

6.6.6 Quality inspection shall be validated by defect, product, material, line and operating condition. Record traceability, false accepts/rejects and human disposition authority. Product release shall remain subject to approved Quality controls.

6.6.7 Robotics and autonomous mobile equipment require machinery/task risk assessment, bounded movement and speed, safe states, obstacle handling and emergency-stop controls. EHS and environmental analytics shall validate sensor integrity and detection limits and preserve statutory monitoring and reporting.

6.6.8 Production planning, supplier selection and energy optimisation shall retain authorised approval for material commitments and remain within safety, environmental permits, equipment, product-quality and continuity constraints. Relevant Engineering/OT, EHS and Quality owners shall approve affected use cases.

6.7 Lifecycle Governance and Third-Party Assurance

6.7.1 Before procurement or development, document the business problem, measurable benefit, alternatives, intended scope and foreseeable failure or misuse. Moderate-risk use requires a documented assessment; high-risk use requires a formal multidisciplinary impact assessment covering people, safety, quality, environment, legal, privacy, security, autonomy and resilience.

6.7.2 Design and build in approved environments with separate development, test and production, version control, secure engineering and change approval. Record requirements, architecture, data lineage, model/version, dependencies, limitations and validation results.

6.7.3 Production release requires recorded approvals and risk acceptance, completed validation, trained users, operating instructions, monitoring thresholds, support ownership, incident processes and tested fallback appropriate to risk. Use phased rollout where necessary. Material unresolved safety or security issues shall prevent release.



Policy on Governance and Responsible Use of AI Systems	Department	: IT Security
	Document No.	: PO-006
	Revision No.	: 1.1
	Rev. Release Date	: 06/10/2026

6.7.4 Procurement shall involve relevant business, IT Security, Data Privacy and Legal reviewers. Assess vendor/model provenance, licensing, security, data use, subprocessors, processing locations, resilience, support and concentration risk. Open-source components require licence, provenance, security and maintenance assessment.

6.7.5 Contracts shall address confidentiality, data ownership and permitted use, provider training, security, incident notification, change notice, audit evidence, service levels, intellectual property, appropriate liability provisions, data return/deletion and exit. Third-party supply does not transfer JSW business accountability.

6.7.6 Monitor material provider and model changes, including embedded AI features. Reassess changed purpose, data, integration, model, autonomy or operating conditions. At retirement revoke access and credentials, preserve required evidence, dispose of data/models according to obligations and maintain continuity for dependent processes.

6.8 Incidents, Exceptions, Training and Enforcement

6.8.1 Report material errors, unsafe outputs, bias, misuse, suspected leaks, anomalous agent actions and near misses promptly through established incident channels. Where continued operation could cause material harm, authorised operators shall contain or suspend affected automation and use a validated fallback while protecting people and stabilising operations.

6.8.2 Preserve relevant logs, model/version, inputs, outputs, approvals and actions securely. Coordinate business owners with Security, Privacy, Legal, HR, Engineering/OT, EHS and Quality as applicable. Assess notification duties, investigate root cause, track corrective actions and revalidate before controlled restoration.

6.8.3 Provide responsible-use training for users and role-specific training for developers, approvers, operators and reviewers. Cover limitations, data handling, verification, fairness, IP, safety, incident reporting and human accountability. Protect good-faith reporting from retaliation.

6.8.4 Exceptions require a documented business justification, risk assessment, compensating controls, named owner, scope, expiry and approval through applicable executive/delegated authority with relevant control-owner concurrence. Review before expiry; maintain a central exception record.

6.8.5 Non-compliance may lead to access withdrawal, suspension, corrective action, contractual remedies or disciplinary action under applicable law and JSW procedures. Retain inventory, assessments, approvals, validation, monitoring, changes, incidents, training, vendor assurance and retirement evidence under applicable records requirements.

7. RACI Matrix

R = Responsible; A = Accountable; C = Consulted; I = Informed. Specialist approvals remain mandatory where applicable. The Group CIO/designated executive authority decides high-risk approvals; the business owner retains accountability for operational outcomes.

Activity	Responsible	Accountable	Consulted	Informed
Intake, inventory and classification	Business owner; Digital/IT	Business owner	Security, Privacy, Legal; OT/EHS/Quality/HR as applicable	CIO
Security/privacy assessment	IT Security; Data Privacy	Head – Information Security (security); Privacy lead (privacy)	Business, Digital/IT, Legal	CIO
Safety and quality assessment	Engineering/OT; EHS; Quality	Relevant functional head for their assessment	Business, Digital/IT, Security	CIO
Build, validation and change	Technical owner; independent validator where required	Technical owner	Business, Security and relevant specialists	Business head

Policy on Governance and Responsible Use of AI Systems

Department : IT Security

Document No. : PO-006

Revision No. : 1.1

Rev. Release Date : 06/10/2026

Activity	Responsible	Accountable	Consulted	Informed
Production approval	Business sponsor; Digital/IT	Business head for low/moderate risk; CIO/designated executive authority for high risk	Security, Privacy, Legal and affected specialist approvers	Operators/users
Supplier due diligence and contract	Procurement; technical owner	Procurement head	Business, Security, Privacy, Legal	CIO
Operation, monitoring and retirement	System/technical owner	Business owner	Digital/IT and relevant control owners	Management
Incident coordination	System owner; relevant incident teams	Business owner for AI outcome; incident lead for response	Security, Privacy, Legal, OT/EHS/Quality/HR	Management as severity requires
Policy exceptions	Business owner; control owner	Applicable executive/delegated authority	Security, Legal and affected specialists	Audit/governance
Independent assurance	Internal Audit	Internal Audit head	Business and control owners	Relevant audit forum

8. Measurements and Monitoring

Report the following measures through existing governance forums. Targets and system-specific thresholds shall be documented by owners according to risk. Measures shall balance value with safety and quality and shall not discourage reporting of incidents or human overrides.

Area	Evidence / measure	Frequency	Owner
Inventory coverage	Registered systems with owners, risk tiers and current approvals; enabled embedded AI coverage	On onboarding/change; quarterly reconciliation	Digital/IT; business owners
Assessment and release	Required assessments, validation, specialist concurrence and approvals completed before release	Before pilot/production; material change	Business/technical owners; reviewers
Access and data	Access and permitted data/integrations reviewed; unauthorised use and exposure tracked	Quarterly; upon material change	System owner; IT Security/Privacy
Performance and safety	Relevant accuracy/drift/bias, false accepts/rejects, overrides, near misses and fallback results	Threshold-based monitoring; monthly high-risk review; quarterly moderate-risk review	System owner; relevant specialists
Agent execution	Out-of-bound actions, approval bypass, tool failures, loops and containment readiness	During operation; periodic control review by risk	Technical owner; IT Security
Incidents and exceptions	Severity, containment, overdue corrective actions and expiring exceptions	Ongoing; quarterly governance summary	Business/incident and control owners
Training and supplier assurance	User training completion, vendor findings and material model/provider changes	Before access/onboarding; annual review and material change	HR; Procurement; Digital/IT
Business value and sustainability	Benefits against approved case; resource consumption where material	Quarterly	Business owner

9. Reference

Applicable laws and contractual obligations in each operating jurisdiction prevail where stricter. Legal shall determine applicability; this policy does not assert certification or substitute for legal advice. External frameworks may guide subordinate standards and assessment methods.

**Policy on Governance and
Responsible Use of AI Systems**

Department : IT Security

Document No. : PO-006

Revision No. : 1.1

Rev. Release Date : 06/10/2026

Category	Reference
JSW framework	PO-001 – Information Security Management System Framework
JSW policy	PO-002 – Information Security Policy
JSW procedure	PR-011 – IT Asset Management Procedure; applicable privacy, incident, change, procurement, EHS, OT and Quality procedures
External guidance	ISO/IEC 27001:2022; ISO/IEC 42001; ISO/IEC 23894; NIST AI Risk Management Framework and Generative AI Profile
Records	AI register; risk/impact assessments; approvals; validation; monitoring; exceptions; incidents; vendor and retirement evidence

10. Annexure

Annexure A – Terms and Definitions

Term	Definition
Artificial intelligence (AI)	A machine-based system that infers from inputs how to generate outputs such as predictions, content, recommendations or decisions that can influence physical or virtual environments.
AI system owner	The accountable business role that owns the purpose, value, risk acceptance, controls and ongoing performance of an AI system.
Model owner	The role responsible for model design or selection, validation, versioning, performance and technical lifecycle.
High-impact use	A use that may materially affect safety, employment, legal rights, financial outcomes, product conformity, critical operations, the environment or vulnerable persons.
Generative AI	AI that creates new text, images, audio, video, code, designs, simulations or other content.
Agentic AI	AI capable of planning or executing multi-step tasks through tools, systems, data sources or other agents with a defined degree of autonomy.
Human-in-the-loop	A human must review and approve the output or action before execution.
OT/ICS	Operational technology and industrial control systems used to monitor or control physical processes, machinery and infrastructure.
Material change	A change to purpose, users, model, training data, data sources, integrations, autonomy, thresholds, operating environment or potential impact that could alter risk.
AI incident	An event or near miss involving unsafe, unintended, unlawful, deceptive, inaccurate, biased, insecure or unauthorized AI behavior or outcome.
Human-on-the-loop	A human supervises execution and can intervene or stop actions.
Human-in-command	Humans retain authority over system purpose, deployment and withdrawal.
Independent validation	Review with sufficient independence from development and appropriate competence for the use case.

Annexure B – Abbreviations

Acronym	Meaning
AI	Artificial Intelligence
CIO	Chief Information Officer
ISMS	Information Security Management System
OT / ICS	Operational Technology / Industrial Control Systems
EHS	Environment, Health and Safety
ERP / MES	Enterprise Resource Planning / Manufacturing Execution System

Policy on Governance and Responsible Use of AI Systems

Department : IT Security
Document No. : PO-006
Revision No. : 1.1
Rev. Release Date : 06/10/2026

Acronym	Meaning
MFA / RBAC	Multi-Factor Authentication / Role-Based Access Control
IP	Intellectual Property
RACI	Responsible, Accountable, Consulted, Informed
NDA	Non-Disclosure Agreement

Annexure C – AI Risk Classification and Approval

Classify by actual impact and access, rather than product label. Apply the highest relevant tier. Sensitive data or write access requires specialist assessment and may raise a use case's tier. Unacceptable use cannot be approved as a routine exception.

Tier	Typical use / impact	Minimum approval and evidence
Low	Assistive, reversible productivity use with non-sensitive data and human verification.	Approved enterprise tool and usage standard; business owner; training and output verification. Centrally registered service permitted.
Moderate	Decision support or limited workflows: forecasting, maintenance prioritisation, approved internal knowledge assistants.	Business-owner approval; documented risk, Security/Privacy and Technology reviews as applicable; validation, monitoring and fallback.
High	Material effect on people, product, finance, environment, compliance or critical operations; privileged agents or process optimisation affecting plant limits.	Formal impact assessment; executive approval and applicable specialist concurrence before pilot and production; independent validation, human oversight, logging and periodic assurance.
Unacceptable / prohibited	Unlawful/deceptive use, unacceptable safety risk, control bypass or unapproved safety-critical autonomous control.	Do not develop, procure, deploy or use. Any proposed safety-critical automation must first satisfy separate competent engineering and safety approval, applicable law and this policy's protective requirements.

Annexure D – Minimum Approval Evidence

The business owner shall retain the applicable evidence below. Mark exclusions as not applicable with a reason and reviewer concurrence. High-risk systems require a formal impact assessment; this checklist alone does not replace it.

1. Business purpose, measurable benefits, alternatives and accountable owners.
2. Inventory record, risk classification, intended users and affected stakeholders.
3. Approved data categories, provenance, rights, quality, retention and privacy assessment.
4. Architecture, integrations, tool permissions, autonomy and prohibited actions.
5. Security review; Legal/HR/EHS/Quality/OT review where affected.
6. Model/version, limitations, realistic validation results and independent validation where required.
7. Human decision authority, operating envelope, challenge channels and escalation.
8. Tested fallback, stop/containment and rollback/recovery appropriate to risk.
9. Vendor due diligence, contractual controls, support and exit arrangements.
10. User training, operating instructions, monitoring thresholds and incident readiness.
11. Recorded release decision, residual-risk acceptance, review dates and material-change triggers.