

Global Artificial Intelligence (AI) Governance Policy v1.0

Executive Summary

This v1.0 policy focuses on the enterprise governance standard with explicit roles, risk tiers, controls for agentic AI, confidentiality rules, lifecycle requirements, and audit-ready documentation expectations.

TABLE OF CONTENTS

1. Policy Governance and Ownership	4
2. Purpose	5
3. Scope	6
4. Responsible AI Governance Operating Model	7
5. AI Risk Classification Framework	7
6. Definitions	9
7. Responsible Artificial Intelligence Principles	10
8. Acceptable Use of AI Tools	11
9. AI Agent Governance Standard	12
10. Confidentiality and Data Protection	14
11. AI System Lifecycle Governance	15
12. AI Security Practices	16
13. AI Incident Response and Reporting	17
14. Model Documentation and Audit Requirements	18
15. Vendor and Third-Party AI Governance	19
16. Complaints and Regulatory Escalation	19
17. Policy Amendments and Exceptions	20
18. Contacts and Review Route	20
19. Roles and Responsibilities	21

DOCUMENT CONTROL

Version History:


Date	Version	Description of change	Owner	Approved by
	1.0	Initial release		

Authorization:

Prepared By	Reviewed By	Approved By
Name:	Name:	Name:
Designation:	Designation:	Designation:
Date:	Date:	Date:

Distribution List:

Sr. No	Department or Function Name
1	



1. Policy Governance and Ownership

1.1 Policy intent

This policy establishes the enterprise governance framework for the responsible use of AI across Lupin. It is written to support business adoption while preserving data protection, information security, quality, and regulatory compliance. The policy is intended to be practical and to guide LUPIN employees on what they may do, what they must not do, and which approvals are required before AI is used in higher-risk settings.

The policy also introduces explicit ownership, a risk-tiered decision model, governance for autonomous and agentic AI, and stronger requirements for documentation, monitoring, and incident response.

- This policy applies enterprise-wide and is binding unless a formally approved exception exists
- The latest approved version must be stored in the company-controlled policy repository
- Local practices and team conventions may add controls, but they may not weaken this policy

1.2 Policy owner and accountable functions

The Office of CIO (OCIO) is the primary owner of this policy. OCIO is responsible for maintaining the document, coordinating periodic reviews, and ensuring that AI governance requirements are embedded into business processes. OCIO works in partnership with Information Security, Data Privacy, Legal & Compliance, Quality & Regulatory Affairs, IT & Digital Transformation and relevant business functions.

Ownership of this policy is distinct from ownership of individual AI systems. Every deployed AI system must have a named Model Owner or System Owner who is responsible for the lifecycle of that system, including intended use, performance, controls, maintenance, monitoring, and retirement. Business sponsors remain accountable for the business outcomes and operational use of AI.

1.3 AI Governance Council

LUPIN shall maintain an AI Governance Council to review material AI use cases, approve high-risk deployments, resolve control questions, and provide cross-functional direction. The Council should meet on a defined cadence and may convene ad hoc for urgent matters such as critical incidents, high-risk approvals, or regulatory questions. The Council provides a formal decision-making body and is a central forum for balancing innovation, risk, quality, and compliance. The Council should also review risk tiering, validation expectations, exception requests, and unresolved policy interpretation issues. All the decisions and conditions of approval should be documented and retained for audit purposes.

1.4 Review cadence and recordkeeping

This policy shall be reviewed at least once a year and sooner when required by material regulatory changes, significant incidents, introduction of new high-risk AI capabilities, or major changes to the company's operating model. Review outputs should capture what changed, why it changed, who approved the changes, and the date of effectivity. All governance records associated with this policy, including Council approvals, exceptions, risk reviews, and evidence of employee communication, must be retained in line with record retention requirements.

- Annual review is the minimum; interim reviews are required when risk conditions materially change.
- All superseded versions must remain archived and traceable.
- Evidence of approval and communication must be retained.

2. Purpose

2.1 Why this policy exists

AI can accelerate research, improve productivity, support decision-making, and reduce administrative burden. At the same time, AI introduces meaningful risks: hallucinated content, biased outputs, privacy violations, cybersecurity weaknesses, unauthorized disclosure of confidential information, regulatory exposure, and loss of accountability when automated systems are poorly governed.

The purpose of this policy is therefore to provide Lupin with a wide framework that enables the safe, secure, and ethical use of AI while ensuring that controls scale with risk. The policy is intended for a regulated environment and is written to support inspection readiness and defensible decision making.

2.2 Control objectives

The policy defines the minimum requirements that AI users, system owners, and control functions must follow. These requirements are intended to ensure that AI systems are used only for an approved purpose, with appropriate oversight, technical controls, validation, and monitoring in place. All the mentioned control objectives are organized around governance, risk classification, acceptable use, privacy, security, lifecycle management, agencies, agent governance, incident response, vendor management, and documentation.

- AI must have a defined intended use
- High-risk users must receive explicit approval before deployment
- AI outputs must not be treated as inherently correct

2.3 Alignment to external frameworks

This policy is designed to align with widely recognized governance and risk-management approaches, including NIST AI RMF 1.0, ISO/IEC 42001, applicable privacy laws, and relevant healthcare and life-science expectations such as Good ML Practice concepts. Alignment to these frameworks does not remove the need to comply with internal policies or local laws; instead, it provides a structure for consistent and mature governance.

- NIST AI RMF informs risk identification, governance, measurement, and management
- ISO/IEC 42001 informs management-system discipline and continuous improvement
- EU AI Act (2024)— including prohibited practices (Article 5), high-risk AI obligations (Annex III), and phased compliance timelines
- FDA AI/ML draft guidance (AMT framework 2024)
- India Digital Personal Data Protection Act (DPDPA)
- GDPR (EU)
- Applicable privacy laws and sector-specific obligations remain mandatory

Policy Requirement:

- All AI systems must be mapped to applicable regulatory obligations before approval
- High-risk systems must document regulatory classification and compliance controls
- AI Governance Council ensures regulatory alignment validation

3. Scope

3.1 People in scope

This policy applies to all employees, contractors, consultants, temporary staff, interns, and approved third parties who use, configure, support, develop, procure, or manage AI systems on behalf of Lupin. The policy applies irrespective of job family or seniority. Responsibilities vary by role, but the core requirements are binding on everyone in scope.

Managers are responsible for ensuring that their teams understand the policy and do not rely on unapproved practices. System administrators and technical teams are additionally responsible for implementing the technical and access controls defined in this policy.

3.2 Technology and activity scope

The policy covers all AI technologies used for business purposes, including Generative AI, algorithmic AI, machine learning models, predictive analytics, computer vision, NLP systems, AI coding assistants, recommendation engines, and autonomous or semi-autonomous AI agents. It also applies to AI features embedded within broader platforms and enterprise software.

Covered activities include experimentation, development, procurement, pilot testing, deployment, monitoring, integration into workflows, content generation, decision support, workflow automation, and external publication or distribution of AI-assisted outputs.

- Both internally built and third-party AI systems are covered
- AI features embedded within business software are covered
- Pilots and proofs of concept are covered, not only production deployments

3.3 Environment and data scope

The policy applies across Lupin-managed infrastructure, cloud services, mobile devices used for work, endpoints, vendor platforms, and any environment where Lupin work is performed. It covers inputs, prompts, training data, reference material, model outputs, logs, and derivative work products created with AI assistance. Any use of company information in AI systems falls within scope. This includes confidential business information, regulated records, intellectual property, source code, employee information, and personal data.

- Company devices and approved work use on personal devices are in scope
- Cloud-hosted and vendor-hosted AI environments are in scope
- Data used to prompt, train, evaluate, or operate AI is in scope

3.4 Regulated and restricted contexts

Special care is required where AI is used in or near regulated activities, including pharmacovigilance, medical information, manufacturing, quality systems, regulatory submissions, safety case processing, and any workflow involving GxP-relevant records. These contexts may trigger enhanced validation,

documentation, and approval requirements. Where local law, customer commitments, or product quality requirements are stricter than this policy, the stricter requirement applies.

4. Responsible AI Governance Operating Model

Lupin's responsible AI governance is anchored by a single, cross-functional governance body – the **Responsible AI (RAI) Committee** – which holds end-to-end accountability for principles, governance, risk, and capability building across all AI initiatives.

4.1 Responsible AI (RAI) Committee

Role	Participants
Chair / Members	CIO, CFO, CHRO, Chief Legal Officer
Standing Invitees	CDAO, CISO

Cadence: Meets every 6–8 weeks, with ad-hoc sessions convened for high-impact or critical-risk use cases, AI incidents, or regulatory developments.

4.2 Mandate and Decision Rights

The RAI Committee is responsible for the following two core mandates:

A. Establish Principles and Governance:

- Set clear principles for ethical, safe, and responsible AI use across Lupin
- Enforce governance through monitoring, compliance, and audit mechanisms
- Approve the AI Policy and all material updates
- Approve the list of permitted AI tools, platforms, and use cases (Annexure 1)
- Define platform-level and use-case-level guardrails
- Approve high-risk and critical-risk use cases prior to deployment
- Adjudicate exceptions, deviations, and policy overrides

B. Manage Risks and Build Capability

- Evaluate legal, regulatory, ethical, security, and reputational risks associated with AI
- Guide the business teams on feasibility, risk mitigation, and remediation
- Oversee AI incident management, root-cause analysis, and corrective actions
- Drive enterprise-wide capability building, training, and awareness on Responsible AI
- Review KPIs and metrics on compliance, risk posture, and adoption
- Stay current on evolving regulations (EU AI Act, DPDP Act, GDPR, HIPAA, RBI, SEBI, FDA) and align internal policy accordingly

5. AI Risk Classification Framework

5.1 Purpose of risk tiering

Risk classification must be performed before a new AI use case is approved. The initial classification should be reevaluated whenever the intended use, system autonomy, data sensitivity, user population, or integration scope changes.

- Risk tiering is mandatory before approval of a new use case

- The tier determines the level of review, documentation, and monitoring required
- Reclassification is required when the risk profile materially changes

Tier 1 - Minimal Risk

Tier 1 covers low-impact, assistive uses where AI supports personal productivity and the risk of material harm is low, provided that confidential information is not disclosed and human review remains in place. Typical examples include summarizing non-sensitive text, drafting internal meeting notes, brainstorming ideas, or improving the readability of general content.

Tier 1 use cases still require users to verify outputs, avoid prohibited data inputs, and comply with approved tool requirements. Minimal risk does not mean no risk.

- Examples: email drafting, summarization of non-confidential material, formatting help, translation of low-risk content
- Minimum controls: approved tool, no confidential or personal data input unless specifically allowed human review before use

Tier 2 - Limited Risk

Tier 2 includes internal analysis, workflow support, and decision-support use cases that can influence internal operations but are not independently making high-impact or regulated decisions. Examples include internal forecasting, knowledge retrieval, coding assistance within approved environments, and workflow triage where a human retains decision authority.

These use cases require clearer intended-use statements, stronger logging, named owners, and documented controls because they can influence downstream decisions even if they do not make those decisions autonomously.

- Examples: internal analytics, knowledge assistants, search and retrieval tools, workflow routing support
- Controls: documented intended use, owner assignment, access control, logging, and documented human oversight

Tier 3 - High Risk

Tier 3 applies to AI systems that could materially affect patient safety, product quality, regulatory compliance, or major business outcomes. These systems may support regulated processes, influence quality or safety decisions, generate content used in formal submissions, or automate actions in sensitive enterprise environments. Tier 3 systems require formal review by the AI Governance Council and may require validation, detailed test evidence, change control, enhanced monitoring, incident playbooks, and explicit go-live approval conditions.

- Examples: pharmacovigilance signal support, manufacturing optimization impacting process decisions, AI-assisted regulatory drafting, clinical analytics with material operational consequences
- Controls: formal approval, validation evidence, traceability, monitoring, rollback and containment procedures, defined owner and sponsor

Tier 4 - Prohibited Use

Tier 4 identifies uses that are not permitted because the residual risk is unacceptable or because the activity conflicts with law, policy, ethics, or company risk appetite. Prohibited uses must not be piloted, tested, or deployed without a specific policy amendment or executive exception approved through the defined governance route. Users must escalate uncertainty rather than assume a questionable use is acceptable. The following are examples of critical use cases:

- Autonomous pharmacovigilance decisions (e.g., adverse event causality)
- Autonomous batch release or quality disposition decisions
- AI-driven modifications to MBR/BMR without human approval
- AI-generated or synthetic data in regulatory submissions
- AI use in hiring, firing, or employee disciplinary decisions
- Processing patient-identifiable data without DPIA approval
- Use of public LLMs for GxP data or regulated records

5.2 Risk factors to consider

Risk classification should consider multiple dimensions such as sensitivity of data used, the consequence of a wrong answer, the level of autonomy, the ability to explain outputs, integration with enterprise systems, degree of human oversight, scale of use, and whether the system is used in a regulated process. If there is a doubt between two tiers, the higher tier should be assigned until the matter is reviewed and classified correctly.

6. Definitions

6.1 Core AI terms

For this policy, Artificial Intelligence means software-based methods and systems designed to perform tasks that mimic or augment aspects of human intelligence, including pattern recognition, reasoning, prediction, content generation, recommendation, and decision support. Generative AI refers to models that generate new content such as text, code, images, audio, or video in response to prompts. Algorithmic AI refers to analytical or predictive systems that classify, rank, score, or forecast based on data.

An AI Tool is any application, system, or embedded feature that uses one or more AI techniques to perform or support a task. A Large Language Model, or LLM, is an advanced language model trained on large amounts of text and capable of understanding and generating human language.

6.2 Governance and lifecycle terms

An AI Agent is an AI system capable of executing multi-step tasks, invoking tools or APIs, or taking actions with limited or no human intervention. Autonomy level refers to the degree of independent action a system may take. Intended Use means the specific approved purpose, context, user population, and operating boundaries for an AI system. Model Owner or System Owner means the person accountable for the lifecycle and control environment of a deployed AI capability.

Human Oversight means the controls through which a qualified human reviews, supervises, approves, or can interrupt an AI system. Model Drift refers to the reduction in model performance or reliability caused by changes in data, process context, or the operating environment over time.

6.3 Data and confidentiality terms

Lupin Confidential Information means non-public information belonging to Lupin that could cause competitive, operational, legal, or reputational harm if disclosed. Personal Data means any information relating to an identified or identifiable natural person. Sensitive or special categories of personal data, health information, or patient-related data require heightened protection.

Approved Enterprise AI Platform means an AI platform that has undergone company review and has been explicitly authorized for defined business uses under specified conditions.

- Public AI tools are not automatically Approved Enterprise AI Platforms
- Approval may be conditional and limited to a specific use or dataset
- Confidentiality obligations apply to prompts, uploads, outputs, logs, and derived content

7. Responsible Artificial Intelligence Principles

7.1 Human oversight and accountability

Humans remain accountable for decisions, actions, and work products assisted by AI. No employee may transfer accountability to a model, a vendor, or an automated workflow simply because an AI system was used. The degree of human oversight must match the risk of the use case and should be defined as part of the intended use and approval process. For high-risk contexts, human oversight must be meaningful. This means the reviewer has sufficient competence, time, information, and authority to question, reject, or override the output.

- AI may support judgment, but it does not replace accountability
- The human reviewer must be able to understand the context and consequences of the output
- Oversight must be designed, not assumed

7.2 Safety, reliability, and fitness for purpose

AI systems must be fit for their intended use. Before deployment, the business users must understand what the system is expected to do, its limitations, the conditions in which it performs well or poorly, and the consequences of error. Validation, testing, and pilot evaluation should be proportionate to risk. AI systems should not be used outside their defined purpose without reassessment.

7.3 Transparency and explainability

Users and reviewers should understand when AI is being used, what role it is playing, and the important limitations that may affect trust and interpretation. The required degree of explainability depends on the use case. Some low-risk uses may only need clear user guidance, while higher-risk uses require stronger documentation, rationale capture, and traceability. Externally shared or published material that is materially generated or transformed using AI should include disclosure where required by law, policy, contract, or business process.

- The use of AI should not be hidden where disclosure is required or material to interpretation
- Documentation should state key assumptions, limitations, and approval boundaries

7.4 Fairness and bias mitigation

AI systems can inherit and amplify biases from data, prompts, rules, or deployment context. Lupin requires teams to identify bias risks relevant to the use case and implement reasonable mitigations before operational use. Bias review is especially important for people-related, medical, safety related, and prioritization use cases. Bias mitigation may include dataset review, prompt design controls, calibration checks, escalation rules, and periodic monitoring for unexpected skew or adverse patterns.

8. Acceptable Use of AI Tools

8.1 Approved tools and approved uses

Employees may use AI tools for company work only when the tool is approved for the specific purpose in question. Approval of a tool does not automatically authorize every possible use of that tool. Conditions of use may differ by data type, role, geography, business unit, or deployment model.

The IT Department and Office of AI Governance shall maintain an Approved AI Tools Register that identifies the approved platform, scope of approval, major limitations, and any prohibited or restricted uses.

- Use only approved tools for business purposes
- Confirm the intended use is within the approved scope, not merely that the tool exists on an approved list
- Follow any platform-specific instructions, banners, and technical restrictions

8.2 User responsibilities

Every user is responsible for verifying AI outputs before relying on them or sharing them. Users must apply professional judgment, confirm facts, check references, and ensure that the content is appropriate for the audience and purpose. AI-generated text that appears polished may still be incorrect, outdated, biased, or misleading.

Users are also responsible for the prompts and data they provide. Treat all inputs to external or vendor hosted AI systems as potentially discoverable or viewable beyond the immediate session unless the approved conditions explicitly state otherwise.

- Verify factual statements before use
- Do not assume citations or references are real without checking them
- Use business judgment and escalate uncertain outputs rather than silently accepting them

8.3 Permitted examples

Permitted uses typically include low-risk drafting, summarization of non-sensitive material, brainstorming, translation of general text, formatting assistance, and internal coding support within approved environments and subject to source-code protections. Some internal analytics and retrieval use cases are also permitted when the approved use and relevant controls are documented.

8.4 Prohibited activities

Certain activities are prohibited regardless of convenience or productivity benefits. These include entering confidential information into unapproved public tools, processing personal data without authorization, making employment decisions using AI, bypassing required human review, or

integrating AI systems into internal applications without the necessary approvals and security controls.

Violations may result in removal of access, disciplinary action, investigation, and other corrective action as appropriate.

- Do not upload confidential business information into public AI tools
- Do not upload personal, patient, or employee data into unapproved AI systems
- Do not represent AI-generated work as wholly original human work where that would be misleading or non-compliant
- Do not use AI to make final employment or disciplinary decisions
- Do not connect AI tools to enterprise systems or IoT devices without written authorization

8.5 Use in regulated content and external communications

AI may only be used in regulated documentation, external communications, or formal deliverables under conditions approved for that context. Additional review, traceability, disclosure, and validation may be required. The mere fact that a document is ultimately reviewed by a human does not make all intermediate uses acceptable. Where publication or external distribution is planned, the appropriate approver must review both the content and the fact that AI was used in its preparation when relevant.

- External-facing material may require explicit disclosure or approval
- Regulated documentation may require documented evidence of review and control
- Do not use AI to shortcut mandatory process steps

9. AI Agent Governance Standard

9.1 Why agent governance is separate

Agentic AI presents a different class of risk than simple assistive AI because the system may plan across multiple steps, call tools, invoke APIs, modify data, trigger workflows, or act at machine speed. This introduces operational, security, and accountability concerns that are not sufficiently addressed by general acceptable-use rules alone.

Lupin therefore applies additional governance requirements to AI agents and any automation that can execute actions beyond passive content generation.

- Agentic capability increases the need for clear boundaries, observability, and containment.
- Approval must consider both the model and the actions it can take through connected tools.

9.2 Autonomy levels

AI systems shall be assigned an autonomy level during review. Level 0 systems provide suggestions only and do not execute actions. Level 1 systems can prepare or queue actions but require human approval before execution. Level 2 systems may execute approved actions while a human remains on the loop for active supervision and intervention. Level 3 systems operate autonomously within explicitly defined boundaries and require the highest level of review and control.

Autonomy level is not a technical label only; it is a governance decision about the system's permitted behavior in production.

- Higher autonomy requires tighter controls, narrower boundaries, and stronger monitoring
- A system may be technically capable of Level 3 behavior but only approved for Level 1 operation

Level	Approval Required
Level 0	System Owner
Level 1	Business Owner + IT
Level 2	AI Governance Council
Level 3	Council + CISO + Quality + Legal

9.3 Identity, access, and boundary control

Every deployed AI agent must have a unique identity and authorization model. Shared accounts, broad administrator privileges, or unmanaged secrets are not permitted. Agents must be granted only the minimum permission necessary for their approved purpose.

Operational boundaries must define which systems, datasets, tools, actions, and decision thresholds the agent may access or invoke. Boundaries should also specify what the agent must not do, when it must escalate to a human, and what conditions cause automatic suspension.

- Use least privilege and service identities
- Document allowed tools, APIs, datasets, and actions
- Prevent agents from expanding their own permissions or chaining into unapproved systems

9.4 Kill switch, rollback, and observability

All autonomous or semi-autonomous agents must support immediate suspension through a kill-switch or equivalent containment control. This is necessary for security events, faulty behavior, unexpected output patterns, integration failures, or governance concerns. Systems must also support rollback or disabling of relevant workflows where practical.

Agent operations must be observable. Logs should capture prompts or triggering conditions where appropriate, actions taken, systems called, outputs produced, exceptions raised, and material decision points. Logging must be sufficient for investigation, audit, and post-incident review.

- Kill-switch capability is mandatory for material autonomous action.
- Monitoring must support real-time or near-real-time detection of abnormal behavior.
- Logging must be retained in accordance with policy and legal requirements.

9.5 Human escalation and guardrails

Agent workflows must define when the system pauses for approval, hands over to a human, or refuses to proceed. Escalation rules should be explicit for ambiguous cases, high-impact outcomes, low-confidence outputs, conflicting sources, and any access to sensitive data or regulated records.

Guardrails may include policy filters, allowlists, deny lists, confidence thresholds, transaction limits, environment segregation, or dual-control approval patterns.

- Humans must remain able to interrupt, override, and investigate the agent
- Guardrails should be documented and tested before go-live

9.6 Agent tool governance

Agents must follow:

- Approved Tool Registry (allowed APIs only)
- MCP / API controls must be documented
- Prompt injection protections required
- Sandbox vs production separation mandatory
- Code execution by agents → restricted & logged

10. Confidentiality and Data Protection

10.1 Core rule for confidential information

To remove ambiguity from the original guideline, this policy states the rule clearly: confidential company information must not be entered into public or unapproved AI tools. Confidential information may only be used in an Approved Enterprise AI Platform and only within the conditions of that platform's approval, the applicable business process, and the data owner's requirements.

Even within an approved enterprise environment, users must apply data minimization and avoid unnecessary disclosure. Approval of the platform does not remove the need to assess whether the specific dataset or document is authorized for that use.

- Default assumption: public AI tools are not approved for confidential inputs
- Enterprise approval is conditional, not blanket permission
- Use the minimum necessary information to accomplish the task

10.2 Personal data and sensitive data

Personal data, employee information, patient data, and any health-related or otherwise sensitive data must only be processed in compliance with applicable privacy laws and internal privacy requirements. Teams must identify the lawful basis, confirm the purpose, apply data minimization, and ensure that appropriate contractual, security, and retention controls are in place.

Where data can be anonymized, pseudonymized, masked, or otherwise de-identified without undermining the use case, those techniques should be considered before using the data in AI workflows.

- Do not use personal data in AI systems without authorization and a valid business purpose
- Consult Data Privacy for cross-border, employee, or patient-related use cases
- Sensitive data requires heightened review and protection

10.3 Intellectual property, copyright, and third-party restrictions

Users must respect intellectual property rights and contractual restrictions when using AI. This applies both to what is entered into AI systems and to what is produced by them. Licensed third-party content, proprietary code, customer information, and partner materials may carry restrictions that make AI ingestion or reuse impermissible.

AI-generated outputs must also be reviewed for possible infringement, inappropriate similarity to third-party material, and implications for patentability, ownership, or publication strategy.

- Do not assume that vendor terms automatically transfer ownership of output in all circumstances
- Consult Legal when publication, patents, licensed material, or external distribution are involved

10.4 Records, retention, and sharing

Prompts, outputs, logs, and derivative work products may themselves become records depending on context. Teams must manage AI-related artifacts in accordance with applicable retention schedules, quality requirements, litigation holds, and information-classification rules.

Sharing of AI outputs should follow the same classification and need-to-know principles that apply to other company documents.

- Store material records in approved repositories
- Do not rely on consumer AI chat history as the system of record for business content
- Apply classification labels where required

10.5 Data Residency Rules

- No patient data / employee data / GxP data may leave approved infrastructure
- Cross-border AI usage requires:
 - DPIA
 - Privacy Office approval
- Approved regions must be defined for:
 - India
 - EU
 - US

11. AI System Lifecycle Governance

11.1 Intake and use-case approval

Every non-trivial AI use case must begin with an intake and approval process. At minimum, the sponsoring team must document the business objective, intended use, user population, data categories, system architecture, proposed autonomy level, and preliminary risk tier. This intake supports early review by the relevant control functions and reduces late-stage rework or deployment delays.

Use-case approval is not a one-time check-box. Significant changes to model behavior, integrations, user scope, or data use may require renewed review.

- Document intended use before procurement or build starts
- Involve control functions early for higher-risk or regulated use cases
- Do not move to production without the required approvals

11.2 Development and design controls

AI solutions must be designed with security, privacy, reliability, and maintainability in mind. Development teams should use secure coding practices, maintain version control, document architecture and dependencies, and ensure that datasets and prompts are appropriate for the intended use. Training data lineage and configuration choices should be traceable.

Teams should also define fallback modes, failure handling, and how the system behaves when confidence is low or when dependencies are unavailable.

- Use trusted datasets and document their source and limitations
- Apply peer review and quality checks during development
- Design for controlled failure, not only for ideal behavior

11.3 Validation and testing

Before deployment, AI systems must undergo testing proportionate to risk. Testing may include functional testing, security testing, bias evaluation, prompt robustness testing, data-quality checks, performance benchmarking, and user acceptance testing. In regulated contexts, additional validation requirements may apply under quality and compliance procedures.

Validation evidence must be retained and linked to the approved intended use. A model that performs well in a general benchmark is not automatically validated for a company-specific regulated use.

- Test the system under realistic operating conditions
- Document limitations, known failure modes, and acceptance criteria
- High-risk uses require stronger validation evidence and formal sign-off

11.4 Deployment, monitoring, and change control

Deployment must occur through controlled infrastructure and approved release processes. Access, secrets, API integrations, and logging must be configured before go-live. Monitoring must be designed to detect drift, degraded performance, abnormal behavior, security events, and policy breaches. Systems should also have review triggers for retraining, recalibration, retirement, or rollback.

Changes to the model, prompts, datasets, integrations, or operating boundaries should follow documented change-control practices proportionate to risk.

- Use approved deployment pipelines and environments
- Define go-live criteria and post-deployment monitoring responsibilities
- Retire or suspend systems that no longer meet performance or control expectations

12. AI Security Practices

12.1 Security governance and baseline controls

AI systems must comply with enterprise cybersecurity requirements and any AI-specific security standards adopted by the company. Security must address the full stack: data, model, application, infrastructure, APIs, identities, secrets, integrations, and monitoring. Baseline controls are mandatory even for internal or pilot deployments because poorly governed pilots can still expose data or create attack surfaces.

Information Security may define additional controls based on data sensitivity, exposure, and autonomy level.

- Apply security-by-design from concept through retirement

- Treat AI components and connected tools as part of the enterprise attack surface
- Do not exempt pilots or experiments from baseline security requirements

12.2 Data, model, and infrastructure security

AI-related data must be protected in transit and at rest using approved cryptographic controls. Training and inference environments must be secured against unauthorized access, tampering, leakage, and misuse. Model artifacts, embeddings, prompts, and configuration files should be treated as sensitive assets where appropriate.

- Where containerization, sandboxing, or environment isolation is practical, it should be used to reduce the blast radius of compromise or malfunction.
- Use encryption, access controls, and approved key-management practices
- Secure datasets, model files, prompt templates, and vector stores as applicable
- Patch and maintain the supporting infrastructure

12.3 Identity, access management, and API security

Access to AI systems must be governed through role-based access controls and the principle of least privilege. Administrative actions should be limited, monitored, and where possible provided through just-intime access patterns. APIs used by AI systems must enforce authentication, authorization, rate limiting, and logging.

Hard-coded credentials, unmanaged tokens, and overly broad service permissions are not permitted.

- Require MFA for administrative or privileged access
- Review access rights periodically and revoke when no longer needed
- Use centralized gateways or equivalent controls for sensitive APIs

12.4 Logging, monitoring, and resilience

AI systems must generate logs sufficient to support operations, security monitoring, incident investigation, and audit. Monitoring should detect not only infrastructure events but also unusual model behavior, agent activity, or output patterns that may indicate malfunction, misuse, or attack. Teams should define alert thresholds and escalation routes aligned to the system's criticality.

Resilience measures may include backup procedures, rollback paths, dependency monitoring, rate controls, and graceful degradation when connected services fail.

- Protect logs from tampering and unauthorized access
- Retain logs according to policy and business need
- Use anomaly detection or equivalent monitoring where justified by risk

13. AI Incident Response and Reporting

13.1 What constitutes an AI incident

An AI incident is any event involving an AI system that creates or could create harm, policy breach, data exposure, security compromise, inaccurate or biased outputs in a material context, unexpected autonomous behavior, operational disruption, or regulatory concern. Not every error is a major incident, but all meaningful anomalies should be assessed and triaged.

Examples include disclosure of confidential information through an AI prompt, an agent taking an unauthorized action, a model generating materially misleading regulated content, or a third-party AI service failing in a way that affects company operations.

- Incidents include security, privacy, quality, and operational failures
- Near misses should also be reported where they could reveal a control weakness

13.2 Reporting obligations and immediate actions

Employees must report suspected AI-related incidents immediately through the defined channels, including information.security@Lupin.com and privacy@Lupin.com, where relevant. Delay can increase harm, particularly when autonomous systems, sensitive data, or regulated workflows are involved.

Initial actions may include suspending the system, stopping workflow, preserving logs and evidence, notifying the owner, and escalating to the appropriate incident-management teams.

13.3 Investigation, containment, and recovery

AI incidents should be investigated by the appropriate cross-functional team based on the nature of the event. Information Security, Data Privacy, Legal & Compliance, Quality, and the system owner may all need to participate. The investigation should establish scope, root cause, affected data or processes, containment steps, remediation actions, and communication requirements. Recovery should include both technical corrections and governance corrections, such as revised approvals, additional monitoring, training, or updated restrictions.

14. Model Documentation and Audit Requirements

14.1 Documentation expectations

Every material AI system must have documentation that explains what the system is for, how it works at an appropriate level, what data it relies on, what controls apply, what testing was performed, and what limitations users must understand. The depth of documentation should be proportionate to risk, but even lower-risk internal systems need enough documentation to support owner continuity and operational clarity.

- Document intended use, owner, business sponsor, autonomy level, and risk tier.
- Describe data sources, assumptions, dependencies, and limitations.
- Retain validation evidence and approval records.

14.2 Model cards, inventories, and traceability

Where appropriate, teams should maintain model cards or equivalent summaries that describe performance, known limitations, ethical considerations, and approved use boundaries in a concise format. Lupin must also maintain a central AI inventory so that governance functions can identify what is deployed, who owns it, and which systems are high risk.

Traceability should link the inventory record, risk assessment, approval, technical documentation, and monitoring or incident records.

14.3 Audit readiness and evidence retention

AI systems used in regulated or sensitive contexts must be supportable during internal review, audit, or regulatory inspection. Teams should be able to produce current documentation, evidence of review and approval, monitoring records, incident history, and relevant change-control history within a reasonable timeframe. Evidence must be stored in approved repositories rather than scattered across informal chats, personal drives, or unmanaged tool histories.

15. Vendor and Third-Party AI Governance

15.1 Third-party risk assessment

External AI vendors can accelerate deployment but also introduce concentration risk, security concerns, unclear data practices, and contractual limitations. Before an external AI service is adopted, Procurement and the relevant control functions must assess the vendor's security posture, privacy controls, reliability, operating model, data-use terms, and support for audit or investigation.

Vendor marketing claims should not be accepted as sufficient proof of control effectiveness.

- Perform due diligence before purchase or onboarding
- Consider whether the vendor uses customer data for training or product improvement
- Assess concentration and exit risks for critical services

15.2 Contractual and operational requirements

Contracts with AI vendors should include clear obligations on security, privacy, confidentiality, permitted data use, incident notification, audit rights where appropriate, subcontracting transparency, service levels, and support for legal or regulatory requirements. Operational teams must also define how the vendor will be monitored after go-live.

Where a vendor platform is approved, the approval should specify what data types and use cases are in scope and what remains prohibited.

16. Complaints and Regulatory Escalation

16.1 Internal complaints and concerns

Employees and other affected people may raise concerns regarding the misuse of AI, privacy violations, inaccurate outputs, unethical behavior, inappropriate automation, or security issues. Concerns should be handled promptly and without retaliation. Users are encouraged to make use of established channels to report concerns.

- Managers should not suppress or informally dismiss AI-related concerns.
- Concerns should be triaged to the right control function.

16.2 External or regulatory escalation

Where law, contract, or regulatory obligation requires notification or escalation, the company will manage the process through the appropriate control functions. Individuals may also have statutory rights to complain to supervisory authorities in certain jurisdictions. Any regulatory communication related to AI should be coordinated through Legal, Compliance, Privacy, Quality, or other designated functions as applicable.

No employee should make unilateral external statements on behalf of Lupin regarding an AI incident or compliance matter unless authorized to do so.

17. Policy Amendments and Exceptions

17.1 Amendments and continuous improvement

Lupin reserves the right to amend this policy in response to legal developments, technology changes, incidents, audit findings, control improvements, and shifts in business strategy. However, amendments must follow formal review and approval rather than ad hoc practice changes. Continuous improvement should be evidence-based. Policy updates should be documented, approved, and communicated to the applicable users. The training and supporting standards should be updated when policy changes materially affect users.

Committing to deploying AI responsibly, covering own and third party as applicable, by minimizing its energy, water and carbon footprint through efficient and sustainable technologies.

17.2 Exceptions

Any exception to this policy must be documented, justified, risk-assessed, approved by the appropriate functions, time-bound, and subject to review. Requests for exception should clearly state what requirement is being waived, why the exception is necessary, what compensating controls will apply, who owns the risk, and the date by which the exception will expire or be re-reviewed.

17.3 References and related documents

This policy should be read alongside relevant corporate policies and standards, including Information Security, Acceptable Use, Data Privacy, Records Management, Quality, Procurement, Third-Party Risk Management, and any future AI standards, templates, or playbooks issued by the company. Related procedures may provide implementation details beyond what is practical to include in a single policy document. The supporting standards may include AI use-case intake templates, model-card templates, agent standards, risk-assessment templates, and incident playbooks.

18. Contacts and Review Route

Questions regarding interpretation of this policy should be directed to the Office of AI Governance in coordination with Information Security, Data Privacy, Legal & Compliance, Quality & Regulatory Affairs, or the relevant business owner as appropriate.

Suspected incidents or urgent concerns should be reported immediately using the established company channels, including information.security@Lupin.com and privacy@Lupin.com, where relevant.

- Policy owner: Office of AI Governance
- Security concerns: information.security@Lupin.com
- Privacy concerns: privacy@Lupin.com
- Approval route: AI Governance Council, with participation from the relevant control functions for in-scope matters

19. Roles and Responsibilities

Role	Responsibilities	Contact Information
AI Governance Team	- Review and approve AI use cases, tool use cases and applications for uses not currently appearing in Annexure 1. - Develop and manage the AI policy - Answer any policy-related questions escalated from managers	<i>Executive Directors, Global CFO, Global CIO</i>
Compliance Team	- Manage ethics and integrity misconduct investigations resulting from the use of AI - Support IT Security incident response for investigations that involve data leaks or breaches resulting from the use of AI.	Compliance Lead, Legal Lead
Information Security Group	- Implement appropriate controls to secure AI access and use in compliance with this policy - Manage access to AI if tools require registration and are owned at the enterprise level - Conduct security risk assessments for AI applications and use cases - Monitor all use of AI	CISO, GRC Lead
Data Privacy Office	Conduct privacy risk assessment for AI applications and use cases	CISO, GRC Lead
All Managers	- Conduct review of specific uses of AI that have been escalated to you for review. - Assist employees with escalating questions on AI to subject matter experts. Escalate questions on policy to AI team when necessary. - Assure that AI-generated content is appropriately labeled as such in all materials in which it appears	-
Employees	- Review all AI output for bias, accuracy, and appropriateness. - Review policies and procedures on an annual basis	-