Artificial Intelligence Governance Policy

Palomar Holdings, Inc. and each of its subsidiaries (collectively, “Palomar”) is committed to the responsible, ethical, and transparent use of Artificial Intelligence Systems (“AIS”) across its insurance operations. This Artificial Intelligence (“AI”) Governance Policy (the “Policy”) aligns with applicable legal and regulatory expectations and is informed by recognized guidance and frameworks, including National Association of Insurance Commissioners (“NAIC”) guidance, including the NAIC Model Bulletin: Use of Artificial Intelligence Systems by Insurers (the “Model Bulletin”), the NIST AI Risk Management Framework, and ISO/IEC 42001, to establish robust governance, accountability, and oversight. The Policy defines clear roles, risk-based processes, and controls to protect consumers, support regulatory compliance, and maintain trust in Palomar’s use of AI technologies, with Board oversight through the Enterprise Risk Management (“ERM”) Committee and day-to-day implementation by senior management.

1. Introduction & Purpose

Mission Statement:
Palomar’s mission is to leverage AI responsibly to enhance insurance services, improve operational efficiency, and deliver value to our customers, while upholding the highest standards of ethical conduct and regulatory compliance.

Regulatory Context:

The Policy is structured to meet the requirements of the NAIC AI Model Bulletin (adopted December 2023), which mandates that insurers implement comprehensive AI governance programs addressing Board and senior management accountability, risk management, transparency, fairness, and ongoing monitoring. Consistent with the Model Bulletin, this Policy is intended to serve as a core description of Palomar’s AIS Program and applies to AI Systems used across the insurance lifecycle, including predictive models used in underwriting, rating/pricing, and claims administration. The Policy also draws on best practices from the NIST AI Risk Management Framework and ISO 42001.

Policy Scope:
This Policy provides a high-level overview of Palomar’s enterprise-wide approach to AI governance, risk management, and oversight. Detailed operational policies and procedures are addressed in the separate AI Usage Policy.

2. Guiding AI Principles

Palomar’s use of AI is governed by the following six principles:

- **Fairness:** AIS must be designed and operated to avoid unfair discrimination and support equitable outcomes for stakeholders.
- **Transparency:** The use, purpose, and output logic of AIS must be clear and understandable to affected parties and regulators, as appropriate.
- **Accountability:** Human oversight is maintained throughout the AIS lifecycle, with clear assignment of roles and responsibilities.
- **Reliability:** AIS must be robust, accurate, and perform as intended under expected conditions.

- **Privacy:** Personal and sensitive data used by or within AIS must be protected in accordance with applicable laws and Palomar's data governance standards.
- **Consumer Protection:** AIS must be used in ways that safeguard consumer rights and prevent harm.

3. Scope & Applicability

The Policy applies to employees, contractors, and other personnel using AI in connection with Palomar business, and to AIS used by or on behalf of Palomar across insurance and corporate functions, including advanced analytical and computational technologies that make, inform, or support decisions impacting consumers, including where AI is vendor-provided, embedded in third-party systems, or used within Palomar's distribution ecosystem (e.g., producers and partners), to the extent applicable based on contractual and oversight rights.

The Policy supplements (and does not replace) Palomar's existing security, acceptable use, vendor management, and incident response controls, and it is supported by the AI Usage Policy and related governance documentation.

4. Governance & Oversight Structure

Role / Area	Responsibilities
Board / ERM Committee oversight:	Consistent with the ERM Committee Charter, the ERM Committee assists the Board in oversight of Palomar's general strategy with respect to enterprise risk management, AI, and cybersecurity; and in assessing and monitoring the Company's risk management framework employed to manage enterprise risks, AI, and cybersecurity.
Executive Accountability and Management Governance:	Palomar's senior management is responsible for implementation and monitoring of the Policy and is accountable to the Board of Directors for AI strategy and oversight.
Decision Rights, Escalation, and Accountability:	Palomar will allocate decision rights and escalation paths for AI use cases based on risk. Higher-risk AIS use cases, material issues, or policy exceptions will be escalated to appropriate senior management governance, and where appropriate, to Board-level oversight through the ERM Committee. Roles and responsibilities will be sufficiently specific to support accountability for: (i) AIS inventory maintenance; (ii) risk assessments and approvals; (iii) monitoring and issue management; and (iv) audit/examination readiness.
Program Owner (Information Security / IT):	Maintains approved tool lists, evaluates AI tools, implements monitoring controls, investigates incidents, and delivers training related to AI. Establishes and maintains standardized AIS intake and risk tiering criteria and provides second-line review for Medium/High-risk AIS (including required controls and go-live readiness), in coordination with Legal/Compliance and relevant stakeholders.
Business Owners / Department Heads:	Own the business use case (including purpose and intended decisions/workflows), ensure staff training, define human-oversight steps, implement required controls, and perform first-line monitoring

Role / Area	Responsibilities
	and issue escalation commensurate with the AIS category and risk tier (including monitoring effectiveness for Category B and Category C AIS in their area).
Legal & Compliance:	Reviews and approves customer-facing or regulator-impacting AI uses prior to launch, including Category C disclosures, regulatory implications, and customer-impact controls (as required by Palomar's Category C requirements).
Three Lines / Audit:	Internal audit (or equivalent) periodically tests adherence to this program's control requirements, including documentation, approvals, and monitoring evidence.
Review Frequency:	At least annually (and more frequently as needed) to reflect changes in applicable laws and regulations, regulatory guidance, business practices, and AI technology risks.

5. Key Definitions & AIS Inventory, Classification, and Change Control

Key Definitions:

- **Artificial Intelligence (AI)**: A branch of computer science that uses data processing systems that perform functions normally associated with human intelligence, such as reasoning, learning, and self-improvement. Machine learning is a subset of artificial intelligence.
- **AI System / AIS**: A machine-based system that can, for a given set of objectives, generate outputs such as predictions, recommendations, content (such as text, images, videos, or sounds), or other outputs influencing decisions made in real or virtual environments, and that is designed to operate with varying levels of autonomy.
- **Algorithm**: A clearly specified mathematical process for computation; a set of rules that, if followed, will give a prescribed result.
- **Predictive Model**: The mining of historic data using algorithms and/or machine learning to identify patterns and predict outcomes that can be used to make or support the making of decisions.
- **Model Drift**: The decay of a model's performance over time arising from underlying changes, including shifts in relationships between training data and deployed data.
- **Palomar AI Use Categories (A/B/C)**: Category A (individual employee assist with human-owned final output); Category B (internal, staff-facing tools); Category C (customer- or external-facing AI outputs or decisions).
- **Sensitive Data & Prohibited Inputs**: Do not enter customer or employee PII, financial account data, PHI, credentials, or claims information containing PII into AI tools—except for approved built-in AI within approved business platforms as described in the AI Usage Policy.
- **AI Usage Policy**: A separate, user-facing policy that is part of this AI governance program and is distributed to personnel within scope.
- **Executive Management**: For purposes of this Policy, "Executive Management" means the CEO/COO, CTO, Legal/CLO, and the affected department head (or their designees), as applicable to the AIS use case.

AIS Inventory (Central Register):
The COO is responsible for the creation and management of a living "AIS Register" to inventory AIS used by or on behalf of Palomar and to support consistent governance, monitoring, and reporting. The AIS Register will be updated on an event-driven basis whenever a new AI use case is approved or a material change occurs, with at least annual review (and more frequent updates as needed) to reflect changes in applicable laws and regulations, regulatory guidance, business practices, and AI technology risks. See **Appendix A** for the standard AIS Register fields.

Category Assignment (A/B/C):
All AIS must be categorized using Palomar's "who sees the output / decision impact / embedded workflow" criteria; if uncertain, classify to the higher-risk category.

Audit and Information Request Readiness:
Palomar will maintain the AIS Register to support efficient responses to audits, risk reviews, and other information requests, including (as needed) a mapping to commonly requested inventory information.

Evidence and Artifact Model:
For each Medium- and High-risk AIS (and as appropriate for other AIS), Palomar will identify expected artifacts (e.g., intake, risk assessment, approvals, vendor due diligence, testing/validation, monitoring, exceptions, minutes where applicable, incident records), including the system of record, owner, access controls, and retention requirements.

Material Change Control:
Any material change (e.g., model, data source, thresholds/rules, customer communication path, vendor feature, or other change increasing risk) triggers re-review at the same approval level as initial deployment and may require re-tiering and updated testing/monitoring.

6. Data Governance, Privacy, and Security Controls (Program Minimums)

Approved Tools Only:
No Palomar information may be entered into unapproved AI tools; only approved AI tools/services may be used for Palomar work.

Data Prohibitions (Hard Stop):
Do not input prohibited data types listed in the AI Usage Policy (including customer PII, PHI, credentials, and claims files containing PII) into any AI tool, even if the tool is approved—except as permitted for approved built-in AI within approved business platforms.

Platform-AI Exception (Controlled Use):
AI features embedded within approved core business platforms (e.g., Majesco Copilot) may be used with PII already inside that platform when the platform has been approved through Palomar's vendor/security evaluation, the AI is built-in (not a separate third party), and the use is within job scope.

Anonymization/De-identification Standard:
When AI use requires data resembling production records, redact direct identifiers and mitigate re-identification risk; contact Security@plmr.com for complex datasets.

Data Use Authorization and Contractual Restrictions:

Ensure data used with AIS is permitted for the intended purpose (including consents, purpose limitations, licensing, and contract terms) and document applicable vendor/third-party restrictions.

Note: Category (A/B/C) reflects exposure context and provides a baseline set of controls; additional requirements may apply based on the AIS risk tier (Low/Medium/High/Prohibited) determined under Section 7.

Category	Security Requirements
A	Approved tools only; prohibited data rules apply; human review of outputs required.
B	Requires IT/Security review and COO and CTO approval before implementation; documented oversight procedures; regular monitoring/auditing; staff training.
C	Requires multi-stage Security + Legal + Compliance review and Executive Management approvals (CEO, COO, CTO, Legal/CLO, and affected department head); comprehensive vendor security assessment; governance framework with monitoring, bias testing, disclosures, human appeal/override, incident response, and audits.

7. AIS Use Case Intake, Risk Tiering, and Control Scaling

Overview:
Palomar applies a risk-based approach to governing AIS use cases. While AIS may be categorized by exposure context (e.g., internal vs. customer-facing), governance expectations are calibrated using horizontal risk tiering across multiple risk dimensions.

Standardized Intake and Initial Risk Tier Assignment:
Before deployment (and upon material change), each AIS use case must have a standardized intake record and an initial risk assessment documented. Palomar assigns an initial risk tier to an AIS use case—**Low**, **Medium**, **High**, or **Prohibited**—based on criteria such as purpose, automation/decision impact, consumer impact, regulatory sensitivity, data types and handling, transparency/auditability, reliability expectations, and potential for unfair or illegal discrimination or other harm.

Tier-to-Controls Mapping (Scaled Diligence):
Low risk use cases follow streamlined review and documentation. Medium-risk use cases require enhanced review, appropriate testing/validation, and defined approvals. High-risk use cases require cross-functional review, senior approval (including Executive Management approval where applicable), robust testing/validation and monitoring, and heightened ongoing oversight. Prohibited use cases are not permitted absent an approved exception pathway (if any).

Risk Tier	Scaled Controls (Summary)
Low	Streamlined review and documentation appropriate to the use case.
Medium	Enhanced review, testing/validation appropriate to the system type, and defined approvals.

Risk Tier	Scaled Controls (Summary)
High	Formal cross-functional review, senior-level approval, robust testing/validation and monitoring, and heightened ongoing oversight/documentation.
Prohibited	Not permitted absent an approved, documented exception pathway (if any) consistent with Palomar's governance processes.

Relationship to A/B/C Categorization:
Where Palomar uses A/B/C categories to describe exposure context, the risk tier determines the level of diligence, approvals, testing, monitoring, and required evidence.

8. AIS Risk Assessment (Required for Category B & C)

Pre-Implementation Risk Assessment:
Before deployment of Category B or C, the business owner and Security/IT will document risks and controls covering at least:

- Data sensitivity and handling approach (including prohibited-data controls).
- Human oversight design (review steps, override authority, escalation).
- Vendor/model risks (training use of inputs, retention/deletion, auditability).
- Customer impact (especially for Category C: disclosures, escalation-to-human, appeal).

Risk Rating:
Rate AIS as Low/Moderate/High based on customer impact, automation level, decision materiality, data sensitivity, and regulatory exposure. Where applicable, align the rating to (or use it to inform) the risk tier assigned under Section 7.

9. Fairness, Outcomes Testing, and Human Oversight

General Standard (Risk-Based):
Palomar will implement human oversight appropriate to the AIS use case's risk tier, decision impact, and consumer/external impact. For decisions that could materially affect customers, applicants, insureds, claimants, employees, or other individuals, Palomar will maintain appropriate human oversight, escalation, and accountability mechanisms, commensurate with the use case.

Testing and Monitoring (Scaled):
For AIS influencing insurance outcomes (e.g., eligibility, pricing, tiering, claims handling, fraud flags, customer communications), Palomar will apply testing, validation, and monitoring standards appropriate to the AIS type and risk tier, including performance/reliability checks and assessments designed to identify indicators of unfair or illegal discrimination, as feasible and appropriate to the use case and regulatory expectations.

Exceptions:
Any exception to applicable governance expectations (including where automation is used in a manner that would otherwise trigger enhanced oversight) must be documented, risk assessed, time-limited where appropriate, and approved through Palomar's exception process.

10. Transparency, Notice, and Human Escalation

Internal Transparency:

- Category B users must know they are interacting with AI; outputs should be labeled in the interface and covered in training.

Customer-Facing Disclosure (Category C):

- Disclose AI interaction up front; provide an option to escalate to a human; identify AI-generated customer communications.

Human Appeal / Override:
Category C AIS must provide a documented process for human review or appeal of customer-impacting outcomes.

11. Vendor / Third-Party AIS Management

Vendor Evaluation:
All AI vendors/tools follow Palomar's vendor approval process with AI-specific criteria including model training/data use, retention/deletion, transparency/explainability, bias/fairness safeguards, auditability/logging, and human oversight capabilities.

Contractual Requirements (Category B & C):
Contracts should support Palomar's governance needs (security, audit, retention/deletion, incident notification, and restrictions on using Palomar inputs for training) because Palomar explicitly evaluates these items for AI tools. As appropriate to the engagement and risk tier, contracts should also support audit rights and/or access to audit reports by qualified entities and vendor cooperation with regulatory inquiries or investigations related to Palomar's use of the AIS.

12. Monitoring, Logging, and Auditability

Ongoing Monitoring:

- Category B: regular monitoring and auditing of tool performance (including model drift/performance decay where applicable) and staff interaction.
- Category C: continuous monitoring (including model drift/performance decay where applicable), regular auditing/compliance reviews, and Executive Management visibility for customer impact.

Logging:
Maintain sufficient logs to reconstruct prompts/inputs (subject to data minimization), outputs, user actions, overrides, version changes, and incidents, to support audit and investigation.

13. Training & Workforce Enablement

Policy Distribution and Acknowledgement:
Palomar will distribute the AI Usage Policy and relevant AI governance guidance to personnel within scope and will maintain evidence of distribution and acknowledgement, consistent with applicable training and compliance practices.

Training Delivery, Employee Acknowledgement and Completion Tracking:
Palomar will provide AI-related training appropriate to job role and risk exposure (including security awareness, approved tools expectations, prohibited data, escalation and reporting, and human

oversight responsibilities). Palomar will require employee acknowledgement of usage rules, track training completion and maintain records in the applicable system(s) of record as part of its governance evidence.

Refresh Cadence:
Training content and delivery cadence will be reviewed and refreshed as needed based on material changes in AI technology, business use, incidents, or legal/regulatory expectations.

14. Incident Response, Issues, and Regulatory Readiness

AI-Related Incidents:
Potential policy violations (e.g., prohibited data entered into AI, use of unapproved tools) must be reported to Security@plmr.com and investigated.

Incident Handling:
AI-related incidents follow Palomar incident reporting/response procedures referenced by the AI Usage Policy, including cooperation with investigations.

Whistleblower Reporting (Code of Conduct):
Concerns related to AI use, compliance, misconduct, or suspected violations of law should be reported in accordance with Palomar's Code of Conduct and Ethics, including available anonymous/confidential reporting channels and the Company's non-retaliation commitment. Nothing in this Policy limits any individual's ability to report possible violations of law or regulation to any governmental agency or entity, consistent with the Code of Conduct and Ethics.

Audit and Examination Support:
Maintain evidence artifacts for each Category B/C AIS (approvals, risk assessments, vendor due diligence, testing results, monitoring reports, training completion, incident records) to support audits and other information requests.

15. Approvals, Exceptions, and Enforcement

Approval Requirements:
Note: Category (A/B/C) provides baseline approval expectations based on exposure context; the AIS risk tier (Low/Medium/High/Prohibited) under Section 7 governs the level of review and may require additional approvals, testing, monitoring, or controls.

Category	Approval / Minimum Controls
A	Approved tools only; no extra approval.
B	IT & Security review, COO and CTO approval + department head; documented oversight procedures; monitoring/audits.
C	Requires multi-stage Security + Legal + Compliance review and Executive Management approvals (CEO/COO, CTO, Legal/CLO, and affected department head); governance framework; monitoring/audits; disclosures and human appeal.

Exceptions:
Exceptions require submission to Security@plmr.com with justification, risk assessment, duration, and approval by COO, CTO + department head; exceptions are documented, time-limited, and reviewed quarterly.

Violations:
Violations may lead to suspension of access, discipline up to termination, and possible legal action; good-faith reporting is protected.

Appendix A: AIS Register Fields

This appendix lists standard fields captured in the AIS Register.

Field	Description
System name	Name/identifier used internally for the AIS.
Vendor / model	Vendor name and model/tool name (including version where applicable).
Category (A/B/C)	Palomar AI Use Category assignment (A/B/C) per the classification criteria in Section 5.
Business purpose	Purpose and intended use case(s).
Business owner	Accountable business leader/department for the use case.
Users	User population (roles, teams) and intended access scope.
Third-party / embedded	Whether vendor-provided, embedded in a third-party system, or otherwise operated by/on behalf of Palomar.
Deployment context	Where/how it is deployed (e.g., internal tool, customer-facing workflow, back-office process).
Data inputs	Key data types/sources used by the AIS (e.g., internal datasets, third-party data, user-provided inputs).
Outputs/decisions influenced	Primary outputs produced by the AIS and the decisions/workflows those outputs support or influence.
Integration points	Systems, tools, or workflows the AIS integrates with (upstream inputs and downstream consumers).
Risk tier	Risk tier per Section 7 (Low/Medium/High/Prohibited) and date assigned/updated.
Approval dates	Dates of required reviews/approvals (e.g., IT/Security, department head, Executive Management, Legal/Compliance as applicable).
Testing evidence	Reference/location for testing/validation artifacts appropriate to the AIS risk and use (e.g., performance checks, reasonableness tests, bias testing where applicable).
Monitoring metrics	Key monitoring measures and where monitoring evidence is maintained (e.g., performance metrics, drift/quality checks, audit logs).
Incident history	Summary or reference to AIS-related issues/incidents, including dates and resolution status (or “none to date”).