CORPORATE POLICY

AI Use and Governance

Responsible use of artificial intelligence at Coastal Construction

Objective

The purpose of this policy is to enable responsible, secure, and productive use of Artificial Intelligence (AI) at Coastal Construction while protecting company, client, employee, and project information and preserving human accountability.

Coastal supports practical adoption of AI where it improves productivity, quality, decision support, communication, and innovation. AI is a tool used by Coastal employees; it does not replace employee judgment, professional responsibility, approval authority, or established business controls.

Scope

This policy applies to all Coastal Construction employees and contractors who use, access, configure, administer, or manage AI for Coastal business or through a Coastal-managed identity, device, application, network, or technology environment.

It applies to standalone AI platforms, AI embedded in approved business applications, AI-enabled integrations and connectors, and agentic or autonomous AI capabilities.

Policy Principles

- **Approved environments.** Coastal business information may be processed only in AI environments authorized under this policy and the related AI Standards.
- **Human accountability.** Employees remain responsible for AI-assisted work and for reviewing outputs before operational or external use.
- **Permission boundaries.** AI may use information the employee is authorized to access, but AI may not be used to bypass, expand, or circumvent application or data permissions.
- **Assistance is different from action.** AI may assist with drafting, analysis, research, summarization, and recommendations. Autonomous or agentic action requires separate authorization.
- **Existing controls still apply.** Use of AI does not remove any review, approval, confidentiality, records, safety, contractual, or other requirement that would apply if the work were performed without AI.

Simple rule

If you could perform the work yourself with the same data and authority, AI may generally assist you in an approved environment. If the AI will act on your behalf, cross system boundaries, change records, or communicate automatically, additional authorization is required.

Key Definitions

AI Tool / System: Any software capability that uses machine learning or generative AI to create, analyze, summarize, classify, recommend, automate, or act on information.

Approved AI Environment: An AI platform, account, or AI-enabled application authorized by Coastal and operated with company-approved security, identity, contractual, and data protection controls.

Embedded AI: AI functionality natively included within a Coastal-approved business application and operating within that application's authorized environment and access controls.

AI Assistance: AI use that creates, summarizes, analyzes, recommends, or prepares work for a human to review and decide whether to use or submit.

AI Action: AI use that changes a system, sends a communication, creates a binding commitment, modifies or deletes records, initiates a transaction, or otherwise performs an operational step on behalf of a user.

Agentic AI: AI that can plan and execute one or more actions with limited direct human intervention, including "work," "cowork," computer-use, multi-agent, autonomous workflow, or similar capabilities.

Third-Party Connector: A plug-in, extension, integration, API connection, MCP server, external model, add-on, or other mechanism that allows an AI service to access or act in another application or data source.

Restricted / Highly Sensitive Data: Information subject to heightened legal, contractual, security, privacy, or operational controls as defined in this policy and the AI Standards.

AI Incident: An event involving unintended disclosure, unauthorized AI use or action, security or privacy exposure, inappropriate permissions, incorrect autonomous action, material reliance on erroneous AI output, or other misuse or loss involving AI.

Approved AI Access Model

Coastal uses a tiered approval model rather than a permanent list of product names in this policy. Current platforms and access requirements are maintained in the AI Standards and may be updated by the Vice President of Emerging Technology without revising this policy, provided the policy requirements remain unchanged.

- **Generally available enterprise AI.** Approved company-managed AI environments may be made broadly available for normal business use. Currently authorized: ChatGPT Enterprise, Egnyte AI, Procore AI, Microsoft Office 365 CoPilot (Free version).
- **Approved by request.** Claude Team, Microsoft 365 Copilot (paid version), and other separately licensed or role-specific AI platforms approved by Emerging Technology, but provisioned only when requested and approved based on business need, licensing, security, or role.
- **Embedded AI in approved applications.** Native AI functionality within a Coastal-approved business application is considered approved for use within that application, subject to this policy and the AI Standards.
- **Agentic or autonomous capability.** Any capability that acts on behalf of a user requires separate authorization from the Vice President of Emerging Technology before use, even when the underlying application or AI platform is otherwise approved.

Data Handling and Security

Use of Company Information in Approved AI

Employees may generally use company, client, project, financial, contractual, and other business information in an approved enterprise AI environment when all of the following are true:

- The employee is authorized to access the information for a legitimate Coastal business purpose.
- The AI environment is approved for the intended use and data type.
- No law, contract, client requirement, project requirement, or data classification specifically prohibits or restricts AI processing.
- The employee follows the same need-to-know, confidentiality, retention, and handling requirements that apply outside AI.

Access to information does not, by itself, mean every AI use is permitted. However, employees are not expected to make technical data-classification judgments beyond what is reasonably known or communicated to them. If information is not marked or known to be Restricted / Highly Sensitive and does not fall into a restricted category below, it may generally be used in an approved enterprise AI environment.

Restricted / Highly Sensitive Data

The following information must not be entered into or processed by AI unless the specific use has been expressly authorized by the Vice President of Emerging Technology and any other required data owner, legal, security, or compliance authority:

- Passwords, MFA codes, authentication secrets, API keys, private keys, access tokens, or other credentials.
- Full Social Security numbers, bank account or routing numbers, payment-card data, unredacted government identification, or similarly sensitive identity or financial data.
- Medical or health information, background-check records, employee investigation or disciplinary records, or other highly sensitive personnel records.
- Attorney-client privileged or litigation material where counsel has restricted processing or external systems.
- Controlled Unclassified Information (CUI), export-controlled information, government-sensitive information, or other regulated information where the applicable requirement restricts AI processing.
- Information subject to a client, owner, project, NDA, contract, or regulatory requirement that specifically prohibits AI use or requires special handling.
- Any information explicitly labeled Restricted, Highly Sensitive, No AI, or equivalent under Coastal or client requirements.

When unsure

Do not guess. If the data is visibly restricted, clearly falls into one of the categories above, or you know the client or contract limits AI use, do not upload it. Ask the data owner, IT, Legal, or Emerging Technology for guidance.

Permission and Access Boundaries

AI must operate within the same authorization boundaries as the employee. Employees may not use AI to retrieve, infer, expose, summarize, or act on information they would not otherwise be permitted to access. AI may not be used to bypass application permissions, security controls, project restrictions, or role-based access.

Third-Party Connectors and Integrations

Employees may use only connectors, plug-ins, extensions, integrations, external models, MCP servers, or similar AI connections that have been explicitly authorized by Coastal. Employees may not build, enable, authorize, or connect third-party integrations themselves for Coastal data or systems. Requests for a connector or integration must be submitted to IT; approved integrations will be configured or implemented by IT and/or Emerging Technology.

Acceptable Use and AI Actions

Permitted AI Assistance

Subject to this policy, approved AI may be used to assist employees with legitimate Coastal business activities, including:

- Drafting and revising internal or external communications for human review.
- Creating email drafts. The employee must review and manually send the message.
- Summarizing meetings, project updates, specifications, contracts, technical documentation, correspondence, or other authorized information.
- Research, ideation, proposal development, planning, analysis, comparison, and decision support.
- Creating checklists, templates, reports, presentations, spreadsheets, code, workflows, and other work product for human review.
- Using native AI features embedded in approved Coastal business applications, provided the use remains within approved permissions and does not become an unauthorized agentic action.

AI Actions Requiring Additional Authorization

Prior written authorization from the Vice President of Emerging Technology is required before enabling or using AI that operates agentically or autonomously, including AI that can independently:

- Send email, Teams messages, text messages, client communications, or other outbound communications.
- Create, modify, move, delete, approve, submit, or transmit records in business systems without a human performing the final action.
- Initiate purchases, payments, contracts, commitments, approvals, change orders, financial transactions, or other consequential actions.
- Operate across multiple systems, applications, repositories, or user accounts to complete work.
- Run persistently or on a schedule to perform actions without direct human review at the point of execution.

At this time, AI is not authorized to automatically send email messages on behalf of Coastal employees. AI may prepare or populate an email draft for employee review, but the employee must perform the send action.

Prohibited Uses

- Use of personal or consumer AI accounts for Coastal confidential or non-public business information.
- Use of unapproved AI platforms, models, connectors, integrations, or agents for Coastal business information.
- Using AI to circumvent security, access controls, records requirements, or business approvals.
- Representing AI output as verified fact when the user has not performed appropriate review.
- Using AI to impersonate another person, fabricate approvals, falsify records, conceal material facts, or perform unlawful, unsafe, discriminatory, deceptive, or unethical activity.

Human Oversight and External Use

Employee Accountability

Use of AI does not transfer responsibility for an employee's work product to the AI system, vendor, or technology provider. The employee using AI remains responsible for the accuracy, completeness, appropriateness, confidentiality, and business consequences of the resulting work.

Employees must review AI-assisted work at a level appropriate to the risk and consequence of the use. AI output should be treated as a draft, recommendation, or analytical aid until the responsible employee has reviewed it.

Review Standard for Communications and Deliverables

AI does not create a new approval hierarchy. The same review and approval rules that would apply if the employee created the work without AI continue to apply when AI assists with the work.

Use	Minimum review expectation
Routine internal work	Employee reviews before relying on or distributing it.
Routine external communication	Employee reviews the content and may send it if the same communication would not otherwise require additional approval.
Work normally requiring supervisor, PM, PX, VP, Legal, Finance, Safety, HR, or other approval	The existing review and approval remains required; AI does not replace it.
Contractual, legal, safety, financial, employment, regulatory, or other high-impact work	Appropriate subject matter and authorized business review is required before use or commitment.

Employees should apply the same professional judgment to AI-assisted external communications that they would to personally drafted communications. If they would normally be authorized to write, review, and send the communication themselves, they may do so with AI assistance after reviewing the final content.

Intellectual Property, Client, and Third-Party Information

- Do not use AI in a manner that violates copyright, license restrictions, NDAs, client agreements, subcontractor confidentiality obligations, or other third-party rights.
- Do not assume that possession of a document eliminates contractual restrictions on how that document may be processed.
- When a client, owner, project, or contract establishes AI-specific requirements, those requirements govern the applicable work and must be followed in addition to this policy.
- AI-generated content that incorporates third-party material must be reviewed for appropriate attribution, licensing, confidentiality, and permitted use when those issues are relevant.

Records and Transparency

AI-assisted work is subject to the same records retention, project documentation, legal hold, and business record requirements as comparable work produced without AI. Employees are not required to label routine AI-assisted content solely because AI was used unless a client, contract, law, project requirement, or Coastal standard requires disclosure.

Roles and Responsibilities

Vice President of Emerging Technology - Policy Owner

- Owns and maintains this policy and the supporting AI Standards.
- Approves AI platforms, material AI capabilities, agentic use cases, and exceptions.
- Coordinates with IT Security, Legal, People & Culture, Finance, Operations, and other stakeholders on higher-risk uses.
- Oversees AI governance reviews, adoption, education, incident coordination, and material policy updates.

IT / IT Security

- Implements identity, access, logging, security, data protection, and other technical controls for approved AI environments.
- Reviews, builds, enables, or administers approved third-party connectors and integrations.
- Supports AI incident containment, investigation, remediation, and access changes.
- Maintains safeguards intended to limit unapproved or shadow AI use where technically practical.

Department Leaders and Subject Matter Experts

- Apply existing business review and approval requirements to AI-assisted work.
- Ensure employees understand applicable project, client, contractual, safety, regulatory, and confidentiality constraints.
- Provide subject matter review when AI-assisted work presents material risk or consequence.

All Employees and Contractors

- Use only AI environments, capabilities, and integrations authorized under this policy and the AI Standards.
- Protect Coastal, client, employee, project, and third-party information.
- Review AI outputs before relying on, submitting, or distributing them.
- Do not delegate authority to AI that the employee does not personally have.
- Report suspected AI incidents, data exposure, misuse, or unauthorized AI actions promptly to IT and Emerging Technology.

Exceptions

Exceptions to this policy must be documented and approved by the Vice President of Emerging Technology. IT Security, Legal, the applicable business owner, or other stakeholders may be required to participate depending on the nature of the request.

An exception request should identify the technology or use case, business purpose, data involved, users or departments affected, requested duration, material risks, proposed controls, and responsible business owner. Approval may be limited by user, project, system, data type, action, duration, transaction threshold, or other condition.

AI Incidents and Reporting

Employees must promptly report suspected AI incidents, including unintended disclosure, use of an unapproved tool or connector, unauthorized agentic activity, security or privacy exposure, incorrect automated changes, unexpected system behavior, or material reliance on erroneous AI output that may affect Coastal, a client, a project, an employee, or a third party.

Reports should be made to IT Help Desk and the Vice President of Emerging Technology. Employees should not conceal, delete, or attempt to independently remediate evidence of a material AI incident unless directed by IT or the incident response team.

Enforcement and Review

Compliance with this policy is mandatory. Violations may result in suspension of AI access, retraining, corrective action, disciplinary action, or other measures consistent with Coastal policies and applicable law.

This policy will be reviewed at least annually and may be revised sooner following material changes in technology, regulation, contractual requirements, security risk, or Coastal business practices. The AI Standards may be updated more frequently by the Policy Owner to reflect approved technologies and operating requirements.

Document History and Version Control

Date	Version	Author	Description
12/2/2025	0.0	PMcCabe / MBandin	Initial draft
8/18/2026	0.1	Emerging Technology	Governance model rewrite; product-independent approvals, data handling, embedded AI, agentic AI, connectors, risk-based review, and exceptions