

## CORPORATE POLICY

# AI Standards

*Operational requirements supporting the AI Use and Governance Policy*

These standards define current approved AI access categories, data handling rules, action boundaries, connector requirements, and request processes. They may be updated by the Vice President of Emerging Technology as technology and business needs change.

## 1. Approved AI Access Categories

| Category                             | Current standard                                                                                                                                                                                                                                                |
|--------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Generally available enterprise AI    | ChatGPT Enterprise, Egnyte AI, Procore AI, Microsoft 365 CoPilot (free version), accessed through the Coastal-managed enterprise environment.                                                                                                                   |
| Approved by request                  | Claude Team, Microsoft 365 Copilot (paid version), and other separately licensed or role-specific AI platforms approved by Emerging Technology. Authorization does not mean automatic provisioning; access is provided based on business need and availability. |
| Embedded AI in approved applications | Native AI functionality in a Coastal-approved application is considered approved when used within that application and its existing permissions. Examples may include Procore, Sage, Egnyte, Microsoft applications, and other approved enterprise software.    |
| Agentic / autonomous AI              | Not approved by default. Requires prior authorization from the Vice President of Emerging Technology for the specific use case.                                                                                                                                 |
| Unapproved / consumer AI             | Not permitted for Coastal confidential or non-public information.                                                                                                                                                                                               |

Product examples above describe the current operating environment and may change without requiring a revision to the governing policy.

## 2. Embedded AI Rule

When Coastal has approved a business application, the application's native AI features are generally approved as part of that application. Separate AI governance approval is not required simply because the vendor adds or exposes an AI-assisted feature, provided the feature:

- Operates within the approved application and its normal authentication and permission model.
- Does not require a user to connect an unapproved third-party model, service, account, plug-in, or external connector.
- Does not independently perform an agentic or autonomous action that requires separate authorization.
- Does not materially change where Coastal data is stored, transmitted, or processed in a way that conflicts with the application's approved security or contractual posture.

### Escalation trigger

If a native feature introduces autonomous actions, external connectors, new third-party data flows, or permissions broader than the underlying application, stop and request review from IT / Emerging Technology before enabling it.

## 3. Data Handling Standard

| Data category     | May use in approved enterprise AI? | Employee rule                                 |
|-------------------|------------------------------------|-----------------------------------------------|
| Public / Internal | Yes                                | Use for legitimate Coastal business purposes. |

| Data category                                                                                                      | May use in approved enterprise AI? | Employee rule                                                                                                                                                  |
|--------------------------------------------------------------------------------------------------------------------|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Confidential business, client, project, bid, estimate, schedule, contract, subcontractor, or financial information | Generally yes                      | You must already be authorized to access the information and there must be no known client, contract, project, legal, or classification restriction on AI use. |
| Restricted / Highly Sensitive                                                                                      | No, unless specifically authorized | Do not upload or process without explicit authorization for that use case.                                                                                     |
| Credentials / secrets                                                                                              | No                                 | Never enter passwords, MFA codes, API keys, tokens, private keys, or authentication secrets into AI.                                                           |
| Information under a known No-AI or special-processing requirement                                                  | No                                 | Follow the stricter client, contract, project, legal, or regulatory requirement.                                                                               |

### How employees identify Restricted / Highly Sensitive data

Employees are expected to recognize obvious or communicated restrictions, not to perform a technical classification exercise for every file. Treat information as Restricted / Highly Sensitive when it is:

- Marked Restricted, Highly Sensitive, No AI, Privileged with an applicable processing restriction, CUI, export controlled, or equivalent.
- A credential, authentication secret, full Social Security number, bank or payment-card number, unredacted government ID, medical record, background-check record, employee investigation/disciplinary record, or similarly sensitive identity/personnel information.
- Covered by a known client, project, contract, NDA, legal, regulatory, or government requirement that limits AI use or external processing.

If none of those conditions is present and the employee is authorized to use the information for the business purpose, the employee may generally use it in an approved enterprise AI environment. When the restriction is unclear, ask the data owner, IT, Legal, or Emerging Technology before uploading.

## 4. Permissions Standard

AI access must not exceed the employee's own access. Do not use AI to search for, infer, summarize, retrieve, or expose information from systems, folders, projects, mailboxes, records, or repositories that you would not normally be authorized to access.

If an AI feature appears to expose information outside expected permissions, stop using the feature and report it to IT.

## 5. Assistance vs. Action Standard

| AI behavior                                                        | Status            | Required control                                                                                       |
|--------------------------------------------------------------------|-------------------|--------------------------------------------------------------------------------------------------------|
| Draft an email or message                                          | Allowed           | Employee reviews the draft and manually sends it.                                                      |
| Automatically send an email or message                             | Not authorized    | No AI system or agent may automatically send outbound communications at this time.                     |
| Draft, summarize, research, analyze, recommend, or prepare content | Allowed           | Human reviews before operational use.                                                                  |
| Populate a draft form, record, or document for human review        | Generally allowed | Human reviews and performs the final submit / save / approval action when the action is consequential. |

| AI behavior                                                                      | Status                                                           | Required control                                                                        |
|----------------------------------------------------------------------------------|------------------------------------------------------------------|-----------------------------------------------------------------------------------------|
| Autonomously change, submit, approve, delete, move, or transmit business records | Requires authorization                                           | Specific agentic use-case approval from VP of Emerging Technology.                      |
| Autonomously purchase, pay, contract, commit, approve, or transact               | Requires authorization; may also be prohibited by other policies | Existing business authority and controls remain mandatory.                              |
| Run a persistent or scheduled agent across systems                               | Requires authorization                                           | Scope, permissions, logging, human checkpoints, and shutdown controls must be approved. |

## 6. External Communications and Deliverables

AI does not change who is allowed to communicate or approve work. The employee must apply the same review requirement that would exist if the employee had written the content without AI.

- If you normally have authority to write, review, and send the communication yourself, you may use AI to draft it, review it, and send it yourself.
- If the communication or deliverable normally requires review by a supervisor, PM, PX, VP, Legal, Safety, Finance, HR, or another approver, that review is still required.
- Do not send AI-generated content without reading and understanding the final content.
- Do not represent unverified AI analysis as established fact; independently verify material facts before relying on them.

### Current email rule

AI may create email drafts. Automatic or agentic sending of email is not authorized.

## 7. Third-Party Connectors, Plug-ins, and Integrations

Only Coastal-authorized third-party connectors and integrations may be used with AI. This includes plug-ins, browser extensions, MCP servers, APIs, add-ons, external models, data sources, and services that connect an AI platform to another application or repository.

### Employee responsibilities

- Do not self-authorize a third-party connector using your Coastal credentials.
- Do not build or deploy a connector that accesses Coastal systems or information unless IT / Emerging Technology has approved the work.
- Do not paste API keys, service credentials, or other secrets into an AI tool in order to create a connection.
- Submit a request to IT when a connector or integration is needed for a business use case.

### IT / Emerging Technology responsibilities

- Review the business need, data flow, permissions, vendor/service, security posture, and operational risk.
- Build, configure, or enable approved integrations using managed credentials and appropriate controls.
- Define allowed users, data sources, permissions, logging, and lifecycle requirements.
- Disable or modify integrations when risk, vendor behavior, business need, or platform capabilities change.

## 8. Agentic AI Authorization

"Work," "cowork," computer-use, autonomous agents, multi-agent systems, or any AI that can take actions on behalf of a user require prior approval from the Vice President of Emerging Technology.

An authorization request should define the user or department, business outcome, systems accessed, data involved, actions permitted, actions prohibited, human approval points, logging, frequency, duration, and responsible business owner.

Approval of the underlying platform or application does not constitute approval of an agentic use case.

## 9. AI Incident Triggers

- You entered Restricted / Highly Sensitive information into an AI service without authorization.
- An AI tool or connector exposed information you did not expect to be able to access.
- An AI agent took, attempted, or queued an unauthorized action.
- AI changed, deleted, submitted, or transmitted business records incorrectly.
- You used an unapproved AI tool, personal AI account, model, connector, or extension with Coastal non-public information.
- Material business action was taken based on AI output later discovered to be substantially incorrect or fabricated.

Report these events promptly to IT Help Desk and the Vice President of Emerging Technology. Preserve relevant prompts, outputs, screenshots, drafts, logs, and records when practical.

## 10. Request and Approval Process

| Request type                                 | Submit to                                                                         | Typical review                                                                                         |
|----------------------------------------------|-----------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| Access to an approved-by-request AI platform | IT / Emerging Technology                                                          | Business need, role, licensing, training, and security.                                                |
| New AI platform or service                   | IT / Emerging Technology                                                          | Vendor, contractual terms, data handling, security, value, and governance fit.                         |
| Third-party connector / integration          | IT                                                                                | Data flow, permissions, credentials, security, operating model, and support ownership.                 |
| Agentic / autonomous use case                | VP of Emerging Technology                                                         | Specific systems, actions, human checkpoints, permissions, logging, business owner, and risk controls. |
| Use of Restricted / Highly Sensitive data    | VP of Emerging Technology plus applicable data owner / Legal / Security as needed | Purpose, data type, platform safeguards, legal/contractual restrictions, and alternatives.             |
| Policy exception                             | VP of Emerging Technology                                                         | Business justification, duration, affected users/data, risk, compensating controls, and owner.         |

## 11. Current Operating Notes

- ChatGPT Enterprise is the generally available enterprise AI environment for normal business use.
- Claude Team is authorized but is provisioned only when requested and approved; it is not provided by default.
- Microsoft 365 Copilot is approved by request and subject to licensing and business need.
- Native AI in approved applications is governed by the embedded AI rule rather than by a permanent product list.
- Egnyte is the correct platform name; references to "Ignite" in prior drafts should be removed.
- Automatic AI sending of email is not authorized. Draft creation is allowed with human review and manual send.
- Employees may not create their own third-party AI connectors for Coastal systems or data.

## 12. Standard Ownership and Review

Owner: Vice President of Emerging Technology. These standards should be reviewed at least quarterly during the initial adoption period and whenever material AI platform, security, contractual, or business-process changes occur.

### Version Control

| Date      | Version | Owner               | Description                                                                  |
|-----------|---------|---------------------|------------------------------------------------------------------------------|
| 8/18/2026 | 0.1     | Emerging Technology | Initial operational standard supporting revised AI Use and Governance Policy |