501CISO
501C CYBERSECURITY LEADERSHIP
Powered by ClearTone Consulting, LLC

CLIENT ADVISORY

CLAUDE AI GOVERNANCE FOR NONPROFITS

Choosing the Right Plan and Controlling Third-Party Data Access

Prepared by ClearTone Consulting LLC | 501CISO Practice | April 2026

EXECUTIVE SUMMARY

As nonprofit organizations adopt Claude AI across their teams, a consistent governance gap is emerging: the tools Anthropic provides to control who can connect to what — and what they can do — vary dramatically between plan tiers. Without a clear understanding of those differences, organizations risk exposing sensitive donor, client, and operational data through unmanaged AI connections to their SaaS systems.

This advisory answers three questions every nonprofit technology or security leader should be asking:

- Which Claude plan actually gives our administrators meaningful control?
- How do we prevent employees from connecting organizational SaaS data to personal Claude accounts?
- What controls must be implemented at both the Anthropic admin layer and the network/endpoint layer to reduce risk?

> **Key Finding:** The Team plan provides connector governance for org-managed accounts but cannot prevent employees from accessing the same SaaS systems through personal Free or Pro Claude accounts. Closing that gap requires either upgrading to Enterprise or implementing network/endpoint controls — or both.

1. HOW CLAUDE CONNECTORS WORK

Claude connectors allow the AI to access third-party SaaS systems — reading files, searching email, creating tasks, and more — on behalf of authenticated users. They are built on the Model Context Protocol (MCP) and are available across Claude's web interface, desktop app, and agentic task tool (Cowork).

The Two-Layer Access Model

All connector access operates through two sequential gates:

Layer	Who Controls It	What It Does
Org-Level Enablement	Owner / Primary Owner	Makes a connector available to the organization. Without this step, no user can access the connector through the managed account.
User-Level Authentication	Individual User	Each user must personally authenticate with the third-party service. Claude inherits only that user's existing permissions — nothing more.

This design has an important security implication: Claude cannot access data that a user cannot already access in the source system. The connector model narrows access — it never expands it.

What Connectors Can Access

The Connectors Directory currently includes over 50 verified integrations. Categories most relevant to nonprofits include:

- Productivity & Files: Google Drive, Microsoft 365 (SharePoint, OneDrive, Outlook, Teams)
- Communication: Gmail, Slack, Zoom
- Project Management: Asana, Linear, Jira
- CRM & Fundraising: Salesforce (via custom MCP connector)
- Custom Internal Tools: Any system with a remote MCP server endpoint

> **Note:** The Microsoft 365 connector is only available on Team and Enterprise plans and requires a one-time setup by a Microsoft Entra ID Global Administrator before individual users can connect.

2. WHAT ADMINISTRATORS CAN CONTROL BY PLAN

Team Plan Admin Controls

The Team plan provides a meaningful first layer of governance for organizations that provision Claude accounts centrally. Owners can:

- Enable or disable specific connectors for the entire organization from Organization Settings > Connectors
- Set per-action restrictions on each connector (e.g., allow read but block write/delete)
- Verify the organization's email domain to block new personal accounts from being created on that domain
- View basic usage data through the admin dashboard

> **Team Plan Limitation:** The Team plan cannot claim or migrate existing personal accounts, cannot enforce SSO, and cannot prevent users from accessing claude.ai through a personal account using a different email address.

Enterprise Plan Admin Controls

Enterprise adds the identity and enforcement layer that most security-conscious organizations require:

- SSO Integration (SAML 2.0 / OIDC): Require all logins to flow through your identity provider (Okta, Azure AD / Entra ID, Google Workspace, Auth0, Ping Identity)
- Domain Capture: Corporate email addresses are automatically routed to the Enterprise workspace — no personal account can be created with that email
- Restrict Organization Creation: Prevents users from creating new personal or team Claude accounts using your verified domain
- Domain Claiming: Discover and migrate all existing personal accounts on your domain into the Enterprise workspace (irreversible)

- SCIM Provisioning: Automate user onboarding and offboarding — when someone leaves, their Claude access is revoked automatically
- RBAC Groups: Assign connector exposure and feature access based on IdP group membership
- Compliance API: Programmatic, real-time access to usage data and content for audit logging and policy enforcement
- OpenTelemetry / SIEM Integration: Cowork emits events for all tool and connector calls, compatible with Splunk, Cribl, and other SIEM pipelines
- Private Inference Deployment: Route prompts through your own AWS Bedrock, Google Vertex AI, or Azure AI Foundry instance — Anthropic never sees the content

Per-Connector Action Controls (Team & Enterprise)

Both Team and Enterprise administrators can configure granular action restrictions for each enabled connector:

Permission Setting	Behavior
Always Allow	Claude can invoke this action automatically when relevant
Needs Approval	Claude must pause and get explicit user confirmation before taking the action
Blocked	Claude cannot invoke this action, regardless of user request

Common governance patterns for nonprofits:

- Email: Allow search and summarize; block send
- Google Drive / SharePoint: Allow read; block create, edit, delete
- Project tools (Asana, Linear): Allow view; block create or status changes
- CRM: Allow read contacts and giving history; block record modification

3. THE SHADOW AI PROBLEM

The Core Risk

Here is the governance gap that most organizations miss: even with a well-configured Team or Enterprise plan, an employee with a personal Free or Pro Claude account can navigate to claude.ai, sign in with their personal email, connect their own Google Drive or Gmail or M365 credentials, and begin working with organizational data — completely outside your administrative visibility or control.

This is not a hypothetical edge case. It is the default behavior for any employee who was using Claude personally before the organization adopted it — and for any employee who simply decides to use their personal account for a work task.

> **Why This Matters for Nonprofits:** Donor records, grant materials, client intake forms, and board communications may all be accessible through a staff member's Google Drive or Outlook account. If Claude can be connected to those accounts outside of organizational control, that data flows through Anthropic's systems on terms that your organization never agreed to.

What the Team Plan Can (and Cannot) Do

The Team plan's domain verification feature can block new personal accounts from being created on your corporate email domain. This is a useful preventive control for new employees, but it has three significant limitations:

- It cannot reach or affect accounts already created before you enabled the control
- It does nothing to stop an employee from using a personal Gmail or other non-corporate email to create a Claude account
- It does not prevent access to claude.ai on a personal account that was already in existence

What the Enterprise Plan Adds

Enterprise's domain capture and SSO enforcement close the gap for corporate email addresses. When fully configured:

- Any login attempt with a corporate email address is automatically routed to the Enterprise workspace
- Existing personal accounts on the corporate domain can be claimed and migrated
- Users with corporate email addresses cannot access personal Free/Pro accounts — those accounts remain but become inaccessible

> **Remaining Gap:** Neither Team nor Enterprise can prevent an employee from using a personal Gmail or other non-corporate email to access claude.ai and connect their work SaaS tools. This requires network and endpoint controls.

4. NETWORK AND ENDPOINT CONTROLS

Closing the shadow AI risk completely requires controls outside of Anthropic's admin console. These operate at the network and device layer and are applicable regardless of which Claude plan you are on.

Tenant Restriction / Header Injection (Most Effective)

Some Secure Web Gateway (SWG) and Cloud Access Security Broker (CASB) vendors support injecting an HTTP header into outbound traffic to claude.ai. Anthropic's infrastructure reads this header and validates the session against an approved organization ID list. Sessions not matching an approved ID are blocked with an IT administrator message.

This means an employee on a corporate device — even using a personal Gmail account — will be blocked from accessing claude.ai unless they are authenticated to your approved Enterprise workspace. This is the same architectural pattern used for Microsoft 365 and Google Workspace tenant restrictions.

Vendor / Approach	Mechanism	What It Blocks
SWG with header injection (e.g., dope.security, Netskope, Zscaler)	Injects anthropic-allowed-org-ids header into outbound traffic	Personal Free, Pro, and Max accounts on any email address
Firewall / DNS blocking	Block claude.ai entirely at the network perimeter	All Claude access on managed network; use allowlists for managed devices only
MDM-managed device policy	Restrict browser access or enforce proxy routing on corporate devices	claude.ai access on managed devices regardless of network
Content DLP (CASB)	Inspect content being uploaded to claude.ai for sensitive data patterns	PHI, PII, donor data, financial data — regardless of which account is in use

DLP as a Complementary Layer

Account-based controls (SSO, domain capture, tenant restrictions) stop unauthorized sessions. Content DLP stops sensitive data from leaving the organization even through authorized sessions. For nonprofits handling health information, donor records, or grant-related financial data, both layers are needed.

Organizations should classify claude.ai as a third-party SaaS application in their DLP tooling and apply the same content inspection policies they would to Google Drive or Dropbox.

Cowork-Specific Considerations

Cowork is Claude's agentic task interface, capable of autonomous, multi-step operations across connected SaaS systems. Because Cowork can take write actions without per-action user confirmation, its governance requirements are more stringent than the standard chat interface.

- Enterprise admins can restrict Cowork connector actions org-wide from the admin console

- Cowork emits OpenTelemetry events for all tool calls, connector invocations, and file operations — compatible with SIEM pipelines
- For organizations requiring data sovereignty, Cowork can be deployed via MDM to route inference through AWS Bedrock, Google Vertex AI, or Azure AI Foundry
- Custom MCP connectors must be reachable over the public internet — they cannot be on a private network without allowlisting Anthropic's IP ranges

5. PLAN COMPARISON AND CONTROL CAPABILITY

The following tables summarize the governance capabilities available at each plan tier and across control layers.

Table 1: Connector Governance by Plan

Control Capability	Team Plan	Enterprise Plan
Enable / disable connectors org-wide	✓ Yes	✓ Yes
Per-action restrictions (read vs. write)	✓ Yes	✓ Yes
User-level auth enforcement	✓ Yes	✓ Yes
Custom MCP connectors (Owners only)	✓ Yes	✓ Yes
Block new personal accounts on corp domain	~ Partial	✓ Yes
Claim / migrate existing personal accounts	✗ No	✓ Yes
Enforce SSO (SAML 2.0 / OIDC)	✗ No	✓ Yes
SCIM provisioning / auto-deprovisioning	✗ No	✓ Yes

Control Capability	Team Plan	Enterprise Plan
Domain capture (corporate email lockdown)	✗ No	✓ Yes
RBAC group-based connector scoping	✗ No	~ Beta
Compliance API (programmatic audit access)	✓ Yes	✓ Yes
OpenTelemetry / SIEM integration	✓ Yes	✓ Yes
Private inference (Bedrock / Vertex / Azure)	✗ No	✓ Yes
MDM-managed Cowork deployment	✗ No	✓ Yes

Table 2: Shadow AI Risk — Control Layer Coverage

Risk Scenario	Team Plan	Enterprise Plan	Network / Endpoint Layer
Block new personal accounts on corp email domain	✓ Yes	✓ Yes	✗ N/A
Prevent corp email from accessing personal accounts	✗ No	✓ Yes (SSO + capture)	~ Partial
Block personal accounts on non-corp email (Gmail, etc.)	✗ No	✗ No	✓ Yes (header inject)
Block claude.ai on non-compliant managed devices	✗ No	✗ No	✓ Yes (MDM + proxy)
Content DLP — prevent sensitive data upload	✗ No	✗ No	✓ Yes (CASB / DLP)
Auto-revoke access when user leaves org	✗ No	✓ Yes (SCIM)	~ Partial

Risk Scenario	Team Plan	Enterprise Plan	Network / Endpoint Layer
Full audit log of connector and tool calls	~ Basic	✓ Full (Compliance API)	~ Supplemental

6. RECOMMENDED GOVERNANCE POSTURE BY ORGANIZATION TYPE

Organization Profile	Recommended Plan	Minimum Additional Controls
Small nonprofit, <50 staff, low data sensitivity, minimal SaaS integration	Team	Domain verification; acceptable use policy; quarterly connector review
Mid-size nonprofit, 50-200 staff, handles donor PII or client records, uses M365 or Google Workspace	Enterprise	SSO enforcement; domain capture; SCIM provisioning; DLP classification of claude.ai
Large nonprofit or health/social services org, 200+ staff, handles PHI or regulated data	Enterprise + Private Inference	Full stack: SSO, SCIM, domain capture, tenant restriction via SWG, CASB with content DLP, SIEM integration, Zero-Data-Retention addendum
Organization currently on Team plan with employees using personal Claude accounts	Evaluate Enterprise upgrade	Immediate: domain membership audit. Short-term: deploy tenant restriction if Enterprise upgrade is delayed

7. ADMINISTRATOR ACTION CHECKLIST

On Team Plan — Immediate Actions

- Navigate to Organization Settings > Connectors and audit which connectors are currently enabled
- For each enabled connector, review and configure per-action permissions (read vs. write vs. blocked)
- Verify your organization's email domain to prevent new personal account creation
- Publish an acceptable use policy that explicitly addresses personal Claude accounts and organizational data
- Identify any employees known to be using personal Claude accounts for work tasks

On Enterprise Plan — Foundational Configuration

- Verify your domain via DNS TXT record in Organization Settings > Organization and Access
- Configure SSO (SAML 2.0 or OIDC) through your identity provider — enable 'Require SSO for Claude'
- Enable SCIM provisioning to automate user lifecycle management
- Run Domain Memberships report to identify all existing personal accounts on your domain
- Enable domain capture and initiate account claiming/migration (note: this is irreversible — communicate to affected staff first)
- Configure RBAC groups in your IdP to assign Claude roles and connector access by function
- Enable OpenTelemetry and connect to your SIEM pipeline before enabling Cowork at scale
- Review and configure Cowork connector action restrictions in the admin console

Network / Endpoint Layer — Regardless of Plan

- Classify claude.ai as a managed third-party SaaS application in your DLP/CASB tooling
- Evaluate SWG vendors that support Anthropic tenant restriction header injection
- Deploy MDM policy to restrict claude.ai access on unmanaged or non-compliant devices
- Define sensitive data categories (PHI, donor PII, financial records) and create DLP rules for claude.ai uploads
- Include Claude usage in your quarterly access review process

8. KEY LIMITATIONS AND HONEST CAVEATS

- Personal devices: If an employee accesses claude.ai from a personal device on a personal network, neither Anthropic's controls nor network-layer controls on your corporate infrastructure can prevent it. Policy, training, and contractual agreements with staff are the only available controls.
- Third-party connector terms: Connectors in the directory are built and maintained by third-party developers. Each has its own terms of service and privacy policy, which are presented during authentication. Your organization should review these terms for connectors accessing sensitive data.
- Custom MCP connectors must be internet-reachable: If you build a custom connector to an internal system, that server must be reachable over the public internet from Anthropic's cloud. Systems behind your firewall require allowlisting Anthropic's IP ranges.
- Domain capture is one-way: Once enabled, it cannot be reversed. Ensure SSO and SCIM provisioning are fully operational before enabling, or affected staff will lose access to Claude entirely.
- RBAC group-based connector scoping is currently in beta: Treat as a supplement to, not a replacement for, per-connector action restrictions.

CONCLUSION

Claude's connector framework is powerful and its admin controls are meaningfully designed — but they are not a turnkey data governance solution. The gap between what a Team plan can enforce and what is needed to actually prevent unauthorized data exposure is significant, and it widens as organizations grow in size, data sensitivity, and SaaS complexity.

For most nonprofits that have moved beyond a pilot phase, the Enterprise plan is not a luxury — it is the minimum viable governance posture for managing AI access to organizational data. Layering network and endpoint controls on top addresses the shadow AI risk that no SaaS admin console can fully close.

ClearTone Consulting is available to assess your current posture, facilitate a plan selection conversation with Anthropic, and help design the technical and policy controls appropriate for your organization's risk profile.

Contact 501CISO / ClearTone Consulting

501CISO Practice | Fractional CISO Services for Nonprofits

501CISO.com | Schedule a discovery call to discuss your organization’s AI governance needs.

This document reflects product capabilities as of April 2026. Claude’s feature set and administrative controls evolve rapidly; readers should verify current functionality directly with Anthropic. This advisory does not constitute legal advice and utilized Claude in its development.