cyberhaven

ChatGPT Enterprise *Documentation*

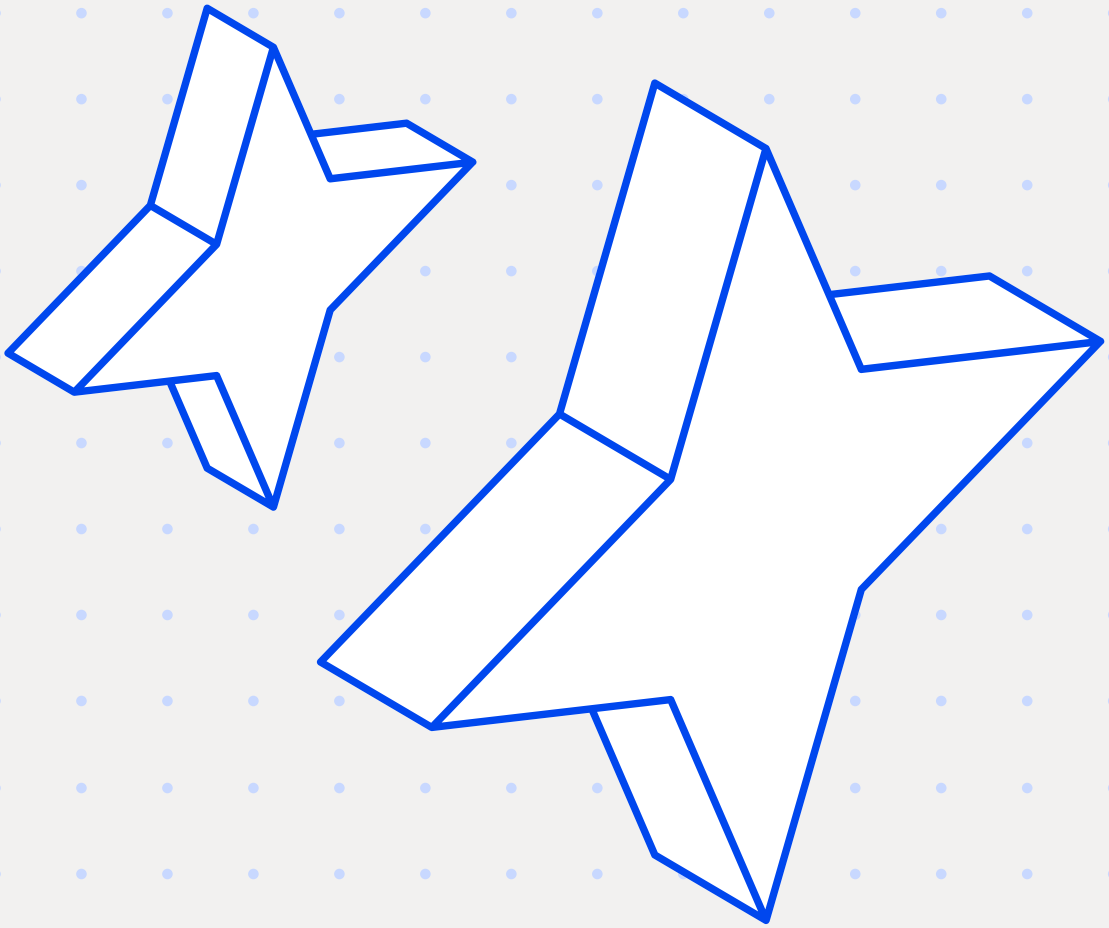

ChatGPT Enterprise

The ChatGPT Enterprise Cloud Connector integrates Cyberhaven with OpenAI's ChatGPT Enterprise to provide visibility into AI usage across your organization. The integration uses OpenAI's Compliance API to monitor conversations, file uploads, custom GPT knowledge files, and tool invocations. The connector performs content inspection on conversation transcripts and attached files to detect sensitive data shared with or generated by ChatGPT.

Requirements

The integration requires the following to function:

Requirement	Details
Eligible plan	Available for ChatGPT Enterprise and Edu workspaces.

OpenAI API key	Must be created with the "Default Project - All Permissions" scope. The key is shown only once at creation, so store it safely.
Compliance API access	Email `support@openai.com` with the last 4 digits of the key, the key name, the creator name, and the requested scope (`read`). OpenAI reviews the request and grants access.

Workspace ID	Found in the OpenAI account under **Profile icon > Workspace settings > General > Workspace identifier**.

IMPORTANT

Use the API key only for Cyberhaven. Do not share or reuse it for other applications or scripts. Rotating or revoking the key in any other context will break the connector.

Coverage

The connector provides visibility into the following activities:

Conversations

- User messages sent to ChatGPT (prompts)
- Assistant responses (completions)
- File uploads within conversations
- Tool invocations logged via `APP_LOG` events (for example, MCP tools)

Custom GPTs

- Knowledge files uploaded to custom GPTs
- GPT configuration and sharing settings
- GPT ownership information

Content inspection

The connector inspects three types of content for sensitive data:

- **Conversation transcripts**: The full conversation text, including user prompts and assistant responses.
- **Conversation file attachments**: Files uploaded within conversations.
- **GPT knowledge files**: Files uploaded as knowledge sources for custom GPTs.

File scanning respects the file type and file size filters configured in scan settings.

Metadata collected

Conversation events

For each conversation event, the connector collects:

- Conversation ID and title
- Message author (user or assistant), model used, and tools invoked
- File references (file ID and file name)
- Custom GPT ID and name (when the conversation uses a custom GPT) ● Temporary chat flag (whether the conversation is saved or ephemeral) ● Workspace ID
- Timestamp of each message

Custom GPTs

For each custom GPT, the connector collects:

- GPT name
- Creator email
- Sharing configuration
- Knowledge file list

Limitations

- Image-only messages (generated images with no text) are not scanned. ● On the first scan, the connector uses the configured scan lookback period when provided; otherwise it defaults to 30 days.
- The connector tracks activities for all users in the workspace by default. Use the Users scope setting to limit scanning to specific users.

- OpenAI enforces rate limits on the Compliance API (typically 50 requests per minute per endpoint). The connector handles this automatically with backoff and retry.
- The API key does not expire but can be manually rotated by your OpenAI admin. If rotated, you must re-authenticate the connector with the new key.

OpenAI-side data retention

The following retention behaviors are enforced by OpenAI and are outside Cyberhaven's control. For the source documentation, see OpenAI Admin API Reference (you must be logged into your Enterprise workspace to access this link).

- **Compliance log files are retained for 30 days.** If the connector is offline for more than 30 days, events from that gap cannot be recovered. OpenAI does not support backfilling historical data.
- **Chat-uploaded files are held for 48 hours only.** The connector polls regularly to capture these before expiration, but files uploaded during extended connector

downtime may be lost.

- **GPT and Project files are retained indefinitely** by OpenAI because active GPTs require them.
- **User-deleted conversations remain available** via the Compliance API for up to 30 days after deletion, even though the user can no longer see them in ChatGPT. ● **Events may take up to 30 minutes** (p99) to appear in OpenAI's compliance log files after the action occurs.

ChatGPT Enterprise Deployment

This guide outlines the steps to deploy the ChatGPT Enterprise Cloud Connector in the Cyberhaven Console. The integration uses OpenAI's Compliance API with an API key to read compliance logs, user information, and custom GPT data from your ChatGPT Enterprise workspace.

Before you begin, review the prerequisites: ChatGPT Enterprise prerequisites.

Set up OpenAI prerequisites

Complete these steps in your OpenAI account before configuring the connector in Cyberhaven.

1. Go to the OpenAI API Platform Portal.
2. Confirm the correct organization (your Enterprise workspace) is selected. Do not use a personal account.

3. Create a new API key with the **Default Project - All Permissions** scope. This must be a fresh key dedicated to Cyberhaven.
4. Copy the API key and store it safely. OpenAI only displays the key once. 5. Email `support@openai.com` to request Compliance API `read` scope. Include the last 4 digits of the key, the key name, the creator name, and the requested scope. 6. Wait for OpenAI to confirm the scope has been granted to the key. 7. Locate your **Workspace ID** in **Profile icon > Workspace settings > General > Workspace identifier**.

Connect Cyberhaven to ChatGPT Enterprise

To connect your ChatGPT Enterprise workspace, log in to your Cyberhaven Console and follow these steps:

1. In the Cyberhaven Console, click **Connectors** in the left navigation bar. 2. Click the **Add connectors** tab and then click **Connect** on the ChatGPT Enterprise card.

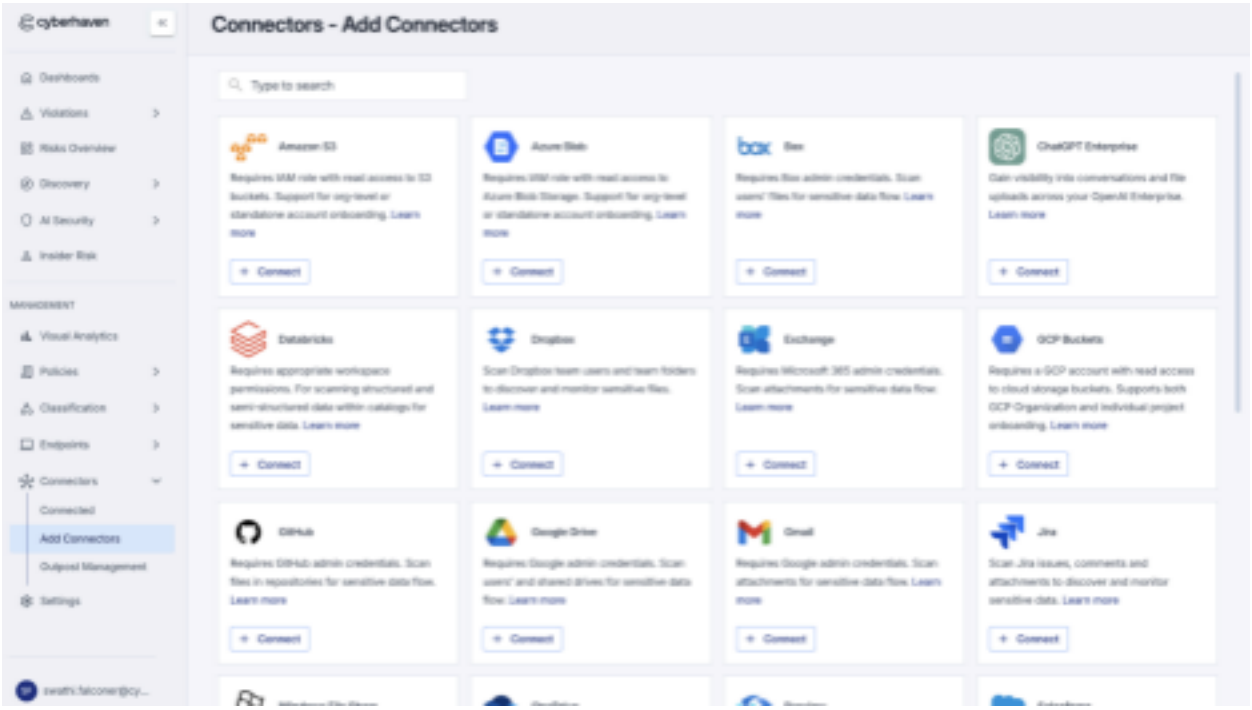

3. On the Add Connector window, follow the instructions in the connection guide under **Connector details**.
4. Enter your **Workspace ID** from the OpenAI prerequisites step.

5. Enter your **API Key**. The key must already have Compliance API `read` scope granted by OpenAI support.

6. Click **Connect**.
7. Review the connector details and click **Save**.

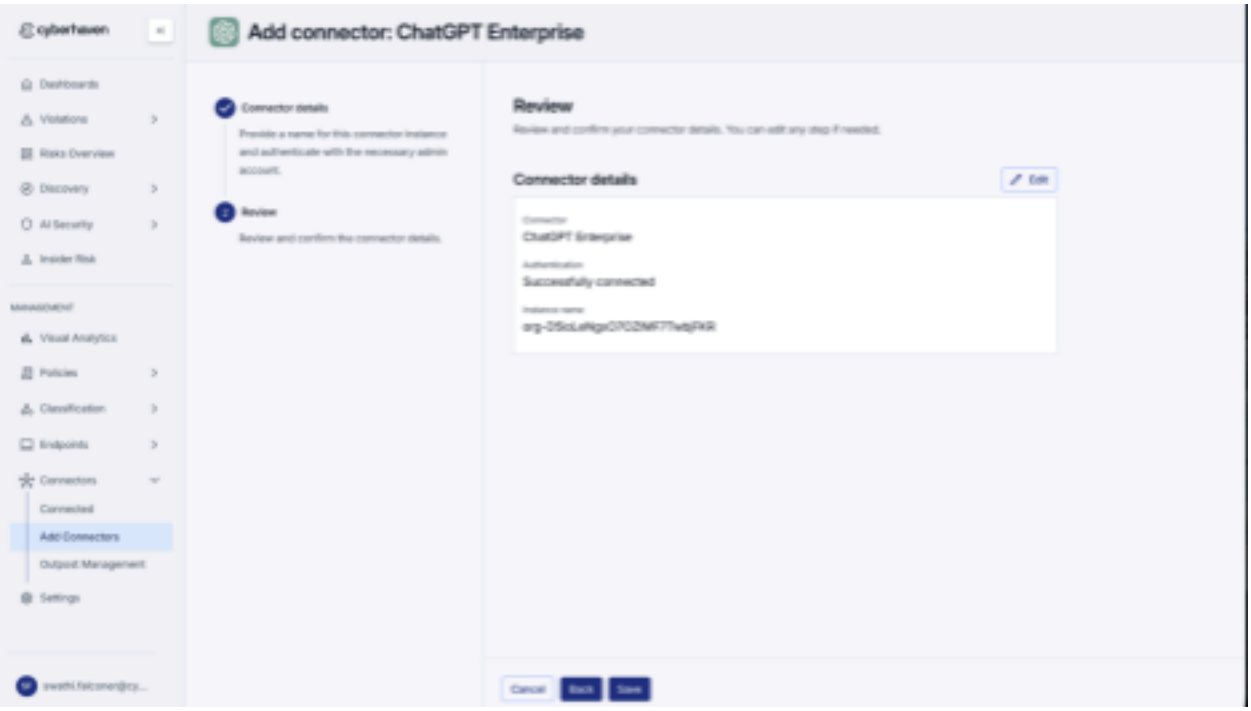

The newly connected ChatGPT Enterprise connector is displayed in the **Connected** tab. You can now click on the connector to configure scope and scan settings.

After connection, Cyberhaven begins retrieving ChatGPT Enterprise compliance logs for up to 30 days of historical data. OpenAI drops any data older than 30 days. Events may take up to 30 minutes to appear in the Console after they occur in ChatGPT, per OpenAI's compliance log delivery SLA.

Configure Scan

Once the ChatGPT Enterprise connector is connected, you can configure historical and forward scans to discover and classify conversations and custom GPT content in your ChatGPT Enterprise workspace. Forward scans continuously discover and classify new conversations and content after scan configuration. Historical scans provide coverage by scanning conversations and content created before the configuration.

1. On the **Connected** tab of the Connectors page, click the gear icon next to the connector.
2. Under **Scope**, choose what content to scan. Click **Next**.
    - **Conversations**: Scan user-assistant chat transcripts. When enabled, choose which users to include:
        - Select **All users** to include all current and future users.

- Use **Select manually** to include or exclude specific users by searching and checking the boxes.
- Select **Advanced** to enter one or more regex patterns to dynamically define the scope of users.

- **Custom GPTs**: Scan knowledge files uploaded to custom GPTs. When enabled, choose which GPTs to include:
  - Select **All GPTs** to include all current and future custom GPTs.
  - Use **Select manually** to include or exclude specific GPTs by searching and checking the boxes.
  - Select **Advanced** to enter one or more regex patterns to dynamically define the scope of GPTs.

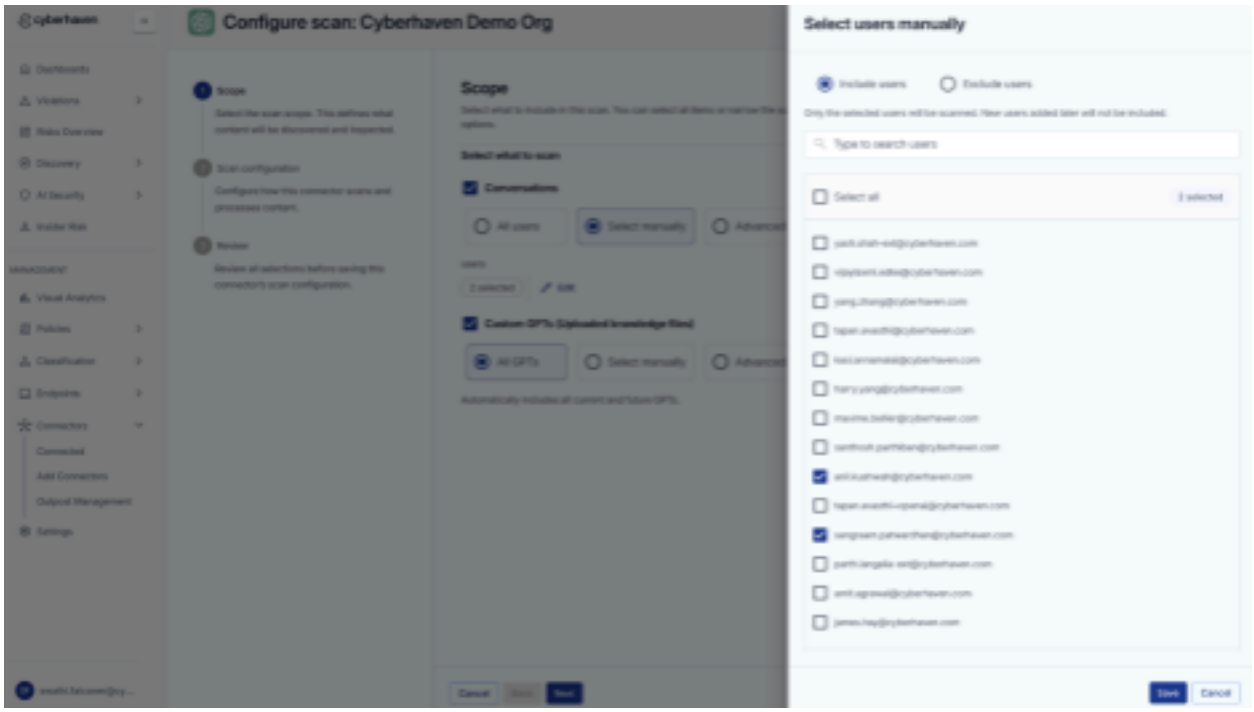

3. In the **Scan configuration** section, enable the scan type and define the scan parameters.
   - Enable **Forward scans** to continuously discover and classify new or modified content.
   - Enable **Historical scans** to scan content created before this configuration.
     - In the **Max last modified days** field, enter the number of days of history you want to scan. OpenAI does not retain compliance log data older than 30 days, so values greater than 30 are capped by

OpenAI.

- Expand **Advanced settings** to configure granular limits and preferences:
  - **Maximum content inspection file size**: Set the maximum file size (in MB) for content inspection.
  - **File types**: Click the Edit icon to adjust the list of file types to be included in the scan.
  - **Capture bucket for generated content**: Choose a configured external storage location to store evidence files, or select None to disable evidence capture for this scan.

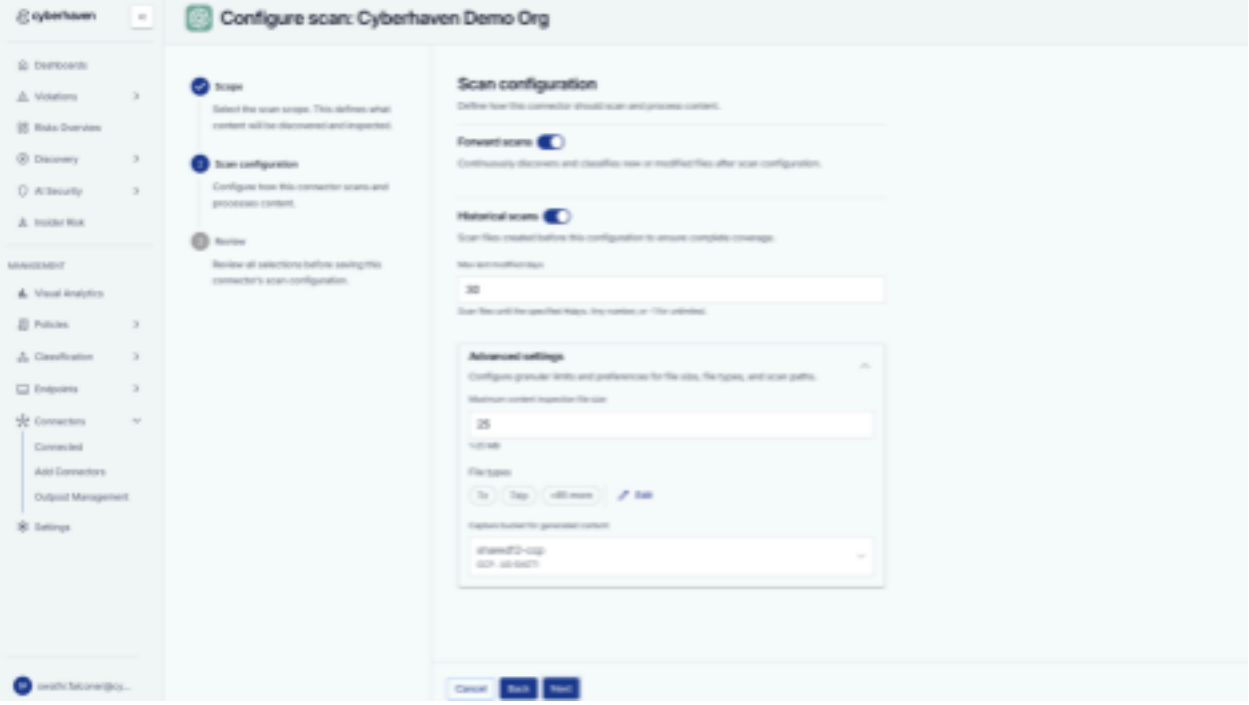

4. Click **Next**.
5. Review your scope and scan configuration and click **Save**.

Re-authenticate Connector

If your OpenAI admin rotates the API key, you need to re-authenticate the connector with the new key:

1. On the **Connected** tab, click the three-dot menu next to the connector and select **Reauthenticate**.
2. Enter the new API key. The Workspace ID should not be changed.
3. Click **Connect** to complete re-authentication.

Delete Connector

To delete the connector, click the three-dot menu next to the connector and click **Delete**.

Documentation

Troubleshooting

- **Authentication failed**: Verify that the API key has Compliance API `read` scope granted by OpenAI support. Keys without this scope cannot access the Compliance API.
- **No events appearing**: It may take 30 minutes to 1 hour for events to appear in the Cyberhaven Console. Verify the connector shows as **Connected** in the Console and that the Workspace ID is correct.
- **Rate limit errors**: OpenAI enforces rate limits on the Compliance API (typically 50 requests per minute per endpoint). The connector handles this automatically with backoff and retry. If errors persist, verify no other applications are using the same API key.
- **Missing files**: Chat-uploaded files are retained by OpenAI for only 48 hours. Files from conversations that occurred during extended connector downtime may be unavailable. GPT knowledge files are retained indefinitely by OpenAI.
- **API key rotation**: If your OpenAI admin rotates the API key, re-authenticate the connector with the new key. The old key stops working immediately.