# Orchestra AI core policy

## Purpose and scope

This policy sets clear, practical rules for how Orchestra employees, freelancers, temporary personnel, and other approved non-employees use AI tools in day-to-day work. The goal is to enable safe, high-quality work while protecting client trust, confidential information, and Orchestra's reputation.

This policy applies to:

- All employees
- Contractors and freelancers or temporary personnel working on Orchestra deliverables or using Orchestra systems

For external vendors and partners who use their own tools or systems, Orchestra's expectations should be addressed through contracts, specifications, and project communication rather than assuming they have access to Orchestra-approved AI tools.

## How to use this policy

For each policy area, this document includes:

- Rule statements that define what is required
- Example workflows that show how work should happen
- An owner role responsible for exceptions and interpretation

## Roles and accountability

- Policy owner: Accountable for the policy text, exceptions, and updates
- AI steering group or cross-functional AI council: Resolves tool approvals, major exceptions, and periodic reviews
- IT and security owner: Approves enterprise tooling, integrations, sandboxing, and access controls
- Legal: Advises on IP, disclosures, client contract requirements, privacy, confidentiality, and vendor terms
- Design leadership owner: Sets practical rules for image generation, image editing and visual manipulation
- Content or copy leadership owner: Sets practical rules for AI-assisted writing, editing, messaging, tone, and client-facing copy standards
- People managers: Coach correct usage and reinforce review norms
- All staff: Remain responsible for policy adherence and the quality, accuracy, and client appropriateness of outputs

# Privacy and data security

## Rules

1. Use only Orchestra-approved AI tools for work, including any work involving confidential internal information or client information.
2. Do not paste or upload confidential or sensitive information into public AI services or personal accounts.
3. Treat client information as confidential by default. If unsure, treat the information as confidential or sensitive.
4. Keep client data separated. Do not combine details from different clients in the same prompt, workspace, or dataset.
5. Any AI tool that requires training on Orchestra data, connects to external apps, or ingests large datasets must go through IT review and sandboxing before use.

## Allowed on approved Orchestra tools and accounts

- Drafting, summarizing, rewriting, and analysis using Orchestra-approved enterprise tools (for example ChatGPT Enterprise), including **client confidential** and **internal confidential** information, as long as you:
    - keep each client's information separate from other clients
    - share outputs only within the correct client workspace or team channel
    - minimize what you paste or upload to what is needed for the task
- Creating internal checklists, outlines, and first-pass drafts that are then reviewed and edited by a human before client delivery

## Requires review and approval

- Any use of **non-approved tools, personal accounts, browser extensions, or unvetted AI services**
- Any request to connect AI tools to Orchestra systems (Google Drive, email, calendar, Slack, ticketing) or to client systems
- Any plan to fine-tune, train, or otherwise adapt a model using Orchestra or client data
- Use of personal information as AI input when it is sensitive, unnecessary for the task, or used at scale.
- Where this policy says "requires review and approval," it means submitting a request through the AI intake form to the policy owner and AI council, who will approve, deny, or approve with guardrails.

## Never allowed

- Mixing two clients' confidential information in the same workspace, thread, project, or dataset
- Copying raw outputs into external-facing work without human review
- Entering credentials, API keys, or access tokens into any AI tool
- Inputting sensitive personal information (numerical identifiers, passwords, financials).

### Example workflows

Workflow A: Using AI with client information

1. Classify the input as public, internal, client confidential, or personal data
2. If internal or client confidential, use only approved enterprise tools
3. Use the appropriate client workspace, project, or approved client GPT
4. Save outputs in the appropriate client workspace. Do not combine confidential information from different clients in the same chat, project, GPT, or knowledge base
5. If the task requires broader data access, request IT review and approval before proceeding

Workflow B: Handling a mistake/accidental disclosure of confidential information

1. Stop and preserve context (what tool, what was shared, when)
2. Notify the IT and security owner and the policy owner
3. Follow internal incident steps and client contract requirements for client notification regarding client confidential information if required, in coordination with Legal
4. Record the incident for quarterly review and process improvement

Owner role: IT and security owner with the policy owner

## Use of AI in writing and assistance

### Rules

1. AI outputs must be reviewed for accuracy, tone, client appropriateness, and relevant risks such as bias, harmful stereotypes, or inappropriate content before external use. Team leads should determine when added review is needed.
2. AI output must be reviewed and edited before it is shared externally or treated as final.
3. AI must not be used as a substitute for core judgment on sensitive decisions, especially where client risk is involved.
4. Do not paste AI output verbatim into client deliverables without verification and editing.
5. Managers should coach usage that strengthens work quality rather than replacing thinking.

6. Verify any inputs to ensure that you have permission to use them (no use of unlicensed third party IP) and that they are appropriate for the project.

### Allowed on approved Orchestra tools and accounts

- First-pass drafting for emails, outlines, press materials, Q&A prep, summaries, and meeting materials
- Rewriting for tone, clarity, structure, and length
- Creating checklists, project plans, agendas, and internal templates
- Building interview questions, internal briefing docs, media lists, and stakeholder messaging frameworks

### Requires review and approval

- Any content that makes product claims, particularly in regulated industries like finance and pharma (review product claims with client)
- Any content containing legal, compliance, regulatory, or contractual statements or claims;
- Any plan to publish or send AI-assisted copy externally without human review and edits

### Example workflows

Workflow A: AI-assisted drafting for a client deliverable

1. Use the approved client workspace, project, or client GPT for any client-confidential material
2. Provide the AI with the source material, context, and constraints needed to do the task well
3. Verify facts, names, dates, claims, quotes, and citations against primary sources
4. Confirm permission to use any third-party material as input to avoid violating license, contract, copyright, or client restrictions
5. Review outputs for accuracy, tone, client appropriateness, and relevant risks such as bias, harmful stereotypes, or inappropriate content
6. Have a second reviewer for high-stakes deliverables
7. Save a brief record of what was AI-assisted if the project requires traceability

Owner role: Policy owner with People managers

## Transparency and disclosure

## Rules

1. Internally, staff should be comfortable noting when AI was used for assistance.
2. Client-facing disclosure should follow client guidelines, contract terms, and the risk posture for the deliverable.
3. If a client prohibits AI use on their account, Orchestra will honor that and document it in account procedures. Client contractual restrictions or guidelines supersede this policy as to all work for the client.
4. 
5. Do not misrepresent AI-generated or AI-assisted work as fully human-created when a disclosure is required for public-facing deliverables.

## Allowed on approved Orchestra tools and accounts

- Team-level transparency during drafting and review
- Using a simple internal note such as “AI-assisted draft reviewed and edited” (disclosure of internal AI use is at the discretion of team leaders as long as it does not mislead or if disclosure is not required by the client)

## Requires review and approval

- Any client request for a formal AI-disclosure label to appear on external deliverables
- Any case where disclosure could affect trust, compliance, or contract terms
- Any public-facing disclosure language (website statements, pitch language, boilerplate) that could be reused broadly

## Example workflows

Workflow A: Client account with AI restriction

1. Record the restriction in the account’s working procedures
2. Confirm the GRL, account lead, and core team are aware of the restriction
3. Use a clear internal flag so teams avoid accidental use
4. If an exception is requested, route through the policy owner, account lead, and any required legal or risk reviewer

Owner role: Policy owner with Account leadership

## Intellectual property and copyright

### Rules

1. Respect copyright, trademarks, licensing terms, and other IP. Do not input third-party copyrighted or other IP-protected materials (e.g., names of artists or celebrities, artworks, trademarks or logos, quotes) unless Orchestra has rights to use them for that purpose. .For example, do not include the names of famous people or movies in prompts (e.g., create a press release in the style of the latest comic movie franchise; create song lyrics in the style of a particular singer)
2. Treat AI outputs as requiring the same review for originality, permissions, and claims as any other work.
3. For externally shared visuals, follow Design's approved guidance on what is permitted.
4. Avoid generating or using outputs that imitate identifiable branded styles, logos, artwork, people, or proprietary assets without authorization.
5. Keep a conservative posture for high-visibility work products.

### Allowed on approved Orchestra tools and accounts

- Summarizing or rewriting materials you have rights to use for the task
- Editing Orchestra-owned or properly licensed images using approved tools and workflows
- Internal concepting and exploration that is not intended for publication or client delivery
- Generating metadata such as headlines, alt text, captions, or keyword tags for assets you have rights to use

### Requires review and approval

- Any AI-generated or AI-manipulated image used in public-facing work or shared with a client, unless it fits Design's approved guidance
- Any use of third-party articles, images, or creative assets beyond confirmed license terms or approved project use
- Any use that could confuse ownership, originality, or attribution for client or public-facing work

### Example workflows

Workflow A: Text content and rights

1. Confirm the source is Orchestra-created, client-provided with rights, licensed, or public domain
2. If unclear, do not upload full text or images. Use summaries and/or seek approval
3. Check for originality and edit for distinctiveness
4. Ensure any claims are supported by primary sources or client data

Workflow B: Images and pitch decks

1. Determine whether the image is AI-generated, AI-assisted edit of a licensed image, or a client-provided asset
2. For internal concepting, use approved tools and treat AI-generated visuals as draft concepts, not final assets
3. For client-facing or public-facing use, follow Design's approved guidance on AI-generated and AI-edited images
4. Seek Legal or risk review when the use involves likeness, endorsement, trademarks, logos, third-party artwork, licensed assets, or unclear ownership
5. Maintain a record of source and license for any external deliverable

Owner role: Legal or risk owner with Design leadership owner

## Approved tools and shadow AI

### Rules

1. For any confidential  internal or client information, use only Orchestra-approved tools and accounts.
2. Unapproved AI tools must not be used with client confidential data or internal confidential data.
3. For client work that does not include confidential inputs, use approved tools by default. Exceptions can be requested when there is a clear productivity benefit and low risk (per discussions with IT and Legal).
4. Browser extensions and plugins that capture typed input or page content are prohibited unless approved by IT.
5. Orchestra may periodically run "shadow AI" checks. The default response is evaluation and standardization, not punishment, but repeated willful disregard of this policy can lead to access blocks.

### Allowed on approved Orchestra tools and accounts

- Routine drafting, summarizing, research, analysis, and editing
- Creating Custom GPTs, Projects, or client/account assistants inside approved enterprise tools
- Uploading approved internal or client materials to the appropriate client workspace, Project, or GPT
- Testing new AI workflows using approved tools and approved data
- Submitting tool requests and evaluations through the AI intake process

### Requires review and approval

- Any use of non-approved tools or personal accounts for work purposes or for client work product, even if inputs seem non-sensitive
- Installing or enabling browser plugins, extensions, desktop AI apps, or local model tooling on managed devices without IT approval

### Example workflows

Workflow A Tool exception request

1. Submit a short request through the AI intake form
2. Specify the use case, data type involved, and risk level
3. IT reviews security and privacy posture
4. Policy owner and AI council decide approve, deny, or pilot with guardrails
5. AI tool terms/contract reviewed by Legal
6. Approved tools are added to the tool list and training materials

Owner role: IT and security owner with the policy owner

## Contractors and third-party partners

### Rules

1. Contractors, freelancers, temporary personnel, and approved third-party partners working on Orchestra deliverables must follow applicable AI rules in this policy.
2. With approval from IT, contractors may be granted access to Orchestra-approved AI tools when needed for their work.
3. Contractors must not use personal AI accounts or unapproved tools with Orchestra or client information.
4. Contractors must disclose AI use in any public facing deliverables, as well as when requested or required by project instructions.
5. Contracts should include confidentiality, tool restrictions, disclosure expectations, and indemnification provisions (work with Legal).

### Allowed on approved Orchestra tools and accounts

- Contractors using Orchestra-approved tools and accounts for permitted tasks, under the same rules as employees

### Requires review and approval

- Any contractor request to use personal tools, personal accounts, or non-approved services with Orchestra or client information
- Any contractor request for access to Orchestra-approved AI tools
- Any vendor that (a) will be inputting Orchestra or client confidential information into an AI tool, (b) uses a workflow that ingests, stores, or processes Orchestra or client materials using the vendor's own AI systems, or (c) will be using AI to create public facing deliverables for clients. Work with Legal to ensure that the vendor contract includes appropriate AI terms. Work with It and design/copy owners to ensure appropriate guidelines are in place covering the types of AI and use cases, including adherence to client contract requirements

Owner role: Policy owner with Legal

## AI notetaker usage

### Rules

1. Consent and trust come first. If a client or external attendee is uncomfortable with an AI notetaker, disable it for that session.
2. The notetaker should clearly announce itself and its purpose.
3. Apply data minimization. Do not record sensitive meetings when notes can be taken manually or summarized afterward.
4. Review notes before sharing internally, and never forward raw notes externally without review.
5. Follow client restrictions on confidentiality and use of AI and document them in account procedures.

### Allowed on approved Orchestra tools and accounts

- Internal meetings with clear consent and non-sensitive topics
- External meetings where consent is explicit and the content is appropriate for recording
- Sharing a cleaned summary internally after review, while keeping raw transcripts limited to the smallest necessary audience

### Requires review and approval

- Meetings covering sensitive client strategy, personnel issues, legal topics, or crisis work
- Any situation where consent is unclear or disputed

### Example workflows

Workflow A: Client call with notetaker

1. Confirm the attendee list and whether external participants will be present
2. Use Orchestra's Legal-approved notetaker disclosure or consent language at the start of the call
3. If any objection, disable for the session
4. After the call, review notes, remove sensitive items, and summarize
5. Store notes in the appropriate client workspace

Owner role: IT and security owner for tool configuration; Account leadership for usage norms

## Governance of AI assistants and custom tools (including the registry)

### Rules

1. Personal prompts and personal custom GPTs on Orchestra-approved AI tools that are used only by an individual do not require registration, but must follow this policy.
2. Any shared assistant intended for multiple team members must be registered with a named owner and basic details.
3. Registration should not become a bottleneck. Low-risk assistants built inside approved enterprise environments with no external integrations can be auto-approved once registered.
4. Higher-risk assistants require review before broad use, including assistants that:
    - Ingest client materials at scale
    - Connect to external apps or data sources
    - Are shared broadly across the company
1. Admin approval remains required for integrations with outside apps.

### Allowed on approved Orchestra tools and accounts

- Personal assistants used only by you for drafting, checklists, summarization, and similar tasks, within approved enterprise tools
- Shared assistants that are low risk, once registered with the minimum required fields and a named owner

### Requires review and approval

- Any assistant that ingests or stores client confidential materials beyond routine, case-by-case prompting
- Any assistant with integrations, automation, external data connections, or broad company-wide sharing

### Example workflows

Workflow A: Low-risk shared assistant

1. Create inside approved enterprise tooling
2. Confirm no external integrations and no cross-client datasets
3. Register with minimum fields and an owner
4. Share with the intended group and include usage boundaries

Workflow B: Higher-risk assistant

1. Submit through AI intake form with the assistant description and data plan
2. IT and security review data flows and permissions
3. Policy owner and AI council approve pilot or deny
4. Document approved scope and required review steps
5. Reassess at quarterly review

Owner role: Policy owner with IT and security owner

## Summary

This policy is meant to help Orchestra use AI with more confidence, consistency, and speed, while protecting client trust and the standards that define the work. As tools, expectations, and use cases continue to change, this framework should give teams a clear foundation for moving forward, making smart decisions, and expanding what's possible inside approved guardrails. It will continue to evolve based on real-world usage, recurring questions, and the needs of the business.