# Claude Usage Rules

**Miroma Group - All staff with Claude Teams access**

These rules apply to everyone with access to Claude through the Miroma Group account. They exist to protect confidential information, client data, personal data, the Group's reputation, and its regulatory obligations under UK GDPR and our client contracts.

Claude is a tool to support human work, not replace human judgement. Every output must be reviewed by a person before it is acted upon or shared.

**If you are in doubt about whether something is permitted, ask the AI team or Legal before you act.**

## Access & Identity

### 1. Use your work email and SSO

Log in with your company email through single sign-on. Do not use any personal Claude accounts connected to your work address.

**Why:** *This is how access is controlled, audited, and removed when someone leaves. It also keeps Group's contractual protections in place. Personal accounts sit under different terms.*

## Data & Privacy

### 2. Incidental B2B contact data is permitted; bulk personal data is not

Minimal, incidental business contact data - such as a client's name and email in a draft email, or signatory details in a contract - is permitted where the context is clearly professional and the data is not the focus of processing.

Do not upload contact lists, CRM exports, spreadsheets of client contacts, or any bulk personal data under this exception.

**Why:** *Incidental use in a professional context is low-risk. Bulk or systematic processing of personal data is not — it requires contractual and regulatory controls that are not in place for Claude.*

### 3. No HR, recruitment, or special category data

Do not upload anything covered by special category data under UK GDPR. This includes health, biometric, racial or ethnic origin, religion, politics, trade union membership, or sexual orientation.

HR and recruitment material is also prohibited: CVs, applications, performance reviews, salary records, disciplinary notes, and absence data.

**Why:** *Special category data is not covered by Anthropic's Data Processing Addendum. Uploading it is a regulatory breach, not just a policy one.*

### 4. No personal identifiers in client's customer data

Anonymise before upload. Strip names, email addresses, phone numbers, postal addresses, account numbers, IP addresses, and any other identifier that links data to a real person. Aggregated and anonymised data is permitted unless a client contract specifically restricts or limits use of AI tools – if you are unsure whether a restriction applies, check with Legal before proceeding. Pseudonymised data, where identifiers have been replaced or removed but users could still be re-identified, remains personal data and is subject to UK GDPR.

**Why:** *Under our client contracts we are a data processor and Claude would be a sub-processor. Contracts require explicit client approval before personal data goes to a third-party AI tool.*

### 5. No social media comments with names, handles or identifiers attached

If you are using Claude to draft responses to social posts, remove the name, handle and timestamp columns first, as well as any tagged users in the post. Work from the post content, then map responses back manually. Before processing this data in Claude, review your Data Processing Agreement with the relevant client. If this use falls outside its scope, ensure the relevant approval process in the client agreement is followed before proceeding. Note the restrictions in rule 4 apply in full here, including in relation to pseudonymised data.

**Why:** *A publicly visible name or handle does not mean it is fair game. The data subject belongs to the client, not to us. Inputting identifiable or pseudonymised data without contractual approval breaches our data processor agreement with clients – see rule 4 for sub-processor obligations that apply.*

### 6. Finance and expense data

Expense claims, invoices, and finance data with incidental personal information attached can be processed where necessary for legitimate business purposes.

### 7. Specific use cases requiring personal data must go through a DPIA first

If you have a use case that requires inputting personal data not otherwise permitted by these rules, you must complete a Data Protection Impact Assessment (DPIA) and obtain sign-off from Legal and the AI team before proceeding.

**Why:** *A DPIA ensures risks are identified and mitigated before processing begins, not after. Early flagging protects the business and the individuals whose data is involved.*

## Integrations & Connectors

### 8. MCP connectors require central approval

Default position: no MCP connectors. To request one, contact the AI team and your IT provider. Requests are assessed case-by-case based on use case, data scope, and risk.

Any approved connector must be scoped to a specific folder or workspace - not a full account.

**Why:** *MCP connectors give Claude live access to other systems. Without scope controls they can pull personal data, confidential client information, or HR records into a session without anyone realising. They also consume tokens at 15–80× the rate of a normal upload.*

### 9. No MCP connection to email systems

Gmail, Outlook, and any other inbox integration is blocked.

**Why:** *Email contains personal data, client confidential information, and HR matters in unpredictable combinations. There is no scope control granular enough to make this safe.*

### 10. Cowork currently not permitted

While we work out the data concerns around Cowork, particularly providing it with folder access, we ask that for now this is not used. It is currently disabled in your agency instance.

**Why:** *Cowork gives Claude access to files on your local machine. Without scope controls it can read, summarise, or surface anything in its reach, including data the user has forgotten is there. Restricting Cowork while we assess provides protection against data leaks.*

## Usage & Cost

### 11. Default to Sonnet; escalate to Opus only when needed

Sonnet handles most work. Use Opus only for deep analysis, strategy, or complex reasoning. Use Haiku for simple classification and batch tasks.

**Why:** *Opus costs 5× Sonnet per token. Defaulting to Opus is the single fastest way to burn through your agency's allocation.*

### Quality, Review & Professional Standards

#### 12. Keep prompts and outputs reviewable

Save prompts used to create client deliverables. Review every output for accuracy, tone, bias, and IP risk before it leaves the agency.

Do not send Claude outputs to clients as-is without human review and adaptation. AI-generated content must be owned by a person before it goes out under the agency's name.

#### 13. Claude does not replace professional or functional review

Claude can help you prepare, draft, or summarise material to support specialist review - but the specialist function (Legal, Finance, Compliance, HR) must do the substantive review itself.

Do not use Claude to perform a review that should be done by Legal, Finance, or another function, and then present the output as if that review has taken place.

**Why:** *Claude can make plausible-sounding errors in specialist domains. For example, a contract reviewed only by Claude has not been legally reviewed - Claude does not carry professional liability and can miss issues a qualified lawyer would catch. Professional sign-off is a control and required for our insurance purposes.*

#### 14. Claude must not be used to assess, rate, or make decisions about individuals

Do not use Claude to evaluate someone's performance, recommend a salary or pay decision, assess conduct or disciplinary matters, or reach any conclusion about an individual that would normally involve HR or a line manager. This applies even where no personal data has been uploaded - the judgement itself must come from a person, not an AI tool.

**Why:** *Claude has no knowledge of the individual, their context, or the organisation's standards. A Claude-assisted assessment is not a fair or defensible one. These decisions carry legal and human consequences that require human accountability.*

### Retention & Transparency

#### 15. Delete individual chats when no longer needed

Delete conversation histories once the work they relate to is complete and any outputs have been saved elsewhere.

Where a conversation has involved any personal data - including incidental contact data permitted under rule 2 - delete it once the work is complete. Do not retain chats containing personal or confidential information beyond the immediate working session

**Why:** *Retaining unnecessary conversations increases the risk of inadvertent data exposure and complicates compliance with data minimisation obligations under UK GDPR.*

#### 16. Consider whether to disclose AI use to clients

Some clients will want to know if AI tools were involved in producing their work. Where there is any doubt, check with your account lead before delivering AI-assisted work externally.

Do not represent AI-generated content as entirely human-authored where that would be misleading.

**Why:** *Transparency builds trust. Misrepresentation, even inadvertently, creates reputational and contractual risk.*

### Reporting

#### 17. Flag concerns, breaches, or unusual behaviour immediately

If something looks wrong, an output seems off, or you suspect data has gone where it shouldn't - tell the AI team immediately.

If personal data or confidential information is involved, also tell Legal.

**Why:** *Early flagging contains issues. Late flagging escalates them.*

**Breach of these rules may, where appropriate, result in disciplinary action under the standard Group or your employer's disciplinary procedure. If in doubt, ask the AI team or Legal before you act.**