LLM Decision Guide for Professional Services Firms

Summit Guard

A practical trust, review, and block matrix for law, accounting, advisory, and consulting firms using ChatGPT, Claude, Microsoft Copilot, Gemini, AI-enabled SaaS, and agentic workflows.

Use this guide to make consistent decisions about when AI use can proceed, when it needs human review, and when it should be blocked until governance controls are clear.

> *Governance-only disclaimer: This guide provides practical AI governance guidance. It is not qualified professional advice, certification, formal assurance, or an audit opinion. Firms should adapt it to their own obligations, client commitments, approved tools, and risk appetite with appropriate qualified review.*

The three decision states

Green: trust with routine controls

Use AI where all of the following are true:

- The tool is firm-approved or part of an approved AI-enabled SaaS platform.
- Inputs are low sensitivity and appropriate for the tool.
- The task supports internal drafting, administration, or low-risk analysis.
- A person remains accountable for the work.
- Vendor settings, retention, access, and logging are understood.
- Evidence can be produced if a client or leader asks how the tool is controlled.

Amber: review before use

Use AI only after explicit review where any of the following apply:

- Client context, confidential information, personal information, financial data, or professional judgement may be involved.
- The output may influence client advice, reports, calculations, correspondence, recommendations, or deliverables.
- The tool is approved, but the use case is new or higher impact.
- Human review expectations need to be written down.
- Approval and review evidence should be retained.

Red: block until resolved

Do not proceed until ownership, approvals, and controls are resolved where any of the following apply:

- Confidential client material would enter an unapproved tool.
- Privileged, sensitive, personal, financial, or client-restricted information may be exposed without approval.
- AI output would be sent externally or treated as final without accountable review.
- A workflow or agent can take action in firm systems without a human approval point.
- Vendor data use, retention, logging, or access controls are unknown.

Decision matrix by work type

| Work type | Green: trust | Amber: review | Red: block |

| --- | --- | --- | --- |

| Internal admin | Draft meeting notes, agendas, internal summaries, and low-risk templates using approved tools. | Summaries involving staff, client, financial, or sensitive operational detail. | Uploading confidential records into an unapproved tool or relying on AI for accountable decisions. |

| Client research | General background research, issue lists, market summaries, and question framing. | Research used to support client advice, reports, proposals, calculations, or recommendations. | Treating AI output as authoritative without source checking, context review, and accountable sign-off. |

| Client-facing work | Formatting, tone improvement, and structure support after content has been reviewed. | Draft reports, emails, analysis, or deliverables that need professional review. | Sending AI-generated content to a client without review, source checks, and responsibility for the final work. |

| Confidential uploads | De-identified examples, synthetic data, or approved firm-managed tools with known settings. | Client files, contracts, matter notes, financial data, or personal information in approved workflows only after review. | Confidential, privileged, sensitive, or client-restricted information in personal or unapproved AI tools. |

| Automated workflow actions | Draft-only automation with no external send, no system change, and clear logs. | Workflow suggestions, data updates, CRM actions, or ticket handling that require human approval. | Autonomous sending, filing, deletion, payment, client update, or record change without approval and rollback. |

| AI agents | Narrow internal helper with limited tool access, logs, and human confirmation. | Agent that retrieves from multiple systems or prepares actions for approval. | Agent that can act across systems, expose data, or complete client-impacting steps without tested controls. |

1. Internal administration

Most firms can allow low-risk administrative use when the tool is approved and the data is appropriate.

Usually green:

- Drafting an internal agenda.
- Turning rough notes into a task list.
- Improving the wording of a non-sensitive policy reminder.
- Summarising a non-confidential training document.

Move to amber when the material includes staff issues, client references, financial information, sensitive business plans, or anything that could be misunderstood as a final decision.

Block the use when staff want to paste confidential records into an unapproved tool, use personal accounts for firm work, or rely on AI to make an accountable decision.

2. Client research

AI can help shape research questions, compare themes, and identify possible issues. It should not be treated as an authority.

Green use:

- Brainstorming questions to ask a client.
- Producing a first-pass topic map.
- Summarising non-confidential background material.

Amber use:

- Research that supports advice, reporting, calculations, strategy, or recommendations.
- Research involving client-specific facts.
- Use of AI-generated citations or source summaries.

Red use:

- Relying on unsourced AI output as fact.
- Using AI to interpret client obligations without qualified review.
- Uploading client information to tools that have not been approved for that data.

3. Client-facing work

Client-facing work requires accountable human review. AI can assist drafting, structure, and clarity, but the firm remains responsible for the final output.

Before AI-assisted material is sent externally, check:

- Accuracy and completeness.
- Source support.
- Client context.
- Confidentiality and data boundaries.
- Tone and professional judgement.
- Whether the output overstates certainty.
- Whether the reviewer can explain the final position.

If the reviewer cannot stand behind the output, it is not ready.

4. Confidential uploads

Confidential uploads are the most common failure point.

Ask five questions before uploading files, text, extracts, screenshots, emails, transcripts, or data:

1. Is the tool approved for this data type?
2. Are firm and client restrictions understood?
3. Are prompts, files, outputs, and logs retained, and for how long?
4. Is data used to train or improve vendor models?
5. Can the firm evidence why this use was approved?

If the answer is unclear, treat the use as amber or red until reviewed.

5. Automated workflow actions

AI-enabled SaaS and workflow tools can move beyond drafting into action.

Higher-risk actions include:

- Sending client communications.
- Updating CRM or matter records.

- Filing documents.
- Creating invoices or payment steps.
- Closing tickets or complaints.
- Triggering approvals.
- Changing access permissions.

For these workflows, policy alone is not enough. The firm needs approval points, logs, rollback steps, and clear ownership.

6. AI agents

Agentic workflows create a different risk profile because the system may retrieve information, reason across context, and call tools.

Before allowing an agent to act, confirm:

- What systems it can access.
- What data it can retrieve.
- What actions it can take.
- Where a human must approve.
- How prompts, plans, actions, and outputs are logged.
- How an action can be paused or reversed.
- Who owns incidents and exceptions.

If these answers are not available, the agent should remain in draft-only or advisory mode.

Leadership evidence checklist

A professional services firm should be able to produce:

- AI tool and use-case register.
- Approved, restricted, and blocked use rules.
- Data handling rules for client and confidential information.
- Vendor notes covering data use, retention, logging, and admin controls.
- Human review expectations for client-facing work.
- Approval records for higher-risk use cases.
- Exception and incident log.
- Review date and owner for the decision matrix.

Practical next step

Review three real workflows:

1. One internal administration workflow.
2. One client research or client-facing workflow.
3. One AI-enabled SaaS or agentic workflow.

Classify each as green, amber, or red. Then record the missing evidence that would move the workflow to an approved state.

Summit Guard helps professional services firms turn AI use into visible governance, clear data boundaries, human review points, and leadership-ready evidence.

Contact Summit Guard: https://summitguard.com.au/contact?source=/insights/llm-decision-guide-professional-services