

AI Risk & Controls Review

ILLUSTRATIVE EXAMPLE

20-person professional services business

Scope:

Five proposed AI uses involving internal and client information

Not client work

Example of how DecisTech assesses proposed AI uses, identifies proportionate controls and recommends where not to proceed.

AI Risk & Controls Review

BUSINESS PROFILE	20-person professional services business
SCOPE	Five proposed uses of AI involving internal and client information
PURPOSE	Determine which uses can proceed, what controls are required, and which proposed uses should not proceed in their current form.

Can these proposed AI uses proceed?

The answer is different for each use.

Some uses involve limited information, straightforward human review and low consequences if an output is imperfect.

Others involve confidential information, important decisions or consequences that require significantly stronger controls.

Two proposed uses should not proceed in their current form.

The assessment therefore does **not** recommend either:

- unrestricted AI use; or
- a blanket prohibition on AI.

It recommends different decisions according to the use.

OVERALL PRINCIPLE

AI risk depends on the use, the information involved, the consequences of error and the controls around it—not on whether a tool is labeled “AI”.

Recommendation by proposed use

PROPOSED USE	RECOMMENDATION	MAIN REASON
Use 1 1 Draft routine internal material from non-sensitive information	→ PROCEED WITH STANDARD CONTROLS	Limited information sensitivity and straightforward human review
Use 2 2 Summarize client meetings containing ordinary confidential information	⇒ PROCEED WITH ADDITIONAL CONTROLS	Useful administrative purpose, but confidentiality, recording and accuracy require stronger controls
Use 3 3 Search approved internal procedures and knowledge	➔ PROCEED WITH CONTROLS	Useful if access, source quality and verification are managed
Use 4 4 Upload highly sensitive client files to an unapproved public AI service	→ DO NOT PROCEED	Information sensitivity and unverified supplier/data handling create disproportionate risk
Use 5 5 Automatically issue client-facing professional recommendations without responsible human review	→ DO NOT PROCEED	Consequences of error and required professional judgment are too high

IMMEDIATE RECOMMENDATION

Allow the first three uses only within defined working rules.

Those rules should establish:

- approved tools and account types
- what information may and may not be entered
- when human review is mandatory
- how important outputs are verified
- who remains responsible for decisions
- when supplier/data-handling review is required
- which uses are prohibited

Do not allow employees to decide individually that a tool is suitable merely because it is publicly available.

Business context

Business profile

The illustrative business has 20 employees and provides professional services to business clients.

Employees regularly work with:

- internal documents
- client correspondence
- meeting notes
- business procedures
- client files
- commercially sensitive information
- professional recommendations

The business is interested in allowing useful AI assistance without creating unnecessary restrictions.

At the same time, management does not want employees to:

- place confidential information into unsuitable services
 - rely on generated content without review
 - automate important professional judgment
 - create inconsistent practices between teams
 - assume that every AI feature in existing software is automatically appropriate
-

Current situation

Employees are increasingly exposed to AI functionality through:

- general-purpose AI services
- office software
- meeting tools
- search and knowledge systems
- features built into existing business applications

Without a clear framework, employees may make different decisions about:

- which tools are acceptable
- what information can be entered
- which outputs need checking
- when AI-generated material can be used externally

The business therefore needs practical controls rather than broad statements about whether AI is "safe" or "unsafe".

Proposed uses

The assessment considers five representative uses.

1 Use 1
Routine internal drafting

Use AI to prepare first drafts of:

- internal agendas
- routine internal summaries
- non-sensitive administrative material
- internal wording based on non-confidential information

No sensitive or confidential client information is required.

2 Use 2
Client meeting summaries

Use AI to:

- transcribe or process meeting content
- prepare draft summaries
- identify actions
- prepare follow-up points

The meetings may contain ordinary confidential client information.

3 Use 3
Internal knowledge search

Use AI to help employees:

- find internal procedures
- locate approved guidance
- summarize internal knowledge
- answer questions using a defined set of approved business information

4 Use 4
Highly sensitive client files in an unapproved public AI service

An employee proposes uploading highly sensitive client material into a publicly available AI service that has not been reviewed or approved by the business.

The employee wants the service to:

- summarize the files
- identify key information
- draft analysis

5 Use 5
Automated client-facing professional recommendations

Use AI to generate and issue professional recommendations directly to clients without responsible human review before the advice is delivered.

How the risks were considered

Each proposed use was assessed against the following factors.

Business purpose Is there a clear and useful reason for using AI?	Information involved What information must the system process?
Confidentiality and privacy Would inappropriate disclosure or processing create a meaningful problem?	Accuracy and reliability How important is it that the output is correct and complete?
Human oversight Can a responsible person review the output before it matters?	Consequence of error What could happen if the output is wrong, incomplete or misleading?
Supplier and data handling Is the service appropriate for the information being processed?	Access Could information become available to people who should not have access?
Retention How long might submitted information or generated outputs be retained?	Accountability Who remains responsible for the final action or decision?
Reversibility Can the use be stopped easily if it proves unsuitable?	

Risk overview

The ratings below are qualitative decision aids. They are not claims of mathematical precision.

CONTROL INTENSITY					
→ PROCEED WITH STANDARD CONTROLS		→ PROCEED WITH CONTROLS	→ PROCEED WITH ADDITIONAL CONTROLS	→ DO NOT PROCEED	
Proposed use	Information sensitivity	Consequence of error	Human review	Supplier/data concern	Overall recommendation
1 Use 1 Routine internal drafting	Low	Low	Straightforward	Low if approved tool used	→ PROCEED WITH STANDARD CONTROLS
2 Use 2 Client meeting summaries	Medium	Medium	Required	Medium	→ PROCEED WITH ADDITIONAL CONTROLS
3 Use 3 Internal knowledge search	Medium	Medium	Required for important use	Medium	→ PROCEED WITH CONTROLS
4 Use 4 Highly sensitive files in unapproved public AI	High	High	Does not resolve data-handling concern	High	→ DO NOT PROCEED
5 Use 5 Automated professional recommendations	Medium–High	High	Absent in proposed use	Medium	→ DO NOT PROCEED

Interpretation

Risk does not increase simply because more AI is involved.

It increases when factors such as these combine:

- sensitive information
- weak supplier controls
- limited verification
- significant consequences
- absent human responsibility

A comparatively sophisticated AI use may be acceptable in one controlled context.

A simple use may be unacceptable if it involves information or decisions that should not be handled that way.

Routine internal drafting

→ PROCEED WITH STANDARD CONTROLS

Proposed use

Employees use an approved AI tool to prepare first drafts of routine internal material using non-sensitive information.

Examples could include:

- agenda drafts
- internal administrative summaries
- meeting preparation notes
- wording for routine internal communications

Information involved

The proposed use does not require:

- confidential client files
- sensitive personal information
- commercially restricted material

PRINCIPAL RISKS

Incorrect or invented content

The tool may produce statements that were not present in the source information.

Scope creep

A low-risk use can gradually expand into employees entering more sensitive information unless boundaries are clear.

Existing controls required

- use an approved account/tool
- restrict the use to information permitted by business policy
- require the employee to review the output
- do not treat generated content as authoritative merely because it is well written

Business value

The use may:

- reduce repetitive drafting
- help employees structure information
- speed up routine administrative work

Over-reliance

Employees may accept generated wording without checking it.

RECOMMENDATION

→ PROCEED WITH STANDARD CONTROLS

This is a proportionate use where:

- information sensitivity is low
- consequences of error are limited
- human review is easy

The control burden should remain proportionate.

Do not create an approval process so heavy that the administrative control costs more than the risk it addresses.

Client meeting summaries

➡ PROCEED WITH ADDITIONAL CONTROLS

Proposed use

Use AI to help create:

- meeting transcripts where appropriate
- draft summaries
- action lists
- follow-up notes

The meetings may contain ordinary confidential client information.

Information involved

Unlike routine internal drafting, this use may involve:

- client names
- business information
- confidential discussions
- commercial plans
- decisions
- meeting participant information

The information is therefore more sensitive.

PRINCIPAL RISKS

Confidentiality

Meeting content may be sent to or retained by a service that is not appropriate for the information.

Inaccurate attribution

A system may incorrectly identify:

- who said something
- who owns an action
- what was agreed

Excessive reliance

A generated summary may be treated as the definitive record even when it contains errors.

Business value

The use may:

- reduce administrative follow-up
- improve consistency
- make actions easier to identify
- provide a useful draft record

Recording or transcription obligations

Recording or transcribing a meeting may create additional requirements depending on the circumstances.

The business should confirm any relevant notice, consent or other requirements before recording or transcription is used.

Loss of nuance

A short summary may omit:

- qualifications
- disagreement
- uncertainty
- contextual details

REQUIRED CONTROLS

Approved service

Only use a service/account that the business has approved for the type of information involved.

Human review

The responsible employee must review the summary before it is:

- sent externally
- treated as an agreed record
- relied upon for important actions

Information minimisation

Do not process more meeting information than is useful for the defined purpose.

Meeting-information rules

Define when:

- transcription is allowed
- recording is allowed
- meeting content may be processed

Action verification

Important actions and commitments should be verified against the meeting context.

Retention awareness

Understand how recordings, transcripts, prompts and generated summaries are retained where relevant.

RECOMMENDATION

➡ PROCEED WITH ADDITIONAL CONTROLS

The administrative value is credible, and responsible human review remains practical.

The use should not proceed through arbitrary public accounts or without clear rules for confidential meeting information.

Internal knowledge search

➔ PROCEED WITH CONTROLS

Proposed use

Employees use an AI-enabled search or question-answering capability to find and summarize approved internal information.

Potential sources include:

- procedures
- internal guidance
- templates
- policies
- approved reference material

PRINCIPAL RISKS

Outdated source material

The system may find or summarize information that is no longer current.

Access problems

A search capability could make information easier to discover by people who should not have access to it.

Loss of source context

A generated summary may omit limitations or conditions contained in the underlying document.

Business value

The use may:

- reduce time spent searching
- make internal knowledge easier to access
- reduce repeated questions
- improve use of existing procedures

Source ambiguity

Employees may not know whether an answer came from:

- an approved procedure
- an old draft
- informal notes
- unrelated material

Confident answers from incomplete information

The system may produce a clear answer even where the source material is:

- incomplete
- inconsistent
- ambiguous

REQUIRED CONTROLS

Approved sources

Define which information sources the system may use.

Do not index everything simply because it is technically possible.

Source visibility

Where important, employees should be able to inspect the underlying source material.

Human verification

Important decisions should not rely solely on an AI-generated answer where the underlying source can and should be checked.

Access controls

The AI search capability should respect the underlying information-access model.

AI should not become a way to bypass existing permissions.

Content ownership

Responsibility should exist for keeping important procedures and guidance current.

RECOMMENDATION

➔ PROCEED WITH CONTROLS

The use has credible value, but its reliability depends heavily on:

- source quality
- permissions
- traceability
- human verification

The main risk is not that the system can search quickly.

It is that employees may trust an answer without understanding the quality of the information behind it.

Highly sensitive client files in an unapproved public AI service

— DO NOT PROCEED

Proposed use

An employee proposes uploading highly sensitive client files to a publicly available AI service that has not been approved by the business.

The intended purpose is to:

- summarize the files
- identify key points
- draft analysis

Information involved

The files may contain highly sensitive:

- client information
- confidential business information
- personal information
- commercially restricted material

PRINCIPAL RISKS

Unknown or unsuitable data handling

The business has not established whether:

- the service is appropriate for this information
- submitted data is retained
- data may be reused
- administrators or subprocessors may have access
- appropriate contractual protections exist

Loss of control

Once highly sensitive information is submitted to an unsuitable service, the business may have limited ability to reverse the disclosure.

Human review does not solve the central risk

Even perfect human review of the generated output would not resolve the risk created by sending highly sensitive information to an unapproved service.

Business value

The proposed use could save some manual review effort.

That potential benefit does not justify the current approach.

Confidentiality

Uploading the material may conflict with obligations the business owes to the client.

Employee-created supplier decision

An individual employee would effectively be choosing a new information-processing service without the business evaluating the implications.

RECOMMENDATION

— DO NOT PROCEED

The use should not be permitted in the proposed form.

The issue is not that highly sensitive information can never be processed using AI.

The issue is that the proposed service and data-handling arrangement have not been established as appropriate for the information.

Possible future route

If there is a strong business case, the business may separately assess whether an approved service can support the use under appropriate contractual, security, privacy, confidentiality, access and retention controls.

That would be a different proposal requiring its own assessment.

Automated client-facing professional recommendations

— DO NOT PROCEED

Proposed use

AI produces professional recommendations and sends or presents them to clients without responsible human review.

PRINCIPAL RISKS

Accuracy

The system may:

- misunderstand facts
- omit relevant considerations
- invent information
- apply inappropriate reasoning

Consequence of error

Incorrect advice may directly affect:

- client decisions
- financial outcomes
- contractual decisions
- operational actions
- the business's professional reputation

Automation bias

Clients or employees may give undue weight to a recommendation because it appears complete, confident or immediate.

Business value

The proposed use could reduce professional staff time and accelerate responses.

Those potential benefits are outweighed by the risk in the current design.

Context

Professional recommendations often depend on:

- client-specific circumstances
- incomplete information
- judgment
- business context
- exceptions

Accountability

The proposed design removes the responsible human review point while still presenting the output as a business recommendation.

Why a disclaimer is not sufficient

A generic statement that “AI may make mistakes” does not replace responsible review where the business is issuing professional recommendations.

RECOMMENDATION

— DO NOT PROCEED

AI may potentially assist a professional with:

- research
- analysis
- first drafts
- identifying questions
- structuring information

The responsible professional should remain accountable for the recommendation ultimately provided to the client.

Acceptable alternative

A better model is:



not:



The same tool can present different risks in different uses

The risk assessment should focus on the use rather than assigning one universal risk level to a technology.

Imagine the same approved AI service being used for two tasks.

	Use A Lower-risk use	Use B Higher-risk use
Use and purpose	Draft an internal meeting agenda using: – public information – non-sensitive internal notes	Analyze highly confidential client records and produce a recommendation that could materially affect the client.
Information sensitivity	Low.	High.
Potential consequence of error	Limited.	High.
Human review	Easy.	Essential.
Supplier/data requirements		Significant.
Likely decision	→ PROCEED WITH STANDARD CONTROLS	Requires substantially stronger controls and may not be acceptable depending on the service and proposed process

PRINCIPLE

The tool is only one part of the risk.

The decision also depends on:

- why it is being used
- what information is involved
- who receives the output
- what happens if the output is wrong
- whether a person remains responsible
- how the supplier handles the information

Therefore:

“Approved AI tool” does not mean “approved for every use”.

Practical control framework

The business does not need a large governance bureaucracy to establish useful controls. It needs a small number of clear rules employees can actually follow.

CONTROL 1

Approved tools and accounts

Define which services and account types employees may use for business information.

Employees should not treat public availability as business approval.

CONTROL 2

Information rules

Define categories such as:

- information that may be used
- information that requires an approved controlled environment
- information that must not be entered into certain tools

The categories should use language employees understand.

CONTROL 3

Human review

Define when generated content must be reviewed before:

- external use
- decision-making
- inclusion in formal records
- reliance on important facts

CONTROL 4

Verification

Require important facts, calculations, quotations, sources, actions and recommendations to be checked where errors could matter.

CONTROL 5

Access

AI-enabled systems should not expand access beyond what employees reasonably need.

CONTROL 6

Supplier/data review

Perform proportionate review before allowing a service to process information where confidentiality, privacy or security materially matters.

CONTROL 7

Prohibited uses

Define a short list of uses the business does not permit.

Examples from this assessment include:

- highly sensitive client information in unapproved public AI services
- automatically issued professional recommendations without responsible human review

CONTROL 8

Responsibility

Employees should understand that:

Using AI does not transfer responsibility for the business decision to the tool.

CONTROL 9

Escalation

Provide a simple route for employees to ask:

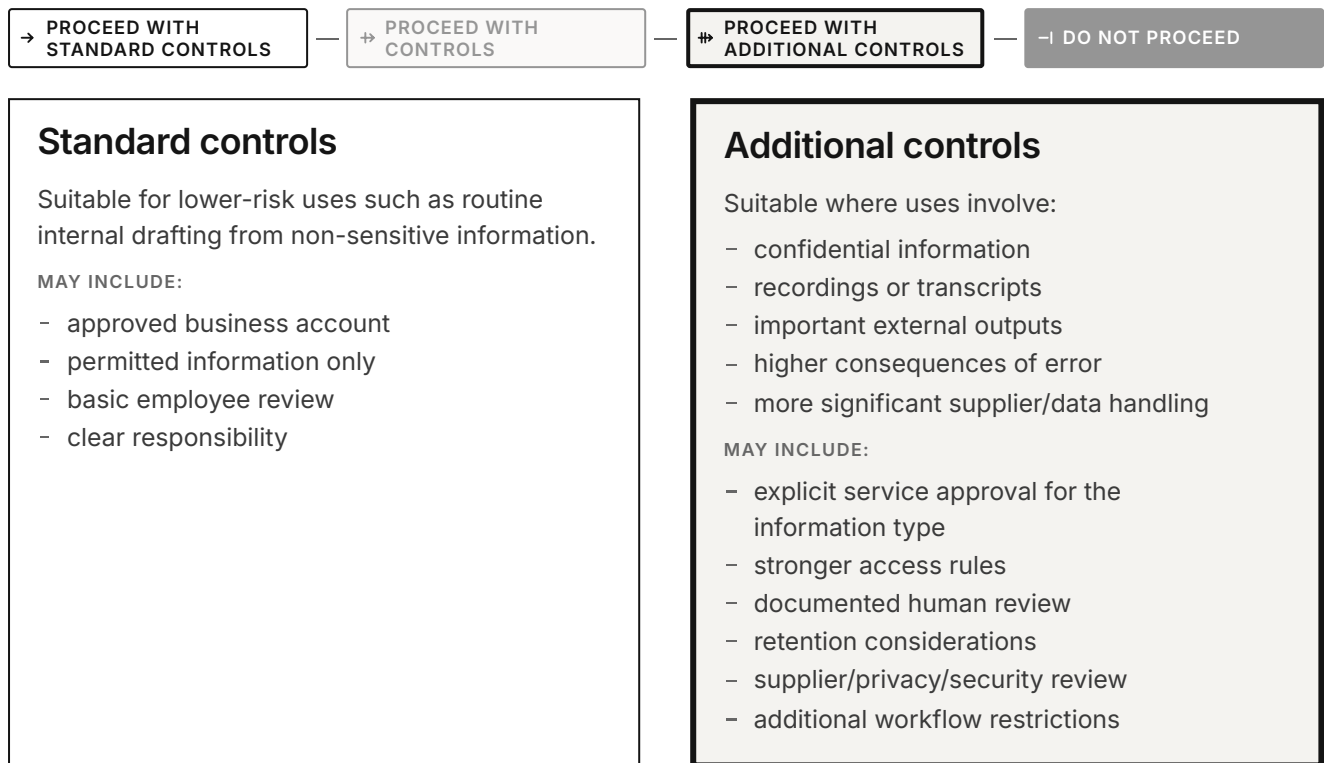
"Can I use this tool for this information and this purpose?"

when the situation does not fit the standard rules.

Standard controls versus additional controls

Not every use needs the same governance burden.

CONTROL INTENSITY



PRINCIPLE

Controls should increase with the risk of the use.

A small business should avoid both extremes:

- no meaningful controls
- controls so burdensome that employees bypass them

What should be prohibited?

The business should avoid a long list of theoretical prohibitions employees cannot remember.

A short list of clear boundaries is more useful.

Based on this assessment:

–I PROHIBITED

Unapproved handling of highly sensitive information

Do not place highly sensitive client or business information into an AI service that has not been approved for that use.

–I PROHIBITED

Intentional permission bypass

Do not use AI search or summarisation to obtain information the employee is not authorised to access.

–I PROHIBITED

Uncontrolled experimentation with client information

Do not use real confidential client information merely to test an unfamiliar AI service.

Use non-sensitive test material until the service and use have been appropriately assessed.

–I PROHIBITED

Unreviewed professional recommendations

Do not automatically issue professional recommendations to clients without responsible human review.

–I PROHIBITED

Pretending generated content is verified

Do not represent generated facts, quotations, summaries or recommendations as checked when they have not been checked.

What should require escalation?

Some proposed uses should not be automatically prohibited, but they deserve review before employees proceed.

Examples include:

⤴ new AI services	⤴ new categories of confidential information
⤴ large-scale document processing	⤴ automated client communications
⤴ automated decisions about people	⤴ AI use in high-consequence professional work
⤴ new recording/transcription practices	⤴ system connections that give AI broad access to internal information

The escalation question should be practical:

Is this materially different from an already approved use?

If yes, assess it before treating it as covered by existing approval.

Supplier and data-handling questions

Before approving an AI service for sensitive business use, the business should understand the questions that matter for that use.

These may include:

Data use

What does the supplier do with submitted information?

Training or service improvement

Can submitted business information be used to improve models or services, and can that use be controlled?

Subprocessors

Which other organizations may process it?

Security

What security controls and documentation are available?

Deletion and exit

What happens to business information when content is deleted, an account is closed or the service relationship ends?

Retention

How long may prompts, files, transcripts or outputs be retained?

Access

Who may access submitted information?

Location

Where may relevant processing occur?

Account controls

Does the business have appropriate administrator controls, authentication, access management and user management?

Contractual position

Are terms appropriate for the type of business information involved?

PRINCIPLE

The depth of supplier review should be proportionate.

A tool used only for non-sensitive internal drafting does not necessarily need the same assessment as a system processing highly confidential client information.

Recommended next actions

Now

1. Define approved AI services

Identify which services/accounts employees may use for business purposes.

Avoid an uncontrolled collection of employee-selected tools.

2. Define information boundaries

Create practical rules covering:

- non-sensitive information
- ordinary confidential business/client information
- highly sensitive information

Employees should be able to understand the categories without needing specialist legal knowledge.

3. Define mandatory human review

State clearly which outputs must be checked before:

- external use
- important decisions
- professional recommendations
- formal record creation

4. Publish the prohibited-use list

Keep it short and understandable.

5. Create an escalation route

Give employees somewhere to ask about unfamiliar uses.

Next

Review higher-risk approved uses

For uses such as meeting transcription or confidential knowledge search, confirm:

- supplier suitability
- permissions
- retention
- source controls
- human-review requirements

Provide practical employee guidance

Focus on:

- what employees may do
- what they may not do
- what requires review
- what requires escalation

Avoid turning guidance into a long theoretical AI manual.

Improve source information

For knowledge-search uses:

- remove obsolete material
- identify authoritative sources
- maintain important procedures

Later

Reassess controls when:

- new tools are introduced
- AI capabilities change materially
- more sensitive information is proposed
- automation begins making or triggering important decisions
- new legal/contractual obligations apply
- the business's use becomes broader or more consequential

Do not assume today's approval automatically covers materially different future uses.

RECOMMENDED IMMEDIATE DECISION

Permit useful lower-risk AI uses within clear business rules.

Apply stronger controls where confidential information or higher-consequence outputs are involved.

Do not allow highly sensitive client information into unapproved public AI services.

Do not automate client-facing professional recommendations without responsible human review.

How AI uses were assessed

DecisTech assesses the proposed business use rather than assigning one generic risk rating to “AI”.

The assessment considers:

Business purpose What useful outcome is the business trying to achieve?	Information involved What information must the system receive or access?	Confidentiality and privacy What would happen if the information were processed or disclosed inappropriately?
Accuracy and reliability How important is output quality?	Human oversight Does a responsible person remain in control before the output matters?	Consequence of error What could happen if the system is wrong?
Supplier and data handling Is the service appropriate for the information and use?	Access Could the use expose information more broadly than intended?	Retention What information may persist and for how long?
Accountability Who remains responsible for the action, output or decision?	Reversibility Can the business stop or change the use if it proves unsuitable?	

ASSESSMENT PRINCIPLE

Proceeding with controls does not mean a use is risk-free.

It means the remaining risk can be managed proportionately for the proposed use.

ASSESSMENT PRINCIPLE

A recommendation not to proceed is not a judgment that AI is inherently unsafe.

It means the particular use, information, process or control environment does not justify proceeding in its proposed form.

The objective is not to eliminate every possible risk.

The objective is to make deliberate, proportionate decisions about where AI can be useful without transferring unacceptable risk to the business or its clients.