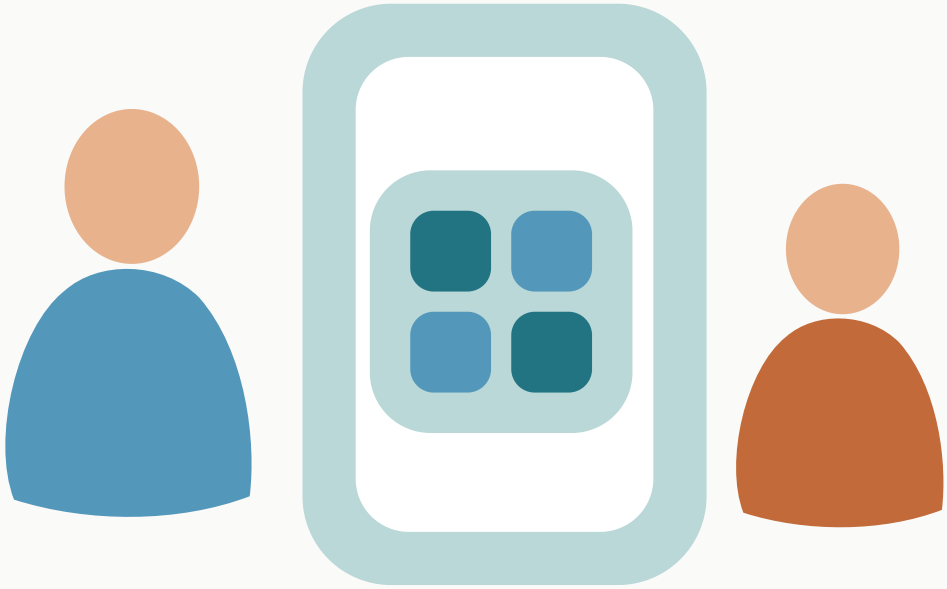
Clowry AI

SHORT GUIDE

Delegating work to AI agents

Six principles



Tom Clowry · Clowry AI

November 2026

Clowry AI is the consultancy trading name of Considered Ideas Ltd

clowryai.com

Introduction

AI agents can pursue a goal over several steps, selecting tools and using the results to decide what to do next. Using agents involves decisions about what to delegate, how to organise the work and whose interests it should serve. These decisions require human judgement, and we remain accountable for their consequences.

That raises practical questions. What should you delegate? How much authority should the agent have? How will you judge work that you did not do yourself? And if agents take over the work through which people learn, how will you and your colleagues develop the expertise you need?

This guide helps you work through these questions. It is intended for people using agents in their own work or responsible for their use across a team or organisation.

Sharing and adapting

You may copy, share and adapt this guide under CC BY 4.0. Credit Tom Clowry, include links to the guide and the licence, and indicate if you have made changes.

WHAT IS AN AGENT?

An AI agent combines an AI model with software tools. Depending on its configuration and permissions, it can search records, update a system or send a message.



The six principles

The guide sets out six principles, and a design record for setting down your decisions.

PRINCIPLE	WHAT IT ASKS	WHAT TO DECIDE
P1 **Purpose**	What should improve for the people who use or do this work?	The improvement you want, who benefits and how you will recognise it.
P2 **Context**	What information and instructions does the agent need, and who can resolve gaps or ambiguities?	The permitted sources, what they cannot show and who resolves questions.
P3 **Authority**	What can the agent do without further approval, and what needs a person's decision?	What the agent can do, what needs approval and how the work can be stopped.
P4 **Review**	How will you judge whether the agent's work is accurate, relevant and appropriate?	Who checks what, with which evidence and how much time.
P5 **Challenge**	What would make you reconsider the answer?	The assumption to test and the evidence that would change the decision.
P6 **Learning**	How will you and your colleagues develop and maintain the expertise to judge the agent's work?	The expertise to retain or develop, and the practice and feedback needed to support it.

Before giving an agent information or permission to act, read 'Information you can supply' under Principle 2 and 'Putting changes into practice' in the complete guide.

What this guide covers

The guide is non-technical. It focuses on how work with agents can be organised and governed, rather than how they are built or configured. Configuring an agent's accounts, system access, connections, logging and testing requires specialist technical expertise.

SCOPE AND LIMITATIONS

This guide provides general information on organising work with AI agents, not advice tailored to a particular organisation, system or decision. It is not legal or regulatory advice.

Its examples, checklists and templates do not establish that a proposed use is safe, lawful or suitable. You need to assess and test the proposed use in your own circumstances. Following the guide does not by itself meet the requirements of any law or standard.

The full notice is on page 32.

Using this guide

Start with an area of work in which you are considering using AI agents. Use the six principles to assess what should change, and record your decisions and unresolved questions in the design record.

Involve colleagues who carry out or oversee the work and, where appropriate, people who use the service or rely on its results. Include those responsible for authorising the work and for the relevant technical, information security, data protection and legal decisions.

The principles are related. Changing what an agent can do, for example, can change the checks needed. Revisit your decisions as you develop the proposal and review the work in practice.

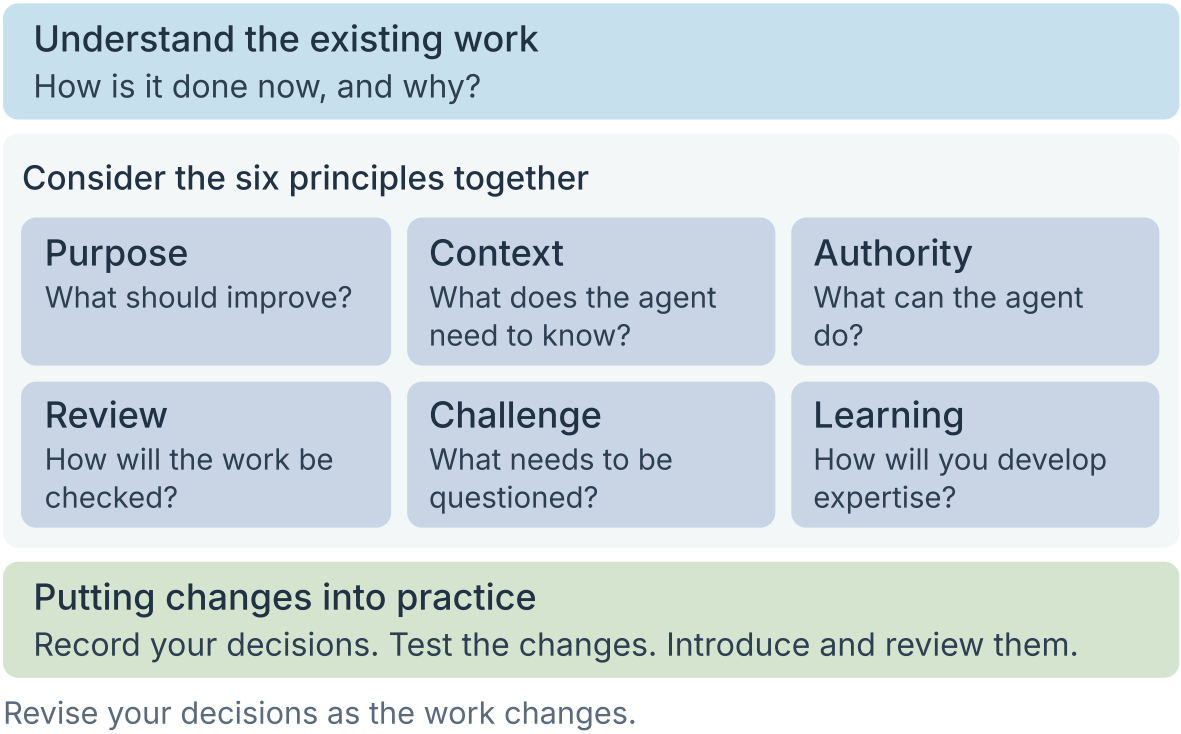


Figure 1 How to use the six principles when reorganising work

PRINCIPLE 1

Purpose

What should improve for the people who use or do this work?

P1

If an agent completed this task in minutes, what would you change about the rest of the work?

Decide what the work should achieve and for whom. Then consider how delegation could improve the way it is done.

WHY IT MATTERS

Has the work as a whole improved?



A task that finishes sooner has not necessarily improved the work. A quicker reply is of little use if it leaves someone unable to act, and more frequent analysis adds little if nobody has time to read it.

So begin with the people who use the work. What should be better for them, and how will you know?

KEY POINT

Judge the improvement across the whole piece of work.

IN PRACTICE

Five steps

01 Define the improvement

Be specific about what you want to change. Fewer errors, shorter waits and more time for difficult cases each require different arrangements. State who should benefit, and how you will recognise the improvement.

02 Check for side effects

Check for adverse effects as well as improvements. Would the change create work or inconvenience elsewhere, including in related processes? An agent writing software changes can produce them more quickly than their outcomes can be properly assessed. An agent updating records can spread a mistake to every system and report that uses those records before anyone notices it.

03 Ask whether the work should be organised differently

If an agent can produce an analysis quickly, the team can discuss early results, explore alternatives and request revisions before deciding.

Agents can also make it practical to examine more cases or adapt explanations for different audiences. Allow time for colleagues to use and discuss the results.

04 Decide whether delegation is worthwhile

Three questions help here. Is the agent likely to meet the standard you need? Will the result arrive in time? How long will it take to review? How long the review takes depends on who carries it out. Understanding the subject helps a reviewer judge whether the result is accurate, relevant and appropriate. Estimate the review time for the person who will actually do it.

Then compare the cost of using the agent with the cost of doing the work without it. Include the time spent briefing the agent, reviewing and correcting its work, and any attempts that fail. Even where the agent would save time or money, you also need to weigh what the improvement is worth and what a mistake would cost.

05 Decide with those involved how the work should change

With the people doing the work, go through what it involves and agree what continues, what changes and what stops. If delegation releases time, agree with them how it will be used.

Likely to succeed	Time the agent takes	Time to review	Outcome
Often falls short	Quick	Quick	Of little use, as the work is redone
Reliable	Too slow for the deadline	Quick	Arrives too late to help
Reliable	Quick	As long as the work itself	The time saved is spent checking
Reliable	In time	Manageable	**A candidate for delegation**

Figure 2 Three practical considerations when assessing delegation. Drawing on Ethan Mollick, Management as AI superpower, One Useful Thing, 27 January 2026.

RISKS

When objectives conflict

An objective also reflects assumptions about what is desirable and fair, and two authorised objectives, such as reducing cost and improving service, may conflict. Agree who will weigh them, and consider whose interests are affected.

EXERCISE

Examine a proposed improvement

1. Choose an activity that an agent could perform more quickly. Explain what that would allow someone else to do better. If the next stage would remain unchanged, consider whether it also needs to be reorganised.
2. Identify an activity that could take place earlier, be done differently or be discontinued. Check the implications with the people who do it and those who depend on it.

IN YOUR DESIGN RECORD · PURPOSE SECTION

Describe the improvement you want, who should benefit and how you will judge the result. Explain the proposed changes to the work and any additional responsibilities they create.

PRINCIPLE 2

Context

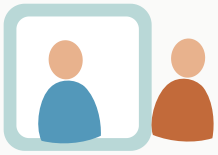
What information and instructions does the agent need, and who can resolve gaps or ambiguities?

What would an experienced colleague notice that the written instructions leave out?

Provide the information and guidance the agent needs to do the work. Arrange for a colleague to answer questions that the available material does not resolve.

WHY IT MATTERS

What the procedure and the records leave out



Procedures and records may not contain all the knowledge that a task requires. An experienced colleague can recognise an unusual request, remember why a rule was introduced or follow a practice that has never been written down.

An agent can draw on its training and the information available to it, but it may not recognise an unwritten practice or understand why it matters. Explain relevant practices and exceptions in the brief you give it. Some cases will still need a person's judgement, such as an unusual request. State which of these cases the agent must pass on, who will respond and what work can continue in the meantime.

KEY POINT

Explain relevant unwritten practices and exceptions in the brief.

IN PRACTICE

Four steps

01 Use only approved information

List the information the task needs. Use only information that your organisation has approved for this AI service, and supply no more than the task requires. Use an account that your organisation provides or has approved, not a personal one.

Do not enter, upload or give an agent access to personal or confidential information unless your organisation has explicitly approved its use with that AI service for that purpose. If you cannot confirm that approval, do not proceed.

02 Examine a case that required judgement

Ask someone to show you a recent case that initially looked routine. Working through it can reveal knowledge or checks that a written description of the process does not include.

What did they notice? Why did it matter? What did they check before acting? What should the brief say? Turn the answers into instructions.

03 Improve the instructions

Use what you learn to improve the instructions. If two policies appear to conflict, clarify which applies before the agent proceeds. If a record covers only part of the service, explain what it leaves out. Include source dates so the agent and the reviewer can distinguish current guidance from an earlier version.

Add each explanation to the brief, noting who supplied it and whether it applies to future work.

04 Allow time for questions

Repeated questions can increase a colleague's workload even when the agent works quickly. Allow time to answer them, group those that can wait and pause work that depends on an unresolved answer.

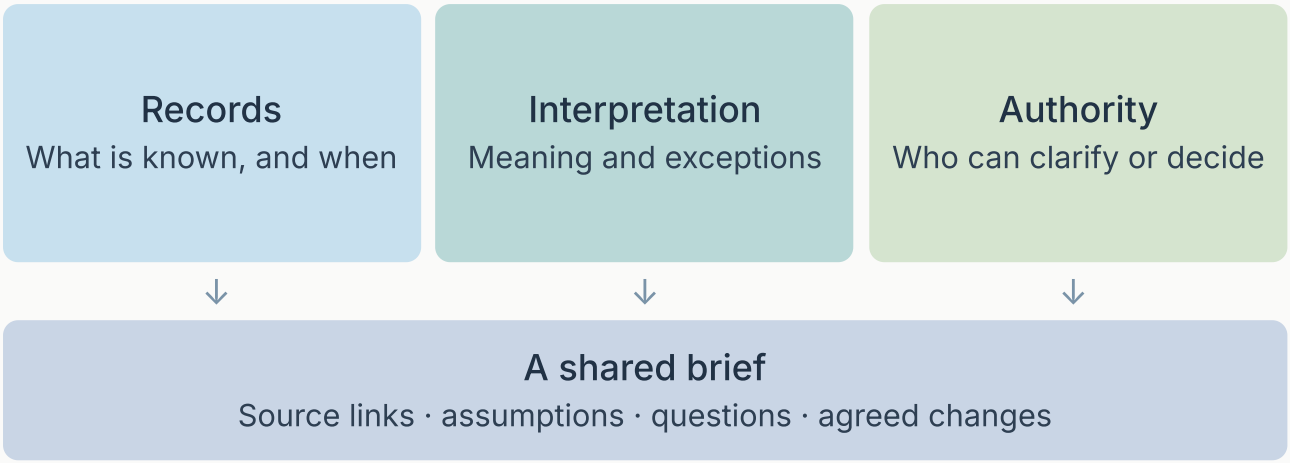


Figure 3 Information and explanations kept together in the brief

RISKS

Check the interpretation before acting

An agent can produce a plausible interpretation from incomplete information. Where a decision depends on knowledge that has not been recorded, someone familiar with the work should check the interpretation before the agent acts on it.

EXERCISE

Check what the instructions leave unexplained

1 Choose a case that required an experienced colleague to depart from the usual process. Identify what could have been explained in advance and what required advice from a colleague.

2 Write one instruction that would help an agent recognise when to ask.

IN YOUR DESIGN RECORD · CONTEXT SECTION

List the permitted sources, who authorised their use and what they cannot tell you. Name who can answer questions, who keeps the guidance current and which questions must be resolved before the work continues.

PRINCIPLE 3

Authority

What can the agent do without further approval, and what needs a person's decision?

Agree what the agent can do without further approval. State when it needs a person's decision and what it should do while it waits.

P3

WHY IT MATTERS

Proposing an action and implementing it have different consequences



A proposal can still be revised. An action that has been implemented can commit the organisation, affect other people and be difficult to reverse. For example, a reply can be revised until it is sent.

KEY POINT

Specify which actions the agent can take without further approval and which need approval first.

IN PRACTICE

Five steps

01 Examine the actions the task involves

Before giving an agent a task, examine the actions it would involve. Permission to prepare a recommendation need not include permission to act on it. Some actions require further approval because of their consequences for other people.

Who could be affected by a mistake? How serious could its effects be? Could they be put right? The actions that need prior approval, and those that require closer checks, should be agreed by those responsible for the work and for the agent.

02 Set out the permissions in the brief

The brief should tell the agent the result required, which information and tools it can use, what it can change, when to ask for help and what records and evidence to provide. Include an approved example of the required standard where you can.

03 Match the agent's access to the work

Being able to use a tool does not mean having permission to use it for every purpose. Give the agent access that matches the work that has been authorised. An agent does not always follow an instruction as intended. Where an error could have serious consequences or be difficult to reverse, do not rely on instructions alone. Enforce limits through access controls or approval requirements.

Before use, test that approval requirements, access restrictions and stopping arrangements work as intended.

04 Permission can be given in advance

A person need not approve every small step. Where testing shows that the agent meets the agreed requirements, it can be given permission to repeat the task under stated conditions. Agree in advance when the agent can act, how many actions or attempts it can make, how the result will be checked and who can stop the work. Establish which actions can be stopped or cancelled, and how completed actions will be checked and errors corrected. Experience can justify extending, retaining or reducing permission.

05 Decide whether one agent can pass work to another

An agent can be configured to pass part of a task to another agent. State in the brief whether this is allowed. Any permission given to the second agent should not exceed what the first agent was given for that task, and the record should show which agent carried out each action.

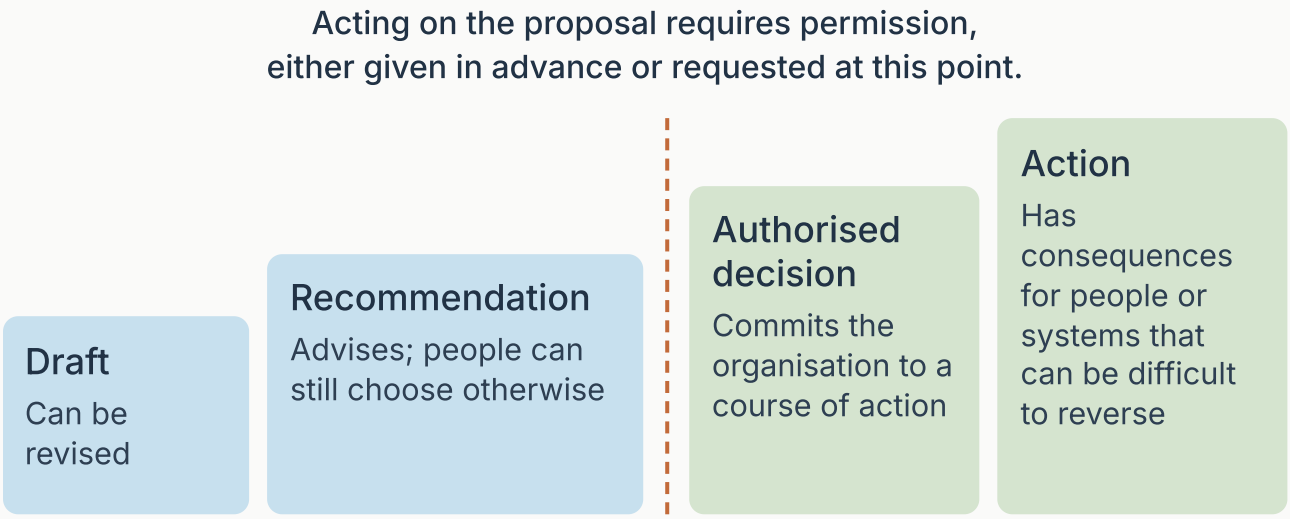


Figure 4 Different permissions for preparation, decisions and action

P3

RISKS

Material that the agent reads can contain instructions

An agent can encounter instructions embedded in emails, documents or web pages that conflict with the instructions it should follow. Attempts to redirect an AI system in this way are a form of prompt injection. The material can come from inside or outside the organisation. The UK National Cyber Security Centre warns that prompt injection may never be fully mitigated. Limit what the agent can access and do if it follows such instructions.

National Cyber Security Centre, Prompt injection is not SQL injection (it may be worse), 8 December 2025.

RISKS

Notes an agent keeps between tasks

An agent can keep notes, lessons or preferences from earlier work and use them in later tasks. It can draw on these as well as its instructions. Keep them visible, decide who can change them and review them with the brief.

EXERCISE

Check when the agent needs approval

1. Choose an action the agent is allowed to take and suppose that a condition for taking it has not been met. State what work can continue, what must stop and who can decide how to proceed.
2. Check the instructions against a case involving incomplete information or conflicting requests. An approaching deadline should not give the agent permission to exceed its authority.

IN YOUR DESIGN RECORD · AUTHORITY SECTION

Describe what the agent can do without asking, what needs approval and what it must do when a condition is not met. Explain who could be affected by an error and whether its effects could be put right. Include how someone can stop the work and check what has already happened.

NOTES

PRINCIPLE 4

Review

How will you judge whether the agent's work is accurate, relevant and appropriate?

Check the agent's work at the points where your judgement is needed. Anyone checking it needs the evidence, time and authority to act on what they find.

WHY IT MATTERS

When approval is not a meaningful check



Approval is not a meaningful check if you cannot examine the evidence needed to make the decision. Agree what needs to be checked, by whom, and allow enough time to do it.

Madeleine Clare Elish describes a "moral crumple zone" in which a person bears responsibility or blame for an automated system despite having limited control over it. Assigning someone responsibility for an agent's work is therefore insufficient unless they also have the information, time and authority to intervene.

Elish, M. C. (2019), Moral crumple zones, Engaging Science, Technology, and Society, 5, 40–60.

KEY POINT

Approval requires access to the evidence needed to judge the work.

IN PRACTICE

Five steps

01 Decide what needs your attention

You cannot check everything, so decide which conclusions and actions need your attention, and allow the time and information to assess those carefully. Decide which situations always need attention, such as conflicting sources or a proposed action outside the brief.

02 Place checks where they can make a difference

Place checks where they can prevent an error or change a decision. A single approval at the end can come too late. Examine the instructions before the agent begins, decide unusual cases as they arise, check calculations and important claims before the result is used, and afterwards inspect a selection of apparently routine work.

03 Provide the evidence for review

You should be able to examine the original evidence rather than depend entirely on the agent's account of its own work. For a recommendation, provide the sources, important assumptions and alternatives considered. For an action, retain what the agent was asked to do, which permission authorised the action and what changed. Checks must also be capable of finding errors that the agent has not flagged.

04 Use a defined standard where one exists

Some work can be checked against a defined standard. Software can be run against its tests, and a calculation can be compared with a known result. Where such a check exists, an agent can run it, correct its work and run it again before you see the result. Set a limit on repeated attempts and require the agent to report unresolved failures. Changes to the agreed acceptance criteria need separate approval. You still need to decide who sets the check and inspect a selection of the results. Other judgements, such as whether a recommendation suits the people affected, cannot be settled by an automated check alone.

05 Account for everyone's time

A task can finish sooner while requiring more human effort. Include time spent briefing, answering questions, checking, correcting and maintaining sources when you judge whether the change is worthwhile. Look at whose time is being used as well as the total.

If the agent handles the easy cases and one colleague receives only the difficult exceptions, their work can become harder even though they handle fewer cases. Discuss with them what support and variety the role needs.

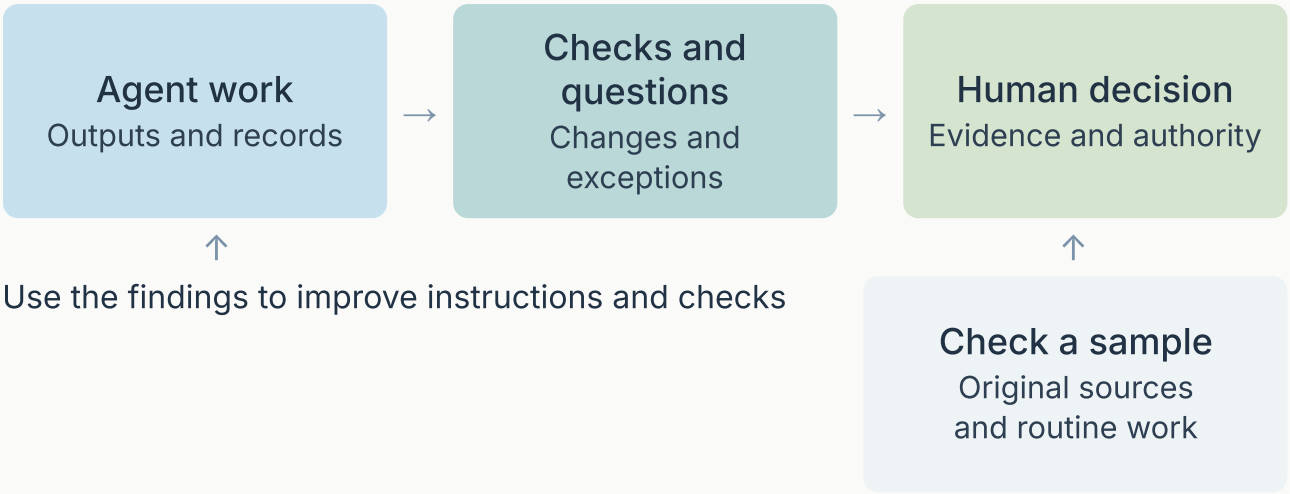


Figure 5 Reviewing the agent's work before a decision

RISKS

What the reviewer does not see

A reviewer who sees only the agent's output may miss relevant information that the agent omitted. From time to time, carry out a piece of the work without the agent and compare the results.

P4

EXERCISE

Ask someone to review a real result

1. Ask a colleague to review a completed piece of the agent's work. Can they establish what it did, inspect the supporting evidence and identify anything that needs correction? Check whether they have enough information and time to make that assessment.
2. If they have to repeat the work to assess it, improve the records and explanations provided. If the review happens after someone has already relied on the result, consider which checks need to take place earlier.

IN YOUR DESIGN RECORD · REVIEW SECTION

State what each reviewer must assess, the evidence they will receive and the time allowed. Explain what happens if they disagree or cannot respond, and how you will check work in which the agent has reported no problems.

PRINCIPLE 5

Challenge

What would make you reconsider the answer?

If two agents agree, what evidence supports their conclusion?

Decide which conclusion or assumption needs testing, and what you would do differently if it proved wrong.

WHY IT MATTERS

Review and challenge



Review establishes whether the work is reliable enough to use. Challenge examines whether its assumptions, interpretation or proposed approach should change. The same assessment can involve both.

Two agents can agree because the evidence is persuasive. They can also agree because both started with the same assumption or missed the same information. Agreement alone does not tell you which has happened.

KEY POINT

Agreement is not enough to establish that a conclusion is justified.

IN PRACTICE

Six steps

01 Decide how much scrutiny the decision deserves

How much scrutiny a decision deserves comes down to the potential cost of a mistake and whether another perspective would meaningfully alter the choice. If neither justifies further analysis, it may be better to decide and move on.

02 Choose a check that could resolve the doubt

Choose a check that could resolve the specific doubt. For a factual claim, check it against its source. For a recommendation, consider another explanation of the evidence or examine whether the proposal is practical.

03 Choose how to examine the work

For a narrow question, one analysis followed by a check may be sufficient. A complex proposal needs examination from several perspectives. If you ask an agent for a separate assessment, withhold the preferred conclusion so that the agent is not instructed to justify it.

04 Add perspectives when they could change the decision

Separate analyses can reveal different assumptions, but they can still rely on the same incomplete evidence. Use them when another assessment could help resolve a question. Where information is missing, gathering evidence or speaking to the people concerned may be more useful.

05 Identify what would change your view

An instruction to be critical gives an agent little direction. A more useful brief asks which assumption the recommendation depends on, what evidence could contradict it, what should be done differently if that evidence appears and how a credible alternative compares.

06 Consider what kind of disagreement it is

A dispute about a cost estimate can be resolved by better evidence. A disagreement about how much extra cost is justified by better access requires someone to weigh those priorities. Counting how many agents favour each option will not resolve it.

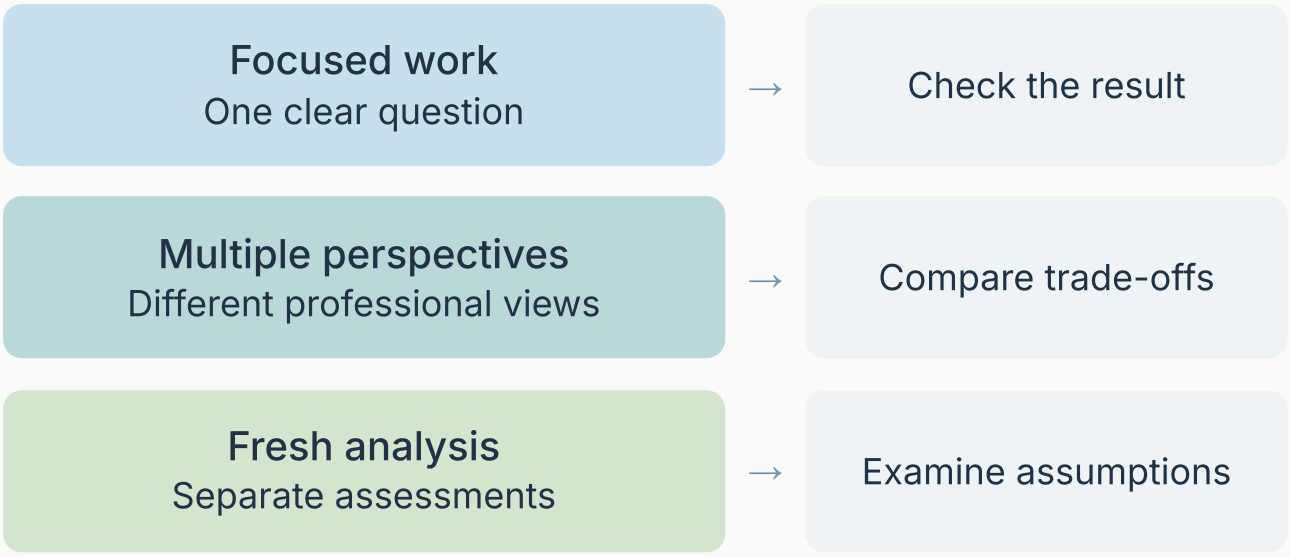


Figure 6 Three ways to examine an analysis

RISKS

Agreement between models may add less evidence than expected

A preprint published in 2026 examined nine AI models from seven model families on benchmark judgement tasks. Their errors were correlated, reducing the panel's effective information to about two independent votes. In these evaluations, the panel offered no meaningful improvement over the best individual model. The finding illustrates why agreement between models should not be treated as independent confirmation.

Kohli, G. (2026), Nine judges, two effective votes, arXiv:2605.29800 (preprint).

P5

EXERCISE

Decide what would change the answer

1 Choose a recommendation and identify one plausible finding that would make you reconsider it. Decide whether to investigate by checking a source, requesting another analysis, speaking to someone or testing the proposal in practice.

2 If further analysis would not resolve the remaining question, explain what evidence you need to gather next.

IN YOUR DESIGN RECORD · CHALLENGE SECTION

Record the assumption you need to test and how you will test it. Explain what evidence is needed before a decision can be made, and record any disagreement for those responsible for the decision to consider.

NOTES

PRINCIPLE 6

Learning

How will you and your colleagues develop and maintain the expertise to judge the agent's work?

When AI performs the work through which expertise has traditionally developed, how do people learn to judge the result?

You and your colleagues need knowledge and experience to assess an agent's work. Plan how you will acquire and maintain both as the division of work changes.

WHY IT MATTERS

The expertise paradox



AI can replace the expertise needed to perform a task while leaving us dependent on that expertise to make good judgements about the result. People have traditionally developed that judgement by doing the work AI is now performing for them.

People often learn through the work itself. They make an initial assessment, discover what they missed and discuss the result with someone more experienced. If an agent takes over that work, the opportunity to learn may change too. This matters at every career stage.

P6

KEY POINT

People need practice to retain the expertise that oversight depends on.

IN PRACTICE

Six steps

01 Agree with each colleague what they need

Someone beginning a role and an experienced colleague can both be asked to assess an agent's work, but they need different opportunities to maintain and extend their knowledge. Someone beginning a role needs experience of ordinary cases. An experienced colleague needs to understand changes in the work and encounter unfamiliar situations.

02 Allow time to form a view first

Relevant expertise helps a colleague assess an agent's draft, but it does not remove the risk of accepting its framing. Forming an initial view before reading the draft can be useful for experienced colleagues as well as those learning the work. Discussing where and why the answers differ can help them understand the subject more fully.

03 Practise on ordinary and difficult cases

You and your colleagues should practise on ordinary and difficult cases, compare different approaches and examine examples in which an agent's work is incomplete or wrong, with feedback from someone who can judge the work.

04 Use agents to support learning

Agents can also support development. They can help someone explore a different approach, prepare practice exercises or rehearse how to explain a difficult decision. Where the learner cannot yet judge the material, an experienced colleague should check it and help interpret the results.

05 Consider what depends on a human relationship

Some parts of the work depend on a human relationship. A service user may need someone to listen, explain a decision or take responsibility for resolving a difficulty. Agent support can help the colleague advising them prepare and follow up. Consider what the service user needs from the person advising them before changing who provides that advice.

06 Make development part of the working day

Agree time for practice and feedback within the working day. Explain to colleagues how their work will change and how they will learn the parts that are new to them. Make clear where colleagues can obtain help and allow time for questions.

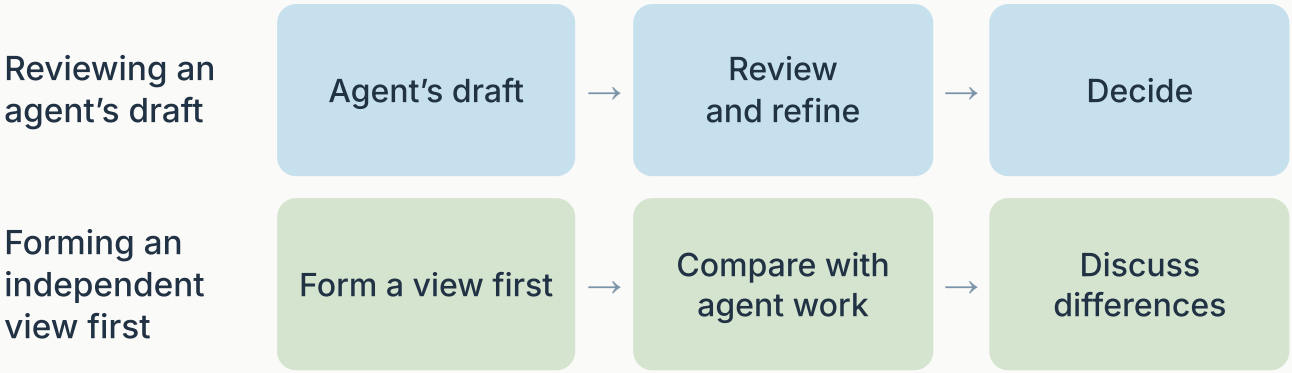


Figure 7 Two approaches to reviewing and learning from an agent's work

Choose the sequence according to what each colleague needs to learn.

RISKS

AI assistance can affect learning during a task

In a randomised study reported in a 2026 preprint, 52 developers learnt an unfamiliar programming library with or without an AI assistant. Both groups then took a test without AI. Those who had used AI during the task scored lower on average, suggesting they had learnt less from the work. Participants who asked conceptual questions or sought explanations tended to score higher than those who delegated the coding. These patterns were observed, rather than assigned experimentally, so they do not establish cause and effect.

The practical question is whether using AI helps people develop the understanding they need to assess its work. This study examined learning during a short task, not long-term expertise.

Shen, J. H. and Tamkin, A. (2026), How AI impacts skill formation, arXiv:2601.20245 (preprint).

EXERCISE

Plan how you and your colleagues will develop your judgement

1 Identify one judgement that you or a colleague will still need to make after delegation. What experience currently helps develop it? Decide how that experience will continue, or what you will put in its place.

2 Agree how you will recognise progress. For example, can someone explain an alternative, identify an unsupported conclusion and say when they need further help?

IN YOUR DESIGN RECORD · LEARNING SECTION

Describe the expertise you and your colleagues need, how you will practise and who will provide feedback. Review the division of work as you gain experience and responsibilities change.

NOTES

Design record template

Use this record to agree what work to delegate to an AI agent and how that work will be organised and overseen. Record your decisions under each principle, note unresolved questions and update the record as the work changes.

Completing this record or checklist does not establish that a proposed use is safe, lawful or suitable. Assess and test the proposed use in your own circumstances.

Use the spaces below for handwritten notes, or download the Word template to complete the record digitally.

Download the Word template

Work being considered	Person responsible
AI service and model used	Version or date of the instructions, permissions and tests
Date	Next review date

Purpose

What should improve, for whom, and what would people do differently? How will you judge the result?

Context

What information can the agent use, and who has authorised this? What relevant practices and exceptions should the brief explain? How will any newly collected information be handled? Who can resolve gaps or ambiguities? Which questions must be resolved before the work continues?

Authority

What can the agent do without further approval? Who could be affected by an error, and can its effects be put right? What must the agent do when a condition is not met? What needs approval? Who can approve those actions? How can the work be stopped?

Review

Who will assess whether the work is complete, accurate, relevant and appropriate, and what evidence will they use? When do they check, and what time and authority do they have? What happens if they disagree or cannot respond? How will you check work in which the agent reported no problems?

Challenge

Which conclusion or assumption needs testing? How will you test it, and what evidence would be sufficient to decide?

Learning

What expertise will you and your colleagues need? How will you practise and receive feedback?

Putting changes into practice

How will you test permissions and stopping arrangements before use, and establish whether the change works? Who will review the arrangements in use? What would lead you to change or stop them, and where will problems and corrective action be recorded?

The complete guide

This short guide contains the six principles and a design record template. In the complete guide, each principle follows the same order: why it matters, background, in practice, examples, risks and an exercise, with a further section where a principle needs one.

The complete guide provides fuller guidance on personal and confidential information and on putting changes into practice. Read both in the complete guide before you give an agent access to information about people or permission to act. The complete guide also covers what an agent is, what it works with, how it carries out work and what it is used for, and includes a checklist, a glossary and references.

The complete guide is free to read at clowryai.com/delegating-to-agents.

About the author

Tom Clowry is an independent AI specialist and adviser. His work focuses on applying AI, with particular expertise in agents and human capability.

If you would like to discuss the guide, please email enquiries@clowryai.com or contact Tom on LinkedIn.

Edition	First edition, November 2026. Revised November 2026.
Licence	© 2026 Considered Ideas Ltd. Licensed under the Creative Commons Attribution 4.0 International licence (CC BY 4.0). You may copy, adapt and share it, provided you credit the source, link to the licence and state any changes you have made.
Suggested credit	Tom Clowry, Delegating work to AI agents: short guide, Clowry AI, 2026. clowryai.com/delegating-to-agents.
Publisher	Clowry AI is the consultancy trading name of Considered Ideas Ltd. Registered in England and Wales, Company No. 17258515. Registered office: 66 Paul Street, London EC2A 4NA.

Scope and limitations of this guide

This guide provides general information on organising work with AI agents, not advice tailored to a particular organisation, system or decision. It is not legal or regulatory advice.

Its examples, checklists and templates do not establish that a proposed use is safe, lawful or suitable. You need to assess and test the proposed use in your own circumstances. Following the guide does not by itself meet the requirements of any law or standard.

The guide is provided in good faith, using information available at the time of publication or revision. AI systems, provider terms and legal requirements can change. Before acting, check the current requirements that apply to your organisation and country, confirm what the system can do and test the proposed use. You are responsible for how you apply the guide and should obtain professional advice where appropriate.

Reading, downloading or using the guide does not by itself create an adviser–client relationship with Tom Clowry or Considered Ideas Ltd.

To the fullest extent permitted by law, Tom Clowry and Considered Ideas Ltd accept no liability for loss or damage arising from use of, or reliance on, this guide. Nothing in this notice excludes or limits liability for death or personal injury caused by negligence, fraud or fraudulent misrepresentation, or any other liability that cannot lawfully be excluded or limited.