



Implementing AI in day-to-day work

Introducing practical AI assistance into everyday work without creating vague expectations or unmanaged shadow use.

WORKS WITH THE ASSISTANTS TEAMS ALREADY USE



EDITION 28 / EMBER & AMBER

EXECUTIVE WORKING EDITION - 08 OCT 2026

28

AT A GLANCE

Introducing practical AI assistance into everyday work without creating vague expectations or unmanaged shadow use.

Most people already use AI at work, often quietly and without guidance. That leaves managers with two problems: vague expectations about what good use looks like, and unmanaged use of tools with company information. This playbook gives individuals and team leads a practical way to bring AI into planning, writing, meeting follow-up, research and communication. It covers the tasks that fit, the information rules, request templates, verification habits, simple measures of time saved and routines for sharing what works.

WHO IT IS FOR Team leads, managers and individual contributors who want everyday AI use to be productive, consistent and safe.

TIME TO USE 50 minutes to read; one team session to apply

What you will have at the end

- A map of everyday tasks that suit AI help and the checkpoint each one needs
- Written expectations your team can follow, replacing guesswork and hidden use
- Request templates for writing, meeting follow-up, research and planning
- A short verification habit that catches the typical errors
- A way to measure time saved and share working methods

Contents

01	Why everyday AI use needs a deliberate approach	03
02	Which everyday tasks suit AI and where you must stay in charge	04
03	Information rules and approved tools for everyone	05
04	Writing, meetings, research and planning, step by step	07
05	A short checking routine for everything that leaves your desk	09
06	Knowing whether it saves time and keeps quality	11
07	Sharing what works and keeping expectations clear	12
08	Where everyday adoption goes wrong, with one illustrative scenario	14
09	A first quarter that produces evidence	15
10	Terms used in this publication	16

EXECUTIVE BRIEF

Why everyday AI use needs a deliberate approach

01



The value of AI at work comes from many small, repeated tasks done a little faster and a little better, which is exactly why it needs shared habits rather than individual improvisation.

Day-to-day work is made of recurring small tasks: drafting an email, preparing for a meeting, turning notes into actions, summarizing a long document, planning a week. Each takes ten to forty minutes. An assistant can often produce a usable first draft in a minute, leaving the person to improve, check and send. Across a team and a quarter, that adds up.

Two things go wrong when nobody sets expectations. First, some people use AI heavily while others avoid it, so quality and speed vary and resentment grows. Second, people use whatever tool is nearest, including personal accounts, and paste in material that should not leave the company. Both problems are management problems, not technology problems, and both are fixable with a plain agreement.

- Say what AI is for on this team: first drafts, summaries, structure, options and checking.
- Say what it is not for: final decisions, confidential material in unapproved tools, anything sent unread.
- Name the approved tools and where to get help when unsure.
- Make clear that you are accountable for anything you send, whoever or whatever drafted it.
- Agree how use is disclosed when it matters, for example in client deliverables.

THE WORKING PRINCIPLE

AI drafts, you decide. Treat every output as a competent but unverified colleague's first pass. You supply context, judgment and the final check.



This playbook is deliberately practical. It assumes you want to get better at using an assistant tomorrow, and that you lead or sit in a team that needs consistent habits more than a policy document nobody reads.

WORKING QUESTION

If a new colleague asked how your team uses AI, could anyone give them a clear, consistent answer?

THE WORK MAP

Which everyday tasks suit AI and where you must stay in charge

02

Start with tasks that are frequent, low in stakes and easy to check. Move toward higher stakes only when your checks are reliable.

TASK	GOOD FIT FOR AI	HUMAN CHECKPOINT
Drafting and tightening emails and updates	Strong: tone, structure, brevity	You read it fully; confirm facts, names, commitments
Meeting follow-up: summary and actions	Strong, with notes or a transcript you own	Attendees confirm owners and dates
Weekly planning and prioritizing	Good: options and sequencing	You decide priorities and trade-offs
Summarizing long documents	Good for orientation	Read the key sections yourself before relying on it
Background research and comparisons	Moderate: finds angles and questions	Verify every claim against a source you open
Preparing questions, agendas and briefing notes	Strong	Check against what you actually know
Performance feedback, sensitive or legal messages	Weak: use for structure only	You write, with HR or legal input where needed

A useful test is the cost of a quiet mistake. A slightly awkward draft costs little. A wrong date in an action list costs a missed deadline. A fabricated fact in a client briefing costs trust. The

higher the cost, the more of the checking must be done by a person who knows the subject.

Tasks to hold back at first

- Anything involving confidential client, employee or financial information until tools and rules are settled.
- Messages that carry legal, contractual or disciplinary weight.
- Decisions about people, such as hiring, promotion or performance ratings.
- Anything you would not be able to explain or defend if asked how it was produced.
- External communication that nobody reviews before it goes out.

These boundaries are not permanent. They exist so that people build confidence on safe tasks first, learn the failure patterns on low-stakes work and bring that judgment to harder tasks later.

WORKING QUESTION

Which three tasks do you repeat every week that sit in the strong rows of this table?

SETTING UP SAFELY

Information rules and approved tools for everyone

03

A one-page rule set that people actually remember beats a long policy they never open.

Shadow use happens when approved tools are missing, slow or unclear, so people improvise. The cure is a clear, easy approved route plus simple rules. Check your organization's policies and any legal or contractual obligations on personal, client and confidential information, and involve your security, privacy or legal contacts for anything unclear.

INFORMATION	RULE	WHAT TO DO INSTEAD
Client names, contracts and pricing	Approved tools only	Use placeholders such as Client A

INFORMATION	RULE	WHAT TO DO INSTEAD
Employee or personal data	Do not paste unless the tool is approved for it	Describe the situation without identifiers
Unreleased plans, financials, source code	Approved tools with the right settings only	Share the structure, not the content
Public information and your own drafts with no sensitive content	Generally fine	Still verify facts
Passwords, keys, access links	Never	Remove before sharing anything

Setup for each person

- Use a company-approved account, not a personal one, for work content
- Know whether your tool retains or learns from inputs, and what the setting is
- Turn off features you do not need, such as connecting to inboxes or drives, until you understand what they can access
- Know who to ask when a task does not clearly fit the rules
- Keep a note of significant uses where disclosure may be expected

A placeholder habit

01 Remove identifiers

Swap names, account numbers and addresses for labels before pasting.

02 Keep the shape

Leave roles, dates and amounts in a form that still lets the assistant help.

03 Put the real details back yourself

Restore names and figures in your own document, and check them as you do.

COMMON MISTAKE

Treating a chat window as private by default. Unless your organization has confirmed the tool's terms and settings, assume anything you paste could be stored or reviewed by someone outside your team.

WORKING QUESTION

Does everyone on your team know which tool to use for work content and what must never go into it?

**TASK PLAYBOOKS**

Writing, meetings, research and planning, step by step

04



A good request has a purpose, an audience, the material to work from, constraints and the form of the answer you want.

Playbook 1: drafting a message

- 01 Give the purpose and audience**
Say who will read it, what you want them to do and what they already know.
- 02 Supply the facts**
Paste the key points yourself. Do not ask the assistant to supply facts it cannot know.
- 03 Set the style**
State length, tone and any phrases to avoid. Share a sample of your own writing if you want it to sound like you.
- 04 Edit and own it**
Cut anything you would not say, check every name, date and commitment, then send.

**REQUEST TEMPLATE: MESSAGE**

Draft a reply to a customer who is unhappy that delivery slipped two weeks. Audience: a long-standing customer, frustrated but reasonable. Facts to use: new date is 14th; cause is a supplier delay; we can offer expedited shipping at no cost. Do not promise a refund. Tone: direct and warm, under 150 words. Do not invent any other facts.



Playbook 2: meeting follow-up

01 Use your own notes or an approved transcript

Check that attendees know a transcript or recording tool was used, as your policies and local rules require.

02 Ask for structure

Request decisions, actions with owner and date, open questions and anything marked unclear.

03 Confirm with the room

Send the draft to attendees for corrections within a day. Owners and dates are the items most often wrong.

REQUEST TEMPLATE: MEETING FOLLOW-UP

From the notes below, list: decisions made; actions with owner and due date, writing 'owner unclear' where the notes do not say; open questions; and any statements you are unsure how to interpret. Do not infer owners or dates that are not in the notes.

Playbook 3: research and summaries

Use the assistant to generate a list of questions, a structure for the topic, and a summary of a document you paste in. Then verify. Ask for claims to be tied to the passage that supports them. Open the sources yourself. If a statistic, quote or citation appears that you did not supply, treat it as unverified until you find it in the original. Assistants can produce references that look real and do not exist.

REQUEST TEMPLATE: SUMMARY

Summarize the document below for a manager with five minutes. Give: the main point in two sentences; five key facts with the paragraph each came from; risks or caveats the author states; and what is not covered. Quote exactly where you rely on specific wording.

Playbook 4: planning your week

Paste your commitments, deadlines and rough effort estimates and ask for two or three possible sequences with trade-offs stated. Then choose. The assistant does not know your priorities, your energy or the politics of your week. Use it to see options quickly, not to decide.

WORKING QUESTION

Which one of these four playbooks will you try this week, and what will you check before sending the result?



VERIFICATION HABITS

A short checking routine for everything that leaves your desk

05



Most AI errors in everyday work are small and findable: a wrong name, an invented fact, a confident tone around a weak claim.

The routine below takes two to five minutes. It is quick because it targets the typical failure patterns rather than rereading everything with equal suspicion. Make it a habit before anything you would be embarrassed to correct later.

Before you send or share

- Names, titles, dates, numbers and amounts match your source
- Every factual claim came from you or a source you opened, not only from the assistant
- Commitments and promises are ones you are able and willing to keep
- Tone fits the relationship, and nothing sounds generic or overconfident
- No confidential detail slipped in that the reader should not have
- You could explain the content to the reader without the draft in front of you



TYPICAL ERROR	WHERE IT HIDES	QUICK CHECK
Invented fact or citation	Research and summaries	Find it in the original; if you cannot, remove it
Wrong owner or date	Meeting actions	Confirm with attendees
Missing nuance	Summaries of long documents	Read the sections that carry the decision



TYPICAL ERROR	WHERE IT HIDES	QUICK CHECK
Generic or off-voice tone	Emails and posts	Read aloud; edit to how you would say it
Agreeable but wrong advice	Planning and options	Ask for the strongest case against the suggestion

Ask the assistant to check itself, then check it

A useful second pass is to ask the assistant to list the assumptions it made, the claims it is least sure of and what it would need to confirm them. This is not proof, because the same system can be wrong about its own uncertainty. It is a prompt for your attention, and it often points to the exact sentence worth checking.

ILLUSTRATIVE EXAMPLE

A manager asks for a summary of a supplier agreement and gets a clean paragraph on payment terms. Opening the contract shows a late-payment clause the summary omitted. The summary was useful for orientation and unsafe as a substitute for reading the clause.

WORKING QUESTION

What is the single most common small error in your own AI-assisted drafts, and what check will catch it?

MEASURING RESULTS

Knowing whether it saves time and keeps quality

06



Measure a handful of tasks before and after. Count the checking time. Ask the people who receive the work.

General impressions of productivity are unreliable. Pick four or five recurring tasks and time them for a couple of weeks without AI help, then with it. Include the time to check and edit. Pair the time numbers with a quality signal: how often the work was returned for correction, or how the recipient rates it. A saving that comes with more rework is not a saving.



01 Select tasks

Choose tasks that recur at least weekly, such as status updates or meeting follow-up.

02 Log the baseline

Note minutes per instance for two weeks. A simple note on a calendar is enough.

03 Apply the templates

Use the request templates and checking routine and log minutes again.

04 Review together

At a team meeting, look at the numbers and the stories behind them. Keep what works.



MEASURE	HOW TO CAPTURE IT	READING IT
Minutes per task, including checking	Personal log or timer	Down without more rework
Returned for correction	Count of drafts sent back	Flat or falling
Recipient view	One question to colleagues or clients	Clarity and usefulness at least as good
Spread of use	Share of team using approved tools	Wider and on approved tools
Use on higher-stakes tasks	Count of tasks moved up the map	Growing slowly with checks




ILLUSTRATIVE NUMBERS ONLY

Suppose weekly status updates take 40 minutes and, with a template and a check, take 25. That is 15 minutes a week for one task and person. The figures are illustrative. Measure your own, and do not count time you simply moved elsewhere.

Resist turning these measures into individual performance scores. The aim is to learn which methods work, not to rank people. When measurement feels like surveillance, people hide their use and the shadow problem returns.


WORKING QUESTION

What will you stop doing if the numbers show no real gain on a task?

TEAM ROUTINES

Sharing what works and keeping expectations clear

07



A few light routines let good methods spread and keep expectations realistic for everyone, including the skeptics.

Adoption tends to split a team into enthusiasts, skeptics and the quietly uncertain. The routines below give all three a place. Enthusiasts share what works. Skeptics supply the checks. The uncertain learn from examples that are close to their own work, not from abstract demos.



- Ten-minute show-and-tell at a regular meeting: one task, the request used, what was changed before sending.
- Shared template folder with request templates, each tagged by task and with its checking step.
- A short team agreement, one page, on approved tools, information rules and disclosure.
- A channel for questions and near misses, with a rule that nobody is criticized for raising one.

- Quarterly review of the agreement, informed by the metrics and the near misses.

Setting expectations fairly

Be explicit about what is and is not expected. Using AI is a way to work, not a duty, and not using it is not a failing. What is expected is quality, accountability and compliance with the information rules. Avoid quiet pressure to produce more simply because tools exist. If output volume rises without a plan, quality and morale usually fall.

WHAT TO NEVER DELEGATE

Decisions about people, commitments made on behalf of the company, anything requiring professional judgment or authority you hold, and the final check on facts. You can use help to prepare. You cannot hand off accountability.

- 01 Run a pilot**
Pick a small group and three tasks for a month.

- 02 Capture templates**
Save the requests that worked, with the edits that were still needed.

- 03 Extend**
Add people and tasks as the checks prove reliable.

WORKING QUESTION

What would make it easier for the most hesitant person on your team to try this safely?

PITFALLS

Where everyday adoption goes wrong, with one illustrative scenario

08

The common pitfalls are about habits and expectations rather than the technology.

- Sending drafts unread because they look polished.
- Using personal accounts and pasting company material because the approved route is awkward.
- Treating a summary as the document and missing the clause or caveat that mattered.
- Producing more output rather than better output, and flooding colleagues with text.
- Leaving expectations unspoken, so some people over-use and others fear using at all.
- Skipping disclosure where clients or colleagues reasonably expect to know.

ILLUSTRATIVE SCENARIO (HYPOTHETICAL)

A project coordinator at a fictional services firm uses an assistant to write meeting actions. One action is assigned to the wrong person because two colleagues share a first name. Nobody confirms. The task is missed. After the incident the team adds a rule: owners and dates are confirmed with attendees within a day, and the template says 'owner unclear' instead of guessing.

The fix is small and cheap, which is typical. A better template and a confirmation step removed the failure without banning the tool. Teams do best when they treat each slip as a design problem for the routine, not as a reason to stop or a reason to blame.

Warning signs to watch

- Colleagues say messages feel generic or oddly formal.
- People cannot explain content they sent.
- Factual corrections after sending are rising.
- Use goes quiet because people fear getting in trouble.

If you see these, adjust the routine first: tighten templates, add a check, clarify expectations. Escalate to policy changes only when the pattern points to a rule gap. Review your own policies and current legal requirements with the right internal contacts as your use widens.

WORKING QUESTION

Which warning sign are you most likely to miss on your team, and who would notice it first?

30 / 60 / 90

A first quarter that produces evidence

09

Begin with clear rules and three safe tasks, prove the checks work, then widen gradually.

First 30 days: ground it

- Publish a one-page team agreement on approved tools and information rules
- Choose three low-risk recurring tasks and log baseline minutes
- Share the message, meeting and summary templates and try them
- Name where people go with questions and near misses

Days 31 to 60: test it

- Hold the first show-and-tell and start a shared template folder
- Introduce the pre-send checklist as a team habit
- Re-measure the three tasks including checking time
- Add disclosure expectations for client-facing work

Days 61 to 90: decide

- Review the metrics and near-miss log and revise the agreement
- Move one or two tasks up the work map with named checks
- Retire templates that did not help
- Decide where to extend, change or stop use

GLOSSARY

Terms used in this publication

10

TERM	MEANING
Assistant	A tool that responds to written requests with drafts, summaries, options or answers.
Prompt or request	The instruction and material you give an assistant to work from.
Shadow use	Use of AI tools outside approved channels or without the knowledge of managers.
Approved tool	A tool your organization has reviewed and permitted for defined kinds of information.
Placeholder	A neutral label such as Client A used in place of a real name or identifier.
Hallucination	A confident statement, quote or citation that is not true or does not exist.
Transcript	A written record of what was said in a meeting, made by a recording or transcription tool.
Disclosure	Telling a reader or client that AI helped prepare the work, where expected or required.
Baseline	A measurement of time and quality before a change, used as the comparison point.
Template	A reusable request with blanks to fill, tested and shared across the team.
Near miss	An error caught before it caused harm, recorded so others can learn.

TERM	MEANING
Accountability	Responsibility for the result stays with the person who sends or acts on it.



A note on use

This playbook offers general working guidance and is not legal advice. Check your organization's policies and current requirements.

TAKE THE NEXT STEP

Take the free Halden assessment at halden.company/assessments, or talk to us at halden.company/contact about turning this into a plan for your organization.

