STRATEGEN AI

Managing AI Projects

A Strategen AI practical companion to PMBOK®

A practical reference for initiating, planning, delivering and transitioning AI projects, including broader projects with an AI component.

For project managers, sponsors, business owners, workstream leads and delivery teams.

Build, buy, configure, integrate and adopt AI capabilities with proportionate project management.

Independent companion

This publication preserves the PMBOK foundation and practical workflow of Project management & governance: A practical companion to PMBOK®. It is not produced, sponsored or endorsed by the Project Management Institute.

Contents

Select an entry to move to the relevant section. Chapter and section headings also appear in the PDF bookmarks and Word navigation pane.

How to use this reference 4

01 Understand the framework 6

1.1 Where PMBOK fits 6
1.2 The edition map: why online explanations differ 7
1.3 Principles as observable behaviour 8
1.4 Performance domains as management questions 10
1.5 Focus areas and the project lifecycle 12
1.6 Traditional knowledge areas: a translation aid 13
1.7 Projects, programmes, portfolios and operations 14
1.8 Tailoring: choose controls for the context 16
1.9 Choose predictive, adaptive or hybrid delivery 18
1.10 The worked case: an internal knowledge-assistant pilot 20

02 Authorise and govern 22

2.1 Govern the investment, manage the work 22
2.2 Build a business case that can survive a change of mind 23
2.3 A charter establishes the project and its boundaries 25
2.4 Give each role a real job 26
2.5 Decision rights are more useful than a crowded RACI 28
2.6 Use responsibility maps without hiding authority 29
2.7 Use gates to make choices, not to collect documents 30
2.8 Assurance increases confidence; authority remains explicit 32
2.9 Keep the reason for a decision, not only its outcome 33

03 Define scope and quality 35

3.1 Turn the ambition into a scope agreement 35
3.2 Make requirements testable and traceable 37
3.3 Build a work breakdown structure that exposes missing work 39
3.4 Design quality and acceptance before testing begins 41
3.5 A small toolkit for diagnosing quality problems 44

04 Plan time, resources and costs 46

4.1 Estimate honestly when knowledge is incomplete 46
4.2 Model the schedule before drawing the timeline 48
4.3 Find the path that controls completion 50
4.4 Plan capacity, responsibility and collaboration together 52
4.5 Resolve competition for shared capacity 53
4.6 Build a budget that distinguishes control from permission 55
4.7 Use earned value to connect spending with progress 57
4.8 Forecast the finish and choose a response 59

05 Lead delivery and change 61

5.1 Run short feedback loops and make work flow visible 61
5.2 Engage stakeholders around decisions and consequences 63
5.3 Communicate around decisions and useful action 64

5.4 Adoption is work inside the plan....66
5.5 Keep the plan integrated when change arrives....67
5.6 Report for a decision, not for reassurance....70
5.7 A light operating rhythm that keeps decisions moving....71
06 Manage risk and assurance....73
6.1 Write risks so that somebody can act on them....73
6.2 Assess exposure and choose a response....74
6.3 Quantify uncertainty when it improves the decision....76
6.4 A RAID log should produce movement....77
6.5 AI projects need evidence about use and behaviour....78
07 Source and sustain....81
7.1 Procurement and sourcing: buying a result you can use....81
7.2 Sustainability and public value: judge the whole consequence....85
08 Transition and realise value....87
8.1 Close the project and transfer continuing responsibility....87
8.2 Benefits need an owner, a baseline and a test....89
8.3 Test the value case with sensitivity....90
8.4 Work through the midpoint governance decision....92
8.5 Turn lessons into changed practice....93
09 Reusable templates....95
T01 Project charter....96
T02 Business case and option decision....98
T03 Delivery baseline and requirements....100
T04 RAID and dependency tracker....102
T05 Decision and change record....104
T06 One-page project status report....106
T07 Release, handover and acceptance....108
T08 Benefits record and review....110
10 Apply and troubleshoot....112
10.1 A 30-day learning and application plan....112
10.2 Diagnosing a distressed project....113
10.3 A one-page project review....115
10.4 Formula desk reference....116
11 Plain-language glossary....118
11.1 Terms used in this guide....118
References and further reading....123

How to use this reference

This Strategen AI handbook helps you manage projects that build, buy, configure, integrate or adopt AI capabilities, including broader projects with an AI component and material changes to existing services. It is designed to improve project decisions, delivery and proportionate governance. It assumes you understand organisations and leadership, while providing the project management vocabulary and AI context needed to apply the guidance. Read the first part for orientation, then use the contents and PDF bookmarks to move directly to the question you are facing.

The organising idea: establish why the work deserves investment, define how success will be recognised, give people clear authority, manage delivery with evidence, and leave someone accountable for the outcome after the project closes.

Three ways to read it

Table F1 | Three routes through the guide

Your immediate need	Suggested route	Useful result
Learn the framework	Start with the edition map, principles, domains and delivery approaches; then follow the worked examples.	A vocabulary that connects to practical decisions.
Start or reset a project	Read the business case, charter, governance, scope and planning sections; complete the relevant templates.	An authorised, testable plan with named owners.
Improve oversight	Read decision rights, risk, forecasting, assurance, status reporting and benefits.	Earlier intervention and clearer continuation decisions.

What the different elements mean

Framework facts are short, sourced descriptions of published frameworks. **Practice guidance** is this handbook's original explanation of broadly used project techniques. **Worked examples** and **templates** are original teaching material. They are suggestions to tailor, rather than official PMI forms or organisational requirements.

The main example is a fictional internal knowledge-assistant pilot. Its twelve-week schedule, AUD 120,000 funding envelope, staffing, measures and decision limits are teaching assumptions. They do not represent a client's or employer's budgets, approvals or policies. Use the example to understand the reasoning, then substitute your actual organisational context. Short contrasting situations show how the same management practices apply to purchased AI, a predictive feature inside a larger project and an assistant permitted to take actions.

Choose the relevant depth

An AI capability can change the work needed for a project without replacing its management approach. Read the guidance on scope, planning, people, cost, quality, sourcing and transition alongside the relevant specialist assurance. For a limited AI component, incorporate that work into the existing project plan and records. For a larger or uncertain

capability, make the additional experiments, dependencies and decisions visible at an appropriate level of detail.

This is a guide to managing projects that deliver or change AI. Using an approved AI tool to draft a status report or organise meeting notes is a separate supporting practice. It does not, by itself, turn the underlying project into an AI implementation. The project manager still checks the resulting record, protects information under the organisation's rules and owns any decision based on it.

The useful minimum

For a small initiative, combine the charter, milestones, RAID items, decisions and benefits measures in one controlled workspace. Add detail when it changes a decision, protects people, clarifies accountability or reduces a material uncertainty. An unused register does not improve control. A short record that the team updates and the sponsor acts on often does.

> **Try this first:** Choose one live project. In five sentences, state its problem, intended outcome, owner, next decision and largest uncertainty. Any sentence you cannot write clearly points to useful work in this handbook.

This is an independent Strategen AI learning and practice companion, adapted from the original guide dated 20 November 2026. It is not produced, sponsored or endorsed by PMI. PMBOK and PMI are trademarks of Project Management Institute, Inc. Published framework concepts remain attributable to their sources; the practical explanations, adaptations, examples, diagrams and templates are original companion material. Consult the official publications for their complete formal content. Strategen AI's website is Strategen-AI.com.

01 Understand the framework

Learn the language and select an approach that fits the work. The edition map separates official framework facts from this handbook's original practice guidance.

1.1 Where PMBOK fits

PMBOK is a body of project management knowledge expressed through PMI's publications. The PMBOK Guide helps organise and apply that knowledge. The associated Standard for Project Management sets out a recognised foundation. A local methodology turns selected concepts into your organisation's roles, decision points, working practices and records. These are related layers, rather than interchangeable names. [1, 5]



Figure 1 | From knowledge to a working project

The layers connect through organisational decisions; a project tool does not supply policy or authority.

Build the missing bridge to your organisation

Imagine the framework says to manage uncertainty. Your methodology specifies who owns risk, how serious exposures are escalated and which decisions require assurance. Your project plan then names an actual owner and trigger: if source access is not approved by the end of week two, the sponsor chooses between narrowing the pilot or moving its dates. A task board makes the agreed action visible. The board is useful because the decision arrangement already exists.

The same distinction applies to a charter. The concept is authorisation and shared direction. The local implementation might be a two-page document or a signed record in an established approval system. Its usefulness comes from the decisions it records, the authority of the approver and whether the team understands it.

Five layers to keep separate

Table 1.1 | Five layers of project management

Layer	Question it answers	Example
Organisational policy and delegations	What are we permitted or required to do?	Spending authority, privacy requirements and procurement rules.
Knowledge and standards	What should competent practice consider?	Value, planning, risk, stakeholders and tailoring.
Methodology and governance	How will this organisation manage this class of work?	Roles, gates, escalation rules and assurance arrangements.
Project-specific plan	How will this initiative be delivered and controlled?	Scope, schedule, costs, owners, quality criteria and baselines.
Tools and records	Where do we coordinate work and retain evidence?	Board, schedule, decision log and accepted deliverables.

A practical test

If you can change the software and still explain who decides, what success means and how exceptions are handled, you have a management approach. If those answers disappear when the board disappears, you probably have a tool configuration with important governance work still missing.

Use this handbook to develop the middle layers. Existing law, policy, contracts and authorised delegations still determine what you can approve. A project role or a RACI letter does not grant additional legal or financial authority.

AI engineering, evaluation and specialist assurance provide work and evidence within these arrangements. For example, a privacy review may establish a condition for using a collection, while an evaluation identifies a limitation in the proposed task. The project manager translates the resulting work into scope, owners, dependencies, acceptance evidence and a decision date. The reviews do not replace the delivery plan.

An optional connection with APIG

Strategen's APIG framework considers Actors, Practices, Infrastructure and Governance. [28] It can help a team ask whether the proposed capability fits the people doing the work, the workflow, the supporting environment and the decision arrangements. Use it where that organisational view is useful, then record the resulting work and decisions in the existing project plan. APIG adds no mandatory phase, gate or questionnaire to this guide and does not replace the PMBOK foundation.

1.2 The edition map: why online explanations differ

PMI's current public catalogue identifies the Eighth Edition as published in November 2025. It combines principles and performance domains with renewed process guidance. The table below is a navigation aid for recognising the vocabulary used in different sources, rather than a claim that older techniques have stopped being useful. [1, 2, 3]

Table 1.2 | PMBOK edition and publication map

Reference	Structure to recognise	How to use it in practice
Sixth Edition, 2017	10 knowledge areas; 5 process groups; 49 processes.	Understand traditional plans, process terminology and many existing workplace templates.
Seventh Edition, 2021	12 principles; 8 performance domains.	Understand the strong emphasis on outcomes, systems thinking and tailoring across delivery approaches.
Process Groups: A Practice Guide, 2022	5 process groups; 49 processes; predictive delivery emphasis.	Consult as a companion when you need further process-based guidance.
Eighth Edition, 2025	6 principles; 7 performance domains; 5 focus areas; 40 tailorable processes.	Use as the current orientation, then select practices for the work.

Historical counts and dates are supported by PMI's edition history; the Process Groups guide is a separate publication. [4, 6] The Seventh Edition FAQ explicitly retains the usefulness of earlier process knowledge. [5]

Do not mix the numbers

When a web article says there are twelve principles and then describes the eighth edition, check which material it is actually using. Likewise, the ten knowledge areas remain useful vocabulary but are not the seven domains of the current edition. Different maps can address the same real work without having identical boundaries.

This guide's chapters use an original practical sequence: understand, authorise, plan, deliver, assure, transition and improve. This is not the official table of contents. It is intended to make the concepts easier to find and use.

What counts do not tell you

Forty processes does not mean forty forms, forty meetings or forty mandatory steps. A small project might combine several needs in one discussion and one record. A large programme may need many more records because its context demands them. The question is whether a decision is supported and an exposure controlled, rather than whether a document exists for every label.

> **Reading habit:** Record the edition beside any course notes or templates you adopt. Verify unfamiliar definitions against that edition before merging them into your own framework.

1.3 Principles as observable behaviour

The current framework identifies six principles. The shorthand below paraphrases their themes; it is not a verbatim list of the official headings. Their full wording can be checked in PMI's public contents preview. [2]

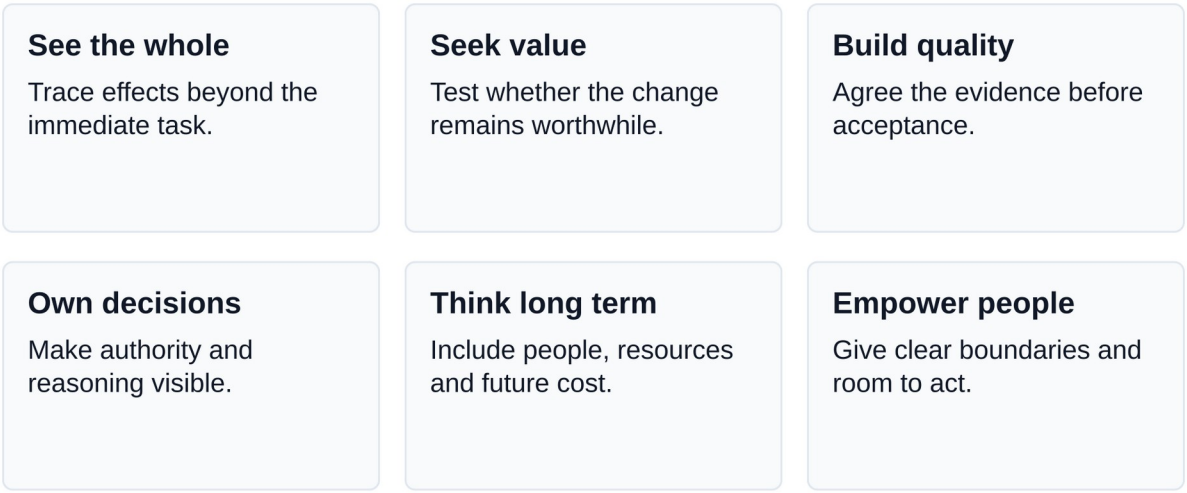


Figure 2 | Six principle themes translated into action

Original paraphrases of the six Eighth Edition principle themes. Consult the official wording in reference 2.

Turn a principle into something you can see

Table 1.3 | Principles translated into observable practice

Principle theme	A practical behaviour	Evidence to look for
Whole-system perspective	Check how a local decision affects people, operations, suppliers and other initiatives.	A change assessment includes downstream ownership and dependencies.
Value	Reconsider the preferred solution when new evidence changes the benefit or cost.	The latest business case explains why continuing still makes sense.
Quality built into the work	Agree what acceptable means before work is presented for sign-off.	Testable requirements, review criteria and useful defect evidence.
Accountable leadership	Make decisions within authority and expose exceptions early.	Named decision owners, honest forecasts and recorded reasons.
Sustainability	Consider lasting social, environmental and economic consequences.	Whole-life cost, accessibility, resource use and operating obligations.
An empowered culture	Give the team clear boundaries and room to solve problems.	People can raise problems, resolve routine work and ask for help safely.

These behaviours are original application examples. They are useful prompts for a sponsor or project manager even when the organisation uses another project method. For AI work, an observable behaviour might be protecting time for representative evaluation, testing an assumed saving against the checking effort it creates, or giving the future service owner a funded role in the decision. The principle becomes useful through that action and its evidence.

One decision can engage all six

Suppose the team proposes broadening a knowledge assistant's content sources. A whole-system view asks who maintains those sources and who may see the answers. Value asks whether broader coverage addresses a meaningful user problem. Quality asks how the new content will be tested. Accountability asks who is authorised to approve the change. Sustainability asks whether the service can be maintained affordably and with acceptable social and environmental effects. An empowered culture lets the team propose options and raise concerns without hiding uncertainty.

Use principles at the point of choice

Do not turn this table into a recurring six-question form for every minor action. Use it when the team faces a trade-off: speed versus confidence, more scope versus adoption, lower price versus exit difficulty, or a visible launch versus a reliable service. One short decision note can record the relevant considerations.

> **Leadership prompt:** What would a reasonable user, future service owner or affected stakeholder wish we had considered before making this decision?

1.4 Performance domains as management questions

The Eighth Edition identifies seven domains: Governance, Scope, Schedule, Finance, Stakeholders, Resources and Risk. They describe related areas of practice. They are not seven sequential stages. [1, 2]

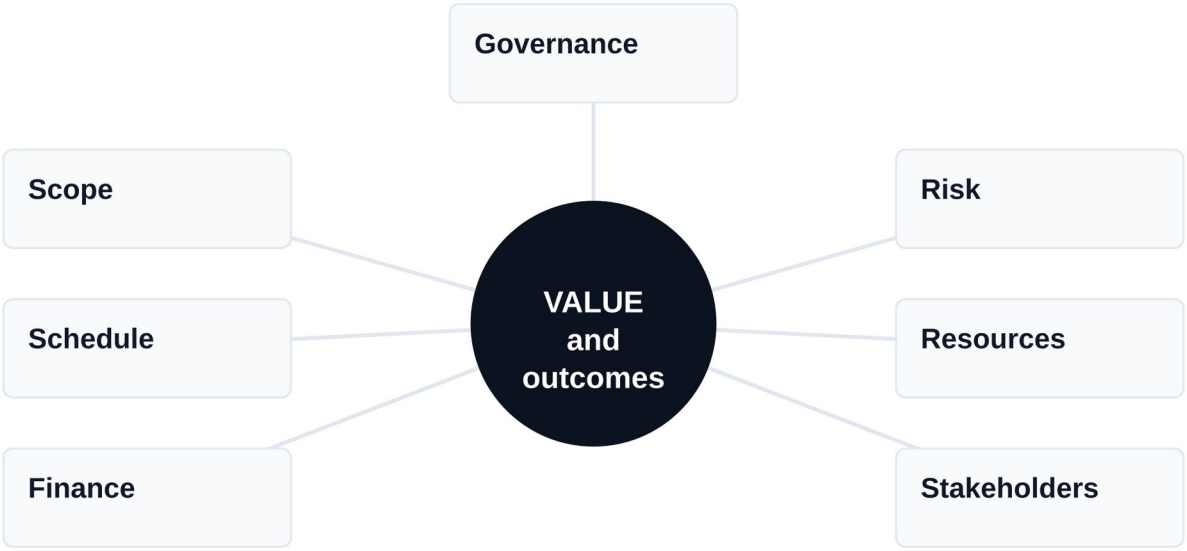


Figure 3 | Seven connected areas of project attention

Domains are concurrent areas of attention. Their positions do not indicate sequence or priority.

A practical weekly scan

Table 1.4 | Performance domains in use

Domain	Question for the project	A useful sign of control
Governance	Are decisions being made by the right people, using current evidence?	Clear next decision, authority and escalation route.
Scope	Are we delivering the agreed outcomes and boundaries?	Traceable requirements and accepted changes.
Schedule	Is the remaining sequence achievable?	Credible data, supplier and evaluation dependencies, capacity and completion forecast.
Finance	Is the expected whole cost understood and affordable?	Actuals, commitments, forecast and controlled reserves, including evaluation and continuing service costs.
Stakeholders	Are affected people involved and able to use the result?	Feedback, observed task performance, adoption evidence and unresolved concerns.
Resources	Can the people and suppliers actually do the work?	Capacity commitments, skills and ownership, including content specialists and reviewers.
Risk	What could alter success, and who is responding?	Owned exposures, triggers and tested responses.

The questions and evidence cues are this handbook's original practice prompts. The seven-domain scan should connect with existing reporting, rather than create a second reporting system. A strong model-performance score is one piece of evidence; it cannot answer whether the remaining work is funded, users can complete their tasks or an owner will sustain the service.

The connections matter

A supplier delay can increase cost, compress testing and reduce adoption time. It is therefore a schedule, finance, quality, stakeholder and risk problem. A scope decision changes the baseline only after its consequences have been considered. A finance forecast that ignores the schedule may understate extended support or contractor costs.

Quality, communications, procurement, integration and benefits remain necessary work even though this list does not contain a separate domain bearing each of those names. This handbook treats them explicitly. Do not infer that a topic disappears because a later framework organises it differently.

Avoid seven disconnected reports

Start with one overall delivery and value judgement. Use the domains to explain the evidence behind it. If an initiative is within its spending limit but cannot show that anyone will use the result, a green finance indicator should not conceal the adoption problem. Equally, an isolated amber item does not automatically require stopping the whole project; the decision depends on consequence, options and authority.

> **Useful question:** Which domain looks healthy only because another team is absorbing its cost or risk?

1.5 Focus areas and the project lifecycle

The five current focus areas use the familiar labels Initiating, Planning, Executing, Monitoring and Controlling, and Closing. PMI's public material distinguishes these management activities from lifecycle phases. [2, 7]

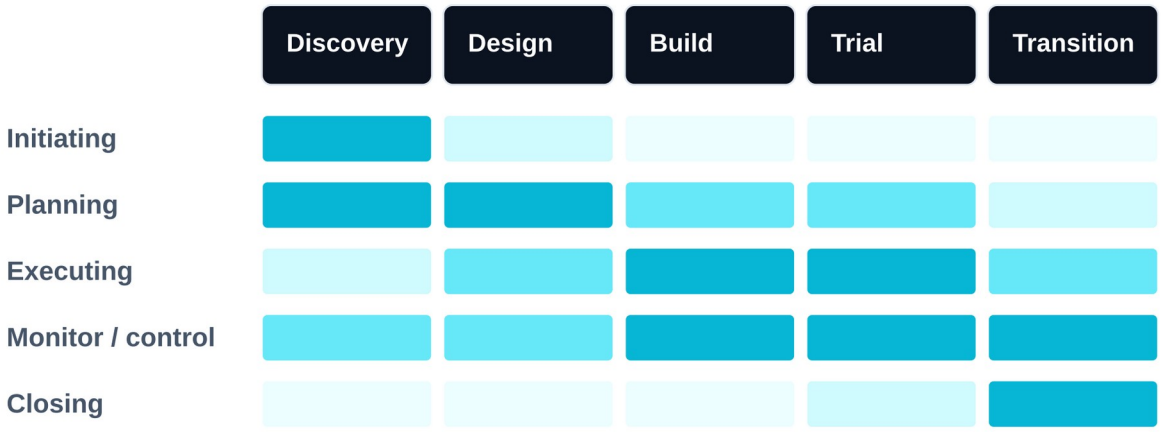


Figure 4 | Lifecycle phases and recurring management work

Management activities recur within and across phases. The top row is an example life cycle, not a mandatory sequence.

Two different views of the same project

A lifecycle describes the progression of the work: for example, discovery, design, build, trial and transition. Management activities occur within and across those phases. The team can plan the trial while executing the build, monitor both and close a completed supplier package. A new phase may need fresh authorisation and a refined plan.

Table 1.5 | Recurring management activities

Management activity	What you are doing	A practical example
Initiating	Establish purpose, boundaries, stakeholders and authority.	Authorise a bounded discovery with a spending limit.
Planning	Decide how work will be delivered and measured.	Set a trial cohort, acceptance criteria and escalation triggers.
Executing	Coordinate and perform the agreed work.	Configure access, curate content and run user sessions.
Monitoring and controlling	Compare evidence with expectations and act.	Investigate missed milestones, forecast effects and implement authorised changes.
Closing	Complete obligations and transfer accountability.	Accept a phase, hand over records and release temporary resources.

Gates are decisions, not dates

A date in a plan tells you when a decision is expected. A gate requires evidence and an authorised judgement. You can hold the meeting on schedule and still decide that the initiative is not ready. Possible outcomes include proceeding, proceeding with explicit conditions, revising, pausing or stopping.

For an adaptive team, a review of an increment provides feedback and learning. It does not automatically grant permission to expand use, spend more or expose a new population to risk. Put those decisions in the governance arrangement without requiring a full committee for every backlog change.

Draw your own lifecycle

Use the phases that reflect the real work and obligations. A procurement-heavy project may need a sourcing phase; an event project has a fixed delivery date; a research project may have experiments and evidence checkpoints. Show when ownership changes and where a commitment becomes difficult to reverse. Those are useful places to consider a decision or assurance review.

A purchased AI capability may require supplier selection, configuration, integration and adoption without model-development work. A project developing a new capability may need several experiments before the team can make a credible delivery commitment. In either case, link the experiment to a question, an effort limit, evidence and a decision. Returning to configuration after evaluation is also a planning and control event: assess the remaining work, update the forecast and check whether the authorised boundaries still hold.

> **Common error:** A five-box diagram running from initiating to closing can make planning appear finished once execution begins. Keep planning and control visibly active as evidence changes.

1.6 Traditional knowledge areas: a translation aid

The traditional ten-area vocabulary is still common in plans, training and supplier documents. The labels below are conventional topic names. Their practical questions and working records are this handbook's original translation, not an official crosswalk to the Eighth Edition. [4]

Table 1.6 | Traditional knowledge areas and working records

Traditional area	Practical question	Typical working record
Integration	Do the parts of the project still fit together?	Integrated plan and decision/change record.
Scope	What is included, excluded and accepted?	Scope statement, requirements and WBS/backlog.
Schedule	What depends on what, and when will it finish?	Dependency-based schedule and milestones.
Cost	What will it cost, and what is our latest forecast?	Estimate, baseline, actuals and forecast.
Quality	What does good enough mean and how is it demonstrated?	Criteria, tests, reviews and acceptance evidence.

Traditional area	Practical question	Typical working record
Resources	Who and what are available to do the work?	Capacity plan, assignments and resource agreements.
Communications	Who needs what information to make which decision?	Audience, cadence, channel and escalation plan.
Risk	Which uncertainties could alter success?	Risk register, response actions and trigger checks.
Procurement	What must be bought and how are obligations managed?	Sourcing plan, contract and supplier record.
Stakeholders	Who affects or is affected by the change?	Stakeholder analysis, engagement and feedback.

Understand the connecting work

Integration means assessing consequences across these topics and keeping the shared plan coherent. For example, an approved change to the data sources may require a different supplier agreement, additional review work, a new test set and revised user guidance. The team should also identify the affected configuration, any acceptance evidence that must be renewed, the capacity for that work and the effect on operating cost. Updating the scope statement alone leaves the project only partly controlled.

An input is information or an artefact needed to do some work. A tool or technique helps perform the work. An output is the resulting decision, record or deliverable. A risk review might use the latest schedule and supplier advice as inputs, scenario discussion as a technique and revised responses as outputs. These terms describe a reasoning chain, not a demand to create a separate form at every step.

Connect plans to registers

Plans state how work will be managed. Baselines preserve an approved reference for comparison. Registers hold changing operational information. A risk management approach may remain stable for several months while the risk register changes each week. A live forecast may change without replacing the authorised baseline. Keep their dates, purposes and owners clear.

> **Control question:** If a change is approved today, which related plans, commitments, tests and communications must also change?

1.7 Projects, programmes, portfolios and operations

Use these terms to locate the decision you are making. A project is temporary work organised to produce a distinct output or outcome and create value in a particular context. A programme coordinates related changes to obtain outcomes and benefits that need joint management. A portfolio selects and balances investments against strategy, capacity and risk. Operations sustain services and ongoing performance. These are plain-language working explanations rather than quoted standard definitions.



Figure 5 | Investment choices change services and create value

Projects contribute outputs; continuing ownership and use are needed to obtain outcomes and value.

The same initiative needs several owners

Imagine a portfolio objective to improve access to organisational knowledge. A programme might combine content ownership, records uplift, staff capability and a new digital service. A project could deliver a bounded knowledge-assistant pilot. Operations then maintains approved sources, supports users, monitors performance and makes routine service decisions.

The programme should manage shared benefits and dependencies rather than duplicate every project's task board. The portfolio should decide which investments deserve scarce capacity rather than direct day-to-day team work. A project should identify the operating obligations it creates before closure, including who will fund and perform them.

The illustration shows related levels of management; it does not require a programme between every portfolio and project. A bounded AI feature can be delivered within an existing project. Conversely, several AI projects may share source remediation, identity integration or staff capability work that needs coordination beyond any one project manager's authority.

Separate output, outcome and benefit

Table 1.7 | Value at project and organisational levels

Concept	Example	Evidence
Output	A knowledge assistant, curated sources, training materials and completed sessions.	Accepted deliverables and readiness checks.
Outcome	Staff find relevant approved information more reliably.	Observed tasks, source accuracy and appropriate escalation.
Benefit	Useful time is released and avoidable rework falls.	Baseline comparison, attributable change and agreed measurement.
Wider impact	Better organisational service or decision quality.	Longer-term evaluation with other contributing factors considered.

An output does not guarantee the next link. Users may reject the workflow, source material may remain outdated or time savings may be absorbed without improving service. Treat the links as hypotheses to test. Use a benefits owner who can influence the working practices that turn a technical capability into value.

Avoid the orphaned service

Projects often have a visible sponsor and budget. The subsequent service can be less visible. Before committing, ask who owns the product roadmap, content, support queue, operating cost, assurance and eventual retirement. A service without these owners may appear cheap only because its obligations are missing from the business case.

For the assistant, continuing work includes maintaining the approved collection, responding to problems, reviewing supplier changes and deciding when further evaluation is required. Transfer those responsibilities and their funding explicitly. A project can close after an agreed handover or retirement even though benefits reviews and service monitoring continue. A later material expansion may justify a new project or another authorised change; routine operation should not leave the original project open indefinitely.

> **Portfolio prompt:** What other work will lose capacity if we approve this initiative, and is this the better use of that capacity?

1.8 Tailoring: choose controls for the context

Tailoring is deliberate selection and adjustment of practices. It requires a reasoned fit between the work and its management approach. PMI's public Eighth Edition material explicitly includes tailoring. [1] The decision rules below are original suggestions for applying that idea.

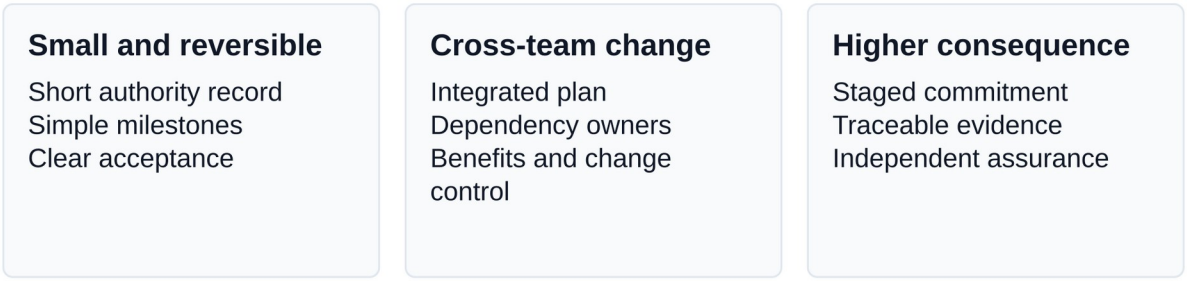


Figure 6 | Increase control with consequence and uncertainty

Consider consequence as well as size

Table 1.8 | Proportionate management arrangements

Context	Start with	Add when justified
Small, familiar, reversible change	Short charter; named owner; milestones; live action/RAID record; acceptance criteria.	Specialist review for a particular exposure or dependency.

Context	Start with	Add when justified
Several teams or material service change	Integrated plan; capacity commitments; change decisions; benefits owner; regular sponsor review.	Dependency coordination and independent readiness checks.
High consequence, difficult reversal or substantial uncertainty	Explicit authorities; staged commitments; traceable evidence; assurance and fallback.	Deeper modelling, external review and stronger release conditions.

A small budget can still create serious privacy or operational consequences. Conversely, a larger routine purchase may not require elaborate discovery. Look at novelty, affected people, uncertainty, regulation, supplier concentration, reversibility, organisational capacity and the cost of failure. Also consider which actions the AI can take and whether a person can detect and reverse an unsuitable result before it has an effect.

Tailor the work as well as the controls

For a purchased ticket-summary feature, the necessary work may be a focused set of permission and task tests, user preparation, a supplier-change arrangement and an identified service owner. A predictive component in a larger warehouse project can use that project's scope, baseline, change process and handover. Building a new model may add data preparation, training and evaluation work; these belong in the relevant work packages and estimates.

Extending an assistant from giving advice to creating purchase requests changes the operating boundary. Plan the transaction integration, permitted actions, confirmations, failure recovery and acceptance evidence. This is more than adding another interface feature, but it does not require a second project methodology. These examples illustrate different management responses, not a classification or scoring system.

Record the tailoring decision once

Write a short note covering the chosen delivery approach, decision structure, records, cadence and assurance. Explain why it fits and who agreed. Identify the triggers that would require more control, such as a wider user group, sensitive data, a new supplier, increased funding or a material service dependency.

Do not use tailoring to waive obligations that the project lacks authority to waive. If a mandated process appears disproportionate, ask the authorised owner for an alternative that satisfies its purpose. Preserve the decision and rationale.

Make the system easier to use

Use existing information where possible. A single source of milestone data can support the team board and executive report. Keep a field only if someone uses it to act or decide. Review records after several cycles: remove duplication, clarify ambiguous labels and archive items that no longer need attention.

Revisit at transitions

Discovery may need rapid experiments; release may need firmer evidence and operating controls. As uncertainty falls or exposure increases, change the approach. A tailoring decision made at initiation is an initial design, not a promise to use the same controls forever.

> **Useful test:** Can each recurring report, meeting and approval be linked to a decision, commitment, learning need or material exposure?

1.9 Choose predictive, adaptive or hybrid delivery

Delivery approaches organise learning and commitment differently. Predictive work develops a substantial plan before implementation and controls departures from that plan. Adaptive work uses short feedback cycles to refine the solution as understanding grows. Hybrid delivery deliberately combines different arrangements for different kinds of uncertainty or constraint. Current PMI guidance accommodates all three. [2, 5]

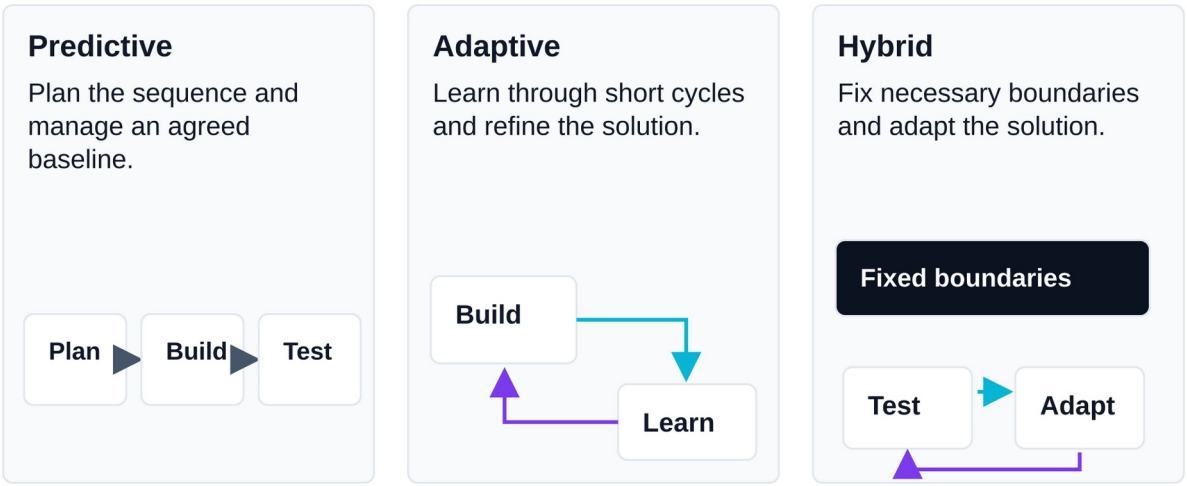


Figure 7 | Three ways to organise delivery and learning

Choose the delivery approach for the work and its uncertainty. Iteration still requires explicit decisions and evidence.

Choose by the characteristics of the work

Table 1.9 | Choosing a delivery approach

Question	Predictive tendency	Adaptive tendency
How stable is the required result?	Requirements are well understood.	Useful requirements must be discovered.
How expensive is change?	Late changes are costly or difficult to reverse.	Small changes are affordable and can be tested.
Can users evaluate partial results?	Meaningful feedback requires a substantial completed component.	Users can test increments early.
What constrains approval?	Formal packages or physical interfaces dominate.	Frequent review can authorise small adjustments.
How much uncertainty remains?	Work resembles proven delivery.	Solution and user response remain uncertain.

These are prompts, not a scoring formula. A physical construction project can use iterative design and a digital project can have tightly fixed obligations. Adaptive delivery still needs

cost control, quality, ownership and governance. Predictive delivery can still release in increments and respond to evidence.

Design a coherent hybrid

For the fictional assistant pilot, use fixed funding and explicit use boundaries, a dependency-based twelve-week schedule, iterative configuration and user testing, and a release decision based on evidence. Establish which backlog changes the product owner may make and which changes need sponsor or specialist approval. The combination is coherent because the boundaries are explicit.

Do not give every work package the same treatment automatically. Identity integration may have known interfaces and a firm delivery sequence, while the most useful response design still needs experiments. Plan the experiment's question, effort and completion condition, and keep its results distinct from acceptance of the service. If an experiment shows that the proposed approach is unsuitable, that can complete the learning work while changing the investment recommendation.

Avoid accidental hybrid delivery

An organisation can ask a team to explore an uncertain solution while treating every early estimate as an immovable commitment. It can also demand formal approval for every minor backlog refinement while expecting weekly release. Both arrangements create friction. Agree where learning is expected, where commitments are firm and how new information changes the plan.

> **Decision prompt:** What uncertainty should we reduce before committing further money, and what can we safely learn by delivering a small increment?

1.10 The worked case: an internal knowledge-assistant pilot

This fictional case gives the techniques a common setting. A medium-sized organisation wants to know whether an internal assistant can help staff find and interpret approved guidance more effectively. The current process involves searching several locations and asking subject specialists. The project purchases evidence about a possible improvement. Expansion would require another decision.

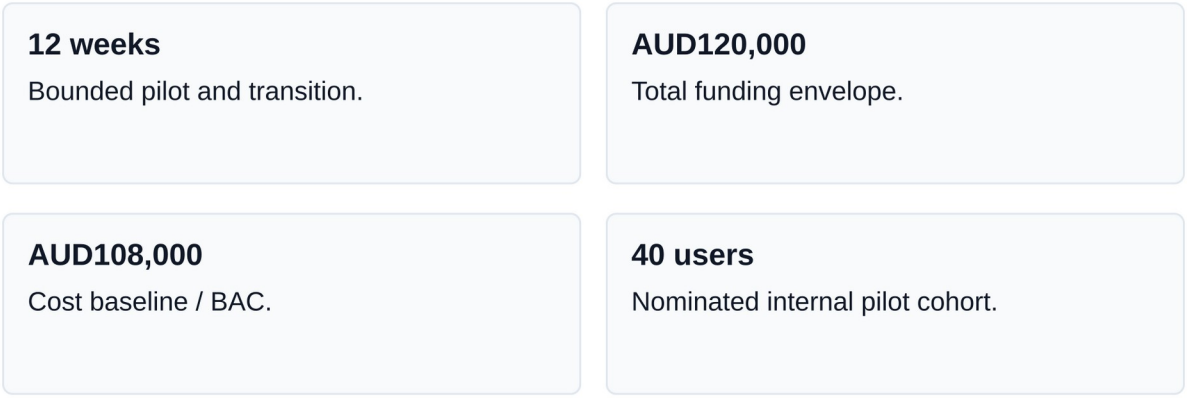


Figure 8 | The fictional pilot at a glance

Shared assumptions

Table 1.10 | Worked pilot assumptions

Element	Teaching assumption
Participants	40 nominated internal staff, with manager-supported participation.
Scope	One approved internal guidance collection, controlled access, source references, evaluations, user guidance and a continuation recommendation.
Exclusions	Public advice, autonomous decisions, personal email ingestion, HR case advice and enterprise rollout.
Duration	12 working weeks under the simplified dependency model. Controlled pre-release evaluation occurs within integrated tests in weeks 8-9. Supervised use occurs in weeks 10-11; the decision pack and handover are completed in week 12.
Funding	AUD 120,000 authorised envelope: AUD 100,000 planned work, AUD 8,000 contingency within the AUD 108,000 baseline, and AUD 12,000 management reserve outside it.
Accounting basis	Illustrative AUD amounts; no tax treatment is modelled. A real plan must state its treatment of internal labour, overheads, accruals and taxes.
Decision	Stop, adapt, continue for a bounded period or propose expansion, using evaluated usefulness, exposure and whole-life cost.
Evaluation	The fictional E01 plan in Chapter 03 specifies the evaluated configuration, tasks, methods and decision conditions. Eight participants from the 40-user cohort join controlled pre-release sessions; the full supervised pilot remains in weeks 10-11.

Element	Teaching assumption
Financial teaching snapshot	At the week-six review, actual cost is AUD 48,000 and the illustrative estimate to complete is AUD 72,000, giving AUD 120,000 at completion. This forecast does not replace the AUD 108,000 baseline or authorise reserve use.

What success must demonstrate

The pilot should produce accepted deliverables and credible learning. The business needs evidence about real task completion, appropriate use, answer support, access boundaries, staff behaviour and operating effort. A working interface is necessary for the experiment, but is not sufficient to establish value or suitability.

The numerical schedule and financial snapshots are teaching assumptions. They are not measurements from a real project, and no historical trend is implied. A later scenario such as a one-week delay is explicitly a what-if change to the baseline case. Evaluation scenario counts, review roles and release conditions are also local teaching choices, not universal performance standards.

Maintain one case record when adapting the example: users and intended tasks; approved sources and permissions; configuration and evaluation versions; schedule and cost assumptions; owners; and the decision currently authorised. If a change alters one of these, check all dependent examples before using them to explain the project's position.

Transfer the management logic

The same management questions apply to a website, service redesign, CRM implementation, event or infrastructure change. What is authorised? Who decides? What depends on what? What evidence supports acceptance? Who inherits the result? What benefit remains plausible? The specialist assurance changes with the subject; the need for coherent project control remains.

> **Transfer the example:** Replace the assistant with a current initiative, then replace the participants, deliverables, constraints, dependencies and measures. Keep the decision logic.

02 Authorise and govern

Build a defensible investment case, establish real authority and decide what evidence is needed before committing further resources.

2.1 Govern the investment, manage the work

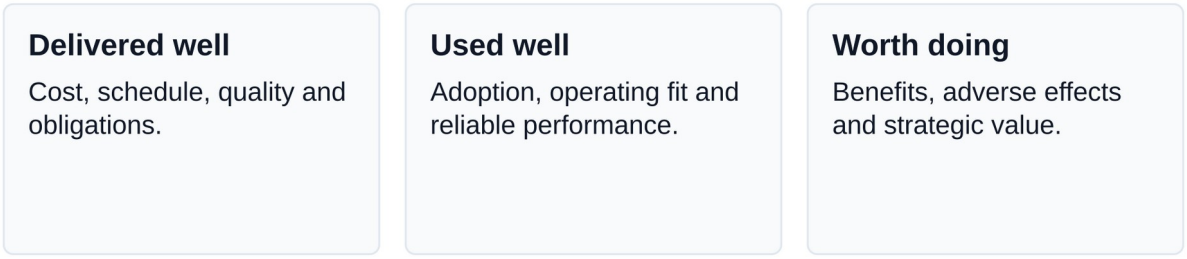


Figure 9 | Delivery performance and realised value are distinct

Purpose. Make decisions about value, authority and exposure while enabling the delivery team to do its work.

A project has two connected conversations. The delivery conversation asks what must happen next, who will do it and whether the plan remains credible. The investment conversation asks whether the intended outcome still matters, whether the proposed course is justified and who has the authority to accept the consequences. A detailed task board can support the first conversation while leaving the second almost untouched.

For the knowledge-assistant pilot, the project manager coordinates document preparation, configuration, testing and training. The sponsor ensures the pilot addresses a worthwhile organisational problem and obtains decisions that cross team boundaries. The service owner decides whether an operational service can be supported. The benefit owner changes the work practices needed to use any time released. These responsibilities need explicit allocation even when one person holds several roles.

Build governance around four questions:

1. **Purpose:** Which outcome are we seeking, for whom, and what evidence would change our mind?
2. **Authority:** Who may authorise spending, scope, release and risk acceptance, within what limits?
3. **Evidence:** What information will those people need before each decision?
4. **Response:** What happens when evidence, forecasts or conditions move outside the approved position?

The project manager should be able to explain the answer to each question in ordinary language. A committee name alone answers none of them. A committee also needs a remit, membership, quorum or decision procedure, access to evidence, and an escalation route for matters outside its authority.

Worked situation. The team proposes finishing on time by removing evaluation of difficult questions. The schedule consequence appears favourable, but the agreed work and evidence have changed. The sponsor must examine whether the remaining evidence can justify the proposed use, and the appropriate authority must decide the scope change. The

team cannot claim completion against the original acceptance conditions merely because the software is finished.

Useful distinction. Delivery success concerns the agreed work and acceptance criteria. Outcome success concerns the change achieved. Investment success concerns whether the total benefits and learning justify the full cost and exposure. Record all three because they may move differently.

An AI experiment may produce a credible finding that the intended capability should not proceed. Accepting that learning deliverable closes the agreed experiment; it does not authorise a service that failed its use requirements. This distinction helps the team report honest results and supports a legitimate stop, change or further-investigation decision.

Common failure: governance becomes a monthly recital of completed tasks. Replace part of that agenda with explicit options and decisions.

Use tomorrow: draw the route from a team issue to the person authorised to resolve it. If a name or authority is missing, resolve that gap before adding another report.

Source boundary. PMI publishes a separate guide covering governance across portfolios, programmes and projects. The four questions and case above are an original working model. [8]

2.2 Build a business case that can survive a change of mind



Conceptual pattern, not data. New evidence can also increase uncertainty.

Figure 10 | Commit in stages as evidence becomes stronger

Purpose. Explain why action is worthwhile and make the comparison between credible options visible.

Start with the problem before describing the preferred solution. For the pilot, “staff struggle to locate current guidance and repeat questions to subject specialists” is a problem statement. “We need an AI assistant” is already a solution. The baseline should establish the frequency, consequence and current cost of the problem, including its variation across teams.

A proportionate business case answers seven questions: What needs to change? What happens if we do little? What alternatives could achieve the change? Which benefits and harms are plausible? What will each option cost over its life? Can we deliver and operate it? What decision is requested now?

Table 2.1 | Business options and evidence

Candidate option	What it might solve	What evidence is needed
Keep present arrangements	Avoids new investment	Cost and consequences of continuing
Improve document ownership and search	Addresses stale material and navigation	Search usability and content effort
Pilot a knowledge assistant	Tests a new retrieval experience	Accuracy, usefulness, cost and staff behaviour
Buy a broader platform	May address several workflows	Fit, integration, lock-in and operating model

Do not make alternatives deliberately weak. Compare them against the same outcome and include future support, licence growth, staff effort, transition and eventual exit. Record ranges when the evidence is thin. A precise cost and an untested benefit assumption do not become reliable simply because they appear in a spreadsheet.

For each credible AI option, make the sourcing assumption clear: build a capability, configure an existing product, integrate a supplier service or enable a feature already owned. Include the work the supplier does not perform, such as preparing information, securing access, evaluating local tasks, introducing the workflow and supporting users. A cheaper licence can still require more project and operating effort.

Test the economics against actual use. Estimate how often the task occurs, the effort under the present method, the assisted task time including checking and correction, and the continuing support and maintenance requirement. State which assumptions are measured and which remain to be tested. These quantities belong in the business case and cost forecast; the guide does not prescribe a universal return threshold.

Worked recommendation. Authorise a bounded 12-week, AUD 120,000 pilot, with AUD 108,000 in the cost baseline and AUD 12,000 of separately controlled management reserve. The decision purchases evidence about usefulness, safe operation and economics. It does not authorise enterprise rollout. The sponsor will decide to stop, adjust, extend or propose expansion using the pilot findings and an updated whole-life case.

Write the assumptions that could reverse the recommendation: sufficient current source documents; staff willing and able to verify answers; viable ongoing support; meaningful improvement over better search. Give each assumption an owner and a date for obtaining evidence. If an experiment leaves a material question unresolved, identify what another bounded test would cost and which decision it could change. Avoid extending a pilot repeatedly without a defined learning purpose.

HM Treasury's public guidance uses strategic, economic, commercial, financial and management perspectives to examine a proposal. Those perspectives are useful questions; this example does not import UK approval requirements into an Australian organisation. [9]

Common failure: treating the business case as a one-time sales document. Revisit it when major assumptions, benefits, costs or strategic priorities change. Continuing expenditure deserves the same disciplined scrutiny as the initial decision.

Use tomorrow: add a short “what would make us stop or change direction?” section to the next proposal.

2.3 A charter establishes the project and its boundaries

Purpose. Create a shared statement of what has been authorised and how the project will be directed.

The business case explains the investment logic. A charter turns an approved intention into a bounded undertaking. It identifies the sponsor and project manager, defines the intended outcome, establishes high-level scope and constraints, and states how decisions will be made. It can be short provided the authority and boundaries are clear.

Table 2.2 | Example charter decisions

Charter field	Fictional pilot entry
Purpose	Test whether staff can locate reliable internal guidance more efficiently
Outcome	Evidence-based decision on a controlled internal knowledge service
Included	Approved source collection, configuration, evaluation, training, 40-person pilot and recommendation
Permitted use	Assistance with internal guidance; users inspect supporting sources. No authority to execute business transactions.
Excluded	Public-facing advice, autonomous decisions, personal email ingestion, HR case advice and enterprise rollout
Delivery limit	12 weeks; approved funding AUD 120,000
Cost control	AUD 108,000 baseline; AUD 12,000 reserve outside baseline, separately controlled
Sponsor	Executive accountable for the intended business outcome
Project manager	Named person coordinating delivery within stated authority
Service owner	Named owner responsible for operational readiness and future support
Decisions	Funding and changes follow actual organisational delegations
Success evidence	Evaluated task quality and effort, acceptable failure behaviour, operational readiness and a justified continuation or stop recommendation

Avoid promises such as “improve productivity by 30%” without defining whose productivity, compared with what, and how it will be observed. A pilot may succeed by establishing that a tempting option should not proceed. Separate the success of the learning exercise from the commercial attractiveness of the eventual service.

Record high-level assumptions and dependencies. If access to a source collection is essential, name its owner and needed date. If the work requires staff from another team, record the agreed capacity rather than assuming their cooperation.

The charter should distinguish approval from consultation. A signature may show that someone has reviewed the proposal without giving them spending or release authority. State what each approval covers and retain the version that was actually approved.

Keep detailed evaluation cases and configuration identifiers in linked controlled records. In this pilot, the charter points to C01 for the configuration and E01 for the planned evaluation. Their purpose is to establish what will be examined and under what conditions, while the

charter retains the durable purpose, permitted use and funding boundary. The E01 plan is developed in Chapter 03.

Worked situation. A manager asks to include external partners in the pilot. The request changes the intended users, information exposure, support arrangements and evaluation conditions. It is therefore assessed through the change process; it cannot be treated as adding a few names to the participant list.

Common failure: the charter becomes a miniature delivery plan and loses its function. Keep the durable purpose and boundaries here; keep frequently changing tasks in the delivery plan.

Use tomorrow: ask the sponsor and project manager to describe the project's outcome and exclusions separately. Any difference is an initiation issue worth resolving.

2.4 Give each role a real job

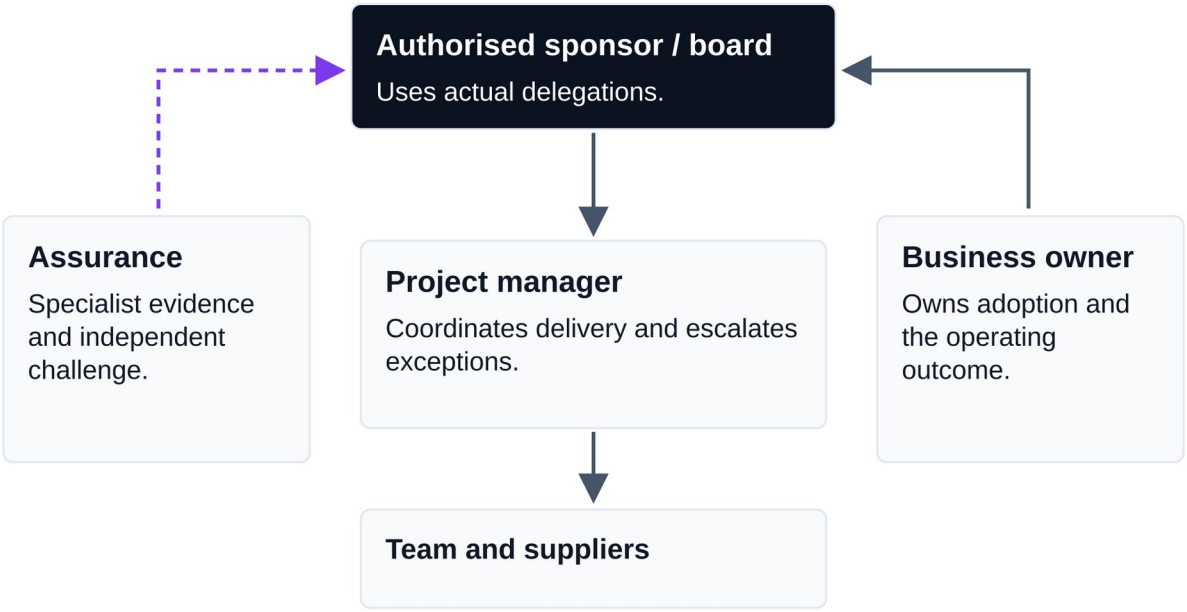


Figure 11 | Separate authority, delivery and independent advice

Purpose. Prevent accountability from disappearing between the people funding, delivering and using the change.

Role titles vary between organisations. Start with the work that must be owned, then map it to actual positions. The following allocation is illustrative and must fit the organisation's governance and delegations.

Table 2.3 | Project roles and contributions

Role	Main contribution	Question the role should answer
Sponsor	Protects the business purpose and obtains organisational decisions	Is this still a worthwhile investment?
Project manager	Integrates plans, people, evidence and delivery controls	What must change to deliver credibly?
Product or business owner	Orders needs and resolves product priorities within authority	Which capability matters most next?

Role	Main contribution	Question the role should answer
Benefit owner	Owns a measurable improvement and the changes needed for it	Is the intended improvement occurring?
Service owner	Accepts ongoing operation, support and lifecycle responsibility	Can we sustain this service?
Delivery team	Designs, creates, checks and improves the outputs	What evidence shows this work is done?
Specialist advisers	Assess relevant issues and recommend controls	What expertise or obligation must inform the decision?
Assurance reviewer	Challenges evidence with appropriate independence	How much confidence should decision-makers place in the case?

Make AI delivery responsibilities visible

The roles above still apply, but generic titles can conceal work that nobody has accepted. Name who will:

- Maintain the source or data inventory, resolve content issues, obtain access decisions and confirm the information available for evaluation.
- Prepare the evaluation method and cases, arrange capable reviewers, preserve results and explain limitations to the acceptance owner.
- Manage supplier evidence, release changes, usage charges and the handover of contractual obligations.

These are responsibilities to allocate, not a requirement to create three new positions. In the fictional case, an evaluation lead coordinates E01; the business owner decides task suitability; the access-control owner decides the permissions requirement; and the service owner accepts operational responsibilities. Delivery, acceptance and specialist decisions remain distinguishable when one person holds more than one role.

A steering group can coordinate decisions across functions, but assigning an issue to “the committee” is often too vague. Record the responsible decision authority and the person following through on the decision. A sponsor may chair the group without having unilateral authority over every topic discussed.

For a small project, the sponsor might also be benefit owner. The business owner might also accept the service. Combining roles can be efficient if conflicts are recognised and evidence still receives suitable challenge. As exposure or complexity increases, separate incompatible responsibilities and strengthen independent review.

Worked situation. The project manager reports that staff do not use the pilot because their managers have not changed the workflow. The team cannot solve that issue merely by sending another training email. The benefit owner and operational managers need to decide what work will change, provide time for adoption and remove conflicting expectations.

Do not let a RACI replace a conversation about authority. “Accountable” on a workshop chart does not create a legal, financial or organisational delegation. Confirm the actual authority for expenditure, contracts, information use, release and acceptance of residual risk.

Common failure: appointing a sponsor who receives reports but cannot resolve priority and capacity conflicts. Escalation needs access to somebody who can obtain the decision.

Use tomorrow: nominate a benefit owner and a service owner before build work accelerates. Ask each what evidence they need to accept their eventual responsibility.

Source note. UK project-delivery guidance explicitly separates project governance and senior responsibility. The allocation above is an original practical design, not a reproduction of its role requirements. [10]

2.5 Decision rights are more useful than a crowded RACI

Purpose. Allow timely decisions while keeping exposure within agreed boundaries.

A RACI identifies who is responsible, accountable, consulted and informed for an activity or output. It can reveal gaps and conflicting expectations. It rarely contains enough detail to operate a decision process on its own. Supplement it with a short decision-rights table for the choices that could materially change the project.

Table 2.4 | Decisions and their evidence

Decision	Prepares the recommendation	Decides, subject to actual delegation	Evidence needed
Change task order within approved boundaries	Delivery lead or project manager	Project manager or delegated lead	Effect on dependencies and commitments
Reprioritise product features	Product or business owner	Product authority within approved scope	Value, effort and effect on acceptance
Increase funding or use controlled reserve	Project manager and finance	Appropriate funding authority	Forecast, options and justification
Add a new user group	Business owner and specialists	Appropriate scope and use authority	Need, impact, safeguards and cost
Make a material configuration or action-boundary change	Technical and evaluation leads with the project manager	Designated change and use authorities	Affected requirements, evidence, integration, capacity, cost and fallback
Permit pilot release	Project manager assembles evidence	Designated release and use authorities	Acceptance, assurance and readiness
Accept ongoing operations	Service transition lead	Service owner with relevant authority	Support, funding, residual issues and exit

The word “appropriate” is a prompt to fill in a real role or delegation reference before using the table. It is not a substitute for an answer.

Set tolerances for matters the team may manage without returning for approval. Describe the boundary, the forecast that triggers escalation and the authority that can reset it. A tolerance is not permission to consume all contingency, ignore an obligation or hide several small changes that cumulatively alter the business case.

Illustrative rule. The project manager may reorder work within the approved scope and milestones, provided acceptance criteria and risk controls remain intact. Any forecast beyond authorised funding, proposed pilot expansion or removal of an agreed safeguard goes to the relevant authority. These examples are not PMI thresholds or organisational policy.

Distinguish routine adjustment from a material change. A correction within an approved configuration tolerance may need a recorded check and updated version, while a new source collection, supplier release or permission to take an action may require broader evaluation and a new use decision. Decide the route from the effect on requirements and authorised use. Do not require a committee for every wording edit, or assume that a technically small edit has no business effect.

Plan for absence. Identify a formally recognised deputy or an escalation route. State the response time needed and what the team should do while waiting. A time-sensitive decision may require a scheduled extra meeting; it does not justify assuming approval from silence.

Common failure: including everyone as consulted on every decision. Ask whose knowledge could materially change the choice, and consult those people. Inform others after the decision where appropriate.

Use tomorrow: list the next five material decisions and their needed-by dates. Check that each has a named decision authority, a recommendation owner and an evidence requirement.

2.6 Use responsibility maps without hiding authority

A responsibility matrix is useful where tasks or decisions cross boundaries. It becomes less useful when every minor activity contains a long list of consulted people. Start with the few deliverables and decisions where responsibility is uncertain.

In RACI, **Responsible** identifies who performs or coordinates the work; **Accountable** identifies the person answerable for the defined result; **Consulted** identifies whose input is needed; and **Informed** identifies who must receive the result. These are working definitions for the example. Always explain how your organisation uses the letters.

Illustrative matrix

Table 2.5 | Responsibility and accountability

Deliverable or decision	Sponsor	Project manager	Business owner	Technical lead	Service owner
Prepare the business case	A	R	R	C	C
Maintain integrated plan	I	A/R	C	C	C
Confirm business acceptance criteria	C	R	A	C	C
Produce and verify configuration	I	C	C	A/R	C
Prepare service handover	I	R	C	R	A
Recommend continuation	A	R	R	C	C

This is a responsibility example, not a delegation schedule. Funding, contracting, security, privacy and release authorities must be identified separately where the organisation requires them. A person may be accountable for preparing a recommendation while another authority makes the investment decision.

Use a linked work-package record for responsibilities that would make the matrix unwieldy. For E01, identify the evaluation lead and content reviewers who do the work, the business

owner who decides R01, R03 and R04, and the access-control owner who decides R02. The service owner separately considers support, monitoring and recovery. This keeps the distinction between doing the work and accepting its result visible without adding every specialist as a column.

Make it operational

Prefer one clearly named accountable person for each defined result. Split a row when it combines several different decisions. Ask each responsible contributor to accept the work and confirm capacity. Set a consultation deadline so silence does not hold the project indefinitely; material unresolved objections must still reach the decision-maker.

When one person holds several roles, preserve the distinctions in the record. A sponsor who is also the business owner may decide on both value and acceptance, but still needs evidence appropriate to each decision. Where independent assurance is required, combining roles should not quietly eliminate it.

> **Check the matrix:** If every cell says accountable, accountability is unclear. If a row has an owner but no capacity, the work remains unplanned.

2.7 Use gates to make choices, not to collect documents



Illustrative decision points. They are not mandated project phases.

Figure 12 | Use gates for decisions and transfer continuing ownership

Gate names and authority are tailored to the organisation. Benefits ownership continues beyond project closure.

Purpose. Reassess whether the next commitment is justified at moments when the project can still change course.

A gate is a decision point. A review, test or assurance activity provides evidence for that decision. A document is one way of recording the evidence. These three things should not be confused.

For the pilot, use a few decisions aligned to increasing commitment:

Table 2.6 | Evidence for continuation decisions

Decision moment	Main question	Minimum useful evidence
Explore	Is the problem worth investigating?	Problem, owner, likely value and bounded exploration cost
Commit to delivery	Is the pilot feasible and justified?	Options, scope, plan, funding, assumptions and risk approach
Begin controlled use	Is this version suitable for these users and purposes?	Evaluation results, safeguards, training, support and rollback
Close the pilot	What have we learned and what is authorised next?	Results, costs, unresolved matters and recommendation
Review benefits	Did the resulting change achieve useful outcomes?	Measured effects, costs, adoption and unintended impacts

These are original example gates, not the PMBOK process groups or a required PMI lifecycle. A project may combine them, repeat them or add decisions for procurement, architecture or deployment where warranted. An adaptive project can deliver increments between gates; governance need not wait for a large final release.

In the fictional schedule, controlled pre-release evaluation is planned within integrated tests in weeks 8-9, including eight participants from the future 40-user cohort. Those sessions take place within the authorised evaluation setting. Their evidence supports the decision about supervised use in weeks 10-11; the existence of a test result does not itself grant that permission.

Benefits review can continue after the project has closed. At closure, transfer the measure, review date and resulting decision to the business or service owner, together with any funded actions. The final card in the illustration therefore represents continuing ownership, not an obligation to keep a temporary project team in place indefinitely.

Every gate should permit a meaningful outcome: proceed; proceed with explicit conditions; revise and return; pause; or stop. Conditions need owners, evidence and deadlines. State whether they must be satisfied before the next activity or may be completed afterwards. “Approved subject to everything being fixed” is too ambiguous to operate.

Worked situation. User evaluation shows faster retrieval but repeated use of outdated sources. The available choices include limiting the source collection, delaying release while ownership is corrected, or ending the pilot if maintenance cannot be sustained. The decision should consider usefulness and operational reality together.

Tailor the intensity by consequence, uncertainty, reversibility, dependency and organisational capacity. A small financial value can still justify strong review if the effect on people is substantial. A large repetitive rollout may be controlled efficiently when the solution and operating context are well understood.

Common failure: scheduling a gate immediately before a promised launch, leaving no practical option to address findings. Allow time for evidence to affect the decision.

Use tomorrow: replace a gate's document checklist with three questions the decision-maker needs answered. Retain only evidence that helps answer those questions or meets a real requirement.

2.8 Assurance increases confidence; authority remains explicit

Purpose. Give decision-makers a credible view of whether the evidence supports the proposed action.

Assurance examines how much confidence should be placed in a claim. The claim might concern delivery feasibility, financial accuracy, security, accessibility, operational readiness or the usefulness of an AI system. Testing contributes evidence; assurance considers whether the right questions were asked, the methods were suitable and the conclusions are justified.

Use a simple claim-evidence-action structure:

Table 2.7 | Testing confidence in a proposal

Claim	Evidence to inspect	If confidence is weak
The pilot is useful for its intended tasks	Representative task results and staff feedback	Narrow the use case or change the design
Only authorised content is accessible	Access tests, configuration and source ownership	Correct access before controlled use
Staff can recognise when to seek help	Observed exercises and support records	Improve workflow and training
The service can be operated	Named owner, support plan, costs and incident procedure	Resolve readiness or limit duration

Make the review proportionate and sufficiently independent of the decision under examination. Independence can mean a peer outside the delivery team for a small initiative, or a formally commissioned independent review for a major programme. Avoid describing a supplier's self-assessment as independent assurance of the purchaser's outcome.

Give reviewers the question, scope, criteria, evidence access and expected output in advance. Ask them to distinguish confirmed findings, unsupported claims and areas they did not examine. Record management's response, including disagreement and the reason. A recommendation without an owner or due date is unlikely to change the project.

Translate the findings into the existing project controls. A limitation may require narrower scope, a changed requirement, another work package, additional review capacity or a later release decision. Record the responsible person, needed evidence and schedule effect. Link the specialist finding rather than copying the same context into another assessment form. The specialist discipline retains its own judgement and authority; the project manager integrates its delivery consequences.

Worked situation. The technical team reports a high average accuracy score. A reviewer asks how often the system answers out-of-scope questions, whether the sample reflects real use and whether severe failures were hidden within the average. The assurance value comes from examining the inference, not simply repeating the score.

A favourable review does not automatically authorise expenditure or use. A decision-maker may also need approvals from separate authorities. Conversely, a project team should not treat an executive's enthusiasm as evidence that unresolved technical findings have been addressed.

Common failure: counting reviews as a measure of assurance quality. Review coverage, evidence quality and completed responses matter more than the number of meetings.

Use tomorrow: select one claim on the dashboard and ask, “What evidence would cause us to withdraw this statement?”

Source boundary. GovS 002 provides a public example of assurance within a wider governance and delivery system. The claim-evidence-action method here is an original learning device. [10]

2.9 Keep the reason for a decision, not only its outcome

Purpose. Preserve enough context to act consistently and revisit a choice when conditions change.

A decision log and a change log answer related questions. The decision log records what was decided, by whom, why and with what conditions. The change log tracks proposed alterations to the approved position, their impact assessment and their disposition. One tool can hold both, provided a proposed change is not mistaken for an approved decision.

Use a compact decision record:

Table 2.8 | Reusable working record

Field	Example
Decision	Limit the first pilot to one approved content collection
Reason	Ownership and update evidence are not yet available for the other proposed collections
Alternatives	Delay all use; proceed with all content; proceed with a narrower scope
Authority	Named person with the applicable use and scope authority
Evidence	Evaluation report and source-owner confirmations, with versions; links to configuration C01 and evaluation plan E01 where relevant
Conditions	Train users on the restricted coverage; monitor unanswered questions
Consequences	Lower breadth; stronger source control; possible additional phase later
Review trigger	Ownership established and evaluation completed for another collection

State the assumptions and dissent that matter. A record can acknowledge that a specialist recommended delay while the authorised decision-maker chose a narrower controlled pilot. Record the actual reasoning and conditions; do not rewrite disagreement into a false consensus.

For a proposed change, assess scope, acceptance, time, cost, resources, benefits, dependencies and risk together. A “free” feature can still create testing, support and maintenance work. An urgent change can use an expedited decision route where one exists, followed by a complete record. Urgency does not create a new delegation.

After approval, update the affected plan, baseline, backlog, contract or acceptance criteria. Retain previous versions and make the effective date clear. Explain the change to the people who must work differently. The authorisation and its implementation are separate steps, both needing completion.

For a material AI configuration change, identify which prior evidence still applies and which requirements must be checked again. Preserve the previous results alongside the new configuration record so a later reviewer can understand the transition. Where the team concludes that a limited change does not affect a requirement, record the reason proportionately. The full evaluation need not be repeated automatically, but continued reliance on earlier evidence should be explainable.

Worked situation. A new requirement for multilingual answers appears mid-pilot. The business owner can rank the request, but adding it changes evaluation coverage, support expectations and potentially the target users. Assess it before deciding whether to trade existing scope, extend the project or propose later work.

Common failure: recording only “approved”. Six months later nobody can determine what was authorised, against which evidence, or whether the conditions were met.

Use tomorrow: take one recent major decision and reconstruct its alternatives, authority, evidence and conditions. Use the gaps to improve the next record.

03 Define scope and quality

Make the intended result observable. Connect the business need, deliverables, requirements and evidence so completion can be recognised.

3.1 Turn the ambition into a scope agreement

Purpose

Scope describes the work and deliverables the project is authorised to provide. It protects the intended outcome by making boundaries visible. The product scope describes what the resulting service or product will do; the project scope includes the work needed to deliver, introduce and hand it over.

A useful scope agreement answers six questions: what problem is being addressed; which people or processes are included; what will be delivered; what is deliberately excluded; what constraints apply; and how acceptance will be decided. Write boundaries before estimating the full project. Otherwise, people can approve the same budget while imagining different products.

For an AI capability, identify the tasks it may support, the information available to it, its permitted actions and the operating setting. Assistance, a recommendation and authority to perform a transaction are different scope choices. State which one is proposed. A broad aim such as “help staff find answers” does not establish those boundaries.

Method

Start with the desired change in business performance. Name the users and operating setting. Break the proposed solution into observable deliverables. Record assumptions separately from commitments. Identify interfaces with other teams, systems and projects. Ask the future service owner to review the boundary: many omissions emerge only when someone considers day-to-day support.

For the fictional pilot, the business question is whether staff can find approved internal guidance more effectively while retaining appropriate human judgement. The pilot tests that proposition with 40 nominated staff and a defined content collection. A positive pilot does not itself authorise organisation-wide deployment.

Table 3.1 | Scope boundaries

In the pilot	Outside the pilot	Boundary evidence
One approved internal guidance collection	Public web crawling or personal email ingestion	Content inventory and permission record
Answers with inspectable source references; users remain responsible for the task	Autonomous decisions, executing business transactions or external publication	Test cases, action permissions and interface design
Supervised use by 40 nominated staff	Access for all staff or the public	User list and access configuration

In the pilot	Outside the pilot	Boundary evidence
Evaluation, training and a continuation recommendation	Permanent operating service and enterprise rollout	Acceptance pack and decision record

A boundary to resolve early

Is source-content remediation included? If the pilot depends on corrected policies but no team owns that correction, the project has a hidden dependency. Either fund the work, secure a dated commitment from its owner, reduce the collection, or change the plan.

Make the completion condition clear. “Content ready” should identify the approved collection, its owner, the material issues resolved and the access behaviour to demonstrate. Uploading documents does not by itself complete that work. If another team owns a prerequisite, obtain its commitment and needed-by date rather than placing an unassigned task in the project schedule.

Tailor the agreement to the capability

Purchased AI inside a service platform. A service team buys a ticket-summarisation feature. The supplier's demonstration informs the choice, but the project still evaluates the organisation's tickets, permissions and workflow. Specify the supplier evidence, evaluation access, release-change arrangements, operating charges and exit obligations. The project may evaluate observable behaviour without inspecting model internals. If the available evidence cannot support the intended use, narrow it or choose another option.

A predictive feature in a larger project. A warehouse upgrade adds demand forecasts that planners review before placing orders. Keep the wider project's baseline and change process. Specify the periods and item groups used to compare the forecasts with the existing method, the business consequence of material errors, and how planners will use the result. Put evaluation and operating ownership in the relevant work packages rather than creating a separate AI administration structure.

An assistant permitted to take actions. Extending an advice-only assistant to create purchase requests changes the authorised capability. Define permitted transactions, authority limits, confirmations, identity, audit records and treatment of failed or duplicate actions. The project manager adds the integration, test, reviewer and recovery work to the plan. Retaining advice-only assistance remains a legitimate option if the additional work and consequences are not justified.

Common failure

“Build an AI assistant” is a solution label, not a scope agreement. It says little about permitted use, evidence quality, service ownership or completion. Keep the scope statement short enough to read, but link it to detailed requirements when ambiguity has material consequences.

Practical test

Ask three stakeholders independently what they expect at completion. Any material difference belongs in the scope conversation now.

3.2 Make requirements testable and traceable

Purpose

Requirements connect a business need to something that can be designed, delivered and checked. Traceability lets you explain why a requirement exists and what evidence demonstrates that it has been met. NASA's engineering guidance illustrates the relationship between identified requirements, their sources and planned verification methods. [11]

Method

Capture the need in the user's language, then make the expected behaviour observable. Give each significant requirement a stable identifier, owner and acceptance method. Record constraints such as permissions, accessibility, response time and recoverability alongside functions. A functional demonstration alone does not establish that the service is suitable for use.

Avoid requirements that combine several unrelated promises. “Fast, accurate, secure and easy to use” needs decomposition. For each part, ask who will judge it, under what conditions, using what evidence and against what agreed threshold. Where the threshold is unknown, commission discovery rather than writing an invented number as a commitment.

For this fictional pilot, the business owner decides task suitability, the access-control owner decides the permissions requirement, and the service owner separately accepts operational responsibilities. These allocations are examples, not universal delegations. The E01 evaluation plan below defines the conditions and method; it does not report completed results.

Table 3.2 | Requirements and acceptance evidence

ID	Need and requirement	Evidence and acceptance condition	Decision owner
R01	A user can inspect the approved source used for an answer	In the agreed answerable cases, source links resolve and reviewers find no materially unsupported statement or omitted material condition under the E01 rubric.	Business owner
R02	Users cannot retrieve content outside their permissions	Positive and negative account tests match the permission matrix. An observed unauthorised disclosure blocks this pilot's release pending correction and re-evaluation.	Access-control owner
R03	Missing or ambiguous evidence receives an appropriate response	Unanswerable cases state the limitation and approved referral. Ambiguous cases seek clarification or state the unresolved limitation, rather than silently choosing between conflicting instructions.	Business owner
R04	Staff can complete the intended task with the service	Representative users complete the agreed tasks to the E01 quality standard. Observe total task time, including checking and correction, for the value decision.	Business owner

Worked example

A stakeholder requests “more accurate answers”. The team asks which questions matter and what a consequential error would look like. The agreed requirement becomes an evaluation across selected task categories, including answer support, omissions and appropriate

refusal. The team records the test data, scoring method, reviewer and intended use. A high overall average cannot compensate for a serious failure in a restricted-use category.

Use traceability in both directions. A requirement without delivery or test coverage is a gap. A feature without a requirement or explicit design rationale may be unnecessary work. When a requirement changes, inspect linked designs, tests, training and operational controls before approving the change.

Link each requirement to the work package that delivers it, the person responsible, its baseline and forecast dates, the configuration evaluated, the evidence and the acceptance owner. A decision owner alone does not tell the team who will do the work. For a small project, use a single linked requirements and delivery record rather than separate registers containing the same fields.

Keep task quality and benefit separate

A usable answer can meet its requirement while offering little time saving. Conversely, a fast answer can be unacceptable because it omits a condition that changes the correct action. Record task quality and total effort separately, then bring both into the business decision. Include support and maintenance outside the individual user's task when assessing the benefit.

A pass applies to the evaluated configuration, tasks and conditions. Before extending the claim to another model release, collection, user group or action, assess what evidence is still relevant and what additional work is needed.

Common failure

Treating a requirement register as an inventory that nobody revisits. Review it at design decisions, demonstrations, acceptance and change assessment. For a small project, one table with links can be sufficient.

3.3 Build a work breakdown structure that exposes missing work

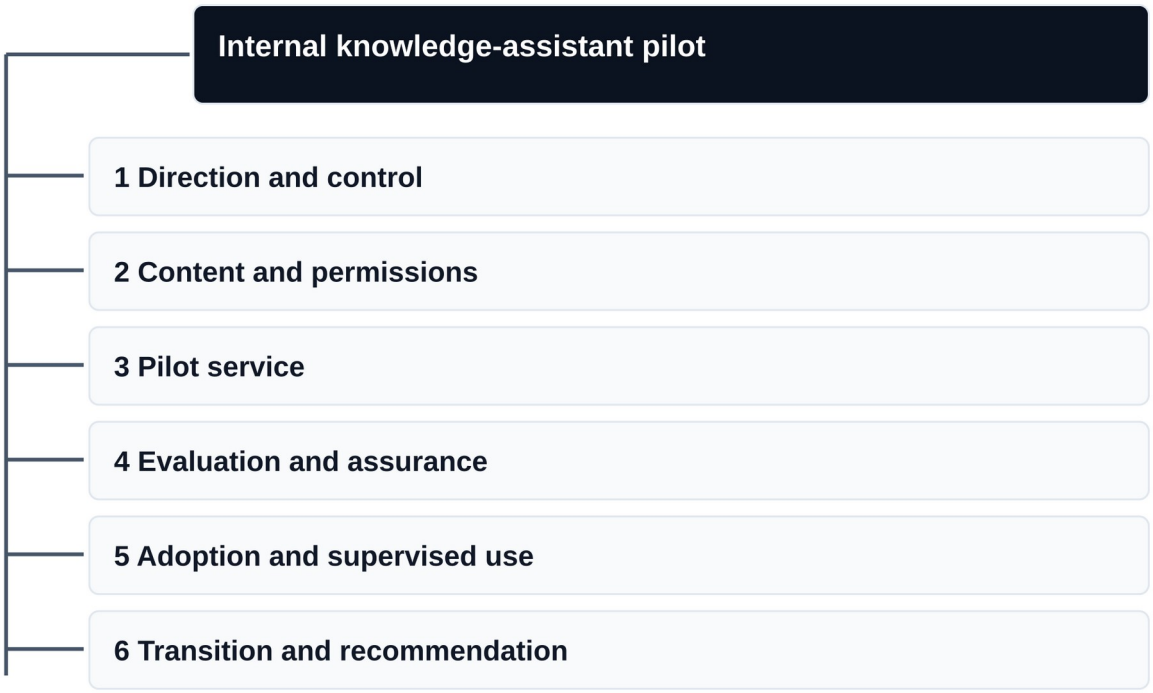


Figure 13 | Decompose all the project work into manageable components

All six scope components remain visible. The work breakdown structure describes outputs and work, not the schedule.

Purpose

A work breakdown structure, or WBS, decomposes the authorised scope into manageable components. A product-oriented hierarchy can connect scope to responsibility, estimating and reporting; NASA provides an accessible example of this approach and its supporting WBS dictionary. [12]

The WBS explains what must be delivered. The schedule explains when activities happen. An organisation chart explains who reports to whom. These are related views, but substituting one for another hides different kinds of information.

Method

Start with the complete project scope. Divide it into major deliverables, then into packages small enough to estimate, assign and accept. Include management, assurance, adoption, transition and closure work where they form part of the authorised scope. Check that the parts cover the whole scope without counting the same work twice. Stop decomposing when extra detail no longer improves ownership or control.

The illustration and worked structure below use the same six components. A diagram of the product alone would omit some of the work required to authorise, introduce and close the pilot.

Worked structure for the pilot.

Table 3.3 | Six components of the pilot work

WBS component	Examples of work packages
1. Project direction and control	Charter; integrated plan; reporting; decision records; closure
2. Content and permissions	Content inventory; owner approvals; remediation; access rules
3. Pilot service	Retrieval configuration; identity and interface integration; source display; monitoring; configuration record
4. Evaluation and assurance	Test design; access testing; response evaluation; controlled user validation; specialist input; evidence pack
5. Adoption and supervised use	User recruitment; guidance; training; feedback support
6. Transition and recommendation	Service runbook; cost model; ownership; continuation or retirement options

Each work package needs a small dictionary entry. For “4.3 response evaluation”, record the output, scope boundary, responsible person, prerequisites, estimate, quality criteria and acceptance authority. A usable output might be an evaluation report with versioned test evidence and unresolved limitations. “Do testing” provides much less control.

In the fictional pilot, the evaluation lead is responsible for that package. It depends on an available C01 configuration, prepared E01 cases, approved access and booked reviewer capacity. Completion means that the planned evidence, scoring decisions and limitations are recorded and presented to the relevant acceptance owners. It does not require the report to recommend release. An experiment can be complete when its credible finding is that the approach is unsuitable.

Locate each piece of work once. The content owner may prepare the permission matrix under component 2 while component 4 verifies the resulting behaviour. Link the dependency and evidence rather than estimating the same test twice. For a purchased capability, configuration, integration and evaluation may replace model-building packages without changing the need for adoption and transition.

A useful workshop

Give each specialist five minutes to identify work missing from the structure. Ask operations about support, users about training, procurement about supplier obligations and assurance about evidence. Then test whether each addition is required for the authorised result or should be proposed as a change.

Common failure

Decomposing to hundreds of tiny actions before the delivery approach is understood. Detailed task lists become expensive to maintain and can create false precision. Plan near-term packages in detail and leave later work at a level appropriate to current knowledge.

Practical test

Can every material budget line, acceptance item and delivery responsibility be connected to the WBS? Unconnected items deserve investigation.

3.4 Design quality and acceptance before testing begins

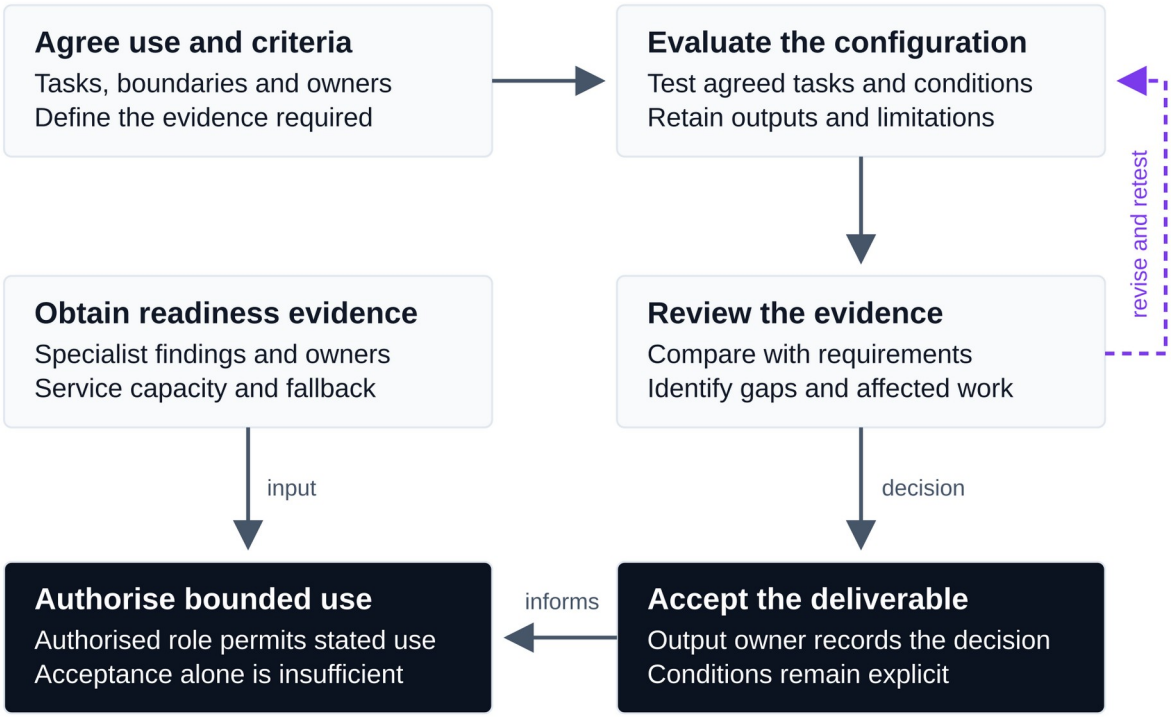


Figure 14 | Acceptance and bounded-use authorisation remain distinct decisions

Acceptance, permission for use and operational handover may be recorded together, but answer different questions.

Purpose

Quality means the delivered result is suitable for its agreed purpose and satisfies the relevant requirements. Define it while designing the work. Testing at the end can reveal defects, but cannot recover every decision that made those defects likely.

Distinguish verification from validation. Verification asks whether specified requirements have been met; validation asks whether the result works for its intended use and users. NASA's public systems-engineering material illustrates both kinds of evidence. [11] Acceptance is the authorised decision to take a deliverable, subject to agreed criteria and any recorded conditions.

Method

Agree criteria and evidence owners. Build preventive practices into the work: peer review, controlled versions, representative test data, clear interfaces and early user feedback. Perform checks during delivery and resolve the causes of defects. Before acceptance, consolidate the evidence, limitations and outstanding items for the person with authority to decide.

Table 3.4 | Dimensions of quality

Dimension	Pilot example	Evidence to retain
Functional behaviour	Questions receive an answer, limitation or referral	Versioned scenarios and observed results
Information quality	The answer is supported by the cited material	Reviewer assessment and source comparison
Access and confidentiality	Restricted material stays restricted	Permission tests, including negative cases
Usability and accessibility	Intended staff can complete representative tasks	Observations and unresolved barriers
Operability	Support can diagnose, disable and restore the pilot	Runbook exercise and named support owner

Worked example

A demonstration answers ten prepared questions correctly. That is encouraging evidence for those questions under those conditions. It does not establish reliability across all staff queries, permission states or future content changes. The acceptance plan therefore includes ambiguous questions, stale material, restricted content, unsupported questions and operational failure scenarios. Sampling and thresholds should reflect the intended use and consequence of failure.

Record defects by impact and urgency, with an owner and retest evidence. A conditionally accepted item needs explicit conditions, a due date and a person responsible for the residual exposure. State whether the condition must be met before use or can be completed afterwards within the authorised boundary. Do not quietly redefine a serious defect as a “future enhancement” to meet a milestone.

E01 evaluation plan for the fictional pilot

The following is a completed fictional evaluation design. Its case counts, repeated runs, participant selection and release conditions are local teaching assumptions. They are not industry thresholds or a claim of statistical assurance, and no test results are invented. The evaluation lead turns the plan into estimated work, committed capacity and an evidence date.

Table 3.5 | Completed evaluation design E01

E01 version 1	Planned evidence
Evaluated configuration	C01 version 1 identifies supplier release S1, prompt P1, retrieval settings R1, collection K1, permission matrix A1 and interface I1. Replace these teaching identifiers with actual recorded versions before applying the plan.
Task set	60 specified scenarios: 30 answerable tasks, 10 ambiguous or unsupported requests, and 20 positive or negative access checks. Include ordinary wording, stale-content challenges and required policy conditions.
Method and variation	Run the 40 response scenarios three times on the same release. Retain all runs and report variation, rather than selecting the best output. Test the 20 access scenarios using the stated accounts. Report scenario counts and run counts separately.

E01 version 1	Planned evidence
Judgement and traceability	Two content reviewers calibrate the rubric on a small separate set. Retain outputs, source passages, judgements and reasons. The evaluation lead resolves scoring disagreements with the relevant content expertise before presenting the results.
User validation	During integrated tests F in weeks 8-9, eight participants from the planned 40-user cohort cover the nominated roles and access profiles in controlled tasks. Compare quality, accessibility and total effort with the current workflow. This precedes the full supervised pilot in weeks 10-11.
Failure and decision rule	Assess R01-R03 against their applicable scenarios and repeated runs, and R04 through user validation. An observed access breach or materially misleading answer blocks this pilot's release. Report other failures by task and severity; correct them or explicitly narrow the permitted use before acceptance and authorisation.

For each case, record the intended task, relevant source and permission state, expected behaviour and the evidence needed to judge it. A materially misleading answer in this example is one whose unsupported statement or omitted condition could change the user's required action or decision. Wording preferences should not be counted as equivalent to that failure. A correct limitation or request for clarification can satisfy R03 when an answer would lack support.

Reviewers should agree the scoring meanings before judging the final results. For R04, observe whether the user completes the specified task correctly and can inspect the supporting information. Record checking, correction and escalation as part of the task effort. Successful technical checks cannot substitute for those observations. End-of-week-nine evidence supports the decision about supervised use; it does not itself grant permission to operate the pilot or expand it.

Choose cases from the intended work, including situations where the appropriate response is to stop or refer. Hold back some cases while improving the service so the final check does not simply replay examples used to tune it. Explain the limits of a small test set: observing no failure does not prove that failure is impossible. NIST's generative-AI profile supports evaluation in relevant conditions and documenting limits on generalisation. [20] The specific E01 design is this guide's practical example, not a NIST prescription.

Separate acceptance from permission to use

Technical verification, business acceptance and permission to operate serve different decisions. The project manager assembles the evidence and makes the route visible; each authority retains its responsibility.

Table 3.6 | Decisions and their evidence

Decision	What it establishes
Accept the deliverable	The nominated owner accepts the agreed output against requirements, recording conditions, exclusions and unresolved limitations.
Authorise bounded use	The authorised role permits specified users, tasks and operating conditions after considering the acceptance evidence, specialist findings and readiness.

Decision	What it establishes
Accept operational handover	The service owner accepts support, monitoring, recovery, fallback and continuing costs. This may be limited to the pilot period.

Permission for a pilot does not extend to enterprise rollout. Record its boundary, duration, monitoring owner and pause trigger. Resolve material conditions or narrow the use before authorisation; an ambiguous “conditional pass” is insufficient. One meeting and a single decision record may capture several decisions on a small project, provided their purposes and accountable roles remain clear.

Plan human review as work. Reviewers need the competence, time, information and authority to recognise relevant errors and stop or override the workflow. Test whether those arrangements are usable, include reviewer capacity in the plan and count the effort in the value calculation.

Re-evaluate when the evidence boundary changes

After a material change, identify the affected requirements and preserve the previous configuration and results. Agree the scope of re-evaluation, its owner and the evidence required. Include the work in the schedule and cost forecast before presenting the changed capability for acceptance. Repeating an unaffected test can waste effort; continuing to rely on evidence that no longer describes the service can leave the decision unsupported. Use the agreed change tolerances and record the reason for the chosen response.

Common failure

Confusing a completed technical test with permission to use the service in every context. The accepted deliverable and authorised pilot use can each be bounded to specified users, data, functions, conditions and period. Expansion is a new decision with new evidence needs.

Practical test

Could somebody outside the delivery team understand why acceptance was justified from the evidence pack alone?

3.5 A small toolkit for diagnosing quality problems

Choose a technique because it clarifies a decision. The following tools can be drawn in a notebook or built in ordinary analysis software. They are useful for investigating patterns, rather than establishing certainty by themselves.

Table 3.7 | Quality investigation techniques

Technique	What it helps you see	Example use and caution
Check sheet	Consistent observations grouped by defined categories.	Record unsupported answers, stale sources, access failures and usability issues. Define categories so reviewers apply them similarly.
Pareto chart	Which categories contribute most to a measured total.	Show defect counts or rework hours. The largest category by count may not be the most consequential.

Technique	What it helps you see	Example use and caution
Cause-and-effect diagram	Several plausible causes grouped for investigation.	Examine content, retrieval, interface, user behaviour and operating process. A diagram proposes causes; testing is still needed.
Five whys	A deeper explanation of a recurring event.	Trace why a source was outdated. Avoid assuming one linear cause or blaming the last person in the chain.
Run chart	How a measure changes over time.	Track task time or defect age consistently. A short sequence can vary by chance.
Control chart	Variation against statistically established process limits.	Use only with suitable process data and assumptions. Control limits are not the same as acceptance specifications.
Sampling	Evidence from a selected subset of work.	Include representative and high-consequence cases. Passing a sample does not prove every future case will pass.

A worked investigation

Suppose reviewers record twenty defective answers. Twelve use outdated sources, five omit an important qualification and three cite unrelated material. The count suggests source freshness deserves attention. But one omitted qualification may have a larger consequence than several stale administrative details. Review severity as well as frequency.

The team investigates why outdated sources remain available. It discovers that the source inventory has no named owner and the update mechanism does not remove superseded versions. The response addresses both responsibilities and technical behaviour. Retesting checks the corrected mechanism and includes previously unseen cases.

Compare observations on a stated basis

Keep the measurement conditions visible. If the task mix, source collection, configuration or reviewers change, investigate that change before interpreting a movement in the figures as improved or deteriorating performance. Report controlled evaluation results separately from observations during supervised use, while using both to guide the project. In the fictional pilot, E01 supports the pre-release decision; weeks 10-11 add evidence about work in practice.

Choose a proportionate response to a pattern. A growing number of unsupported requests may call for clearer user guidance, a scope decision or a change to the source collection. It does not automatically justify adding new content, changing the model or accepting a larger operating burden.

Close the learning loop

Retain the defect, assessed consequence, cause, corrective action, owner, affected versions and retest result. Monitor recurrence after release. If the correction adds substantial operating work, update the forecast and benefits case. Quality improvement can create cost or resource effects that must remain visible.

> **Avoid the shortcut:** A root-cause meeting without evidence and follow-through can produce a plausible explanation that never changes performance.

04 Plan time, resources and costs

Build a feasible delivery model. The worked network and financial examples show how to analyse dependencies, capacity, progress and forecast cost.

4.1 Estimate honestly when knowledge is incomplete

Purpose

An estimate is a reasoned expectation under stated assumptions. It is useful when it helps a decision; it becomes misleading when a tentative number is presented as an unconditional promise. Separate the estimate, the confidence attached to it and the commitment an authorised person is willing to make.

Method

Select an estimating approach suited to the available evidence. Analogy uses a comparable piece of work. A parametric estimate uses an observed relationship, such as effort per validated content item. Bottom-up estimating combines work-package estimates. Expert judgement helps where data is limited, but should expose its assumptions and be challenged by another perspective.

Table 4.1 | Selecting an estimating method

Situation	Useful starting approach	Question to challenge
Early concept	Comparable projects and a range	Is the comparator genuinely similar in complexity and constraints?
Repeated work	Unit rate multiplied by quantity	Is the rate stable at this volume and quality level?
Defined packages	Estimates from the people doing the work	Are integration, review, rework and handover included?
High uncertainty	Scenarios, experiments and staged commitment	Which unknown could most change the decision?

Worked example

Reviewing 200 content items at an observed average of 12 minutes each suggests 40 hours of review effort. That calculation excludes remediation, content-owner decisions and access changes. If 20% of items require 30 additional minutes, add 20 hours. The useful estimate is now at least 60 hours plus the separately identified review and coordination work, with uncertainty around the rework rate.

A simple three-point teaching example uses optimistic, most likely and pessimistic durations of 2, 4 and 8 days. The triangular mean is (2 + 4 + 8) / 3 = 4.67 days. The conventional PERT weighted average is (2 + 4×4 + 8) / 6 = 4.33 days. Neither is a confidence bound or a delivery guarantee. The formulas encode different assumptions and cannot compensate for poor inputs or correlated risks.

Use the next decision to improve the estimate

A small content sample, supplier discovery session or integration prototype may reduce uncertainty more efficiently than a longer estimation meeting. Record what evidence would cause the estimate to change.

Estimate the work that creates credible AI evidence

For a configured or purchased capability, estimate the configuration, data or knowledge preparation, integration, evaluation, adoption and operational preparation that the organisation must still perform. Do not add model-training effort unless training is actually in scope. For a predictive model that does require development, separate obtaining representative data, preparing it, developing the model and evaluating it in the intended workflow. The estimate should reveal those activities rather than bury them in a single “AI development” allowance.

Estimate evaluation from the agreed tasks, configurations and review method. Include preparing expected outcomes, running tests, checking results, resolving disagreements and repeating affected tests after changes. The approved pilot evaluation design E01 v1 provides a quantity basis; a change to its coverage or repeat runs changes the estimate. Confirm the evaluator's available hours separately from the computing time needed to run the tests.

An experiment can have a bounded time and spending commitment without a promise that its hypothesis will succeed. Agree the question, stopping condition, evidence to produce and decision to inform. At the boundary, use the results to refine the next estimate, change approach or stop. Do not schedule repeated experimentation indefinitely until an attractive demonstration appears.

The 200-item content calculation and the three-point duration exercise above are independent teaching examples. They explain estimating techniques and are not additional amounts to be added to the connected pilot's budget.

Common failure

Adding an unexplained percentage everywhere and then adding another project contingency on top. Record risk allowance explicitly so reserves are visible and not counted twice. Update estimates as evidence improves, retaining the reason for each material revision.

4.2 Model the schedule before drawing the timeline

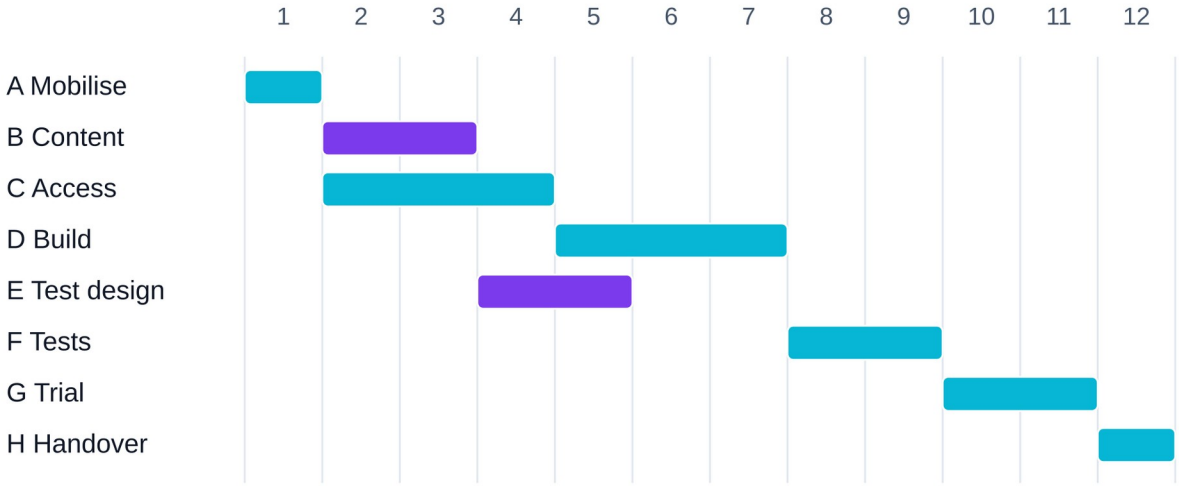


Figure 15 | The twelve-week network shown against time

Purpose

A useful schedule is a model of work, dependencies, resource availability and uncertainty. The Gantt chart is one way to display that model. GAO's schedule guidance explains why reliable schedule logic matters to both forecasting and the assessment of cost consequences. [13]

Method

Derive activities from work packages. Define completion criteria, estimate duration, identify predecessors and confirm available resources. Add real decision, procurement and external dependencies. Review the resulting completion date before committing to it. A deadline can constrain a decision, but entering that deadline in software does not make the work feasible.

Distinguish **effort** from **duration**. Two days of specialist effort may take two weeks of elapsed time if the specialist is available for only one day each week. Waiting for an external approval can also consume duration without much project effort. Record these facts instead of hiding them inside a padded task estimate.

Table 4.2 | Dependency relationships

Dependency type	Meaning	Example
Finish to start	The successor begins after the predecessor finishes	Final access testing starts after configuration is ready
Start to start	The successor can begin once the predecessor begins	Draft user guidance starts as supervised testing begins
Finish to finish	The successor cannot finish before its predecessor finishes	Evaluation summary cannot finish before test execution finishes
Start to finish	The successor cannot finish until its predecessor starts	Rare; use only when the operational logic requires it

Use explicit activities for meaningful waiting, review or mobilisation periods. A hidden lag is harder to explain, own and update. Avoid connecting every task directly to a milestone simply to produce the desired picture; the dependency must describe how work actually becomes possible.

Worked example

The pilot service cannot be built on the assumption that permissions will eventually be resolved. Content readiness and access controls both precede the working build. Evaluation design can progress separately once the content scope is known. The parallel work shortens the project while preserving the prerequisites for integrated testing.

Make AI dependencies usable in the plan

Name the evidence that makes each dependency complete. “Data ready” might require a usable approved extract, known limitations, access for the delivery team and an owner for correction. “Supplier ready” might require access to the evaluation environment, the supported release version, documentation and a named response commitment. A promise to supply these later is an unresolved dependency, not permission to start work that relies on them.

Separate work that can proceed in parallel from work that would invalidate earlier effort. Evaluation design can begin before the working service is finished; final execution needs the configuration, sources and permissions being proposed for use. Reserve time to correct defects and rerun affected tasks. Place specialist reviews at the point when their evidence is available and their findings can still change the design, with a response date before the relevant release decision.

For an existing platform feature, check when the organisation can enable, restrict, test or disable it and which dates the platform owner or supplier controls. For an action-taking capability, include the work to configure tool permissions, test approval boundaries and rehearse recovery before exposing it to live actions. These are activities and dependencies within the chosen delivery approach, not compulsory new phases for every project.

Common failure

Asking people for a percentage complete without asking what remains. A task can stay “90% complete” while a dependency prevents acceptance. Update actual dates, remaining duration, blockers and forecast finish dates at the same reporting cut-off.

Practical test

If one dependency slips, does the forecast change logically? If nothing changes, inspect the model before trusting the green status report.

4.3 Find the path that controls completion

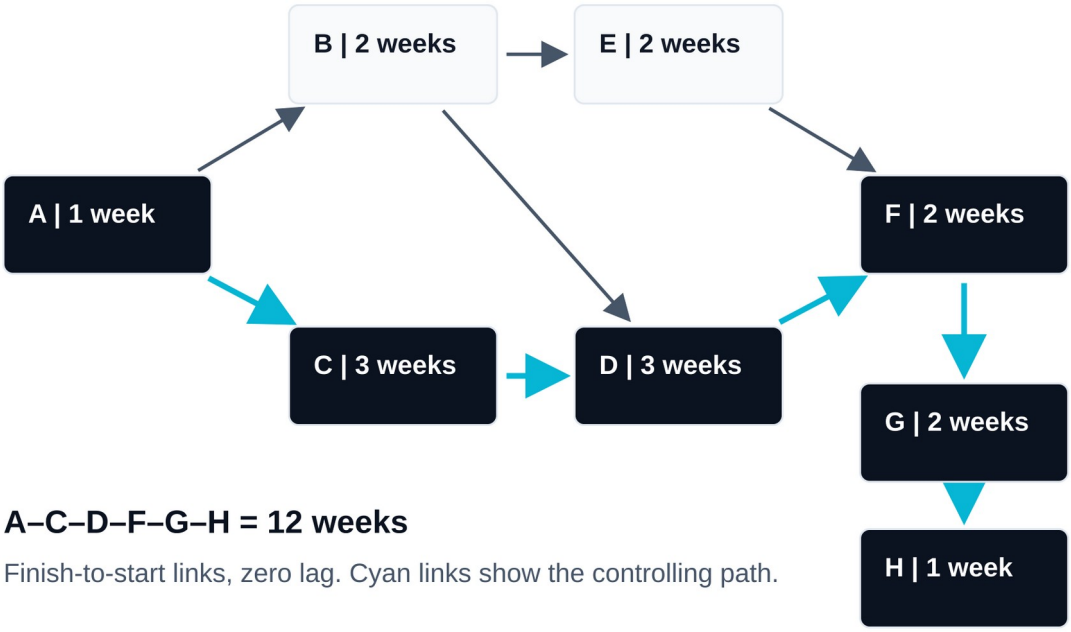


Figure 16 | A twelve-week dependency network

Activity-on-node example with finish-to-start links, zero lag and assumed resource availability. Times are elapsed working weeks.

Purpose

Critical-path analysis identifies the chain controlling the earliest finish under the schedule's assumptions. Float is calculated scheduling flexibility, not spare staff capacity. GAO discusses the importance of valid logic and attention to paths close to becoming controlling. [13]

Worked example

Time is measured in whole working weeks from time zero. Activities use finish-to-start dependencies, have no lags, and have the durations shown. Resources are assumed available without conflicts; calendars and duration uncertainty are deliberately simplified. Under those assumptions the project finishes at time 12, at the end of week 12.

Table 4.3 | Pilot activities and dependencies

ID	Activity	Duration	Predecessor
A	Mobilise and authorise pilot	1	None
B	Prepare content collection	2	A
C	Establish access controls	3	A
D	Build working service	3	B, C
E	Design acceptance evaluation	2	B
F	Run integrated tests	2	D, E
G	Run supervised pilot	2	F
H	Complete handover and decision pack	1	G

Forward and backward pass

Table 4.4 | Forward and backward pass results

ID	ES	EF	LS	LF	Total float
A	0	1	0	1	0
B	1	3	2	4	1
C	1	4	1	4	0
D	4	7	4	7	0
E	3	5	5	7	2
F	7	9	7	9	0
G	9	11	9	11	0
H	11	12	11	12	0

ES and EF mean earliest start and finish; LS and LF mean latest start and finish without moving the modelled project completion. For these finish-to-start links, an activity's ES is the maximum EF of its predecessors. Its LF is the minimum LS of its successors. Total float equals LS minus ES, equivalently LF minus EF.

The controlling path is **A → C → D → F → G → H**, lasting 12 weeks. B has one week of total float; E has two. That does not mean every non-controlling task can consume its float independently without examining shared paths and resources.

Management response

If C needs an extra week and nothing else changes, completion moves to week 13. Finishing B a week earlier does not recover that delay, because D must still wait for C. Recovery must address a controlling activity, a dependency, capacity or authorised scope. Parallelising dependent work may introduce rework or risk; adding people may provide no benefit where the bottleneck is an approval.

Interpret the AI work inside this network

The network describes this configured assistant, not a universal AI life cycle. Activity D includes the service configuration and integration needed for the agreed collection. E prepares the acceptance evaluation; F executes integrated checks and controlled validation against the proposed configuration. The eight controlled validation participants in E01 v1 work within F in weeks 8–9. G is the distinct supervised use by the 40-user cohort in weeks 10–11, and H completes handover in week 12.

In a purchase or a project with a limited AI feature, retain only the activities needed for that work and its consequences. The dependencies still need to express what must be true before a successor can begin. The separate supplier-release scenario later in this chapter adds an activity after F; it does not amend this baseline network.

Common failure

Treating the first calculated path as permanent. Recalculate with actual progress, remaining work and current constraints. The route controlling completion can change.

4.4 Plan capacity, responsibility and collaboration together

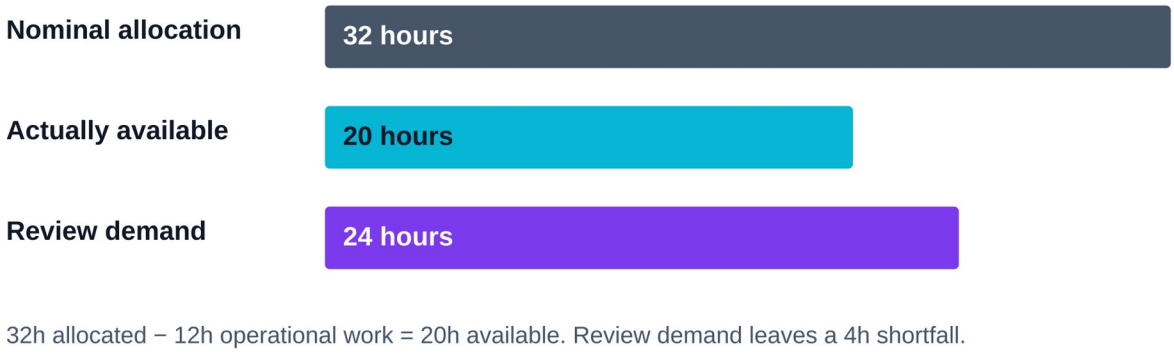


Figure 17 | Resource demand must fit actual availability

Purpose

A named person on a plan is not evidence that capacity has been secured. Resource planning connects the work to people with the time, authority, skills and support needed to complete it. Resource conflicts can invalidate an otherwise sound dependency schedule.

Method

Estimate demand by skill and period. Confirm actual availability with the people and managers involved. Account for existing responsibilities, leave, operational interruptions and review obligations. Identify scarce specialists and their substitutes. Then compare demand with capacity and resolve the gaps through sequencing, scope, additional support or a revised commitment.

Worked example

The pilot needs a subject-matter expert for 24 hours of content review across two weeks. A nominal 0.4 full-time allocation on a 40-hour week appears to provide 32 hours. However, the expert has 12 hours of committed operational work within that allocation, leaving 20 hours available. There is a four-hour shortfall before allowing for unexpected rework. Clarify the allocation instead of assuming overtime will absorb the difference.

Table 4.5 | Capability and capacity commitments

Role or capability	Required contribution	Capacity commitment to confirm
Business owner	Resolve use-case and acceptance decisions	Decision dates and review time
Delivery lead	Coordinate dependencies, forecasts and exceptions	Availability across the full pilot
Technical specialist	Configure, integrate and diagnose defects	Protected build and test periods
Content expert	Validate guidance and assess answer support	Review blocks and turnaround times
Service representative	Prepare support and transition	Early involvement and handover acceptance

A responsibility matrix can clarify who performs work, who accepts accountability, who is consulted and who needs information. Keep it focused on decisions and deliverables that are genuinely ambiguous. It does not replace financial delegations, role descriptions, technical authority or capacity agreements.

Provide the capability to evaluate and decide

Identify who designs the evaluation, who supplies domain judgement, who operates the test environment and who accepts the relevant result. The same person may perform several tasks in a small project, but record each responsibility and address any conflict that would weaken the evidence. A supplier's evaluator can explain product behaviour; the business owner still needs evidence about the organisation's intended tasks. Platform, data and content owners need capacity to resolve faults rather than only attend meetings.

Treat human review as real work. Reviewers need time, relevant competence, access to the supporting material, clear escalation routes and authority to reject or stop unsuitable use. If a reviewer cannot establish whether an answer or proposed action is acceptable, routing it to that person does not resolve the problem. Test the review process during evaluation and make the capacity commitment visible in the delivery and operating plans.

In a separate, fictional evaluation exercise, 120 outputs at six minutes each require 12 review hours. Two hours of preparation and three hours to resolve disagreements bring demand to 17 hours. A ten-hour allocation leaves a seven-hour gap. The team can reduce the agreed test scope only if the resulting evidence remains adequate for the decision; otherwise it must secure capacity or revise the date. These hours are not additions to the connected pilot ledger.

For action-taking AI, distinguish the person who approves an action from the person who later checks its result. Estimate both tasks where both are required, including periods when neither person is available. If the operating model relies on rapid intervention, confirm that the roster and access arrangements can provide it.

Working agreement

Define where decisions are recorded, how blockers are raised, what a reasonable response time is and when an issue needs escalation. Give people permission to report uncertainty and failure early. A team that must first explain why the problem is somebody else's fault will surface risks too late.

Common failure

Making the same person accountable, responsible and consulted for everything, while treating participation by other teams as free. Another is allocating every available hour: a plan with no capacity for variation tends to convert small disruptions into queues.

Practical test

Ask each contributor to confirm both their deliverable and the time they can actually provide. Resolve conflicting answers before setting the baseline.

4.5 Resolve competition for shared capacity

Several individually reasonable project plans can add up to an impossible workload. A portfolio conversation makes that conflict visible and decides which commitments take priority. Project managers should present the implications, while the person authorised to allocate the shared resource makes the trade-off.

An illustrative capacity conflict

Table 4.6 | Shared specialist demand

Initiative	Specialist demand next fortnight	Constraint	Consequence of delay
Knowledge-assistant pilot	4 days	Access review precedes integration tests.	Pilot completion may move unless another activity can absorb the delay.
Service migration	5 days	External supplier window has been booked.	Rebooking may add cost and prolong the old service.
Reporting improvement	3 days	Users want an early release; no fixed external date.	Benefit is deferred, with limited immediate disruption.

The shared specialist has eight available days. Demand is twelve. A priority score does not create four additional days. The decision could defer the reporting improvement and still need to move or reduce one day elsewhere; secure a capable substitute; or change a milestone with its consequences made explicit.

Prepare a feasible recommendation

Confirm whether each activity truly requires that specialist and whether work can be prepared by somebody else. Examine dependencies and the cost of interruption. Identify any obligation or exposure that cannot be traded away by the local group. Present a small number of feasible combinations, including which work is displaced by each.

Do not ask all three project managers to keep their original dates after allocating less capacity. Update their forecasts and communications together. A portfolio decision should flow into the actual plans, supplier discussions and user expectations.

Include recurring AI support in the demand picture. The same platform or domain specialist may be needed for an ongoing service's release checks, an incident and a new project's evaluation in the same fortnight. Those operating obligations are not spare time available to the project. Ask the resource owner to agree the allocation, escalation route and any substitute, then reflect that decision in each affected forecast.

In a smaller organisation, one person may combine product, evaluation and support responsibilities. That can be workable when the workload is visible and decisions remain clear. Combining role names does not remove the hours or create independent review capacity. The twelve-day demand example above remains a standalone portfolio exercise, not a further requirement in the pilot schedule.

Review the systemic cause

If the same specialist is repeatedly overcommitted, the organisation may need earlier demand visibility, shared service capacity, skills development or fewer concurrent initiatives. Treat the pattern as an organisational decision, rather than asking each project to solve it independently.

> **Governance question:** What are we explicitly choosing not to do now, and have the affected owners accepted the consequences?

4.6 Build a budget that distinguishes control from permission

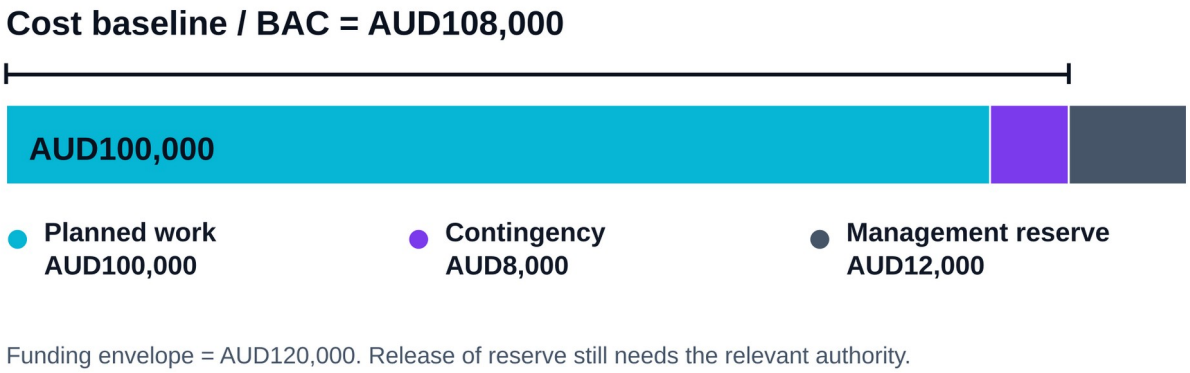


Figure 18 | Funding baseline and reserves are different quantities

Purpose

A funding envelope, cost baseline and cash-flow forecast answer different questions. Funding sets the authorised financial boundary. The baseline supports comparison of planned and achieved work. Cash flow describes when payments are expected. Treating these as the same number hides commitments and reserve decisions.

Worked budget for the fictional pilot.

Table 4.7 | Pilot cost baseline and funding

Budget element	AUD	Treatment in this teaching example
People and specialist support	72,000	Planned delivery work
Cloud, tools and environments	12,000	Planned delivery work
Evaluation activities	10,000	Planned delivery work
Training and transition	6,000	Planned delivery work
Identified-risk contingency	8,000	Within the cost baseline, governed and allocated through work-package controls
Cost baseline / BAC	**108,000**	Budget at completion for the approved baseline
Management reserve	12,000	Outside the baseline; release requires the appropriate authority
Total funding envelope	**120,000**	Total approved funding assumed for this example

This is an illustrative control structure, not a prescription for an organisation's accounting policy. Define the local treatment of contingency, reserves, internal labour, tax, capitalisation and overheads with finance. The distinction between a performance baseline and management reserve is also discussed in PMI's public earned-value material. [14]

Method

Connect estimates to scope and schedule. Identify the basis of each quantity and rate. Include mobilisation, integration, assurance, adoption, decommissioning and ongoing service costs where relevant. Time-phase the approved work, record assumptions and make reserve

ownership explicit. An estimate of project cost alone is insufficient where the investment creates substantial continuing expenditure.

Track actuals, commitments and the forecast cost of remaining work. An invoice not yet received may still represent a liability or obligation. Avoid adding all purchase orders to all actuals indiscriminately: part of a committed amount may already have been invoiced. Use a reconciliation that prevents double counting.

Worked control question

If $48,000 has been spent and the remaining work is forecast at $72,000, the expected total is $120,000. That equals the funding envelope but exceeds the $108,000 baseline. It needs an explanation and an authorised response. The existence of reserve does not automatically approve its use. A forecast overrun does not itself justify management-reserve release or resetting past variances; examine the cause and apply the agreed reserve policy.

Separate delivery cost from the cost of continuing the service

The pilot's AUD 100,000 of planned work, AUD 8,000 contingency, AUD 108,000 BAC and AUD 120,000 funding envelope remain the connected project budget. Within those amounts, make AI-related effort visible in the appropriate work packages: source preparation, configuration, evaluation, human review, integration, adoption and transition. Do not add a second generic AI allowance for costs already included.

Prepare an operating-cost estimate with the receiving service owner and finance. Identify who pays, the charging unit, expected volume, minimum commitments, review effort, support and re-evaluation. Connect usage assumptions to the expected business workload, including retries or multiple service calls where these are part of the proposed design. A supplier's price per request cannot by itself establish the cost of a completed business task. Preserve the distinction between a resource cost used for comparison and an additional cash payment.

Independent worked example: a purchased AI feature

This separate fictional purchase illustrates the calculation; none of its figures changes the knowledge-assistant pilot. Assume 1,200 completed tasks per month, two chargeable service calls per task, and human review of every completed task. All prices and effort rates below are illustrative assumptions, not market quotations.

Table 4.8 | Independent purchased-feature operating cost example

Monthly element	Assumed quantity and rate	AUD
Subscription	Fixed monthly charge	600
Usage	1,200 tasks × 2 calls × AUD 0.05	120
Human review	1,200 tasks × 2 minutes = 40 hours; AUD 80/hour	3,200
Maintenance and support	5 hours at AUD 80/hour	400
Monthly operating resource cost	**Subscription, usage and labour**	**4,320**

If one-off configuration costs AUD 8,000, the first three months cost AUD 8,000 + 3 × AUD 4,320 = AUD 20,960 on these assumptions. This is a resource-cost comparison before tax

and other organisational accounting adjustments; it is not a benefit calculation. Finance should distinguish existing internal capacity from additional expenditure.

If task volume doubles and the per-task review remains necessary, usage rises to AUD 240 and review effort to 80 hours costing AUD 6,400. Keeping subscription and maintenance unchanged for this sensitivity only gives AUD 7,640 per month. The increase is AUD 3,320, not merely the AUD 120 increase in usage charges. In a real forecast, test whether maintenance, licensing tiers and review performance also change. The practical decision is whether the operating model remains affordable and staffed at the proposed scale.

Common failure

Reporting underspend as success while deliverables are late. Unspent money can reflect delayed work. Examine cost alongside accepted progress and the forecast to finish.

4.7 Use earned value to connect spending with progress



Figure 19 | Week-six performance connects spending and completed work

Purpose

Earned value management compares planned work, completed work and actual cost on the same budget basis. It is useful when scope, work-package budgets and objective progress evidence are reliable. GAO's cost guide includes EVM as a way to integrate cost assessment with delivery performance. [15]

At a reporting date, **planned value (PV)** is the budget associated with work scheduled to be complete; **earned value (EV)** is the budget associated with work actually completed; **actual cost (AC)** is the cost incurred for that work. Earned value is a measure of budgeted work, not business benefits, revenue or the market value of the product.

Worked snapshot: end of week 6

The fictional pilot has BAC = $108,000, PV = $54,000, EV = $43,200 and AC = $48,000. These values are stipulated for the teaching example; the fact that week 6 is halfway through the schedule does not itself imply half the budget should have been earned.

Table 4.9 | Measures and interpretations

Measure	Calculation	Result and interpretation
Cost variance, CV	EV − AC	−$4,800: completed work cost more than its budget
Schedule variance, SV	EV − PV	−$10,800: less budgeted work completed than planned
Cost performance index, CPI	EV / AC	0.90: $0.90 of budgeted work for each $1 spent
Schedule performance index, SPI	EV / PV	0.80: 80% of the planned budgeted work achieved
Budget-weighted completion	EV / BAC	40% of the baseline budget represented by completed work

A negative schedule variance is expressed in money, not weeks. SPI = 0.80 does not prove that the project will finish 20% late, or that four-fifths of its benefits have been achieved. Inspect the dependency schedule to understand the date implication.

Make progress credible

Agree earning rules before reporting: accepted milestones, measurable quantities, or clear completion rules for small packages. A package does not earn value merely because money was allocated to contingency. In this example, risk-related budgets remain subject to controlled work-package allocation and objective completion evidence.

Earn progress against the agreed type of work

An authorised experiment can earn its budgeted work when the agreed test is completed, evidence and limitations are recorded, and its decision output is delivered. An unfavourable finding does not make that work disappear. It also does not earn the value allocated to a deployable service. Keep experiment completion, deliverable acceptance and realised business benefits separate in both the work breakdown and reporting.

For the pilot, successful test execution alone does not complete a requirement if its acceptance condition has not been met. Record the configuration and evidence supporting accepted work. When a change makes earlier evidence insufficient, assess the remaining work and its reporting treatment openly; do not claim completion for the revised deliverable solely because an earlier version passed. Agree the treatment through the project's control arrangements rather than changing earning rules to improve the ratio.

Common failure

Using subjective percentages, inconsistent cut-off dates or an obsolete baseline. The resulting ratios look precise while describing incompatible information. Reconcile progress, actual costs and approved changes before interpreting them.

Tailoring

A small pilot may need only accepted milestones, actual spend and a forecast to finish. Use formal EVM when its additional insight warrants the effort.

4.8 Forecast the finish and choose a response

Purpose

A variance explains what has happened. A forecast informs what to do next. The most useful estimate at completion combines actual costs, the remaining work, current risks and the intended recovery plan. Formula-based estimates provide useful comparisons, but their assumptions need to match the situation.

For the fictional pilot, AC = $48,000, EV = $43,200, BAC = $108,000 and CPI = 0.90. Public PMI guidance discusses several EAC formulas and their different assumptions. [16]

Table 4.10 | Alternative cost forecasts

Forecast assumption	Calculation	EAC
Remaining work achieves its original budget rate	AC + (BAC − EV)	$112,800
The current cost efficiency continues	BAC / CPI	$120,000
Re-estimated remaining work costs $72,000	AC + bottom-up ETC	$120,000

ETC means estimate to complete; EAC means estimate at completion. For EAC = $120,000, variance at completion is BAC − EAC = −$12,000. The original $120,000 funding envelope would be fully consumed. Reserve release still requires authority, and spending the full envelope does not establish that continuing is worthwhile.

The required efficiency to finish within the original baseline can also be examined. TCPI to BAC = (BAC − EV) / (BAC − AC) = 64,800 / 60,000 = **1.08**. That requires $1.08 of budgeted work per future dollar, compared with 0.90 achieved so far. It is a challenge to justify, not proof that recovery is impossible.

Management response

First establish the cause: underestimated work, a temporary start-up cost, rework, changed scope, a supplier issue or unreliable measurement. Then compare options. Resolve the bottleneck, remove low-value authorised work through change control, revise the completion commitment, seek additional funding or stop. Show each option's effects on benefits, exposure and operational consequences.

For the pilot, a forecast AUD 12,000 above the baseline, within the existing funding envelope, may be reasonable if evidence supports the remaining value and the response addresses the cause. It may be poor use of funds if the intended use case remains unsuitable. A favourable cost ratio cannot answer that investment question.

Separate what-if: a supplier release changes the remaining work

At the week-six review, assume a required supplier-release change means that integrated tests F must be followed by one week of correction and re-evaluation before supervised use G. Call the additional activity X. It needs 25 evaluation hours and 15 platform-support hours, each at an illustrative AUD 80 per hour. These 40 hours are not included in the existing

estimate to complete. Assume all other costs remain unchanged; an actual forecast must also examine any additional management, support, licence or waiting costs.

Table 4.11 | Separate supplier-change scenario

Measure	Existing baseline or forecast	Separate change scenario
Completion	Baseline end of week 12	Forecast end of week 13
Controlling sequence	A → C → D → F → G → H	A → C → D → F → X → G → H
Actual cost at week six	AUD 48,000	Unchanged at this review
Estimate to complete	AUD 72,000	AUD 75,200
Estimate at completion	AUD 120,000	AUD 123,200
Cost baseline / BAC	AUD 108,000	Unchanged unless reauthorised
Funding envelope	AUD 120,000	Unchanged unless reauthorised

The added cost is 25 × AUD 80 + 15 × AUD 80 = AUD 3,200. New ETC is AUD 75,200 and EAC is AUD 48,000 + AUD 75,200 = AUD 123,200. That is AUD 15,200 above BAC and AUD 3,200 beyond the funding envelope. Reserve already inside that envelope cannot fund expenditure beyond it.

Under the original finish-to-start assumptions, F finishes at the end of week 9. X occupies week 10, G moves to weeks 11–12 and H to week 13. Confirm the evaluators, support staff and participants are available on those dates before treating the new forecast as feasible. Completing content preparation earlier does not remove the new dependency.

Take a specific choice to the authorised decision-maker: provide the time and funding for the added work; retain an acceptable supported configuration if that option exists; narrow the use; or stop. Explain the evidence and operational consequences of each option and the latest useful decision date. Preserve the original baseline and update the forecast while the decision is pending. Do not omit necessary evaluation simply to retain the baseline date.

Common failure

Applying EAC = BAC/CPI mechanically to very early or volatile data, or assuming money can recover time. Review the remaining scope and schedule. As work finishes, conventional SPI tends towards 1 even if delivery was late, so it cannot replace date-based schedule analysis.

05 Lead delivery and change

Coordinate people and feedback, manage changes coherently and report the decisions that will keep delivery and value on course.

5.1 Run short feedback loops and make work flow visible

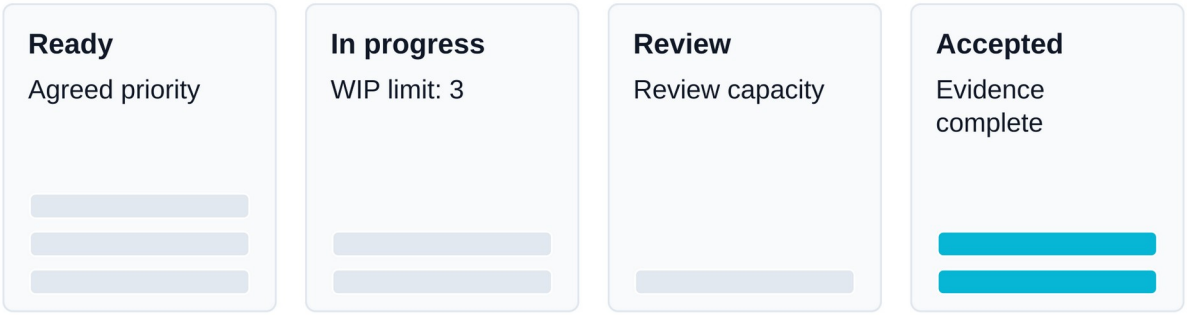


Figure 20 | Make workflow and completion evidence visible

Purpose

Frequent inspection helps teams discover whether their work is useful and whether the way they work needs to change. A visible board can support this, but the board alone does not determine priorities, limit overload or establish completion.

Scrum provides a defined framework with accountabilities, events and artefacts, including a Product Owner, Scrum Master and Developers. Its Product Goal, Sprint Goal and Definition of Done connect direction and completion. [17] Kanban emphasises an explicit workflow and controls on work in progress; its flow measures include throughput, cycle time and the age of unfinished work. [18]

A lightweight pilot rhythm

Agree the next usable result and why it matters. Select only work that is understood well enough to start. Review blockers frequently. Demonstrate completed work to the users and evidence owners. Adjust priorities from what was learned, then improve one aspect of the team's working process. This is a suggested rhythm, not a claim to reproduce the full Scrum framework.

Table 5.1 | Work states and completion policies

Board state	Entry or exit question	Example policy
Ready	Is the need clear and the work small enough?	Required inputs and acceptance evidence identified
In progress	Does the team have capacity to start?	No more than three active items in this fictional team
Review	Is evidence available for an independent check?	Reviewer named; blocked reviews raised promptly
Accepted	Has the agreed completion standard been met for this work item?	Evidence linked, acceptance recorded and dependent records updated

Worked example

The team starts six items but completes only one because the same reviewer is needed for all six. Starting four more items makes the board busier and the queue longer. The response is to finish or unblock existing work, arrange review capacity and adjust the entry policy. A limit of three is an experiment for this team, not a universal rule.

Measure finished work and elapsed time consistently. A service-level expectation can state a probability and elapsed time using relevant historical cycle-time data. With little history, label the estimate as provisional. Comparing teams by story points is unreliable because their scales and work differ; assess their outcomes and local forecasting performance.

Make learning and accepted delivery visible

Label work by its intended result. An experiment card may finish with evidence that an approach is unsuitable. A service-deliverable card requires its agreed acceptance evidence. Both can reach the board's Accepted state, but acceptance of the experiment does not authorise the service for use. Link each card to the relevant requirement or decision, tested configuration and result. A useful demonstration can inform the next step while further integration or evaluation remains open.

Include data and knowledge preparation, evaluation design, integration, user guidance and service preparation in the flow of work. Work hidden outside the board can become the real constraint. Before integrating an AI component, confirm its input/output arrangements, permissions, supported version and expected failure response with the teams responsible. Test the joined workflow as well as the component: a satisfactory model result is insufficient if the interface presents the wrong source, a downstream system rejects the output or the recovery path cannot be used.

For a purchased service, the delivery team may configure and integrate rather than develop or train a model. Track the work it can perform and the supplier commitments it depends on. If a supplier issue blocks completion, assign the internal follow-up owner and a decision date. Do not keep starting unrelated changes simply because the blocked item is outside the team's direct control.

Common failure

Demonstrating unfinished work as if it were accepted, or letting every urgent request bypass the workflow. Define exceptional handling and record its effect on existing commitments.

5.2 Engage stakeholders around decisions and consequences

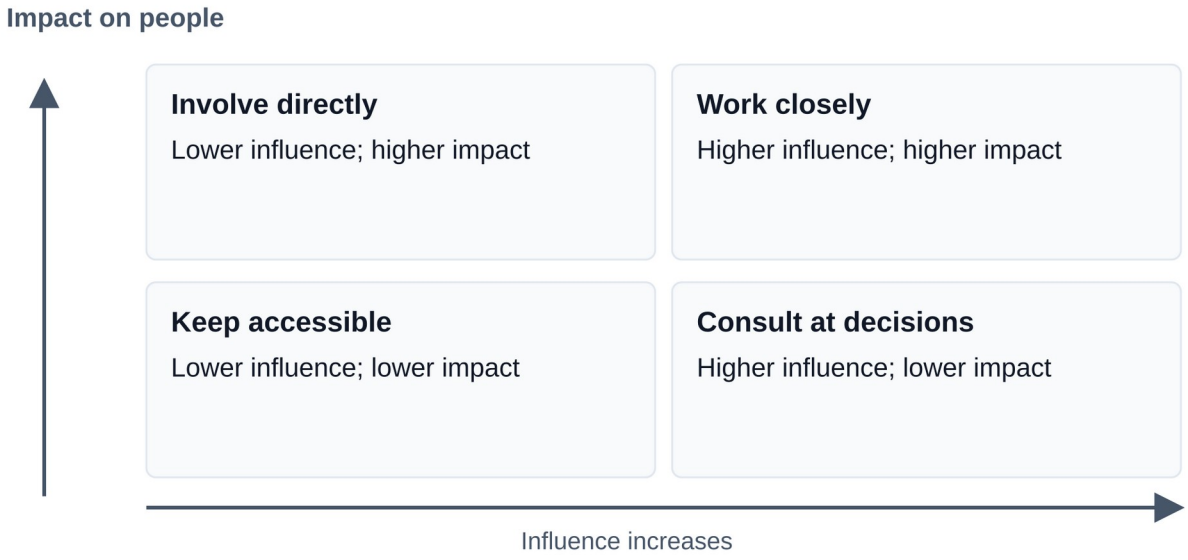


Figure 21 | Influence and impact shape engagement

Purpose

Stakeholder engagement helps the project understand who is affected, whose knowledge is needed and who can enable or constrain the result. Influence is relevant, but so are exposure, accessibility, dependence and the consequences of being overlooked.

Method

Identify people and groups across the project and future service. Ask what changes for each group, what they need, what concerns they have and what they can decide. Include users who are unlikely to attend steering meetings, operational staff who inherit the result and groups that may experience unintended effects. Record an engagement objective and a practical next action.

Table 5.2 | Stakeholder engagement

Stakeholder	What matters to them	Engagement action	Evidence of progress
Sponsor	Value, affordability and continuation choices	Short decision brief at agreed milestones	Timely decisions with conditions recorded
Pilot users	Useful answers and manageable effort	Observe real tasks and test user guidance	Usability findings and resolved barriers
Content owners	Accurate representation and maintenance	Agree source ownership and correction process	Current approved content and clear responsibilities
Service team	Support demand and recoverability	Involve them in design and runbook exercises	Explicit operational acceptance
Assurance specialists	Evidence proportionate to exposure	Agree review scope and checkpoints early	Findings closed or residual issues assigned

Worked example

Pilot users support the concept but rarely participate. Do not immediately conclude that they resist change. The problem may be workload, inaccessible sessions, unclear instructions,

concerns about monitoring or a use case that does not help them. Test these explanations with a few short conversations and observations. Then alter the engagement method or the project proposition.

Differentiate consultation from consent and decision authority. Listening to a team does not automatically authorise a change affecting that team. Likewise, a consulted specialist may provide advice without owning the business decision. Clarify these relationships in the project governance arrangement.

Revisit the stakeholder map as the scope and delivery stage change. Procurement may be central before contract award; support staff become more involved before transition. A static map prepared at initiation quickly loses practical value.

Engage the people affected by the AI boundary

An AI feature can change who prepares information, who checks it and who is held responsible for acting on it. Map those changes even when only a small component of a wider project uses AI. Include the people whose work supplies the data or knowledge, those who review outputs and anyone whose service or decision is affected. For a predictive feature, seek the experience of groups whose cases may be less well represented in the evaluation. For an action-taking capability, involve the owners of the systems and resources on which it may act.

Explain what participation involves, what use or feedback data will be recorded, why it is needed and who can access it through the organisation's agreed arrangements. Resolve these questions with the relevant business and specialist owners before asking users to participate. Give affected people a usable route to challenge an output, report a problem or continue through an alternative process. The project manager's action is to turn those needs into requirements, engagement work and named decisions.

Common failure

Colour-coding people as supporters or blockers and treating the label as a fixed personal trait. Describe observable concerns and their causes instead. A strong objection may reveal a design flaw, an unowned burden or a legitimate constraint.

Practical test

For each significant decision, can you name the people whose knowledge, authority and lived experience must be represented?

5.3 Communicate around decisions and useful action

Communication management concerns the information people need to understand the work and act. Stakeholder engagement concerns the relationship, participation and response of the people involved. A report can be delivered perfectly while a stakeholder remains unconvinced or unable to participate. Use both perspectives.

A small communications plan

Table 5.3 | Communications by audience

Audience	Information or conversation needed	Channel and timing	Owner and feedback
Delivery team	Next work, blockers, dependencies and changes	Shared board and short checks suited to the work	Project manager; confirm actions in the record.
Sponsor	Forecast, value, material exceptions and choices	Brief decision note before the needed-by date	Project manager; record the decision and conditions.
Pilot users	Purpose, permitted use, support and how feedback is used	Accessible introduction, practice and short feedback sessions	Adoption lead; observe understanding and behaviour.
Content and service owners	Upcoming changes, obligations and readiness evidence	Working sessions at design and transition points	Relevant owner; confirm acceptance of responsibilities.
Assurance specialists	Claims, review scope, evidence and unresolved findings	Early agreement and versioned evidence pack	Assurance coordinator; track responses and retesting.

Handle disagreement productively

Describe the decision and the evidence on which people disagree. Separate uncertainty about facts from disagreement about priorities. A technical lead may favour a fast configuration change while a content owner questions whether the underlying guidance is current. Both views can be reasonable, and neither is resolved by telling people to collaborate better.

Ask what would change each person's recommendation. Identify a short test or source of evidence where possible. If the disagreement concerns an accepted trade-off, take the options to the person authorised to make it. Record the rationale and communicate the effect on work, rather than simply naming the winning opinion.

Communicate changes in capability and permission

When a release changes the service, tell each audience what it means for their work. Users need the changed behaviour, permitted tasks, limitations and support route. Reviewers need revised judgement guidance and escalation conditions. Service staff need the effective version, monitoring and recovery arrangements. The sponsor needs any effect on cost, acceptance, timing or value. A supplier's release note is an input to this communication, not the complete message for every audience.

State what evidence supports a performance claim and where its limits matter. Results from the controlled evaluation and feedback from supervised use answer different questions. Present neither as proof of organisation-wide benefit. Keep the current user guidance linked to the version actually in use, and confirm that people relying on a changed rule have received and understood it.

Make information usable

Provide the decision first, context second and evidence behind links. Use plain language, explain acronyms and make materials accessible. Keep sensitive material in an appropriately

controlled location. Confirm the receiver has understood the required action; sending a message is not proof that the action was understood or agreed.

> **Practical prompt:** After every significant meeting, can participants identify the decisions, actions, owners and dates without reconstructing the conversation?

5.4 Adoption is work inside the plan

Purpose. Make the organisational change needed for benefits visible, resourced and owned.

A new capability produces value only when people use it appropriately and the surrounding work changes. Training is one part of adoption. Managers may also need to change procedures, allocate time, adjust measures, clarify responsibilities and retire old practices. Include that work in the project scope and capacity plan.

For the knowledge assistant, define the intended behaviour precisely: staff use approved sources, inspect the supporting material where appropriate, recognise coverage limits and raise uncertain questions through a known route. An instruction to “use AI more” is too vague to evaluate and may encourage unsuitable use.

Start with affected groups and their tasks. Ask what each group will do differently, what they may lose, what they need to learn and who can remove obstacles. People may reasonably resist a change that increases verification work, removes familiar support or creates unclear accountability. Treat that information as design evidence.

Table 5.4 | Adoption and readiness

Adoption need	Pilot action	Evidence of readiness
Understand purpose	Explain the bounded use case and expected benefit	Participants can describe suitable and unsuitable tasks
Have time and support	Managers allocate evaluation time and a support contact	Capacity confirmed and help route exercised
Use the service well	Practice realistic tasks, including failure cases	Observed behaviour meets agreed expectations
Sustain content	Assign source owners and update routines	Owners accept the work and demonstrate an update
Change the workflow	Agree when to use search, assistant or specialist support	Updated workflow is understood and usable

Measure adoption and useful performance separately. Login counts show access, not whether people solved a task or avoided error. A well-used tool can create extra rework. A lightly used tool may be valuable for infrequent complex tasks. Link each measure to the intended benefit and examine unintended effects.

Worked situation. Participants report faster answers but spend longer checking them. The response is to measure total task time and examine the causes, rather than celebrate response speed. The product design, source quality or use case may need adjustment.

Rehearse the changed work

Practice the full task, including checking supporting evidence, rejecting an unsuitable output, using a fallback and obtaining help. Ask participants to explain why they would rely on an answer or decline to act. For an AI capability permitted to take actions, rehearse the point at

which a person must approve, amend or stop the proposed action and what happens when that person is unavailable. Provide a conventional route for work that falls outside the permitted use.

Observe the time and competence required. If reviewers routinely approve outputs without the information or time needed to judge them, revisit the design, workload and use boundary rather than recording training as complete. The manager who owns the work needs to accept the continuing capacity requirement. In a small, reversible feature, short task-based practice may be sufficient; the management obligation is to establish usable behaviour, not to impose an elaborate training programme.

For an established enterprise platform, do not assume familiar access means the new AI feature has already been adopted. Update the relevant procedure, support material and responsibility at the point of work. Retire outdated guidance deliberately and check the effect during supervised use. This also applies when an existing AI service expands to new tasks or users.

Common failure: assigning all adoption to the project manager, who lacks authority over operational practices. The benefit owner and managers must own the changes within their teams.

Use tomorrow: add one acceptance condition for observed user behaviour and one for operational ownership to the readiness review.

5.5 Keep the plan integrated when change arrives

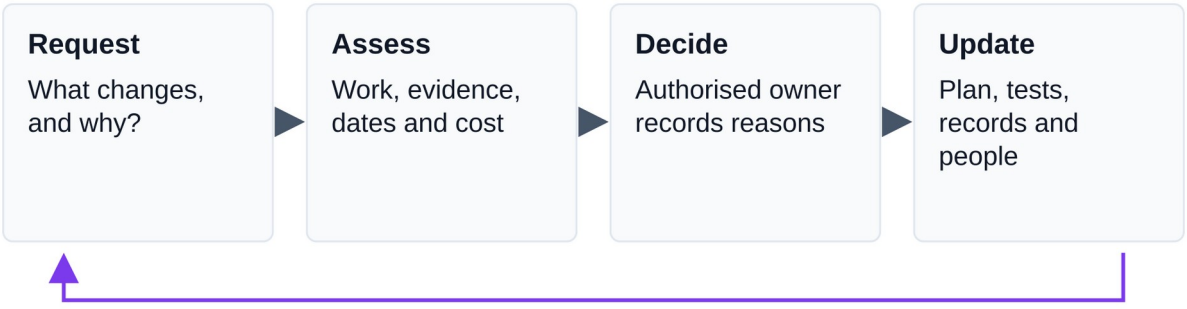


Figure 22 | One change decision updates the connected plan

Purpose

Integration means making connected decisions across scope, time, money, quality, people, risk and benefits. A local improvement can make the overall project worse when its wider effects are ignored. A baseline is an approved reference for comparison; a forecast is the current expectation of what will happen.

Maintain both. If a milestone was approved for week 12 and is now forecast for week 14, show both dates. Updating a forecast does not require pretending the original commitment was different. Rebaselining should follow an authorised decision with its reason and history preserved.

Method for a material change

Describe the request and intended benefit. Compare options, including deferral or a smaller alternative. Assess effects on deliverables, dependencies, cost, schedule, assurance,

operational support and benefits. Identify who has authority to decide. Record the outcome, conditions and rationale, then update the affected records and tell the people who need to act.

Worked example

In week 5, a stakeholder asks to add a second document collection. The request appears to need only configuration. Assessment reveals additional permissions analysis, content-owner review, evaluation coverage and user guidance. The team offers three choices: keep the original pilot; replace some lower-priority material within the current envelope; or extend scope with a revised cost and completion forecast.

Table 5.5 | Example change record

Change record field	Example entry
Request and reason	Add a second collection to test a broader staff use case
Benefits and trade-offs	Wider relevance; less time to validate each collection
Impact evidence	Added tasks, resource demand, evaluation changes and revised forecast
Recommendation	Defer until the first collection meets acceptance criteria
Decision	Named approver, date, conditions, effective version and follow-up owner

Not every backlog adjustment needs a steering committee. Establish a decision boundary: the delivery team may adjust implementation within agreed scope and tolerances; the authorised sponsor or governing body handles changes that exceed those bounds or alter the investment case. The appropriate threshold depends on organisational delegations and project consequences.

Assess the change to the tested service

Keep a record of the configuration covered by acceptance evidence. For the main pilot, C01 v1 identifies the service configuration and E01 v1 identifies the evaluation design. The record should identify the model or service release where that is knowable, relevant prompt instructions, source collection, retrieval settings, permissions, enabled tools and intended use. Link these items to existing configuration and project records instead of copying sensitive details into every form.

Table 5.6 | Effects of an AI configuration or use change

Change being considered	Management consequence to assess	Evidence or record to update
Model or supplier release	Earlier behaviour evidence may not cover the proposed release; support and cost assumptions may also change	Supplier notice, affected requirements, evaluation scope and release decision
Prompt instructions or retrieval settings	The service may answer, refuse or cite material differently	Configuration record, affected task results and known limitations
Data or knowledge sources	Coverage, currency, access and maintenance work may change	Source ownership, permission tests, representative tasks and operating effort

Change being considered	Management consequence to assess	Evidence or record to update
Permissions or connected tools	The capability may reach different information or perform different actions	Access and action boundaries, approval/recovery tests and authorised use
Users, purpose or operating conditions	Previous evidence may not represent the proposed work or affected people	Scope, requirements, specialist advice, adoption plan and benefit assumptions

The table identifies change impacts, not a new scoring framework. The delivery team proposes the affected tests and reviews; relevant evidence owners confirm that the coverage is sufficient for their decisions. A small, well-understood correction may need targeted checks. A material change in purpose, access or autonomy may require broader evaluation and fresh authorisation. Record the reason for the chosen depth, including relevant checks retained from the earlier configuration.

A supplier can change the operating service without a local feature request. Assign someone to monitor release notices and confirm whether changes can be tested, deferred or reversed under the actual service arrangement. If the organisation cannot preserve an accepted configuration, assess the supported alternative and its implications before continuing use. The supplier-release scenario in Chapter 04 shows how this can add work, move the forecast and exceed funding without changing the original baseline.

After the decision, update the implementation instructions, evaluation record, user guidance, financial forecast and operational arrangements that actually changed. Retain the relationship between old and new evidence. Do not record a different configuration under the original version identifier merely because the interface looks the same.

Common failure

Approving changes in meetings but leaving the schedule, cost forecast, acceptance criteria and supplier instructions unchanged. Another failure is accepting “free” enhancements that create ongoing service costs or additional exposure.

Practical test

After approval, can the team point to one coherent current plan and a clear record of what changed? A folder of competing versions is a control problem.

5.6 Report for a decision, not for reassurance



Figure 23 | Report evidence and the next useful decision

Purpose

Good reporting lets the recipient understand the current position, its implications and any action required. A green dashboard is not the objective. The objective is a shared, evidence-based view that supports timely intervention and avoids surprises.

Method

Begin with the reporting date and the decision required. Show the baseline, current forecast and change since the previous report. Summarise accepted progress, cost, delivery confidence, material risks and benefits evidence. Keep detail behind links. Use a consistent cut-off so cost, progress and schedule refer to the same period.

A one-page pilot dashboard.

Table 5.7 | Pilot dashboard contents

Element	What the reader should see
Overall position	A short explanation of delivery confidence and its main cause
Outcome and scope	What remains achievable; any proposed scope decision
Milestones	Baseline date, forecast date, movement and reason
Financials	Actual cost, commitments, forecast to finish, reserve position
Progress	Accepted deliverables and the next evidence-bearing milestone
Exposure	The few risks or issues requiring attention, with owners
Decision request	Options, recommendation, authority needed and latest useful date

Worked exception statement

“Access configuration is forecast to finish one week late. With the present dependencies, the final decision pack moves from week 12 to week 13. The team recommends retaining the evaluation scope and accepting the revised date. A decision is needed by Friday to confirm pilot-user availability.” This tells the sponsor much more than “schedule amber”.

Define traffic-light thresholds in advance, including who can override a mechanical score with an explained judgement. A project can be within budget and still require attention because its

benefit case has weakened. Do not average severe exposure away inside a composite green score. Where confidence is uncertain, say so and identify the evidence needed to improve it.

Choose the reporting rhythm to fit decisions. A short team check may be daily; a sponsor discussion may be weekly or fortnightly; a governing body may review formal milestones. These are examples, not universal requirements. Urgent exceptions should not wait for the next scheduled report.

Report AI delivery evidence at the right level

Separate four statements in the progress narrative: what was explored, what deliverable was accepted, what use was authorised and what business outcome has been observed. Link a material acceptance claim to the configuration and requirements it covers. Where evaluation is incomplete, state the remaining coverage, owner and planned decision date. Explain significant failures and limitations even when an overall result appears favourable.

Keep this within the existing status report. A small feature may need one clear paragraph and links to its evidence. A more complex service may need separate reporting of evaluation, integration and operational readiness because they have different owners and dependencies. Include the cost of review and rework in the forecast, and show unresolved supplier-release or capacity decisions alongside other delivery exceptions.

For the separate week-six supplier-change scenario, report EAC of AUD 123,200 against BAC of AUD 108,000 and funding of AUD 120,000, with completion forecast for week 13. Label it a pending change scenario until an authorised response is recorded. Do not overwrite the original pilot's AUD 120,000 forecast or report the additional funding as already available.

Common failure

Reporting activity volume instead of results: meetings held, emails sent and tasks started. Use accepted outputs and evidence of change wherever possible. Add leading indicators such as ageing blockers, unconfirmed dependencies and delayed decisions before they become missed milestones.

Practical test

Can the recipient identify the required decision within 30 seconds and understand its rationale within two minutes?

5.7 A light operating rhythm that keeps decisions moving

Purpose. Turn the framework into a repeatable way of working that remains proportionate.

Begin with a small set of maintained records: charter and business case; integrated delivery plan; RAID log; decision and change record; short dashboard; and benefits/transition record. They may live in one workspace. Add a separate artefact when it has an owner, a decision purpose and a clear update trigger.

Use a rhythm suited to the project rather than adopting meetings by habit:

Table 5.8 | Integrated monitoring cadence

Cadence or trigger	Focus	Expected result
Short team check-in	Immediate work, blockers and dependencies	Coordinated next actions
Weekly delivery review	Evidence, forecast, risks and upcoming decisions	Updated plan and escalation
Fortnightly demonstration or review	Working outputs and user feedback	Acceptance evidence and learning
Sponsor review at agreed frequency	Value, exceptions, capacity and choices	Decisions within authority
Before release or material change	Assurance and readiness	Explicit permission or further actions
After transition	Benefits, operating cost and unintended effects	Continue, improve, expand or retire

These frequencies are illustrative. Combine meetings when the participants and purpose align. Shorten routine reporting by keeping live evidence accessible; do not use the dashboard meeting to rediscover basic facts.

Use the weekly review to look ahead to source changes, supplier releases, evaluation capacity and the decisions that enable the next work. Bring a material incident or change in permitted use directly to the relevant authority instead of waiting for the routine meeting. Keep the normal project controls in use: the AI component does not need a parallel reporting calendar when the existing one can carry its evidence.

The “After transition” row becomes part of the receiving service and benefit owners' operating rhythm. Agree those owners, triggers and review dates before project closure. The temporary project need not remain open to run every future evaluation; the service does need an accountable way to notice change, reassess the affected evidence and decide whether to continue, restrict or suspend use.

When a project is in difficulty, perform a bounded reset. Establish the current scope, accepted work, cost, commitments and realistic forecast. Identify immediate exposure and stabilise it. Recheck the business need. Develop options with consequences, obtain an authorised decision, then publish the revised position and responsibilities. A reset is an explicit decision, not an opportunity to erase the old baseline or obscure poor performance.

A practical first month. Use the 30-day learning plan in the final reference section to apply this rhythm to one real project. Do not wait for a scheduled review to address urgent exposure.

Use three review questions regularly: What decision became easier because we maintained this record? What information arrived too late? What work are we doing only because the template contains a field? Improve the management system as deliberately as the product.

Common failure: mistaking consistency for identical administration. Standardise the questions and decision expectations while tailoring evidence and effort to the project.

Use tomorrow: select one real initiative and implement the smallest complete control loop: outcome, owner, plan, evidence, decision and follow-through.

06 Manage risk and assurance

Identify uncertainty, test the strength of controls and bring appropriate evidence to the person authorised to decide.

6.1 Write risks so that somebody can act on them

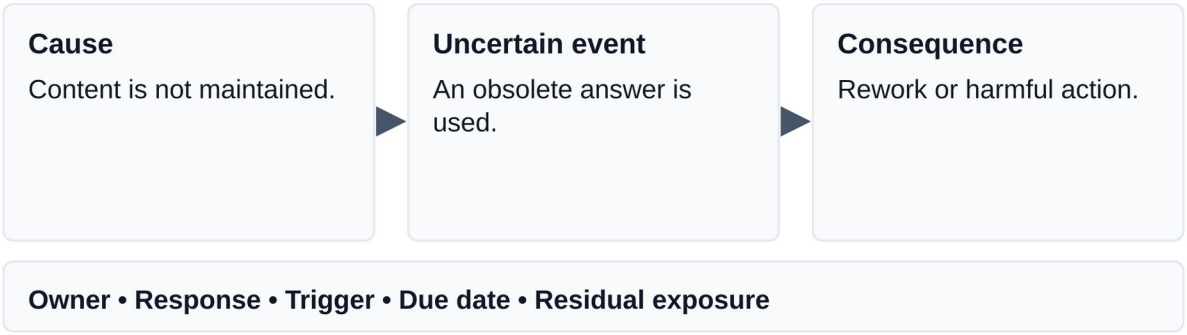


Figure 24 | A risk record connects cause event and consequence

Purpose. Make uncertainty visible early enough to change the plan.

A risk is an uncertain event or condition that could affect objectives. An issue is already present. An assumption is a belief used in planning that has not yet been sufficiently established. A dependency is something the work needs from another activity, party or decision. Treating all four as “risks” obscures the next action.

A useful risk statement links cause, uncertain event and consequence. For the pilot: “Because several documents within the approved collection lack a current owner, outdated documents may remain available to the assistant, causing staff to rely on superseded instructions.” This is more actionable than “data quality risk”. It identifies the condition to investigate, the event to prevent and the consequence to evaluate.

Identify risks from several directions:

- Outcomes: what could prevent the intended improvement?
- Delivery: which estimates, skills, dependencies or suppliers could disrupt the plan?
- People: who might be harmed, excluded, overloaded or misunderstood?
- Operation: what could fail after the project team leaves?
- Opportunity: what favourable uncertainty could increase value or reduce effort?

Ask people who do the work as well as those who approve it. A workshop with only senior participants may miss document maintenance, service-desk pressure or the practical cost of verifying an answer. Include external dependencies and cumulative effects across projects where relevant.

Sample record. R-03; cause: unowned content; event: outdated material used in answers; consequence: incorrect internal guidance and rework; owner: content lead; treatment: confirm ownership, exclude unverified material and test update propagation; trigger: documents beyond review date or contradictory answers; review: weekly during preparation and after source changes.

Keep risk ownership separate from action ownership. The person accountable for monitoring a risk may need several people to complete treatments. Assign one owner for each action, its due date and the evidence of completion.

Positive uncertainty also deserves attention. If an existing search capability performs well, the project could avoid unnecessary integration. Assign somebody to investigate that possibility and compare the options fairly. An opportunity is not simply a benefit already assumed in the business case.

Make the risk statement specific to the proposed AI use. A purchased assistant with limited read access, a prediction supporting a human decision and a capability permitted to change records have different possible consequences. Identify the work affected and the decision the owner may need to make. Do not copy a generic list of AI concerns into the register without a project-specific cause and action.

Connect the response to the delivery plan. If the treatment is additional evaluation, source remediation or a recovery exercise, identify its effort, predecessor and acceptance evidence. If the treatment reduces scope or changes the human-review arrangement, update the requirements, forecast and operating responsibilities. Once a failure is observed, manage the present problem as an issue while retaining any remaining uncertainty as a linked risk.

Common failure: populating a long generic register and then never changing delivery decisions. Keep the most decision-relevant risks prominent and archive resolved items without losing their history.

Use tomorrow: rewrite the top three risks with a cause, event and consequence. If you cannot name a practical next action, refine the statement.

6.2 Assess exposure and choose a response

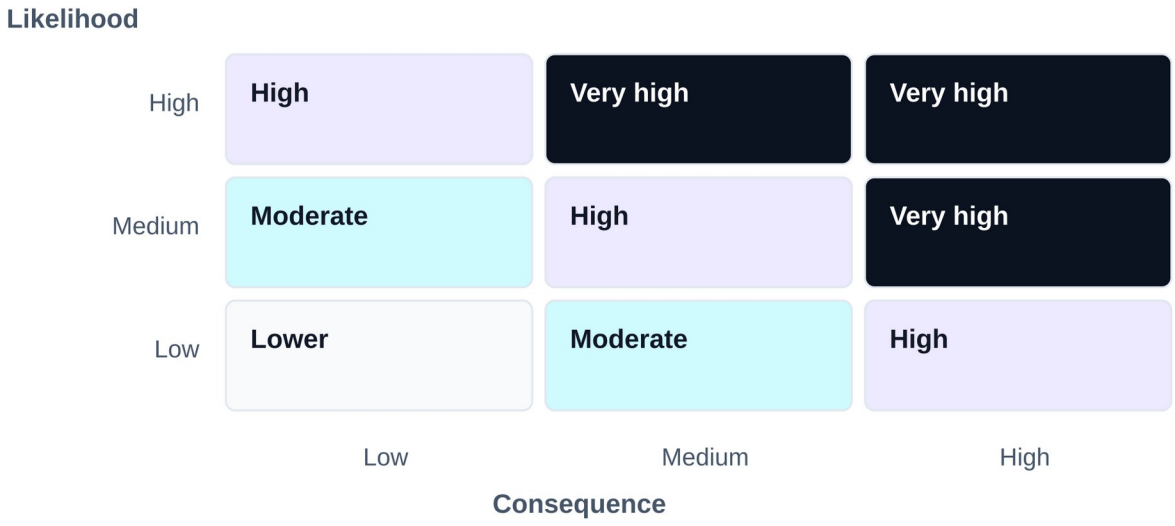


Figure 25 | A qualitative matrix supports judgement

Illustrative ordinal categories for discussion. The diagram does not calculate probability, loss or organisational approval authority.

Purpose. Decide which uncertainties need attention and what to do about them.

Assess likelihood and consequence using agreed definitions. A three- or five-level scale can support prioritisation, but its meaning matters more than its colour. Define consequence across relevant dimensions: cost, time, service, people, information, reputation and obligations. Decide how multiple dimensions will be handled before scoring the register.

A numerical rating from ordinal labels is an aid to sorting; it is not a precise measure of loss. Two risks with the same score can require different responses. A rare, severe harm may

need a specific control even when its calculated score resembles a common, minor inconvenience. Keep the reasoning and any overriding escalation rule visible.

For an adverse risk, consider avoiding the exposure, reducing its likelihood or impact, sharing an appropriate part with another party, or accepting the remaining exposure through the authorised process. Transfer through a contract or insurance rarely removes all accountability, operational disruption or reputational consequence. For opportunities, consider increasing the chance, collaborating to capture the opportunity, or deliberately pursuing it.

Distinguish three views: exposure before existing controls, the current exposure with controls operating, and the expected exposure after planned treatments. Definitions vary between organisations, so label the views explicitly. Do not present a future, hoped-for reduction as though the control already exists.

Worked record. Access leakage is assessed as unacceptable for the proposed pilot. Planned actions include permission verification, negative access tests and source restrictions. Until those actions are evidenced, the current rating remains unchanged. The use decision can be reconsidered after the results are available.

When estimates are justified, expected monetary value can help compare options. A 25% chance of an AUD 8,000 rework event produces an expected value of AUD 2,000. That does not mean the actual event will cost AUD 2,000, nor that adding expected values automatically creates an adequate reserve. Dependence between risks, severe consequences and funding constraints require further judgement.

Reserve should connect to uncertainty rather than operate as an unallocated buffer that hides poor planning. In the case project, AUD 8,000 of contingency is included in the baseline; the separate AUD 12,000 management reserve is subject to its own control. Update the forecast when exposure changes and obtain any required approvals.

Use the organisation's established risk criteria and escalation arrangements. AI-specific evidence can change an assessment or response without requiring a new project scoring system. Keep the proposed use, configuration and consequences visible so the owner can understand what the assessment covers.

A planned human-review or fallback control should remain a planned treatment until its practical operation is evidenced. Confirm the relevant capacity, access, competence and recovery time. If the control cannot work in the intended conditions, revise the use boundary or delivery plan and refer the remaining exposure to the appropriate authority. A favourable average evaluation result does not override an agreed exclusion or an unacceptable severe failure.

The 25% / AUD 8,000 expected-value calculation above is an independent teaching exercise. It does not establish the probability of a pilot event or explain the pilot's separately stipulated AUD 8,000 contingency.

Common failure: accepting a risk because treatment is inconvenient, without checking authority or alternatives.

Use tomorrow: review the highest-consequence risk and ask what evidence shows its existing controls actually work.

6.3 Quantify uncertainty when it improves the decision

Qualitative risk assessment helps prioritise attention. Quantitative analysis can help compare choices, forecast ranges or size a response when credible data and an appropriate model exist. More arithmetic does not automatically create more confidence. Preserve the assumptions and distinguish a modelled result from observed performance.

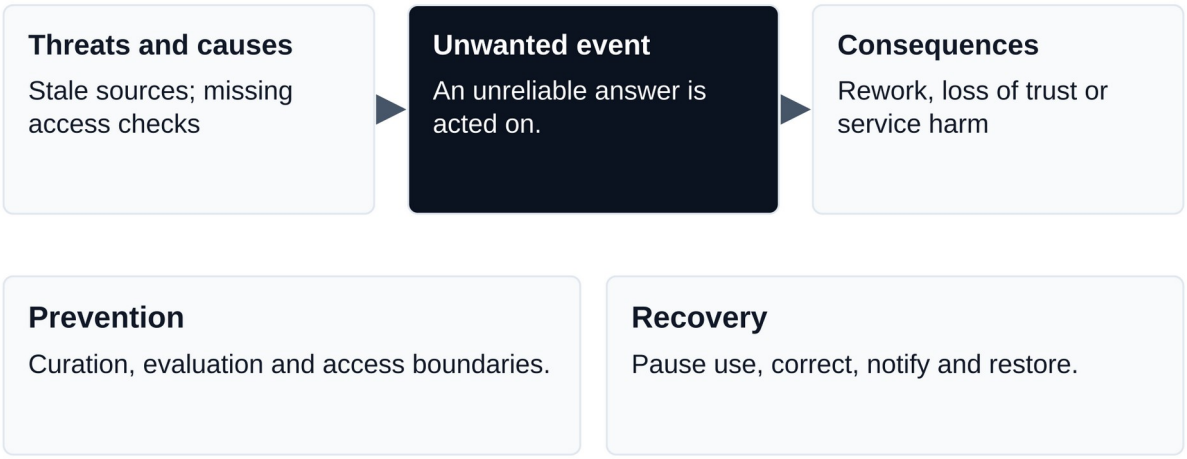


Figure 26 | Prevention and recovery address different parts of a risk

Expected monetary value: a worked comparison

Consider an independent, illustrative source-conversion task. Without an additional check, assume a 30% chance of AUD 20,000 in rework. Expected rework cost is 0.30 x 20,000 = AUD 6,000. A proposed AUD 2,000 check is estimated to reduce the chance to 10%, leaving expected rework of AUD 2,000. The combined expected cost becomes AUD 4,000, a reduction of AUD 2,000 against the no-check option.

This comparison depends on the probabilities and cost estimate being credible. It does not say the actual project will lose AUD 6,000 or that AUD 6,000 is sufficient reserve. It also excludes effects not represented in dollars. A severe harm or binding obligation cannot be settled by a favourable average cost.

Sensitivity before complexity

Ask which assumption most changes the choice. In the example, if the check barely changes the probability, the apparent value falls. If the consequence includes substantial disruption, avoiding it may be more valuable than the direct rework estimate suggests. A small sensitivity table can reveal this more clearly than a single expected value.

Simulation and correlation

A Monte Carlo model repeatedly samples uncertain inputs and calculates outcomes, producing a distribution under the model's assumptions. For a schedule, it can explore completion dates based on activity ranges and dependency logic. A modelled P80 date is one that 80% of those simulations finish by; it is conditional on the model, not a guarantee from real-world evidence.

Shared causes matter. Supplier delay, specialist unavailability and rework can affect several tasks together. Treating them as independent may understate the combined exposure. Review the model with people who understand the work, validate it against experience where possible and explain important omissions.

For AI-related uncertainty, separate observed evidence from assumptions about future use. A controlled task set may inform a comparison while leaving uncertainty about different users, sources, operating conditions or supplier changes. State those limits when presenting a forecast range. Do not turn a small or convenient sample into a precise claim about every future task.

Where the evidence is thin, compare a few plausible cases for workload, review effort, rework and service cost. Identify which assumption would reverse the decision and fund a bounded test of that assumption if the result could justify the next commitment. The standalone source-conversion calculation above remains an illustration of the technique; it is not an extra cost or reserve allocation in the connected pilot.

> **Decision rule:** Use the least complex analysis that can materially improve the choice. Where evidence is too thin, a bounded experiment may be more useful than an elaborate model.

6.4 A RAID log should produce movement

Purpose. Coordinate uncertainty and obstacles without creating a second task-management system.

RAID commonly groups risks, assumptions, issues and dependencies. Some organisations use a different meaning for a letter; state the meaning in the register. The value is a shared record of conditions that could affect delivery and the next action needed.

Table 6.1 | Risks assumptions issues and dependencies

Type	Pilot example	Next useful action
Risk	Source maintenance may be insufficient	Confirm owners and test update routines
Assumption	Staff can spend one hour each week in evaluation	Validate with their managers before planning capacity
Issue	A source owner has declined access	Resolve authority or remove the source from scope
Dependency	Identity configuration is needed before access testing	Obtain a named provider, date and acceptance evidence

For each record, include an identifier, clear description, impact, owner, next action, due date, status and links to relevant evidence. Add fields only when somebody uses them to make a decision. Separate the full history from the short list needing attention now.

Dependency agreement. “ICT will configure access” is too loose. Record what is required, who will supply it, the needed-by date, the delivery conditions and the effect of delay. Confirm the provider has committed capacity. A dependency owner inside the project coordinates the need; the supplying party remains responsible for its delivery commitment.

As evidence changes, convert records rather than leaving contradictory entries. An assumption that proves false may become an issue, a change request or a new risk. A missed dependency may create an issue now and a further risk to the final milestone. Link the records so the causal story remains understandable.

Review the log in a working meeting, not only when reporting upwards. Ask: What has changed? What action is overdue? What decision is needed? Which record can close with

evidence? Escalate records that the current owner cannot resolve, including organisational conflicts and missing external decisions.

For a supplier or platform dependency, record what the project can observe and influence. A version notice, evaluation-access date or permission-change request needs a named external contact and an internal owner who will follow up. Include the latest useful response date and the consequence if it is missed. “Awaiting vendor” is not an adequate action plan.

Link the dependency to the affected requirement, configuration and planned evidence where that relationship matters. If an assumed supplier release changes, link the issue to the change record and revised test activities instead of maintaining conflicting versions in separate logs. Close the original assumption only with the evidence and explanation needed to understand the decision later.

Common failure: “in progress” remains the status for weeks. A named next action and date are more informative than a vague status word.

Use tomorrow: filter the log to overdue actions and dependencies due in the next fortnight. Resolve the ownership gaps before discussing all other entries.

6.5 AI projects need evidence about use and behaviour

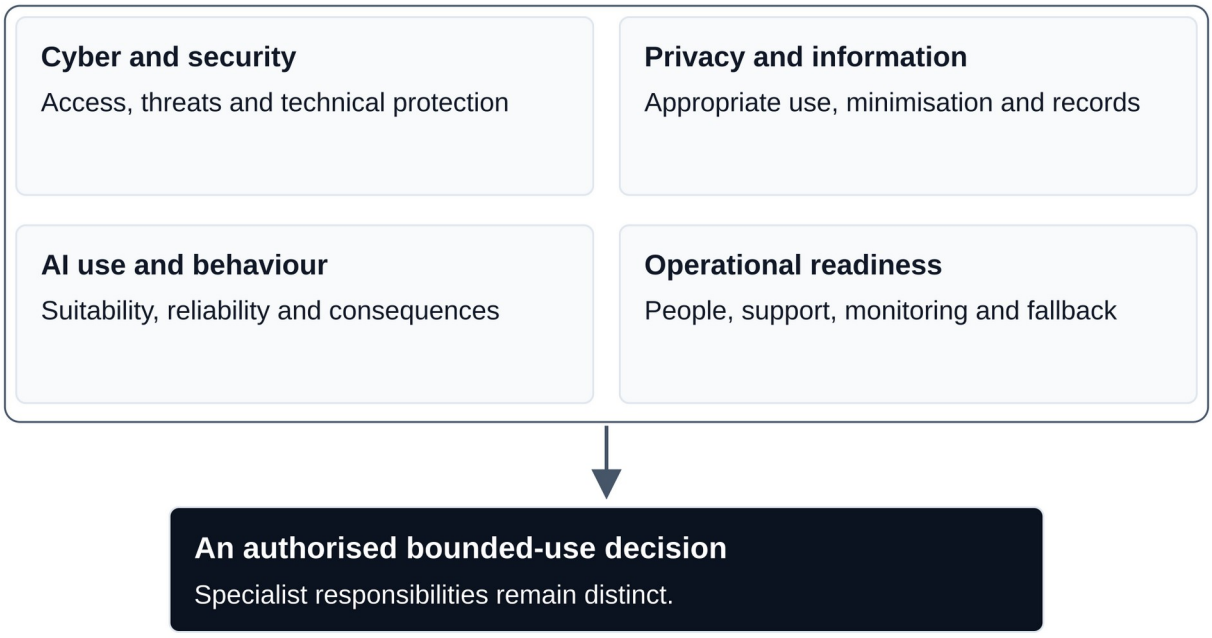


Figure 27 | Different evidence informs an authorised use decision

Purpose. Integrate project control with assessment of whether the AI capability is suitable for its intended use.

Project management controls delivery. Cybersecurity, privacy, information governance, accessibility, legal review and AI assessment examine different aspects of the proposed service. Their findings should meet in the use decision, while each discipline retains its specialist questions and accountable authority. No single completed assessment proves the entire service is suitable.

NIST's AI RMF addresses risks to people, organisations and society, and its 2024 generative-AI profile provides a companion treatment of generative-AI concerns. The framework is voluntary; its public page says AI RMF 1.0 is under revision as at this guide's

research date. [19, 20] Australia's national government assurance framework similarly treats AI assurance as part of broader governance. [21]

For the pilot, assess the actual combination of users, purpose, data, model, retrieval sources, interface and workflow. An approved software platform does not answer whether every possible use is appropriate. Nor does a successful proof of concept establish safe performance at broader scale.

Link existing project and specialist records to the decisions they support. The following is an evidence map, not an assessment questionnaire or a substitute for a specialist process:

Table 6.2 | Project assurance interfaces

Decision or boundary	Example evidence
Intended use and responsibility	Intended use, users, exclusions and human responsibilities
Information access and handling	Source inventory, permission tests, data handling and retention decisions
Behaviour and acceptance	Representative and challenging tasks, unsupported answers, refusal and recovery behaviour
Effects on people	Use-case impact assessment, accessibility and differential performance observations
Usable human responsibilities	Observed user testing, clear limits, escalation and correction routes
Continuing operation and change	Monitoring, incident response, version tracking, change triggers and rollback

Evaluate failure severity as well as averages. A high overall score can conceal a small number of unacceptable failures. Preserve the tested version and configuration, record known limitations, and explain which conditions invalidate earlier evidence. Model changes, new sources, expanded users or new purposes may require reassessment.

The disciplines overlap and should share evidence; a completed cyber assessment alone does not establish that a particular AI use is suitable.

Worked situation. The assistant retrieves current documents but confidently combines unrelated instructions. Access controls may work while the answer remains unsuitable. Add behavioural evaluation and appropriate design changes; do not close the finding merely because the cyber review is complete.

Keep the assurance process usable. Begin with a short description of the use and its consequences, then route to specialist depth when the answers justify it. Maintain one linked evidence record rather than forcing staff to repeat the same context in several forms.

Manage the specialist interface as project work

Seek specialist advice early when the proposed work creates unfamiliar information use, access, external actions, affected groups or obligations, or when existing advice may not cover the intended use. The relevant organisational policy and accountable function determine the review required. A small embedded feature may use an existing review route

with a narrow additional evidence requirement; an expanded service with different consequences may need a broader review.

Before scheduling the review, agree its scope, inputs, reviewer capacity, needed-by date and expected output. The project provides a clear description of the proposed use and configuration, relevant data or source arrangements, evaluation evidence and unresolved limitations. The specialist provides findings, conditions, recommendations or a decision within that function's authority. Record who may approve the use, who accepts deliverables and who accepts operational responsibility. These decisions may involve different people.

Translate findings into work. A privacy finding may require a different data flow before integration can proceed. An accessibility finding may require interface changes and renewed user testing. A behavioural finding may require narrower tasks, different sources or further evaluation. The project manager connects those responses to scope, owners, dependencies, costs and release conditions; the project manager does not replace the specialist's judgement or approve an exception outside their delegation.

Agree which findings prevent the proposed release and which can be managed under an explicitly accepted, time-limited condition. Record the condition, owner, evidence and expiry or review trigger. If the responsible authority cannot accept the remaining position, narrow, defer or stop the proposed use. A project deadline is not evidence that a specialist condition has been satisfied.

Retain the result with the tested configuration and pass continuing obligations to the service owner at transition. The operating team needs to know which changes or incidents require fresh specialist input and who will obtain it. Project closure transfers this work; it does not remove it.

Common failure: treating human review as a control without testing whether reviewers have time, competence, authority and usable evidence.

Use tomorrow: nominate the person authorised to approve the specific use, and define what would trigger suspension or reassessment.

07 Source and sustain

Procurement, commercial arrangements and whole-life consequences belong in the project plan and investment case from the beginning.

7.1 Procurement and sourcing: buying a result you can use

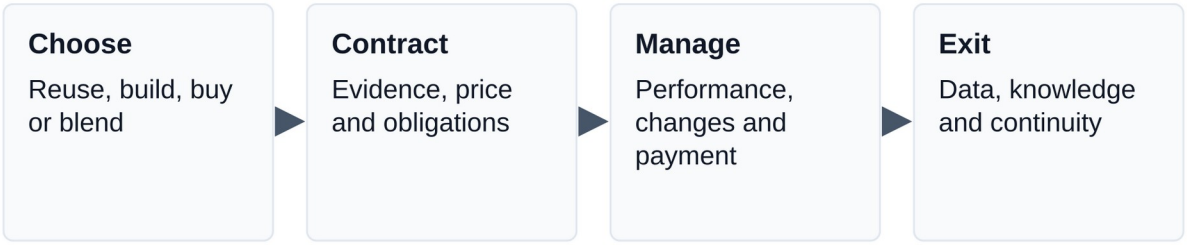


Figure 28 | Commercial choices shape delivery and exit

A supplier can provide specialist capability, capacity or a finished service. The project still needs an internal owner who can explain the need, evaluate delivery and arrange ongoing operation. Purchasing a product does not purchase accountability for the organisation's outcome.

The Australian Department of Finance's Commonwealth guidance treats value for money as broader than purchase price, including fitness for purpose, supplier capability, sustainability and whole-of-life costs. It also emphasises fair dealing, transparent decisions and proportionate procurement effort. These are useful reference principles; Commonwealth rules do not automatically define the requirements of a state agency or private organisation. [22, 23]

Decide how to obtain the capability

Before asking suppliers for proposals, compare reuse, internal delivery, purchase and a mixed arrangement. Start with the business problem rather than the features of a preferred product. Include a credible do-minimum option so the decision has a real comparator.

Table 7.1 | Options and decision evidence

Option	Strong reason to consider it	Question that could change the decision
Reuse an existing service	An approved capability already addresses most of the need	Can it support the workload, users and required controls?
Build internally	The capability is distinctive and the organisation can sustain it	Who will maintain it when the original team moves on?
Buy a product or service	The market offers a mature solution	What adaptation, integration and operating work remains ours?
Combine internal and external delivery	Specialist input complements internal ownership	Are responsibilities at the boundary sufficiently explicit?

Write a short sourcing decision that records the options, assumptions, whole-of-life estimate, capacity implications, major dependencies and recommendation. For a digital service, include implementation, integration, migration, training, support, licences, consumption charges, exit and decommissioning. Estimate the internal effort as well as supplier fees. A low-priced product that requires substantial operational work may have a higher total cost than a more complete service.

For AI work, distinguish acquiring a capability from developing a model. Buying a product may leave the organisation responsible for configuring instructions, preparing knowledge, integrating access and workflows, evaluating intended tasks and supporting users. An existing enterprise platform may offer the feature under a current agreement, but enabling it can still create delivery, data, support and cost decisions. Ask the platform and commercial owners to confirm what the current arrangement covers before treating the feature as ready to use.

If internal development or training is proposed, establish why reuse or configuration is insufficient and who can sustain the resulting capability. If purchase is proposed, establish what is being bought: a defined output, implementation effort, access to a product or a continuing service. That distinction changes the acceptance evidence, cost model and exit plan.

Select against evidence

Decide evaluation criteria and any mandatory conditions before evaluating bids. Separate a mandatory requirement from a scored preference: a high aggregate score should not disguise failure of a mandatory condition. Define what different scores mean, retain individual assessments and record the reasons for the moderated decision.

Use comparable evidence. A scripted demonstration using the same tasks and sample inputs is usually easier to compare than unrelated sales demonstrations. Ask who will actually perform the work and what assumptions underpin the proposal. Test the supplier's claimed delivery approach against a realistic change, failure or handover scenario.

An illustrative evaluation could consider functional fit, delivery capability, implementation effort, ongoing support and whole-of-life cost. Its weights should reflect the particular purchase. Record declared conflicts, control confidential bid information and provide material clarifications through the approved procurement channel. Do not invent additional evaluation criteria after seeing the bids.

Obtain evidence for the intended use

Ask for evaluation access that allows the organisation to test relevant tasks, permissions, limits and failure handling under agreed conditions. Confirm the product version, configuration, test-data arrangements, usage charges and time available for evaluation. Use approved representative material; agree any restrictions with the relevant data and specialist owners before providing it. Supplier benchmarks can inform selection, but the project still needs evidence that addresses its own users, tasks and operating conditions.

For a separate fictional purchase of a customer-correspondence assistant, the supplier's polished demonstration is followed by agreed sample cases, including incomplete information and requests outside the permitted scope. The buyer observes draft quality, correction effort and the ability to keep sending authority with staff. The evaluation informs the implementation estimate and acceptance criteria before a wider commitment. It does not promise that every future case will behave identically.

For an AI feature within an established enterprise platform, start with a proportionate evaluation of the changed workflow and permissions. Reuse relevant platform evidence, identifying what it actually covers and what the new use adds. Existing platform approval and a new feature's task suitability answer different questions.

Choose a commercial arrangement that fits uncertainty

UK government sourcing guidance links risk allocation with pricing and recommends placing risks with those best able to manage them. It cautions against fixed-price arrangements when the scope itself is not fixed. These are transferable design considerations, rather than contractual rules for another jurisdiction. [24]

Table 7.2 | Commercial arrangements

Arrangement	Basic commercial logic	Management attention
Fixed price	An agreed amount purchases a defined result	Make boundaries, acceptance and change pricing clear. The buyer still carries risks outside the supplier's commitment.
Time and materials	Payment reflects agreed effort and inputs	Set rates, reporting, spending limits and regular demonstrations of useful progress. Track remaining work, not only hours consumed.
Cost reimbursable	Eligible costs are reimbursed under agreed terms, with any fee arrangement	Define allowable costs, evidence, oversight and incentives. The buyer retains substantial exposure to cost growth.
Subscription or consumption	Payment follows access, seats, usage or another service measure	Model growth, renewal, overage, service levels, portability and the cost of leaving.

These categories can be combined. A short discovery engagement may reduce uncertainty before a separately authorised implementation commitment. A spending cap controls commitment; it does not guarantee that every desired feature will be delivered within that cap.

Make service dependencies explicit in the agreement

The project manager brings the intended use, delivery plan and evidence needs to procurement and the relevant commercial, legal, technical and specialist advisers. Those functions determine the appropriate terms and approval route. Translate the agreed obligations into project activities, acceptance conditions and continuing ownership rather than treating the contract as a document reviewed only at purchase.

Table 7.3 | AI procurement decisions

Commercial topic	Project decision it supports
Data and permitted use	Confirm what information may be supplied, where relevant processing occurs, how it may be retained or used, and which responsibilities remain with the buyer
Performance and support	Distinguish service availability and response commitments from evidence about answer or task quality; identify support access and incident escalation

Commercial topic	Project decision it supports
Release and feature changes	Confirm notice, evaluation access, version choices, feature-disable options and any ability to defer or reverse a change
Charging and capacity	Identify billable units, minimums, limits, price-change arrangements and likely workload; agree who monitors consumption and responds to exceptions
Deliverables and rights	Confirm access to the configuration, documentation, evaluation records and other material needed to operate and change the service
Exit and continuity	Establish export, transition assistance, continuing access, end-of-service obligations and the time needed to move or retire the service

Do not imply that a supplier has accepted terms that have not been agreed. Where a standard service offers little negotiation, show the remaining dependency and viable alternatives to the authorised buyer. A service may be unsuitable even when its purchase price fits the budget. Conversely, a limited feature may be manageable within existing terms when its boundaries and fallback are clear.

Connect acceptance, payment and ownership

For each purchased deliverable, identify the specification, evidence required, reviewer, acceptance authority, decision period and treatment of defects. Align payment milestones with the contractual agreement and evidence of performance. A demonstration is evidence to examine; it is not automatically acceptance.

Example: a supplier presents a configured knowledge service. Technical checks pass, but operations has no support procedure and the required export has not been demonstrated. Record these as distinct gaps, assess their contractual and operational significance, and use the agreed acceptance mechanism. The project manager should not imply that commercial acceptance, security authorisation and permission to launch are interchangeable decisions.

Maintain a small contract obligations register: obligation, owner, due date, evidence and status. Monitor buyer obligations too, such as providing access, decisions and test data. Escalate emerging disputes while evidence is current. Changes to a project plan do not themselves amend a contract; use the authorised commercial route.

Design the exit at the beginning

Agree who owns and can retrieve data, configuration, documentation and other deliverables. Identify the format and practical usability of exports, required transfer assistance, notice periods, ongoing access during transition and evidence of completed deletion where applicable. Record renewal dates early enough to permit a genuine decision.

Test portability as a practical transition task. Select a representative export, restore or import it into the proposed destination where feasible, and identify what cannot transfer. Depending on the service, the needed material may include source content, metadata, configuration, prompts, access mappings, evaluation cases and operational records, subject to the agreed rights and information controls. Record the effort and restrictions rather than assuming that an export button provides a usable replacement service.

A replacement may require new configuration, integration, evaluation and user guidance even when the data can be moved. Include those tasks and a fallback in the exit estimate. For an embedded platform feature, the fallback may be disabling the AI function and

continuing the conventional workflow; test that this remains workable. For a capability that performs actions, define how outstanding actions, authorisations and audit records will be reconciled before the service is changed or withdrawn.

At closure, reconcile accepted work, outstanding defects, invoices, warranties, support arrangements and access removal. Assign unresolved items to a continuing owner. Test a meaningful export before reliance on the service becomes difficult to reverse.

7.2 Sustainability and public value: judge the whole consequence

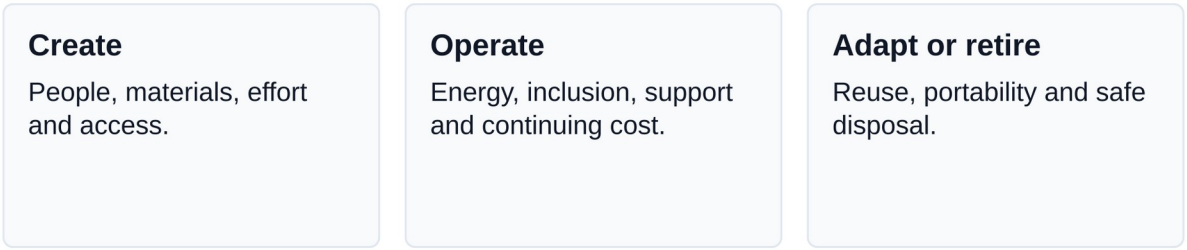


Figure 29 | Consider consequences across the useful life

A useful project can improve one measure while imposing costs elsewhere. A faster service might increase exclusion; a cheaper platform might shift work to frontline staff; an automated process might make decisions harder to understand. Include these effects in the decision, rather than treating them as incidental observations after delivery.

Use three lenses in the business case and at major changes:

Table 7.4 | Sustainability considerations

Lens	Questions for the project	Possible evidence
People and community	Who benefits, who carries the burden, and who could be excluded?	User research, accessibility checks, workload observations, complaints and distribution of outcomes
Environment	What resources are consumed during creation, use and retirement?	Travel, energy, materials, equipment life, waste and disposal assumptions
Enduring organisational value	Can the service be funded, supported, trusted and adapted?	Operating costs, skills ownership, supplier dependence, maintainability and benefits evidence

Select measures proportionate to the decision. A small internal improvement might need a short paragraph, an accessibility check and a named support owner. A large infrastructure or public-facing digital initiative may require specialist assessment and formal evidence. More questions are useful only when they improve a decision or the quality of control.

Keep different measures visible instead of forcing everything into a dollar estimate. Financial return, service access, workforce effects and environmental impacts can be presented together with their evidence and uncertainty. Do not label staff time released as cash savings unless an approved action converts that capacity into an actual budget reduction. Equally, improved service quality can matter even when it does not produce revenue.

For an illustrative staff knowledge service, measure whether answers help people complete real tasks, whether staff can challenge an answer, and whether content ownership is sustainable. Compare any time saved with the effort required to review outputs, maintain sources and support users. Record adverse effects and unintended benefits as evidence emerges.

Make AI consequences measurable without overstating certainty

Choose measures that can change the decision. For people, observe the distribution of task quality, review burden, access to alternatives and the ability to question or correct an outcome. A reduction in one team's workload may be accompanied by extra content maintenance or exception handling elsewhere. Agree who will collect these observations and who can act on them.

For environmental considerations, use relevant supplier or organisational evidence and state its boundary. Establish whether a measure covers the proposed service and period, and whether alternatives have been compared on a consistent basis. Do not convert an unsupported estimate of requests, energy or emissions into a precise sustainability claim. Where the consequence is material and evidence is inadequate, seek specialist input or preserve the uncertainty in the decision.

For continuing economic value, revisit the volume and human-effort assumptions as use grows. The independent purchased-feature example in Chapter 04 shows why a low usage price can coexist with a substantial review cost. A larger service may justify that cost through quality or capacity benefits, but the decision needs the complete comparison. For a small, reversible feature, a brief record of these effects may be sufficient; a separate assessment document is useful only when the decision warrants it.

Before continuation, ask: would we still choose this option knowing what we know now? Give the sponsor a clear choice between continuing, adapting, pausing and stopping. Past expenditure informs the record but should not by itself justify future expenditure. The decision concerns future consequences and available alternatives.

08 Transition and realise value

Close the temporary work responsibly, establish what value has occurred and retain the learning that improves future decisions.

8.1 Close the project and transfer continuing responsibility

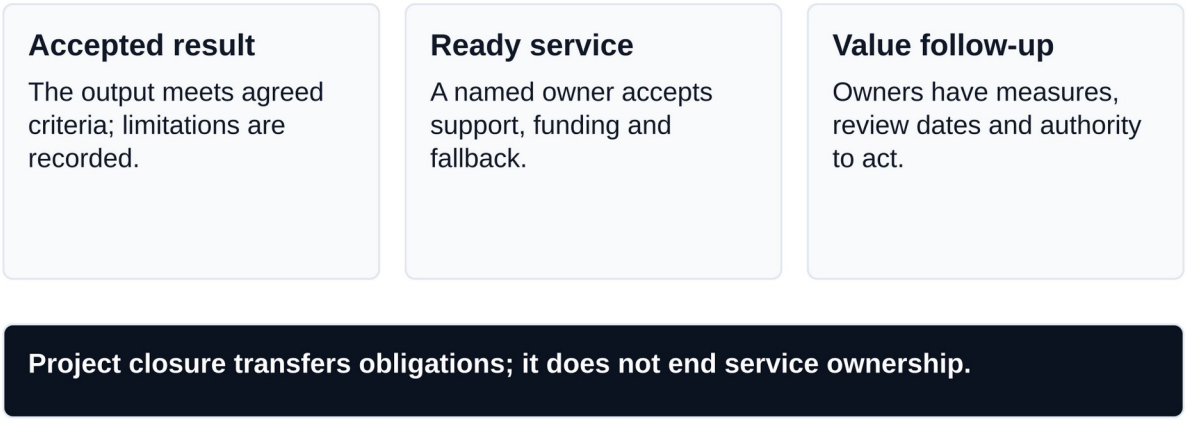


Figure 30 | Close the project and transfer continuing responsibility

Purpose. Finish the temporary undertaking without abandoning the service, obligations or benefit measurement it created.

Closure is an organised decision that the project has completed its authorised work or is ending for a stated reason. A project can close successfully, close with agreed residual work, or stop before the original scope is complete. Record which situation applies and why.

Separate four decisions: acceptance of deliverables; permission to operate or use the service; transfer of ownership and support; and closure of the project. They may occur together, but one does not automatically establish the others. A signed acceptance form is insufficient if no team has the capacity or funding to operate the service.

Table 8.1 | Project closure evidence

Closure area	Evidence to retain
Deliverables	Acceptance against agreed criteria, including exceptions
Finances and contracts	Actuals, remaining commitments, invoices and contract obligations
Operations	Owner, support, access administration, monitoring and maintenance
Residual matters	Open risks and issues with receiving owner and agreed actions
Benefits	Baseline, measures, owner and future review dates
Records	Decisions, configuration, source documentation and retention arrangements
People	Handover completed, resource release and recognition
Learning	What changed, why, and where the lesson will be applied

For the pilot, the end decision may be to disable the environment after capturing results, continue a limited service under an operational owner, or seek approval for another phase.

Do not allow a time-limited pilot to become a permanent production service by inertia. Confirm how users will be informed and how information and contracts will be handled if it ends.

A useful handover includes a rehearsal. Ask the receiving team to demonstrate access administration, a source update, an incident response and a recovery or rollback where relevant. This reveals practical gaps that a document pack may conceal. Confirm how much support the project team will provide after transition and when that support ends.

Hand over the evidence and the ability to act

The release record should identify the service configuration and use boundary that were accepted, the evaluation evidence supporting them, known limitations and any conditions. For the pilot, reference the relevant C01 configuration and E01 evaluation versions and the actual completed results. A planned evaluation design is not a completed result. Link the requirements, open exceptions and separate decisions on acceptance, authorised use and operational readiness.

Agree the continuing work before closure. Record an individual or accountable team for each responsibility, confirm capacity and funding, and obtain the receiving owner's acceptance. A role label in an organisation chart is not evidence that the handover has occurred.

Table 8.2 | Continuing operational ownership

Continuing responsibility	Receiving responsibility to confirm	Operating record or trigger
Sources and data	Content or data owner	Approved material, correction process, update checks and withdrawal of unsuitable material
Service configuration and releases	Service owner with the relevant technical or platform owner	Current version, supplier notices, change assessment and affected re-evaluation
Support and incidents	Service support owner	Help route, triage, escalation, suspension and recovery arrangements
Use and specialist conditions	Authorised use owner and relevant specialist functions	Permitted users/tasks, decision conditions and triggers for renewed advice or permission
Continuing costs	Receiving budget owner with finance	Usage, licences, review effort, maintenance and an exception threshold
Benefits and adverse effects	Business or benefit owner	Baseline, measures, review dates and decisions to improve, expand or retire

These are responsibilities to allocate, not mandatory new job titles. Several may sit with one person in a small service. Resolve conflicts where that person lacks time, authority or the independent judgement needed for a particular decision.

Rehearse a proportionate failure scenario. If the service produces unsuitable results, confirm who can restrict access, disable the feature or restore a supported configuration, and how users continue their work. Do not assume an earlier model or release can be restored when the supplier does not offer that option. In that case the agreed fallback may be conventional

search, manual processing or temporary suspension. For an action-taking capability, include stopping further actions and reconciling actions already performed; where reversal is impossible, define containment and resolution with the affected owner.

Ongoing evaluation, source maintenance, usage monitoring and response to supplier changes continue as service work after the project closes. If the next material expansion requires a new project or investment decision, state the trigger and owner. Closure should leave a controlled service or a deliberate shutdown, not an unowned pilot.

Worked situation. The team has completed delivery, but content owners have not accepted maintenance responsibility. The sponsor must decide whether to delay transition, narrow the service or fund an explicit interim arrangement. Marking the task “complete” does not solve the ownership problem.

Common failure: lessons are recorded as slogans such as “communicate better”. Capture the observed event, its cause, its effect and a specific change to a future practice.

Use tomorrow: name the person who will own each unresolved matter after closure and obtain their acceptance before releasing the project team.

8.2 Benefits need an owner, a baseline and a test



Figure 31 | Delivery enables benefits and operations sustains them

Purpose. Establish whether the project creates a worthwhile improvement and who will sustain it.

Distinguish outputs, outcomes and benefits. An output is the delivered knowledge assistant and its supporting materials. An outcome is staff completing certain information tasks differently. A benefit is a measured improvement that stakeholders value, such as reduced task time without poorer quality. The causal links need evidence; they are not guaranteed by launch.

A benefit record should contain the description, owner, baseline, target, measure, data source, expected timing, necessary operational changes, costs and relevant adverse effects. Include the assumptions linking the project to the benefit and describe how other influences will be considered.

Illustrative calculation. Forty staff perform two relevant searches each day for five days: 400 searches per week. Assuming all 40 participants are active, and total task time including user verification falls from eight minutes to five, gross time released is 400 x 3 minutes = 1,200 minutes, or 20 hours per week. If additional central source maintenance, service support and other work not already included in task time consume six hours, net capacity released is 14 hours. This is a modelled week of active use, not a benefit earned throughout all twelve project weeks. At an illustrative loaded rate of AUD 80 per hour, its nominal

capacity value is AUD 1,120 per week. This is not automatically a cash saving. A budget saving requires expenditure actually to be avoided or removed, and the organisation must decide how released capacity will be used.

Check the assumptions before annualising. Volumes, adoption, seasonality, quality and support effort may change. A short pilot sample does not prove steady performance across all teams. Compare like tasks, record the sampling method, and use an appropriate comparison or baseline where feasible. Report uncertainty and sample limitations.

Separate cashable savings, cost avoidance, service improvement, risk reduction and learning. Avoid adding them when they overlap. For instance, the same hours cannot be counted both as staffing savings and as additional output without explaining the allocation.

If the benefits do not justify the full cost, act on the evidence. The pilot's learning may still be useful, while an expanded service may be a poor investment. A credible recommendation can be to improve existing search, narrow the use case or stop.

Measure the complete task and its conditions

Before supervised use, agree how the team will observe task completion, quality, adoption and additional work. Use the same task boundary for the current process and the proposed service. The pilot's five-minute task time includes the user's verification; the six central hours cover work outside that boundary. Counting the same review time in both would understate the benefit, while omitting it from both would overstate it.

Break out results where they would change the decision, such as answerable versus unsupported tasks, different work groups or tasks requiring specialist escalation. Record unsuccessful tasks and adverse effects as well as completed ones. The business owner should judge whether the improvement remains worthwhile at an acceptable quality, with the required people and operating resources.

Keep capability, adoption and value separate. Passing agreed evaluation criteria can support a use decision. Active use can show whether the workflow is being adopted. Measured improvement against the baseline, after the required effort and costs, supports a benefits claim. None automatically proves the next. If results differ from the business case, update the forecast and the owner's next action rather than quietly changing the original target.

Source note. DTA's public benefits guidance emphasises measurable improvement, ownership, proportionate measurement and validation of forecasts. Its policy scope does not automatically apply to every Australian organisation. [25] PMI also publishes a dedicated benefits-realisation practice guide. [26]

Common failure: closing the project before assigning benefit measurement and operational follow-through.

Use tomorrow: choose the two most consequential benefits and confirm who will measure them, from which baseline, and on which date.

8.3 Test the value case with sensitivity

A continuation decision often depends more on a few uncertain assumptions than on the precision of a large spreadsheet. Vary those assumptions explicitly. The purpose is to discover what must be true for the preferred option to remain attractive.

Use the knowledge-assistant case as a teaching example. Assume 40 possible users, two relevant tasks per working day, five working days per week and an observed three-minute

reduction in total task time, including the user's verification. Assume six hours of additional central maintenance and support each week, outside that task time.

Table 8.3 | Weekly capacity sensitivity

Active users	Weekly tasks	Gross hours released	Less central overhead	Net hours released
10	100	5	6	-1
20	200	10	6	4
30	300	15	6	9
40	400	20	6	14

Each active user contributes half an hour of gross weekly capacity under these assumptions. Twelve active users therefore match six hours of overhead. This is an operational time break-even point only. It excludes project investment, licences, other operating costs and benefits such as improved service quality. It is not a financial break-even or a recommendation to proceed.

Identify the decision-changing assumption

At twenty active users, reducing the measured saving from three minutes to one changes gross released time from ten hours to about 3.3 hours. After six hours of central effort, the net position becomes about negative 2.7 hours. The value claim is sensitive to both adoption and actual task improvement.

Now compare the assistant with better document ownership and search. Apply the same baseline, task definitions and observation method. If the simpler option achieves similar improvements with lower ongoing effort, that is relevant evidence for the next investment decision.

Test the overhead assumption as well as adoption. The six central hours are held constant in this illustration so the effect of active users is visible. A larger source set, more exceptions or a supplier-release review may change that workload. Estimate the additional work from its cause and show the revised comparison, retaining the original case for explanation.

Keep operating resource costs beside the time result. A positive net capacity figure does not establish affordability after licences, usage, maintenance and remaining investment. The independent purchased-feature cost example in Chapter 04 illustrates that separate calculation; its volumes and prices must not be inserted into this pilot without an explicit change of assumptions.

Keep the measures honest

Use observed or clearly labelled assumed inputs. Measure representative tasks, variation and unintended effects. Avoid annualising a short trial without considering changing adoption, holidays, source maintenance and other demand. A nominal hourly value can help compare resource use, but released time becomes cash savings only through an actual spending change.

> **Business-case question:** Which assumption, if wrong, would change our preferred option, and how can we test it before committing more?

8.4 Work through the midpoint governance decision

At the end of week six, the fictional pilot has completed less budgeted work than planned. The original cost baseline is AUD 108,000; the funding envelope is AUD 120,000. PV is AUD 54,000, EV is AUD 43,200 and AC is AUD 48,000. These figures support a performance conversation; they do not explain its cause.

Step 1: establish what is happening

Reconcile the reporting cut-off, completion evidence and actual costs. Inspect which work packages are late and which activities control completion. Ask whether the problem is temporary mobilisation, underestimated content effort, rework, a new requirement or a measurement error. Do not select a recovery action before the cause is sufficiently understood.

Step 2: compare the remaining choices

Table 8.4 | Options and decision evidence

Option	Evidence needed	Trade-off to explain
Continue the authorised scope	Credible remaining-work estimate and cause-specific response.	May consume reserve and still require a schedule change.
Narrow scope through change control	The smaller pilot can still answer a useful business question.	Lower coverage, possibly lower cost or exposure; some learning deferred.
Move the decision date	Revised dependencies, participant availability and cost effects.	More time may protect evaluation quality but delay value or add cost.
Pause to resolve a prerequisite	A defined issue, owner, resolution plan and review date.	Avoids unsafe or wasteful work while retaining restart obligations.
Stop and close responsibly	Updated case showing continuation is no longer justified.	Retains learning and avoids remaining cost; manages contracts and transition.

Step 3: make an explicit decision

If the sponsor recommends continuing to a forecast AUD 120,000, the appropriate authority must decide whether reserve use is permitted and justified, or whether scope, funding or another commitment should change. Record any approved reserve transfer and baseline change separately, preserving the original performance history. The existing funding envelope is not a reason to skip that decision. Record any conditions, such as evidence that the reduced source set is supportable, and a point to test whether recovery is working.

If the separate supplier-change scenario applies

The supplier-release scenario in Chapter 04 is an additional what-if at this same review point. It adds activity X after F, 25 evaluation hours and 15 platform-support hours at AUD 80, with all other costs unchanged for the illustration. ETC becomes AUD 75,200, EAC becomes AUD 123,200 and completion moves to week 13. The original pilot's baseline and AUD 120,000 funding envelope remain unchanged until an authorised decision.

In that scenario, approval to use reserve inside the existing envelope would not cover the additional AUD 3,200 beyond funding. The sponsor needs a feasible choice about additional

funding, an acceptable supported configuration, a narrower scope or stopping, together with the timing and evidence consequences. Record the decision and test resource availability before committing the revised dates. This scenario does not replace the unamended midpoint example above.

Step 4: preserve the evidence

Update the forecast, action owners and communications. Amend baselines only through the agreed authority, retaining the original reference and reason. Inform suppliers through the appropriate commercial process where their obligations change. At the next review, evaluate whether the intervention improved the forecast and the value case.

> **Sponsor question:** Knowing what we know today, which remaining option offers the best justified use of our resources?

8.5 Turn lessons into changed practice

Capture useful lessons while the details and evidence are available. Waiting until closure makes it harder to distinguish the original cause from later explanations. A lesson is valuable when somebody changes a practice and checks the effect.

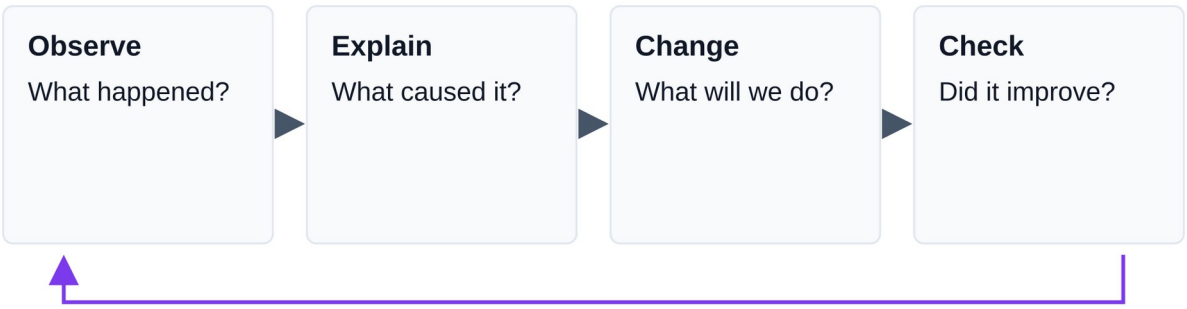


Figure 32 | Turn observed experience into changed practice

A usable lesson record

Table 8.5 | Reusable working record

Field	Example
Observed event	Content approval arrived after configuration work had started.
Evidence	Approval dates, rework tasks and version history.
Contributing cause	Source ownership and approval turnaround were assumed rather than confirmed.
Effect	The team reconfigured the collection and repeated tests.
Practice change	Confirm the content owner and response commitment before baselining dependent work.
Owner and application	Project practice owner adds a dependency check to the next initiation review.
Follow-up	Review the next two projects to see whether late approvals and rework fall.

Improve the system, not only the team

A repeated problem across projects may belong to a portfolio process, procurement arrangement, shared platform or organisational capacity decision. A project manager cannot solve every systemic constraint with a better task list. Escalate recurring patterns with evidence and an owner for the organisational response.

Keep positive lessons too. A source-owner workshop may have prevented weeks of rework. Record what made it useful and when the same approach is likely to help again. Avoid turning one success into a mandatory meeting for every project regardless of context.

Retain AI lessons with their context

Record the relevant use, configuration, source conditions and evaluation method when they explain the result. A lesson from one supplier version or task set may need retesting before it is applied elsewhere. Preserve enough evidence to support the lesson without unnecessarily copying sensitive source material into a general project repository.

Useful practice changes might include obtaining evaluation access before commitment, securing reviewer time before scheduling tests, or agreeing a conventional fallback before enabling a feature. Name the project, procurement or service owner who will change the practice and the next occasion on which its effect can be observed. Share positive evidence about a simpler option or a stopped experiment as well as successful deployment; both can improve future investment decisions.

A short retrospective

Ask what helped, what hindered, what surprised the team and what should change next. Choose one or two improvements that can be tested in the next cycle. Record the expected effect. At the following review, check whether the change actually helped before adding more process.

> **Test of a lesson:** Can you name the next situation where it will be applied and the person who will make that happen?

09 Reusable templates

Copy the fields that support your decisions and tailor them to existing systems. These are original Strategen working templates, not mandatory PMI forms. Replace the bracketed prompts, retain evidence links and identify actual decision authorities. A project with a small AI component can use these fields inside its existing records; it does not need a parallel set of AI forms.

Each template includes a reusable record and a completed fictional example. The examples describe decisions or planning snapshots at stated points in the same knowledge-assistant pilot. A completed planning record is not a claim that a test passed, a service was authorised or a benefit occurred. Record those results when evidence exists.

The pilot baseline remains 12 weeks, 40 nominated staff, one approved internal collection, an AUD 108,000 cost baseline and AUD 120,000 total funding. Controlled validation with eight participants occurs during integrated tests in weeks 8–9; the supervised pilot is planned for weeks 10–11, followed by handover and the continuation decision in week 12. The additional one-week, AUD 3,200 scenario in Chapter 04 is a separate what-if and does not change these completed baseline examples.

Choose the record for the decision

Table 9.1 | Template finder

Record	When to use it	Decision or control it supports
T01 Project charter	Initiate or reset the undertaking	Authorisation, boundaries, owners and funding
T02 Business case	Compare alternatives and revisit value	The next justified commitment, including deferral or stopping
T03 Delivery baseline	Agree and maintain delivery commitments	Scope, requirements, acceptance, dates and forecasts
T04 RAID tracker	Manage emerging uncertainty and dependencies	Owned actions and timely escalation
T05 Decision and change record	Consider a change or record a decision	Authority, consequences, evidence and follow-through
T06 Status report	Review delivery at a stated cut-off	Progress, forecast, exposure and decisions needed
T07 Release and handover	Prepare for use and transfer responsibility	Bounded release, operational acceptance and closure
T08 Benefits record	Plan and assess the intended outcome	Baseline comparison, total effort, ownership and value

T01 Project charter

Discuss the charter before seeking approval. Authorisation establishes a bounded undertaking; it does not turn unresolved assumptions into facts. The AI use boundary tells the sponsor what is being committed to, while the evidence and stop conditions explain when that commitment must be reconsidered. Keep the signed decision and current charter together.

Copy and complete

Project: [Name] **Version/date:** [Version and date] **Decision sought:** [Authorise, revise or close]

Table 9.2 | Project charter | reusable record

Field	Complete in plain language
Purpose and outcome	[Problem, affected people, current performance and improvement sought; identify the question the project must resolve]
Sponsor / delivery lead / business owner	[Names and responsibilities; identify who accepts deliverables, who permits use and who will own the service]
Deliverables and boundaries	[Outputs; intended users, tasks, sources, operating conditions and permitted actions; exclusions; links to detailed requirements]
Success and evidence	[Observable completion and value measures; evidence owner; assessment dates; limits on the conclusion]
Time and funding	[Milestones; funding envelope; cost baseline; reserves and release authority; continuing costs requiring a separate commitment]
Dependencies and assumptions	[Data or content access, supplier capability, interfaces, reviewer capacity and other conditions that could change feasibility]
Decision authority	[Team authority; escalation route; applicable delegations; decisions reserved to specialist or operational owners]
Main exposure and stop conditions	[Material risks; findings or changes that require pause, narrowing, another approach or reconsideration]
Authorisation	[Approver, decision, date, conditions and next review; state exactly what this authorisation permits]

Completed fictional example

Project: Internal knowledge-assistant pilot. **Record:** CH01, version 1, mobilisation. **Decision sought:** Authorise a bounded pilot project.

Table 9.3 | Project charter | completed example

Field	Completed entry
Purpose and outcome	Test whether staff can find and verify approved guidance with less total effort. Compare the proposed assistant with the current search method; establish evidence before claiming an improvement.
Roles	Sponsor: project authorisation and continuation. Project manager: integrated delivery and forecasts. Business owner: task suitability and acceptance. Access-control owner: permission evidence. Service owner: operational readiness and support. These are fictional allocations, subject to real delegations when applied.
Deliverables and boundaries	Configured pilot, prepared collection and permissions, E01 evaluation evidence, user preparation, supervised use, handover and recommendation. Forty nominated staff; one approved collection; assistance only. Exclude public use, autonomous decisions, personal-email ingestion, HR case advice and enterprise rollout.
Success and evidence	R01–R04 and E01 version 1 define acceptance evidence. The benefits record compares correct task completion and total effort with the current method. A useful learning result can support a decision to stop without establishing that the service is acceptable.
Time and funding	Integrated tests complete in week 9; supervised use in weeks 10–11; handover and decision pack in week 12. Planned work AUD 100,000 plus contingency AUD 8,000 gives BAC AUD 108,000. Management reserve AUD 12,000 sits outside BAC but inside total funding of AUD 120,000. Permanent operation is outside this approval.
Dependencies and assumptions	Content-owner approval and access arrangements are required before their dependent work. Evaluation reviewers and pilot users need booked capacity. The supplier must provide a supported configuration suitable for the intended environment.
Decision authority	The team can adjust implementation within the agreed scope. The sponsor decides scope, funding and continuation within the assumed delegation. Access and operational decisions retain their nominated owners. Reserve access is an explicit decision.
Stop conditions	Do not begin supervised use without the required acceptance, access and operational evidence. If an observed access breach or materially misleading answer remains unresolved, correct and re-evaluate or narrow the permitted use. Reconsider the investment if the intended task remains unsuitable.
Authorisation	Illustrative mobilisation decision: authorise the bounded project and its stated funding; retain the week-12 continuation decision. This does not authorise live use, enterprise expansion or permanent operating expenditure.

T02 Business case and option decision

Record significant disadvantages as well as benefits. Separate benefits already evidenced from those still hypothesised. Use a range or an explicit unresolved estimate where uncertainty remains; do not write zero because an operating cost is unknown. Update the case when facts change. An earlier approval is not a reason to continue an unattractive investment.

Copy and complete

Decision: [What must be decided, by whom and by when?]

Need: [Current performance, affected people and reason for intervention. State whether AI is necessary or one possible response.]

Table 9.4 | Business case and option decision | reusable record

Option	Expected benefit or learning	Whole-life cost and effort	Main uncertainty / evidence needed
Continue current arrangements	[Current performance and consequences]	[Ongoing expenditure and effort]	[What persists or worsens?]
Smallest useful intervention	[Improvement or question it could resolve]	[Delivery, adoption and operation estimate or range]	[Assumptions to test before further commitment]
Broader alternative	[Additional benefit]	[Delivery, integration, transition and ongoing estimate]	[Additional dependencies, obligations and evidence]

Recommendation and reason: [Explain the trade-off and the next bounded commitment.]

Benefit owner and measurement: [Owner; baseline; quality and adoption measures; additional review effort; review date.]

Affordability and dependencies: [Funding source, capacity, supplier terms, usage assumptions, continuing costs and confidence in estimates.]

Conditions and decision: [Approve, reject, defer or request evidence; authority; date; conditions; next review and stop criteria.]

Completed fictional example

Record: BC01, version 1, mobilisation. **Decision:** Whether the sponsor should fund a 12-week learning investment before considering wider use.

Need: Staff must locate and verify internal guidance. The amount of avoidable effort has not yet been established. Faster answer generation alone would not demonstrate value if checking, correction or support consumes the apparent saving.

Table 9.5 | Business case and option decision | completed example

Option	Expected benefit or learning	Cost and effort position	Main uncertainty
Continue current search	Provides the comparison method and avoids a new service commitment	Existing operating effort continues; measure it during baseline observations	Current task time, quality and recurring difficulty require observation
Improve source material and search	Could reduce avoidable effort without an AI capability	Obtain a proportionate estimate from content and search owners before treating it as the preferred alternative	Whether content quality or navigation is the main cause of difficulty
Bounded assistant pilot	Tests usefulness, limitations, adoption and operating demands for one collection	AUD 120,000 project funding, including BAC AUD 108,000 and management reserve AUD 12,000; no permanent service funding approved	Suitability, data access, support effort, usage costs and value at larger scale
Enterprise rollout	Could extend a proven improvement to more work	Not estimated or authorised by this case; needs a separate whole-life proposal	Wider tasks, permissions, support demand and consequences are not covered by pilot evidence

Recommendation: Authorise the bounded pilot, keep the current method as the comparator and explicitly test whether source or search improvements are sufficient. A larger commitment depends on the resulting evidence, including legitimate options to narrow, defer or stop.

Benefit ownership: The business owner owns correct task completion and effort measures. The service owner supplies support and continuing-cost evidence. E01 controls the quality comparison; T08 records the value assessment.

Affordability and conditions: The project fits the illustrative AUD 120,000 funding envelope at authorisation. Actual usage and review effort must be recorded; a pilot budget is not a forecast of enterprise operating cost. Content permissions, supplier suitability and reviewer availability remain dependencies. At week 12 the sponsor decides continuation, supported by the business and service owners. Pilot acceptance does not authorise expansion.

T03 Delivery baseline and requirements

Keep delivery responsibility, acceptance authority, baseline dates and forecasts visible. Configuration identifies what was evaluated; evaluation conditions identify the limits of the conclusion. Updating the forecast records current expectations. Changing the baseline requires the agreed authority. Link detailed schedules, WBS records, requirements and cost plans instead of duplicating them here.

Task-time evidence under R04 informs the benefits decision separately. The service owner retains operational acceptance. End-of-week-nine evidence supports a bounded-use decision; it does not itself authorise supervised use, expansion or closure.

Copy and complete

Record/version: [Identifier] **Baseline:** [Effective date and approving decision] **Forecast date:** [Reporting cut-off] **Record owner:** [Name]

Table 9.6 | Delivery baseline and requirements | reusable record

Baseline element	Agreement or reference
Scope and outputs	[Included users, tasks, data or knowledge sources, permitted actions, operating conditions, deliverables and main exclusions]
Completion standard	[Evidence needed to claim completion; identify experiment findings separately from accepted service capability]
Configuration and evaluation	[Version or record identifying the model or supplier release, prompts, retrieval, content, permissions, tools and interfaces that matter; linked evaluation plan]
Milestones and dependencies	[Milestones, external inputs, provider, needed-by dates, capacity and commitment status]
Cost and funding	[Cost baseline; contingency; management reserve; total funding; current forecast and operating-cost boundary]
Change route	[Triggers for review and re-evaluation; who can approve each change; link to T05]

Table 9.7 | Delivery baseline and requirements | reusable record

Deliverable / requirement ID	Acceptance evidence and decision owner	Delivery responsibility	Baseline date	Current forecast
[ID and observable result]	[Condition, evidence and person who decides]	[Person who produces the result]	[Date]	[Date]
[ID and observable result]	[Condition, evidence and person who decides]	[Person who produces the result]	[Date]	[Date]

Completed fictional example

Record: DB01, version 1. **Baseline:** Set at mobilisation under CH01. **Forecast date:** End of week 6. **Owner:** Project manager. This records planned controls; it does not claim completed test results.

Table 9.8 | Delivery baseline and requirements | completed example

Baseline element	Completed entry
Scope and outputs	Forty staff; one approved collection; assistance only. Deliver the configured pilot, content and permissions, E01 evidence, user preparation, supervised use, handover and recommendation. Exclude public use, autonomous decisions, personal-email ingestion, HR case advice and enterprise rollout.
Completion standard	Agree completion for each work package. R01–R04 govern service requirements. Completed learning work does not establish service acceptance or realised benefits.
Configuration and evaluation	C01 version 1: supplier S1, prompt P1, retrieval R1, collection K1, permissions A1, interface I1. E01 version 1 holds the Chapter 03 evaluation design. Match the actual deployed versions to C01 before using results.
Milestones and dependencies	Access controls: end week 4. Integrated tests: end week 9, including eight-person controlled validation in weeks 8–9. Supervised use: weeks 10–11. Handover and decision pack: week 12. Book reviewers and users against these dates.
Cost and funding	BAC AUD 108,000: work AUD 100,000 plus contingency AUD 8,000. Management reserve AUD 12,000; total funding AUD 120,000. Week 6: AC AUD 48,000 + ETC AUD 72,000 = EAC AUD 120,000. Authorise reserve use; fund permanent operation separately.
Change route	Review model or supplier, prompt, retrieval, collection, permission, tool, interface and intended-use changes. The project manager identifies affected evidence, work and forecasts for the authorised decision. Preserve earlier baselines and configurations.

Table 9.9 | Delivery baseline and requirements | completed example

Requirement	Acceptance evidence and decision owner	Delivery responsibility	Baseline	Forecast
R01: Inspect supporting guidance	E01 answerable cases and source rubric; business owner	Evaluation lead, with content reviewers	End week 9	End week 9
R02: Respect access boundaries	E01 positive and negative access cases against A1; access-control owner	Technical lead, with evaluation lead	End week 9	End week 9
R03: Handle missing or ambiguous evidence	E01 unsupported and ambiguous cases; business owner	Evaluation lead, with content reviewers	End week 9	End week 9
R04: Complete the intended staff task	Observed controlled user tasks against the agreed quality rubric; business owner	Evaluation lead, with user representatives	End week 9	End week 9

T04 RAID and dependency tracker

Review actions and changed exposure, not every row mechanically. A risk that occurs may become an issue; retain its history. Where a specialist review raises a finding, link it to the affected requirement, owner, action and decision date. Closing a row requires evidence or explicit acceptance by the appropriate authority.

Copy and complete

Use one row per item. A **risk** is an uncertain future event. An **assumption** is a belief to verify. An **issue** is a present problem. A **dependency** is something needed from another activity or party.

Table 9.10 | RAID and dependency tracker | reusable record

Field	Fill-in prompt
ID, type and raised date	[Unique ID; R, A, I or D; date]
Statement	[What is uncertain, believed, already happening or required?]
Consequence	[Effect on outcome, time, cost, quality, people, acceptance or operating obligations]
Owner	[One person accountable for follow-through; identify the relevant decision owner separately where needed]
Next action and due date	[Specific action, actor and deadline; include required input or evidence]
Trigger / review	[What changes the response; next review date; affected configuration or use if relevant]
Status and evidence	[Open, active, resolved or accepted; dated evidence or decision link]

For a risk, add likelihood and impact using the organisation's agreed approach, response and residual exposure. For a dependency, add the provider, needed-by date, provider commitment and escalation contact. For an assumption, state how it will be tested. This is a delivery record, not a separate AI risk scoring system.

Completed fictional example

Record: D03, dependency, raised at mobilisation; review point before completion of access controls.

Table 9.11 | RAID and dependency tracker | completed example

Field	Completed entry
Statement	Access controls and their configuration evidence must be ready at the end of week 4 for the working-service build.
Provider and commitment	Access-control owner provides the permission matrix and confirms implementation readiness. The week-4 date is the plan; the project manager must obtain and retain the provider's dated commitment.
Consequence	A one-week delay on this dependency moves the modelled finish to week 13 if no recovery action is taken. This is a sensitivity of the original network, not a change to the baseline.
Owner	Project manager owns follow-through; the access-control owner owns the control evidence and related decision.
Next action	Before the week-4 completion point, the technical lead checks configuration readiness with the access-control owner; the project manager records the evidence date and confirms the next work can start.
Trigger and escalation	Escalate a forecast miss as soon as it becomes apparent. Prepare options and their effect on acceptance, resource availability and the finish date. Do not remove permission tests to recover time.
Status and evidence	Open at this planning snapshot. Closure requires the dated permission record and readiness evidence, not a verbal assurance that the work is nearly finished.

An assumption about source quality belongs in its own row, with a content owner and a sampling action. An observed stale policy is an issue requiring correction or exclusion. Keeping these distinct helps the team choose the right response without creating a larger register.

T05 Decision and change record

Record the reason as well as the result. A changed prompt or supplier release can affect evidence and permissions even when its technical implementation is small. Agree the extent of re-evaluation through the existing change route. Approval is incomplete operationally until affected requirements, forecasts, work instructions and commitments have been updated. The separate AUD 3,200 what-if in Chapter 04 illustrates another change assessment; it is not the decision in CH04.

Copy and complete

Record ID: [ID] **Raised by/date:** [Person and date] **Decision required by:** [Latest useful date]

Table 9.12 | Decision and change record | reusable record

Field	Fill-in prompt
Question or proposed change	[What needs a decision, why now and which configuration or intended-use boundary is affected?]
Options considered	[Include retaining current arrangements, narrowing or deferring where feasible]
Expected value	[Benefit, learning or obligation addressed]
Impact	[Scope, schedule, cost, quality, resources, risk, benefits and operation; identify affected evidence and re-evaluation work]
Recommendation	[Preferred option and supporting evidence; state uncertainties]
Decision and authority	[Approve, reject or defer; approver; authority; date; conditions; permission actually granted]
Implementation	[Owner, actions, dates and affected requirements, configuration, plans, forecasts and records]
Communication and review	[Who needs to know; when implementation and consequences will be checked]

Evidence links: [Analysis, estimate, requirement, supplier advice, evaluation plan and meeting record.]

Completed fictional example

Record: CH04, second-collection request. **Raised by:** Requesting business workstream during the pilot. **Decision point:** Before any second-collection preparation begins.

Table 9.13 | Decision and change record | completed example

Field	Completed entry
Proposed change	Add a second internal content collection to the assistant. The approved boundary is one collection.
Options	Continue with one collection; add the second now after estimating and approving the work; or defer it for the continuation decision.
Expected value	Potentially covers additional staff questions. That value has not been demonstrated by the first collection's evaluation.
Impact	Additional content ownership, permissions analysis, remediation, evaluation, user guidance and support work. C01 and affected R01–R04 evidence would need review. No supported schedule or cost estimate for this added scope is yet available.
Recommendation	Defer until the first collection has suitable evidence and a whole-life case for expansion can be assessed. Protect the current evaluation and week-12 decision point.
Decision and authority	Illustrative sponsor decision: defer the request; retain the existing scope, baseline and funding. This decision assumes the sponsor holds the relevant delegation and does not override access or service-owner authority.
Implementation	Project manager records the future option and keeps the current collection boundary in DB01. No new collection is connected and no baseline change is made.
Communication and review	Tell the requester, technical lead, content owner and evaluation lead. Revisit the option with the week-12 recommendation or sooner if the sponsor requests a properly estimated proposal.

Evidence: Current scope in CH01/DB01, collection K1 in C01 and the E01 evaluation plan. The request does not provide evidence that the second collection is ready or valuable.

T06 One-page project status report

Lead with the exception or choice that matters. Distinguish facts, estimates and missing evidence, and explain changes since the previous report. Apply agreed status thresholds with judgement. Keep detailed logs behind links, include ongoing-cost implications when they change the decision and raise urgent exceptions between reporting dates.

Copy and complete

Project / reporting date: [Name and cut-off] **Overall confidence:** [Status plus one-sentence reason]

Decision required: [Decision, recommendation, authorised recipient and latest useful date. Write “none” when no decision is needed.]

Table 9.14 | One-page project status report | reusable record

Area	Current position
Outcome and scope	[What remains achievable; changed assumptions, intended-use boundaries or benefit concerns]
Accepted progress	[Completed work against agreed conditions; evidence links; distinguish learning completed from service acceptance]
Milestones	[Baseline; forecast; movement; cause; next evaluation, adoption or transition dependency]
Financials	[Baseline; actuals; commitments; ETC; EAC; reserve position; usage and review-cost assumptions]
Main risks / issues	[Up to three material items with consequence, owner and action]
Next period	[Next deliverables, decisions, reviewer capacity and dependencies]

Completed fictional example

Project: Internal knowledge-assistant pilot. **Reporting cut-off:** End of week 6. **Overall confidence:** Continuation remains conditional on affordable completion and the planned acceptance evidence; the cost forecast requires a decision.

Decision required: Sponsor to consider the AUD 12,000 forecast excess over BAC and the authorised funding response before further commitments depend on reserve access. Retain the original scope while the cause and remaining estimate are reviewed.

Table 9.15 | One-page project status report | completed example

Area	Completed entry
Outcome and scope	The target remains a bounded 40-person pilot using one collection and a week-12 continuation recommendation. Enterprise rollout and a second collection remain outside scope. No benefit result is claimed at this cut-off.
Accepted progress	The illustrative earned-value record reports EV AUD 43,200 against PV AUD 54,000. This reflects the agreed work-completion rules; it does not assert that R01–R04 have passed or that supervised use is authorised. The project manager must retain the supporting completion evidence.
Milestones	Baseline and current planning forecast remain: integrated tests end week 9; supervised use weeks 10–11; handover and decision pack week 12. Confirm the remaining dependency and resource assumptions when reforecasting. SPI does not calculate a completion date.
Financials	BAC AUD 108,000; AC AUD 48,000; ETC AUD 72,000; EAC AUD 120,000. Forecast variance is AUD 12,000 above BAC and reaches the total funding envelope. Management reserve is AUD 12,000, with access subject to the nominated authority. Check commitments and accruals are included consistently in the estimate.
Exposure requiring attention	Cost pressure needs a cause-based response. The evaluation lead must confirm reviewer capacity before integrated tests. The service owner must confirm support and fallback readiness before supervised use. Each item remains linked to its live record.
Next period	Continue planned integration, prepare E01 evidence, confirm C01 identifiers and book the eight-person controlled validation in weeks 8–9. Return the financial decision and any change to the sponsor with a current forecast.

For the stated inputs, CPI = 43,200 / 48,000 = 0.90 and SPI = 43,200 / 54,000 = 0.80. EAC = 48,000 + 72,000 = AUD 120,000. The separate supplier-change what-if is not included in these figures.

T07 Release, handover and acceptance

Once a decision is made, update the decision fields and link the evidence; preserve the earlier version. Conditional acceptance needs owned and dated conditions, and must not disguise a failed prerequisite. Name who monitors future supplier, model, content, permission or use changes and who decides whether re-evaluation is required. The project can close after transferring those ongoing responsibilities; it should not remain open indefinitely because the service continues operating.

Copy and complete

Deliverable / configuration: [Exact item and version] **Record date:** [Date] **Decision sought:** [Deliverable acceptance, permission to use, operational acceptance or closure]

Table 9.16 | Release, handover and acceptance | reusable record

Field	Fill-in prompt
Intended or authorised use	[Users, tasks, sources, permitted actions, environment, duration and limits; distinguish proposed from authorised use]
Acceptance evidence	[Criteria, configuration, results, limitations, reviewers and evidence links; identify missing prerequisites]
Open items and limitations	[Consequence, owner, action, date and authority for any retained exposure]
Operational ownership	[Service owner; support and escalation; owners of content, access, costs and future changes]
Readiness	[Runbook, access, training, monitoring, recovery, fallback and disable/rollback evidence]
Transfer	[Assets, configuration, records, supplier obligations, continuing expenditure and capacity]
Decisions and authority	[Separate decision, authority, date and conditions for deliverable acceptance, use and operational responsibility]
Follow-up and closure	[Review date; benefits owner; unresolved project obligations; closure recommendation and continuing responsibilities]

Completed fictional example

Record: RH01, version 1, prepared during week 8. **Configuration:** C01 version 1. **Decision sought:** Readiness for the end-of-week-nine bounded-use decision. This is a completed pre-release record, not a report of test results.

Table 9.17 | Release, handover and acceptance | completed example

Field	Completed entry
Proposed use	Forty nominated staff, one approved collection, assistance only, during supervised use in weeks 10–11. No public use, autonomous decisions or enterprise rollout. This use is planned; it is not yet authorised at this snapshot.
Evidence position	E01 version 1 defines the tests and controlled validation against R01–R04. Integrated tests run in weeks 8–9, including eight validation participants. Results and limitations must be reviewed at the end of week 9; this record claims no pass result.
Open prerequisites	Evaluation lead supplies the results; access-control owner decides the permission evidence; service owner decides support readiness. An observed access breach or materially misleading answer blocks this pilot's release pending correction and re-evaluation or a justified narrowing of use.
Operational ownership	Service owner is the nominated recipient of the bounded pilot's support responsibility; formal operational acceptance is still required. Content owner maintains collection K1; access-control owner maintains permissions. Project manager coordinates transfer without taking those owners' decisions.
Readiness to demonstrate	Support must be able to identify the deployed configuration, receive user reports, disable the assistant and direct staff to the current guidance-search method. The runbook exercise must establish who acts and how access and service are restored.
Transfer and costs	Transfer C01, E01 evidence, the user guidance, support contacts, content ownership and supplier commitments. Pilot-period costs remain in the project forecast. Any service beyond the pilot needs an owner, budget and separate authorisation; none is implied here.
Decision status	At week 8, no use permission is recorded. Project manager to present the evidence at the planned week-9 decision. Business acceptance, access approval, operational acceptance and permission for bounded use must each be explicit, even if captured at one meeting.
Follow-up and closure	Week-12 handover and decision pack will record findings, obligations, benefits evidence and the recommendation to expand, modify, pause or stop. Closing the project requires the agreed closure criteria and an owned transfer or retirement decision.

T08 Benefits record and review

Keep output completion separate from benefit achievement. Training attendance and access granted can support adoption, but neither establishes improved performance. Count additional effort once and avoid overlapping benefits. Record other changes that may explain the result. If the service continues, transfer the measurement method, review dates, costs and action authority to named operational owners.

Copy and complete

Benefit ID / description: [A change in business performance or stakeholder experience, not merely an output] **Review date:** [Date and decision supported]

Table 9.18 | Benefits record and review | reusable record

Field	Fill-in prompt
Owner	[Person accountable for the operational change and measurement, including after closure]
Measure and baseline	[Metric, unit, quality standard, method, baseline period and evidence]
Target and timing	[Expected improvement or hypothesis; review date; confidence and assumptions]
Delivery contribution	[Outputs that enable the improvement; avoid attributing all change to the project]
Adoption and dependencies	[Actual use, workflow change, permissions, content currency, skills and support]
Costs and possible disbenefits	[Usage, checking, correction, rework, support, maintenance and adverse effects; relevant social, environmental or economic consequences]
Review evidence	[Observed result, comparison, sample limits, data quality and other influences; distinguish missing evidence from zero benefit]
Decision and next action	[Continue, adjust or stop; owner; authority; date; continuing review responsibility]

Completed fictional example

Record: B01, version 1, measurement plan at the week-six review. **Benefit:** Staff find and verify approved guidance with less total effort while meeting the agreed task-quality standard.

Table 9.19 | Benefits record and review | completed example

Field	Completed entry
Owner	Business owner owns the improvement and measurement. Service owner supplies continuing-cost and support evidence; evaluation lead supplies comparable task-quality evidence.
Measure and baseline	Observe representative tasks using the current method and the pilot. Record correct completion, elapsed task time and active effort, including source inspection and correction. Record support and maintenance effort separately. Baseline observations must be retained before an improvement is claimed.
Target and timing	Hypothesis: the pilot can reduce effort without reducing task quality. No numerical saving is committed before observation. Review controlled validation in weeks 8–9, supervised-use evidence in weeks 10–11 and the overall value case at week 12.
Delivery contribution	Prepared content, source display, permission controls, user guidance and a supported workflow enable the proposed change. Faster generation alone is insufficient.
Adoption and dependencies	Forty nominated users form the bounded pilot population. Record actual participation and task coverage. Current source material, working permissions and adequate checking time are required. Eight validation participants do not establish organisation-wide adoption.
Costs and disbenefits	Count checking, corrections, unsuccessful attempts, support, content maintenance and usage charges. Record accessibility barriers, misleading responses, displaced work or extra burden. Translate released capacity into money only with a clear rate and interpretation; cash savings need an actual expenditure reduction.
Evidence at this review	No realised benefit is asserted at week 6. The measurement plan is established; observations and their limitations will support the week-12 decision. Pilot findings will not automatically generalise to other tasks, users or collections.
Decision and next action	Continue the agreed evidence collection. Business owner prepares the week-12 benefit assessment; service owner supplies operating implications. Sponsor chooses the next commitment using quality, adoption, net value and remaining uncertainty together.

10 Apply and troubleshoot

Turn learning into a usable management approach. The diagnostic, review questions and formula sheet provide a quick route back to the techniques. Use them for the whole project, including its AI component. The purpose is to identify the next useful action or decision, not to generate an AI score.

10.1 A 30-day learning and application plan

Choose one real project you understand. Use it for the whole month so each concept adds to one coherent management approach. Allow approximately 30 minutes of focused learning on most working days and a longer weekly review. Adjust the pace to the project's timing; do not wait for day 30 to escalate an issue that needs attention now.

Table 10.1 | Learning and application plan

Period	Learn and apply	Evidence to produce	Question to ask yourself
Days 1–3	Distinguish project, outcome and benefit. Describe the problem, current method and alternatives, including a response without AI.	A one-page purpose statement with a baseline measure and assumptions	Have I described a worthwhile change, or only a proposed capability?
Days 4–7	Establish sponsorship, authority, use boundaries and a proportionate delivery approach.	A draft charter, responsibility map and next decision point	Who authorises the project, accepts its outputs, permits use and owns future value?
Days 8–10	Break the result into deliverables. Define testable requirements, acceptance evidence and exclusions.	A small scope structure or backlog with configuration and evaluation references	Could another person recognise completion and the limits of the evidence?
Days 11–14	Sequence work, test capacity and estimate costs, including data preparation, evaluation, human review and operation.	A milestone plan, estimate ranges, reviewer commitments and dependencies	Is the plan feasible with the people and supplier inputs actually available?
Days 15–17	Identify risks, issues, assumptions and dependencies. Connect specialist findings to delivery actions.	A short live register with owners, actions and triggers	Which uncertainty or unresolved finding could change the next decision?
Days 18–21	Establish progress evidence, reporting and change control, including changes to suppliers, configuration or intended use.	A one-page status report and an impact assessment with re-evaluation effort	Does the report separate accepted work, learning, forecasts and benefits?
Days 22–24	Review stakeholders, communications, adoption and quality. Observe the intended workflow and human review.	Engagement actions, representative evaluation tasks and readiness evidence	Who must work differently, and do they have time, skills and authority?

Period	Learn and apply	Evidence to produce	Question to ask yourself
Days 25–27	Examine sourcing, operations, fallback, handover and benefits ownership.	A transition outline, continuing-cost owner and benefit review dates	Who will own the service and respond to changes after the project team leaves?
Days 28–30	Run a short governance review and simplify controls while preserving required decisions.	An agreed next-step decision and a personal improvement list	Which practices improved decisions, and which created little value?

At each weekly review, explain one concept aloud without using its acronym. Then show an example from the project. If you can define a baseline but cannot identify which version is authorised, revisit the practice rather than the definition.

Keep a learning log with three columns: concept, observed example and change to my practice. For example: “dependency; content owner approval controls pilot readiness; confirm the decision date and escalation contact.” This makes learning visible through changed behaviour. For an AI component, add a concrete consequence: “evaluation evidence applies to C01 version 1; assess the supplier release change before relying on that evidence again.” A definition is useful when it changes what the team does.

By day 30, aim to explain the project in five minutes: why it exists, what result is authorised, how work will proceed, what could change the plan and which decision comes next. The goal is a usable management system that colleagues can understand. For a small purchased feature, the same learning can produce a few fields in the existing project plan. The number of new documents is not a measure of progress.

10.2 Diagnosing a distressed project

Begin with evidence, current obligations and decision authority. A red status is a signal to investigate. It is not proof that a project should stop, nor a reason to conceal uncertainty behind a new plan. The following table offers starting hypotheses, not automatic diagnoses.

Table 10.2 | Troubleshooting management problems

Observed signal	Investigate first	Immediate management response	Governance choice to prepare
Dates move repeatedly	Dependency logic, data/content readiness, reviewer availability, supplier inputs and estimate assumptions	Reforecast remaining work using demonstrated capacity; make evaluation and rework visible	Change scope, timing or capacity; assess whether value still justifies continuation
Spending rises while accepted work does not	Rework, repeated experiments, usage charges, unclear acceptance and unrecorded scope	Reconcile cost with completed deliverables, accepted learning and remaining obligations	Approve a bounded recovery plan or reconsider the investment
Most items are almost complete	Missing acceptance evidence, blocked reviews, unresolved limitations and work in progress	Inspect a sample and finish a small number of items fully; separate a demonstration from readiness	Resolve acceptance authority or change the delivery sequence

Observed signal	Investigate first	Immediate management response	Governance choice to prepare
Reports stay green despite unresolved issues	Status definitions, stale evidence, changed configurations and escalation behaviour	Report facts, forecast impact and required decisions separately; identify evidence that no longer covers the intended use	Confirm tolerances and require an explicit recovery decision
Sponsor decisions are absent	Authority, competing priorities and decision quality	Present concise options with a decision date and consequence of delay	Reconfirm sponsorship, delegate authority or pause affected work
Users reject a technically finished result	Problem assumptions, workflow fit, accessibility, trust, checking effort and support	Observe real tasks and test the reasons for rejection; measure total effort and quality	Adapt the solution, fund adoption work or reassess expected benefits
Supplier delivery is failing	Contract commitments, buyer dependencies, release changes, evaluation access and evidence of breach	Establish facts with commercial support; preserve evidence and assess integration and exit implications	Choose remediation, replacement or termination through authorised processes
Risks exceed agreed tolerance	Actual exposure, control effectiveness and consequences	Contain immediate exposure within authority and escalate	Accept with conditions, reduce exposure, pause or stop
The original need has changed	Strategy, demand, alternatives and remaining costs	Update the business case using current evidence	Continue, redirect or close, even when delivery performance is good

A recovery proposal should state the cause being addressed, the proposed intervention, the accountable owner, the resources needed and a date to evaluate whether it worked. Retain the previous baseline and decision history. If a new baseline is approved, show what changed and why; otherwise the recovery can erase the evidence needed to learn.

Apply the diagnosis to an AI delivery problem

Suppose a convincing demonstration is followed by poor results on staff tasks. First check whether the tasks, users, source collection and configuration match the demonstration. Then examine requirements, source quality, variation, usability and review effort. The response may be a corrected collection, narrower use, changed configuration or a different solution. “More model training” is not an automatic recovery action, particularly for a purchased capability.

If evaluation is repeatedly postponed because the same specialist must review every workstream, the immediate management problem is capacity and sequencing. Book the reviewer, reduce concurrent work, agree evidence in smaller batches or change the date. Calling the delay an AI assurance issue does not resolve it.

If a supplier change invalidates earlier evidence, establish the affected configuration and requirements, estimate proportionate re-evaluation and update the forecast. Use the existing change authority to choose whether to proceed, retain a supported earlier version, narrow the use or stop. Do not represent the previous acceptance decision as covering a different capability. These are management responses within the original workflow, not a separate risk assessment.

10.3 A one-page project review

Use these questions for a short sponsor conversation, a project reset or your own weekly review. They are an original prompt set, not an audit checklist. Concentrate on exceptions and choices; link to existing evidence rather than recreating it.

Table 10.3 | Eight questions for a project review

Question	Evidence to inspect	Decision if the answer is weak
Why does the project still matter?	Current need, alternatives and benefit hypothesis, including a solution without AI.	Reconfirm, redirect or stop.
What is authorised?	Scope, funding, decision rights, intended users and permitted actions.	Clarify authority before further commitment or use.
What is actually complete?	Accepted outputs and learning evidence, with the configuration and conditions evaluated.	Resolve acceptance, stale evidence or unreliable reporting.
Is the remaining plan credible?	Dependencies, data/content readiness, reviewer capacity, estimates and supplier commitments.	Replan, change scope or secure capacity.
What will it cost and when will it finish?	Baseline, actuals, ETC, reserves and usage/review assumptions; continuing costs separately.	Address cause and authorise any change.
What could materially change success?	Risks, issues, assumptions, dependencies and changes affecting prior evidence.	Assign responses, re-evaluation and escalation.
Will people use and support the result?	Adoption, checking effort, operating ownership, support, monitoring and fallback readiness.	Improve design, transition or the value case.
What decision is needed next?	Options, recommendation, evidence and needed-by date.	Decide or name the missing evidence and owner.

Finish with action

Record the decision, owner, conditions and next review point. If no decision is needed, state what evidence supports continuing within the existing authority. If the answer remains unknown, assign a bounded investigation and a date to return. An explicit unknown is more useful than an unsupported green indicator.

> **Your learning test:** Explain the project to a colleague in five minutes without relying on acronyms. They should understand its purpose, authorised result, delivery approach, largest uncertainty and next decision.

10.4 Formula desk reference

Use consistent units, an agreed accounting basis and reliable inputs. Working-day estimates need a working calendar; accepted scope does not establish realised benefits. For an AI experiment, agree what completed learning work means before assigning earned value. A higher evaluation score does not itself earn the budgeted value of an unaccepted service. The scheduling and EVM chapters explain assumptions and worked calculations; public GAO and PMI guidance provides further reading. [13, 15, 16]

Table 10.4 | Formula quick reference

Measure	Formula	Interpretation or limit
Cost variance	CV = EV − AC	Positive means earned value exceeds actual cost for the reported work; it is not a measure of business benefit.
Schedule variance	SV = EV − PV	Budget-valued progress difference, not elapsed time. It cannot determine the pilot finish date.
Cost performance index	CPI = EV / AC	Budgeted work achieved per cost unit; AC must be non-zero.
Schedule performance index	SPI = EV / PV	Budgeted progress relative to plan; PV must be non-zero.
Estimate at completion	EAC = AC + ETC	Actual cost plus estimated remaining cost.
EAC: original remaining rate	AC + (BAC − EV)	Assumes remaining work performs at its original budget rate.
EAC: current efficiency persists	BAC / CPI	Uses cumulative cost efficiency; unstable early data can mislead.
Variance at completion	VAC = BAC − EAC	Negative indicates forecast overrun against the baseline.
Required remaining efficiency	TCPI = (BAC − EV) / (BAC − AC)	Efficiency needed to meet BAC; a zero or negative denominator requires interpretation, not mechanical use.
Total float, simplified CPM	LS − ES = LF − EF	Assumes the stated dependency model, calendars and constraints.
Triangular estimate	(O + M + P) / 3	Mean under a triangular model; not a confidence limit.
Conventional PERT mean	(O + 4M + P) / 6	Weighted estimate; input quality and assumptions matter.
Expected monetary value	Sum of probability × consequence	Use a clear sign convention and avoid double counting scenarios.
Net capacity released	Gross time released − additional operating effort	Include checking, correction, support and maintenance on a consistent basis. Monetised capacity is not automatically cash saved.

O, M and P mean optimistic, most likely and pessimistic. ES/EF and LS/LF mean early and late start/finish. BAC is budget at completion; ETC is estimate to complete; PV/EV/AC are planned value, earned value and actual cost.

Apply the formula without losing the management question

The baseline pilot reports AC AUD 48,000 and ETC AUD 72,000 at week 6, so EAC is AUD 120,000. Against BAC AUD 108,000, VAC is negative AUD 12,000. This reaches the AUD 120,000 funding envelope; it does not establish authority to draw reserve. The distinct supplier-change what-if adds AUD 3,200 to ETC, giving EAC AUD 123,200. Keep that scenario separate from the baseline example. If the change becomes the expected outcome, update the live forecast promptly while preserving the approved baseline.

For capacity, define the comparison boundary. If task-time observations already include user checking and correction, do not subtract that effort a second time. Add support, maintenance or review work outside that task boundary on the same period and unit basis. Evaluate acceptable quality alongside the time result. A calculation can be arithmetically correct while answering the wrong business question.

11 Plain-language glossary

The definitions explain how terms are used in this guide. They are original learning aids, not quotations from PMI terminology or a replacement for an organisation's defined terms. The AI entries focus on management decisions; NIST provides further technical and evaluation context. [19, 20]

11.1 Terms used in this guide

Table 11.1 | Plain-language glossary

Term	Plain-language meaning and useful distinction
Acceptance criteria	Conditions used to decide whether a particular deliverable is acceptable. Specify evidence, relevant configuration, operating conditions and decision authority, not only an aspiration such as “good quality”.
Actual cost	Cost recorded for work performed during the stated period or to date. Be clear whether reporting includes accruals and how commitments are treated.
Adaptive approach	A delivery approach that uses frequent learning and feedback to adjust detailed work. It still needs purpose, accountability and limits.
Agentic AI	An AI capability that can select and carry out steps or use tools within a delegated task. Specify its permitted actions, authority limits, confirmations, failure handling and stop arrangements. The label does not establish how much autonomy is appropriate.
AI capability	A service or feature that uses AI to support a task, provide an output or take a permitted action. The project may build, buy, configure, integrate or adopt that capability; it does not necessarily train a model.
AI model	A component that uses patterns learned from data to produce outputs such as predictions or generated content. Project acceptance concerns the complete capability in its intended workflow, including data, interfaces, permissions and people, rather than the model alone.
Assumption	Something treated as true for planning before it is sufficiently confirmed. Give material assumptions an owner and a validation point.
Assurance	Examination intended to provide justified confidence that a project, control or result meets the required expectations. Independence should match the review's purpose. Specialist assurance informs project decisions without automatically granting business acceptance or permission to operate.
Backlog	An ordered collection of possible work. Inclusion does not necessarily mean authorisation, funding or commitment to deliver every item.
Baseline	An authorised reference against which change or performance is assessed. A current forecast can differ from the baseline without automatically replacing it. A configuration record identifies what was tested or deployed; it is not a substitute for the scope, schedule or cost baseline.
Benefit	A measurable improvement valued by stakeholders. A delivered AI feature or improved evaluation score may enable a benefit; it does not establish that the improvement occurred in operational work.
Benefits owner	The person accountable for enabling, measuring and sustaining a specified benefit, including after project closure.

Term	Plain-language meaning and useful distinction
Business case	The rationale for committing resources, including options, expected value, costs, risks and assumptions. It should remain relevant as circumstances change.
Change control	The agreed way to assess and authorise changes to controlled commitments. For an AI capability, changes to configuration, information sources, permissions or intended use may affect previous acceptance evidence. Change control differs from helping people adopt a new way of working.
Charter	The authorising statement of the project's purpose, boundaries, sponsorship and management authority. It is usually shorter than the delivery plan.
Configuration record	A versioned record of the components and settings needed to identify the capability being evaluated or used. It can include the supplier or model release, prompts, retrieval settings, source collection, permissions, tools and interfaces. It supports traceability and change decisions.
Constraint	A limitation within which the project must work, such as a mandatory date, available funding or a required standard.
Contingency reserve	Time or funding allowed for identified uncertainty under agreed control. It should relate to exposure rather than being an unexplained extra percentage.
Critical path	The dependency path that determines the earliest finish in the schedule model. Several paths may be critical, and resource constraints may alter the feasible schedule.
Deliverable	A verifiable result produced by work, such as an approved design, configured service or training package. A documented experiment finding can be a deliverable without proving that the proposed service is suitable. Deliverables can be intermediate or final.
Dependency	A relationship in which work or a decision depends on something else. External dependencies need named contacts and agreed dates.
Earned value	The budgeted value assigned to work completed against agreed rules. It is a project performance measure, not revenue, money received, a model score or realised business benefit.
Embedded AI	AI functionality supplied within a wider product or enterprise platform. Its presence may change requirements, evaluation, cost or operating responsibilities even when the project does not select or develop the underlying model.
Estimate	A reasoned assessment of an uncertain quantity. State its basis and range; a precise number can still be highly uncertain.
Estimate at completion	The current estimate of total final project cost, combining cost already incurred with expected remaining cost. Explain the forecasting assumptions.
Evaluation	Planned examination of a capability against stated tasks, criteria and operating conditions. It produces evidence and limitations for a decision. Running tests is not the same as meeting acceptance conditions or receiving permission to use the result.
Fallback	An agreed alternative way to continue or safely stop the work when the AI capability is unavailable or unsuitable. It needs an owner, usable procedure and sufficient capacity; naming a manual process does not prove it will work.

Term	Plain-language meaning and useful distinction
Float	Scheduling flexibility before a stated consequence occurs. Total float is delay available before affecting project finish or an imposed schedule constraint; free float is delay available before affecting a successor's earliest start.
Forecast	The current prediction of an outcome, cost or date using available evidence. Forecasts support decisions even when the approved target stays unchanged.
Gate	A defined decision point that authorises, conditions, redirects or stops further work. A meeting without a decision is not necessarily a gate.
Generative AI	AI that produces content such as text, images or code in response to inputs. Project planning should define how generated content will be checked and used, and what variation or unsupported output means for acceptance.
Governance	The arrangements for direction, authority, accountability and oversight. Governance determines who may decide; management prepares and implements decisions.
Human review	Work performed by a person to examine an AI output or proposed action and decide what should happen next. It requires relevant competence, time, information and authority. Include the effort in the operating plan and value calculation.
Hybrid approach	A deliberate combination of delivery approaches suited to different work and constraints. Explain how their plans, decisions and evidence connect.
Issue	A current problem or condition requiring action. A risk concerns uncertainty; an issue has already arisen or is already present.
Management reserve	Budget held outside the cost baseline for unforeseen work within scope, controlled under agreed authority. Its use requires appropriate authorisation.
Milestone	A significant event or decision in the schedule, usually represented with zero duration. The work required to reach it has duration.
Organisational change management	Work that helps people adopt and sustain changed processes, roles or behaviours. It differs from controlling changes to the project baseline.
Outcome	The change produced when outputs are used, such as a changed service or behaviour. A benefit is the valued improvement arising from that change.
Portability	The practical ability to move relevant data, configuration, records and service work to another arrangement. Contractual export rights alone do not establish that a replacement service can operate; examine formats, dependencies, effort and exit support.
Portfolio	A grouping of investments managed to support strategic objectives. Components need not share the same deliverable or be operationally dependent.
Predictive AI	AI used to estimate an outcome, classify an item or detect a pattern from data. Define the decision it supports, relevant comparison method and consequences of error. It is distinct from a predictive project delivery approach.

Term	Plain-language meaning and useful distinction
Predictive approach	A delivery approach that establishes substantial detail in advance when the work is sufficiently understood. Plans still require monitoring and controlled adjustment.
Programme	Related projects and other work coordinated to obtain benefits that separate management would make harder to achieve.
Project	A temporary initiative organised to create value in a particular context. Its outputs can continue operating long after the temporary work ends.
Prompt	Instructions or other input supplied to a generative AI capability to guide a response. A prompt can affect behaviour and may therefore need version control and proportionate re-evaluation when it changes.
Quality	The degree to which a result meets relevant requirements and is fit for its intended purpose. Additional features do not automatically improve quality.
RACI	A responsibility map showing who performs work, answers for it, is consulted and is informed. It does not replace formal delegations.
RAID	A convenient grouping of risks, assumptions, issues and dependencies. Keep their meanings distinct even if they share one register.
Re-evaluation	Repeating or extending relevant evaluation after a change or new finding. Identify which earlier evidence remains applicable and which requirements need fresh evidence. The appropriate extent depends on the effect of the change.
Requirement	A need or condition that a result must satisfy. Good requirements can be tested or otherwise verified. For AI work, state intended tasks, operating conditions and significant failure conditions so the evidence supports a clear decision.
Residual risk	The uncertainty remaining after the chosen responses and controls. A control being installed does not prove that the residual exposure is acceptable.
Retrieval	Finding information from specified sources for use in a response or task. In a knowledge assistant, retrieval settings, source quality and permissions affect what information is available. Finding a passage does not establish that the resulting answer represents it correctly.
Risk	Uncertainty that could affect objectives favourably or unfavourably. Describe its cause, possible event and consequence sufficiently to choose a response.
Risk appetite	The kind and amount of uncertainty an organisation is prepared to pursue or retain in seeking its objectives. A project cannot assume that its sponsor can waive every organisational limit.
Rollback	Returning a service to an earlier supported configuration after a change. Confirm that the earlier version, data and dependencies remain usable. Rollback may not reverse actions already taken; recovery may also be required.
Scope	The boundaries of the result and the work needed to produce it. Distinguish product scope from project work; include meaningful exclusions, intended users, information sources and any permitted actions.

Term	Plain-language meaning and useful distinction
Sponsor	The person providing organisational backing and accountability for the initiative, with authority defined by local arrangements. The role requires decisions, not only advocacy.
Stakeholder	A person or group that can affect the work, is affected by it or has a relevant interest. Influence and impact are different considerations.
Tailoring	Deliberate adaptation of management practices to the work and context. Record significant choices and retain mandatory organisational controls.
Tolerance	An agreed boundary within which a role may act without further escalation. Specify the measure, authority and response to a predicted breach.
Training and retraining	Work that adjusts a model using data so its learned behaviour changes. Include it only where the project requires it. Configuring a purchased feature, changing a prompt or updating a source collection is not automatically model training.
Usage unit	The quantity on which a service charge is based, such as a request, transaction or amount of processed content. Relate chargeable units, retries and multiple calls to completed business tasks before forecasting operating costs.
Validation	Checking whether the result is suitable for its intended use and users in the relevant setting. Controlled validation informs acceptance but does not itself authorise live use.
Verification	Checking whether a result meets its specified requirements using appropriate evidence. Record the configuration and conditions so the conclusion can be interpreted correctly.
Work breakdown structure	A structured decomposition of project scope into manageable deliverables and work packages. It describes the work; the schedule adds sequence and timing.

References and further reading

Framework facts are attributed to the relevant publications. The examples, practical explanations and AI project adaptations are original companion guidance. Numbered citations link to the entries below. Sources were checked for this edition on 21 November 2026; publication dates identify historical material.

PMI sources establish terminology and framework claims. Technical and specialist sources support the stated management implications; they do not replace organisational advice, decision authority or applicable requirements. Consult the full publication when applying its formal provisions.

[1] Project Management Institute. PMBOK Guide: current edition overview.

Eighth Edition; November 2025.

Open the source publication

Edition, publication date, principle/domain counts and tailoring orientation.

[2] Project Management Institute. PMBOK Guide, Eighth Edition: public table of contents.

2025.

Open the source publication

Official public preview confirming principle, domain and focus-area organisation.

[3] Kelly Heuer, PMI Vice President of Learning. Evolution of the PMBOK Guide.

January 2026.

Open the source publication

PMI-authored presentation hosted by NTT Data; confirms 40 non-prescriptive, tailorable processes.

[4] Project Management Institute. MENA Leadership Institute Meeting: PMBOK edition-history infographic.

October 2024.

Open the source publication

Historical Sixth/Seventh Edition years and structural counts, in the slide titled “Knowledge is Power”. The infographic's broad delivery-approach labels are not adopted as an edition-to-method rule.

[5] Project Management Institute. PMBOK Guide, Seventh Edition: frequently asked questions.

2021.

Open the source publication

Historical 2021 FAQ explaining the Standard/Guide distinction and continuing relevance of process knowledge. Purchasing, subscription and examination details are not current advice.

[6] Project Management Institute. Process Groups: A Practice Guide, publisher catalogue.

2022.

Open the source

Predictive focus, five process groups and 49 processes; distinct from a PMBOK edition.

[7] Project Management Institute. What is a Project, Examples and the Project Lifecycle.

Open the source publication

Public explanation distinguishing lifecycle phases and management process groups.

[8] Project Management Institute. Governance of Portfolios, Programs, and Projects: A Practice Guide.

January 2016.

Open the source publication

Further reading on portfolio, programme and project governance. This guide role and decision examples are original.

[9] HM Treasury. Guidance on developing business cases for projects and programmes.

Updated June 2026.

Open the source

Strategic, economic, commercial, financial and management perspectives; UK requirements are not presented as local policy.

[10] UK Government Project Delivery. Government Functional Standard GovS 002: Project delivery.

Version 2.1; 2025.

Open the source

Governance, roles, assurance and transition context. Used for transferable practice, not jurisdictional applicability.

[11] NASA. Systems Engineering Handbook: Appendix.

Open the source

Appendix A explains verification against requirements and validation against intended use; Appendices D/E illustrate traceability and evidence. NASA-specific institutional obligations are not imported.

[12] NASA. NASA Work Breakdown Structure Handbook.

2021 catalogue record.

Open the source

Product-oriented scope hierarchy and dictionary context. The pilot WBS is original.

[13] US Government Accountability Office. Schedule Assessment Guide: Best Practices for Project Schedules.

GAO-16-89G; 2015.

Open the source

Scheduling logic, model integrity and cost implications of delay. The A-H network and calculations are original.

[14] PMI public learning library. Taking step four with earned value: establish the project baseline.

1995.

Open the source publication

Historical further reading on the performance baseline and reserve separation. Its loose contingency/reserve terminology and assumed role authority are not adopted; the guide retains explicit local control of management reserve outside BAC.

[15] US Government Accountability Office. Cost Estimating and Assessment Guide.

GAO-20-195G; 2020.

Open the source

Reliable estimates and EVM context. Worked budget, arithmetic and interpretations are original.

[16] PMI public learning library. How to make earned value work on your project.

2012 conference paper.

Open the source publication

Historical conference paper supporting CV/SV/CPI/SPI and alternative EAC assumptions. The alternatives are not interchangeable. Reserve terminology and early-data stability claims are not adopted as universal rules.

[17] Ken Schwaber and Jeff Sutherland. The Scrum Guide.

November 2020.

Open the source

Named accountabilities, goals and Definition of Done. This guide rhythm is an original example, not a full Scrum description.

[18] Kanban Guides. The Kanban Guide.

May 2025.

Open the source

Workflow, work-in-progress control and flow-measure context. Board policies and limits are illustrative.

[19] National Institute of Standards and Technology. AI Risk Management Framework.

AI RMF 1.0; 2023.

Open the source

Voluntary AI RMF 1.0, released January 2023. At the 21 November 2026 source check, the public page states that it is being revised. Specialist context only; this guide does not reproduce an AI assessment method.

[20] Autio and colleagues; NIST. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile.

NIST AI 600-1; July 2024.

Open the source

Voluntary generative-AI profile. MEASURE 2.3 and 2.5, printed pp. 30–31, support evaluation in deployment-relevant conditions and limits on generalisation. MAP 4.1, printed p. 26, addresses re-evaluation of modified third-party models or new domains. It does not prescribe this guide's fictional case counts, thresholds or project approvals.

[21] Australian, state and territory governments; Department of Finance host. National framework for the assurance of artificial intelligence in government.

June 2024.

Open the source

Australian government assurance context, including lifecycle and organisational considerations. Applicability is determined by the relevant government and organisation; no all-sector mandate is inferred.

[22] Australian Government Department of Finance. Commonwealth Procurement Rules: Value for money.

Open the source

Whole-life value and non-price considerations. Commonwealth requirements apply only within their defined scope; other organisations use their own applicable policies and delegations.

[23] Australian Government Department of Finance. Efficient, effective, economical and ethical procurement.

Open the source

Fair dealing, transparency and ethical procurement context; local requirements still govern.

[24] UK Government Commercial Function. The Sourcing Playbook.

Updated November 2026.

Open the source

Relationship between scope certainty, commercial risk and pricing. Used as informative practice guidance.

[25] Australian Government Digital Transformation Agency. Benefits Management Policy: Guidance.

Open the source

Benefits measurement, ownership and validation; Commonwealth policy scope is not automatically applied elsewhere.

[26] Project Management Institute. Benefits Realization Management: A Practice Guide.

December 2018.

Open the source publication

Further reading on benefits and organisational value.

[27] Project Management Institute. PMBOK Guide, Eighth Edition: second-printing errata.

Second-printing errata; linked by PMI as updated April 2026, checked 21 November 2026.

Open the source publication

Further reading for owners of the official guide; includes corrected dependency and metric descriptions.

[28] Strategen AI. Strategen AI: current public positioning and APIG.

Accessed 21 November 2026.

Open the source

Current public brand positioning and APIG: Actors, Practices, Infrastructure and Governance. Used for an optional connection to project management, not as a replacement for the PMBOK structure.

Published by Strategen AI. Version 1.0, 21 November 2026. Strategen-AI.com.