AI4BCM

RELEASE 2026.11.1

AI4BCM: Guideline for using AI by BCM professionals

Framework-neutral guidance for BCM practitioners

Nobody taught BCM professionals how to use AI. AI4BCM Guidance does.

AI FOR BCM, NOT BCM FOR AI.

Will AI replace BCM managers?

When tasks become cheaper, people often do more of them, for as long as demand keeps growing (Bessen, 2018). The increased efficiency of steam engines led to greater coal consumption (Jevons, 1865). Similarly, when writing reports becomes cheaper, more BIAs, refreshers and exercises are requested. The role of the BCM will change and grow for as long as there is a significant amount of BCM work that AI cannot do.

EDITORS

Willem A. Hoekstra · Konstantin Gerner

CONTRIBUTORS

Matthias Rosenberg · Charlie Maclean-Bristol · Simon Contini · Totti Karpela · Sachin Kumar · Jason Hoss

Hoekstra, W., Gerner, K., et al. (2026). *AI4BCM: Guideline for using AI by BCM professionals*. CC BY 4.0.

Contents

AI4BCM Guidance

01 Where do I stand?

Ways of working and maturity levels 5

02 The rules that always apply

Data rules 7 · Five principles 8

03 AI across the BCM lifecycle

The lifecycle, applied 10

1 Govern 11 · 2 Embed 13 · 3 Analyse 15

4 Design 18 · 5 Implement 20 · 6 Validate 22

04 Choosing a tool

Which tool for which task 24 · Before you approve a tool 25

05 Prompts that work

Which task first? 26 · The prompt pattern 27 · Build your own skill 28

06 Looking after your prompts and tools

The record card 29 · Keep it up to date 30

References 31 Glossary 34

Short version: the *Quick Start* page, ai4bcm.org/quick-start

PURPOSE

Ask AI to prepare, compare, summarise, retrieve, and challenge information.

Who this guidance is for

This guidance is for BCM professionals, for managers and governance owners who oversee BCM work, and for the people in resilience, risk, crisis management, IT, security, facilities and operations who work beside BCM teams. It follows common BCM lifecycle activities, from governance and embedding to analysis, design, implementation and validation. These are generic on purpose, so the guidance works with any BCM framework, a formal management system or an internal method.

- **BCM professionals**
- **Managers and governance owners**
- **Resilience, risk, crisis management, IT, security, facilities**

This guidance is about how AI can help BCM professionals draft, analyse, review, communicate, maintain and validate their work. It does not cover the resilience of AI systems, or of services that depend on AI. That is a different topic. The full guidance covers where you stand today and the rules that should be applied everywhere when using AI. It explains how to select the right tool for your BCM processes, offering prompts for each task and explaining how to maintain the workflow.

JEVONS PARADOX



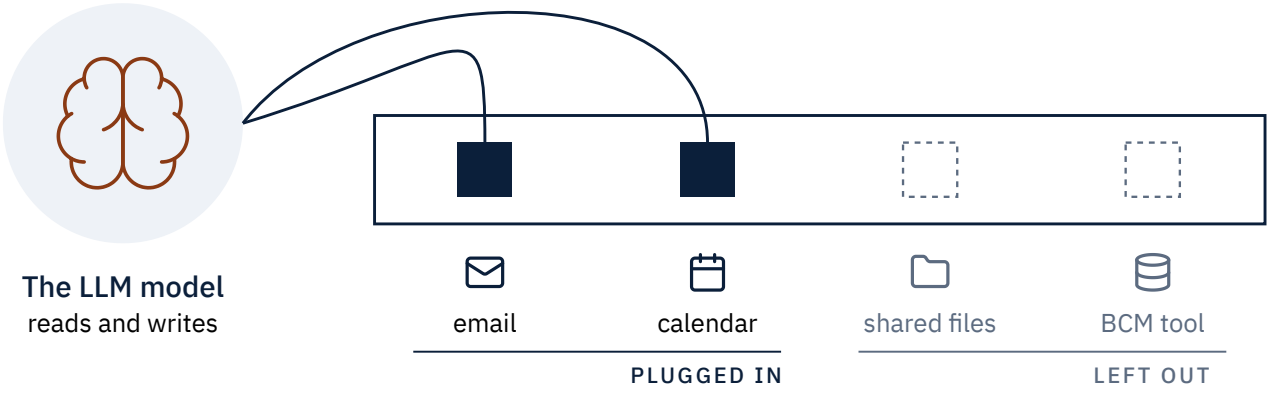
Introduction

AI assists BCM managers with written tasks at each of the six stages. For instance, during the 'Govern and Embed' stages, it drafts policy wording from your approved framework for review and adapts awareness messages for each audience. During the Analyse and Design stages, it prepares business impact analysis (BIA) interview guides and questions the assumptions behind the strategies. During the implementation and validation stages, it drafts plans, writes exercise scenarios and categorises debriefing notes into actions.

The brain in a jar

Imagine the language model (LLM) is a brain in a jar. On its own, it can only read what you paste into the chat box and generate a response. Its usefulness depends entirely on the questions you ask and the context you provide. What you connect it to determines its capabilities. For example, if you connect it to your email or calendar, it can scan your inbox, draft replies independently or update a record.

EXHIBIT 1



Chat with AI4BCM

You do not have to read this guidance from start to finish before you use it.

The AI4BCM chatbot at ai4bcm.org/chat has full access to this guidance. Describe your task as you would to a colleague. For example, you could say, “I have to run BIA interviews next month” or “My plan needs updating after a reorganisation”. The chatbot will then explain how to best use AI for this activity.

EXHIBIT 2

Ask it ai4bcm.org/chat

01

Where do I stand?

First, work out which of the four things you are currently using.

EXHIBIT 3

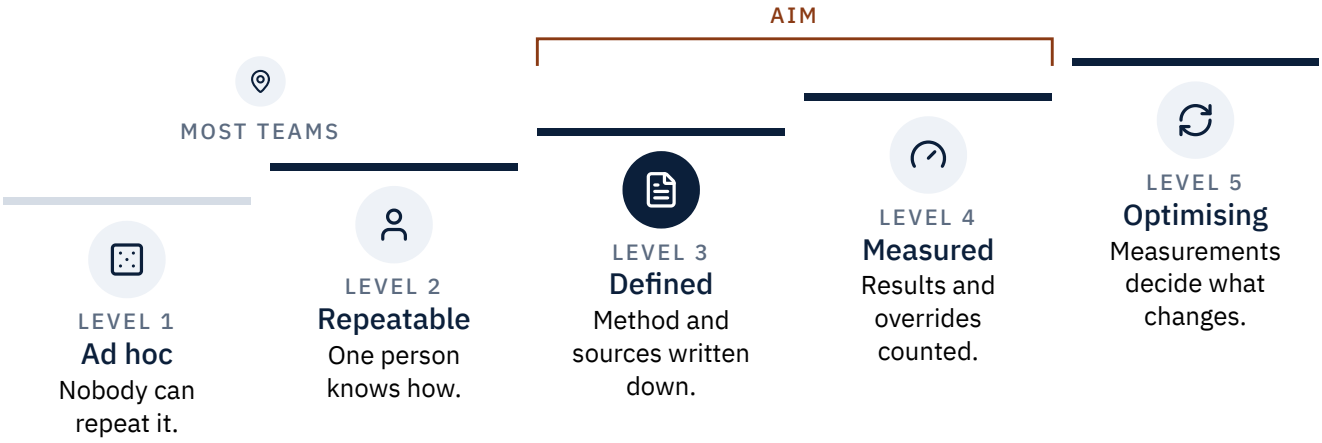
	Chatbot	Skill	Workflow	Agentic AI
What it is	a conversation	your best prompt, with sources	fixed steps, checked by rules	a model working towards a goal
How you use it	ask, and ask again	save it once, reuse it	run it step by step	set a goal and permissions
Who decides what comes next	you, in each prompt	you; the saved prompt guides	the set process	the model, within permissions
Results	vary each time	**similar, still checked**	**comparable, if inputs are fixed**	vary with the path

For most BCM work, a skill or a workflow is the useful choice.

Every important action should be approved by a knowledgeable individual. Results can only be compared if the inputs are controlled, the method is documented and the output is verified. A workflow is the fastest way to achieve this.

The five AI maturity levels

EXHIBIT 4



LEVEL	TYPICAL CHARACTERISTICS	COMMON MISTAKE AT THIS LEVEL
LEVEL 1 Ad hoc	Approved tools used by individuals, each in their own way. Nothing is written down, so no result can be repeated or traced.	Accidental disclosure of sensitive information, and nobody can say which output came from AI.
LEVEL 2 Repeatable	Saved prompts and defined tools let a team repeat what worked. It still depends on the people who know how.	Poor prompt patterns get repeated, and the method lives in one person's head.
LEVEL 3 Defined	Use cases, prompts, sources, review points and limits are written down, given a version number and are owned.	Too much trust in AI (retrieval quality) or in data that comes through connectors.
LEVEL 4 Measured	Quality, overrides, failures and source dates of each workflow are documented. Each workflow has a named owner, approval gates and a tested fallback.	Hidden dependency on workflows and approval bottlenecks, as well as a low override rate, are all signs of poor quality.
LEVEL 5 Optimising	Data measurements determine which workflows change, grow or stop. People set the limits of AI usage and review them afterwards.	Complexity exceeds the capability of governance, and there is unclear accountability across connected systems.

02

The rules that always apply

Data rules

If you paste something into the wrong tool, it will stay there forever.

Data classification decides the tool

The more sensitive the data and the more serious the consequences of an error, the more controlled the environment must be. Public, free online tools are never used for BCM data. They are only useful for generic brainstorming that discloses nothing about your organisation.

Classification levels

Below is a list of data assets of the highest classification level.
Paste it in an approved corporate or private environment only.

- BIA data and risk assessment detail
- incident records and live crisis information
- personal data, contact lists and personnel records
- vulnerabilities and control weaknesses
- recovery strategies and design assumptions
- supplier weaknesses and contractual dependencies
- security-related architecture or site detail
- board papers and executive-session material
- audit findings not yet cleared by the audit owner

Classify before entering; how sensitive the data is decides which tools you may use.

Passwords, keys and access tokens go into no tool, approved or not.

Prompt injection

The text in a document, webpage, email, picture or file can be written in such a way as to make an AI perform an action that you did not ask for, such as ignoring its instructions, revealing data or taking action. AI tools cannot reliably distinguish such text from its context, and providing an additional prompt will not help (OWASP, 2025, ASI01 and ASI06; IMDA, 2026, sections 2.3.1 and 2.3.2). Give a tool that reads data outside of your controlled environment only the permissions it needs for the task (principle 3). If you notice that a document is instructing the AI in unintended ways, inform the owner and your information security team immediately.

Five principles

At every AI maturity level and in every stage of your BCM lifecycle, these five principles apply:

PRINCIPLE 1

Know AI and its limits

Anyone using AI for BCM work is familiar with its strengths and weaknesses. It can invent facts and sources, overlook information that is not included in its sources, and appear confident when it is incorrect. Your organisation should help those using AI to understand these limitations (European Union, 2024, Article 4, as amended in 2026).

PRINCIPLE 2

Match the tool to the classification

Classify the material before you paste it into a tool. The more sensitive it is, and the more serious an error would be, the more controlled the tool must be. Free public tools never get BCM data. The classification levels are listed under *Data rules*, p. 7.

PRINCIPLE 3

Checking AI output and approving what it does

AI-generated output should be reviewed by a competent person before it is included in a BCMS document, report or decision. This person must have the necessary resources and time to do so.

AI can only act independently within a predefined environment owned by a specific individual, with permissions set in the tool itself and every action logged. It must always be possible to stop AI.

Each organisation determines which actions require approval on a case-by-case basis. The more mature your company is, and the more guardrails you have in place throughout your BCM lifecycle, the more you can hand over to AI. It is worth tracking the number of approvals and the rate at which AI output is overridden, but these alone do not demonstrate effective review. A very low override rate can also mean that nobody checks any more (IMDA, 2026, section 2.2.2).

PRINCIPLE 4

AI answers only from approved sources, and cites them

Link your AI tool to approved internal repositories and official external sources only. Ensure that it only provides answers from these sources. Regularly check the cited passage against the claim it supports.

PRINCIPLE 5

Quality over impressing

A long, polished document is not necessarily a good one. AI can make things look finished quickly. Read everything before sending it out, and only keep what will be used and makes sense.

Minimum control checklist

Before any AI-assisted BCM task, check that:

- [ ] the tool is approved for the type of data involved
- [ ] the information is suitable for that environment
- [ ] the output is based on approved, current sources, citations checked
- [ ] reviewer has the necessary knowledge, access to the relevant sources, and time to check
- [ ] any actions or record changes have approval gates
- [ ] you can identify afterwards who did what, and from which sources
- [ ] fallback methods exist if the tool is unavailable

03

AI across the BCM lifecycle

The lifecycle, applied

Each stage chapter names one organisation that uses AI for that stage's work today, cited in References, p. 31.

EXHIBIT 5

THE SIX STAGES



Not sure which stage your task is in? Ask the AI4BCM Guidance chatbot. **ai4bcm.org/chat**

STAGE 1 · GOVERN

Establishing and governing BCM

What this stage covers

This stage covers the establishment and governance of BCM, including policy, scope, roles and responsibilities, governance arrangements, management reporting, and the management framework needed to operate and improve continuity capabilities over time.

Where does AI help, and where are its limits?

Much of this stage is drafting, reviewing and summarising, and AI does that well. It writes first drafts, finds unclear wording and makes reports easier to read. It also helps monitor changes in official guidance, regulation, or organisational context that could affect the BCMS.

This stage also decides who is accountable. AI helps prepare and challenge material. Policy, management intent, accountability and any statement that a BCMS is compliant stay with named people, through the organisation's established governance process.

Used today

CITED PUBLIC EXAMPLE, SOURCE IN REFERENCES, P. 31.

ClearBank, a UK clearing bank, uses generative AI to check its internal policies against each other so that key policies do not fall out of line; this is described by the bank in a UK government collection of case studies (2024), not reported as a result.

What AI can do, by activity

Give the AI the authoritative material for the task: internal governance frameworks, strategy statements, policy templates, organisation structure, prior management review outputs, approved role descriptions, or internally prepared requirement summaries. If external requirements are relevant, use official sources and handle any licensed material in accordance with permission and policy (*Data rules*, p. 7).

Context and requirements

- Summarising official regulatory or guidance developments for review L1
- A simple workflow monitors approved official sources and prepares a short note on their relevance for you. L4

Scope and policy

AI can help in two ways: it can write a first draft and review subsequent drafts for inconsistencies and mistakes. AI is usually more useful for the latter. Ask the AI to highlight unclear wording, unsubstantiated statements and governance gaps, such as missing owners or approval steps. Before the draft goes into formal approval, have the relevant departments check it, such as Legal, Compliance, Information Security, Records Management, and the document owners.

- **Grounded drafting.** Use an enterprise AI assistant alongside approved internal documents to create a first draft of a policy, scope statement or governance note. You can then customise the draft using your own terms and structure. L2
- **Red-team review.** Ask the AI to identify any clauses that are unclear, unrealistic, inconsistent or ambiguous. L2

Roles, responsibilities and competence

- AI can be used to help draft or refine role summaries, such as RACI-style mappings or governance descriptions. Top management can then agree on who is responsible for what and who has which authority. L2
- Drafting interview questions for hiring BCM staff L2
- **Onboarding support.** A dedicated internal assistant helps new practitioners to understand the structure, terminology and expected outputs, based on approved BCMS governance material. L3

Objectives and programme

- Preparing the plan to achieve the business continuity objectives: what needs to be done, what resources are required, who will do it and by when? L2

Reporting and management review

- **Structured management reporting.** Use AI to summarise BCMS metrics, open actions, validation results and review data for management. L2
- Reports that combine several BCMS data sources. They are written with the AI functions of your reporting tools. L3

L1 Ad hoc · L2 Repeatable · L3 Defined · L4 Measured · L5 Optimising

STAGE 2 · EMBED

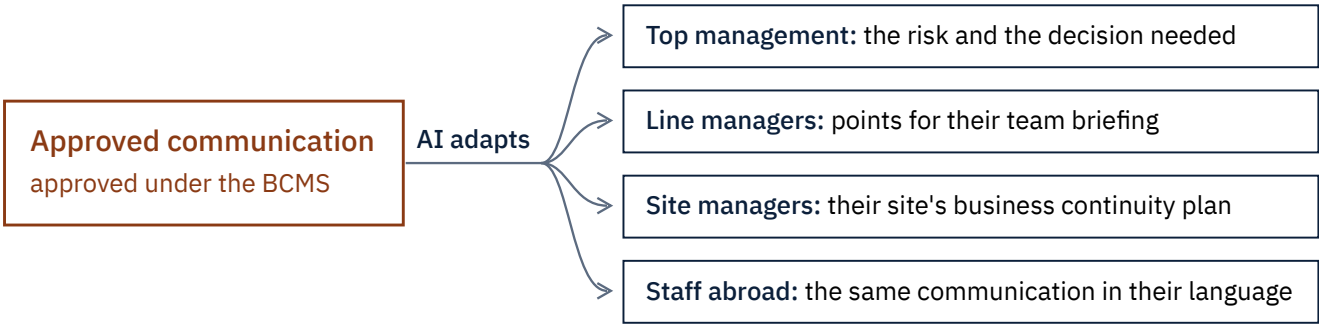
Embedding BCM in the organisation

What this stage covers

This stage addresses the embedding of business continuity in organisational culture and practice, including awareness, engagement, education, communication, and the development of behaviours that support resilience in day-to-day work.

EXHIBIT 6

AI adapts one approved communication for each audience



ILLUSTRATIVE. AWARENESS AND COMMUNICATION AS IN ISO 22301:2019, 7.3 AND 7.4

Where does AI help, and where are its limits?

Much of this stage is explaining, adapting messages for each audience and reading feedback, and AI does that well. It tailors messages by role, simplifies guidance, writes quiz questions, finds themes in staff feedback and answers routine staff questions.

Culture still changes through people. AI makes messages more relevant and quicker to produce. The success of this stage should be judged by whether staff understand their role and participate, rather than by how much content is produced.

Used today

CITED PUBLIC EXAMPLE, SOURCE IN REFERENCES, P. 31.

Fairfax County's 911 call centre in Virginia collaborated with George Mason University to build an AI voicebot that plays the caller in practice scenarios for new and experienced call-takers. The university reports this as a tested prototype for 2025, not reported as a result.

What AI can do, by activity

Firstly, gather together the approved business continuity materials that staff can access. These may include policy extracts, awareness notes, role descriptions, onboarding content, exercise expectations and approved FAQs (*Data rules*, p. 7).

Understanding the culture

Analyse surveys, workshop comments and free-text feedback in an approved, aggregated environment. The aim is to identify patterns, trends, areas of confusion and recurring concerns.

- **Anonymised feedback analysis.** Use a corporate AI or analytics tool to group survey comments into themes such as 'confusion', 'ownership gaps', 'confidence', and 'positive examples'. L2

Awareness and communication

Start from the behaviour you want: for example, more people taking part in BIA interviews, plan owners keeping their plans up to date, and staff familiarising themselves with their response roles. Adapt the wording for each audience, but keep the approved version.

- **Simplification and translation.** Use AI to convert formal business continuity language into shorter, clearer wording or into other languages used within the organisation, with review before publication. L1
- **Current case-study creation.** Search-enabled AI can help you to identify recent public incidents that are relevant to your sector. You must verify the facts before use. L1
- Drafting an awareness campaign: creating messages, choosing formats and deciding on the timing for a specific audience and desired outcome. L2
- **Audience-specific message drafting.** Ask the AI to explain a business continuity topic to executives, line managers, site managers and support staff in plain language and in a specified format. L2
- **Scheduled awareness workflows.** A workflow can be set up to draft recurring awareness messages or manager reminders at regular intervals. Don't forget to approve each one before it is sent out. L3

Training and competence

- Drafting FAQs, quiz questions, short learning texts, manager talking points, intranet pages and posters L2
- **Internal chatbot support.** A virtual assistant can answer routine questions such as where to find a plan, what a person's business continuity role involves, and why an exercise is important. However, it must be based only on approved content. L3

Measuring the culture

Use participation data, survey trends, manager feedback, and validation evidence to assess whether awareness and engagement are improving.

- The themes in staff feedback and questions for the internal assistant show which awareness content should be changed, extended or stopped. L4

L1 Ad hoc · L2 Repeatable · L3 Defined · L4 Measured · L5 Optimising

STAGE 3 · ANALYSE

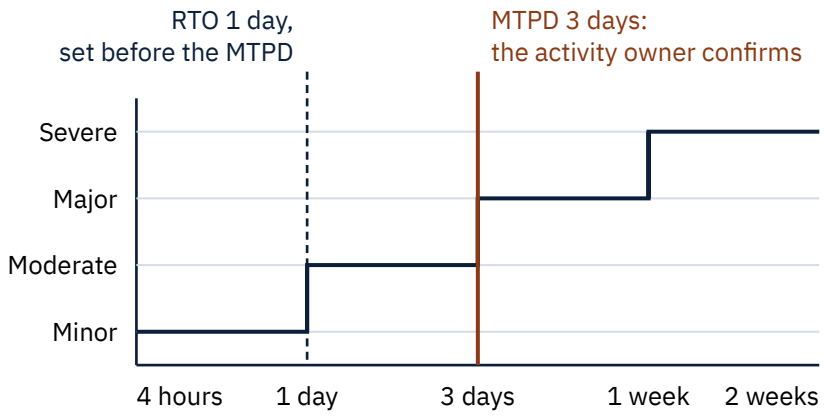
Analysis: business impact and risk

What this stage covers

This stage covers BCM analysis activities, especially business impact analysis (BIA), risk assessment, and the identification of dependencies, priorities, and recovery requirements. These activities draw on many people and systems, and their inputs have to be compared and checked.

EXHIBIT 7

AI shows where impact becomes unacceptable; the activity owner decides where the line sits



Impact of stopping one activity, on your impact criteria

EXAMPLE. AI DRAWS THE CURVE FROM THE INTERVIEWS AND MARKS WHERE IMPACT TURNS UNACCEPTABLE.

Where does AI help, and where are its limits?

AI can structure interviews, summarise workshops, spot inconsistencies and dependencies, and compare data sets. It can also flag records for review when something changes. It can also link information from internal systems, such as HR, application inventories, supplier records, location data and previous incident reports.

This stage involves handling some of the most sensitive information in BCM. The BIA and risk assessment data includes information on financial impacts, staffing dependencies, concentration risks, vulnerabilities, supplier weaknesses and internal operational constraints. Always keep live BIA and risk data in approved corporate or private environments, never in public tools.

Used today

CITED PUBLIC EXAMPLE, SOURCE IN REFERENCES, P. 31.

The US Defense Logistics Agency uses AI risk-assessment models to identify unreliable suppliers and those providing counterfeit or non-conforming parts (2025).

What AI can do, by activity

Gather the approved source material for the analysis. This should include activity lists, organisational structures, system inventories, supplier records, location data, impact criteria, prior incident material and existing risk information (*Data rules*, p. 7).

Preparing the BIA

- Explaining BIA terms such as maximum tolerable period of disruption (MTPD) and recovery time objective (RTO) in easier language L1
- **BIA interview support.** Use an enterprise AI assistant to prepare interview guides and workshop questions, explain definitions, and structure the expected outputs of the discussion. L2

BIA interviews and workshops

- Transcribing and summarising workshops and interviews. The summary includes impacts, dependencies, assumptions, gaps and open questions, rather than just providing a well organised summary. L2

Dependencies and consistency

- **Consistency checking across responses.** Ask the AI to compare BIA outputs from related departments and identify where one team's assumptions or recovery expectations conflict with another's. L2
- **Dependency mapping.** Compare activities with systems, people, suppliers, sites and data to identify shared dependencies, single points of failure and areas of concentration risk. L3

Recovery requirements

AI verifies the proposed recovery requirements, but does not set them.

- **Recovery target review support.** Use AI to test whether a proposed RTO is consistent with impact thresholds, dependencies, and comparable activities. The activity owner and management set the target. L2

Risk assessment

- **Risk scenario support.** Search-enabled tools identify current public threat developments. AI then writes a relevance note for each threat, naming the source and the risk record affected. This is followed by questions for the risk owner. The risk owner then decides on the next steps. L4
- Clustering risks, issues, and recurring themes L3

Keeping the analysis up to date

Where possible, a connected workflow can detect changes such as application retirement, restructuring, changes to suppliers, or site closure.

- Prompting a review when the organisation or its systems change L4
- **Continuous review workflows.** A workflow tool monitors changes in approved source systems and flags any potentially affected BIA or risk assessment records for human review. L5

L1 Ad hoc · L2 Repeatable · L3 Defined · L4 Measured · L5 Optimising

STAGE 4 · DESIGN

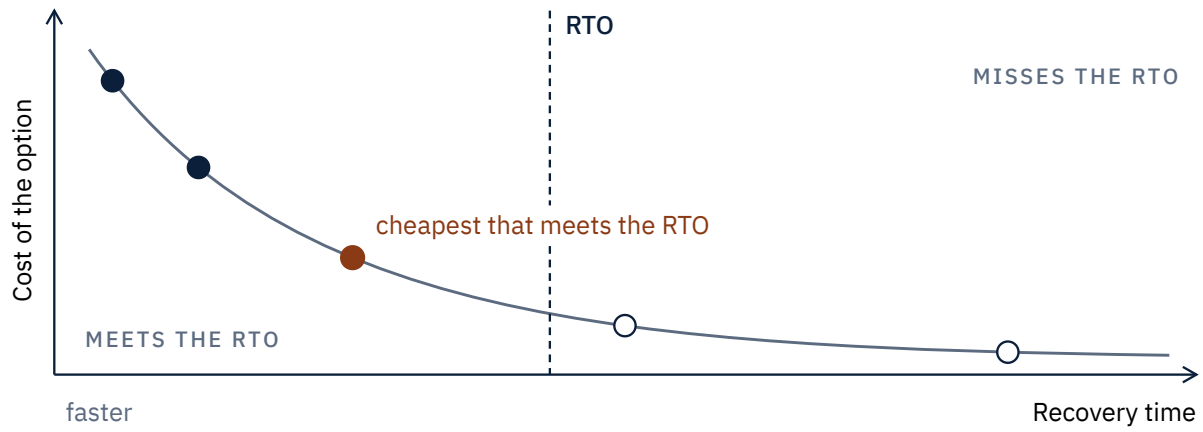
Continuity strategies and solutions design

What this stage covers

This stage covers the design and selection of continuity strategies and solutions to meet business requirements. It is where the organisation moves from analysis to action by considering how it will protect critical activities, recover within required timeframes, and manage trade-offs between cost, resilience, complexity, and practicality.

EXHIBIT 8

AI finds the cheapest option that still meets the RTO



ILLUSTRATIVE. AI PLACES THE OPTIONS BY COST AND RECOVERY TIME AND FINDS THE CHEAPEST THAT MEETS THE RTO.

Where does AI help, and where are its limits?

AI compares options, shows gaps between what you have and what you need, challenges assumptions and organises the facts for a decision. It also brings together information from BIA outputs, dependency maps, location constraints, supplier arrangements, technical architecture, and indicative cost assumptions.

AI should not be allowed to decide what the organisation invests in, accepts or prioritises without effective safeguards and quality checks in place. Management should make decisions after consulting those affected.

Used today

CITED PUBLIC EXAMPLE, SOURCE IN REFERENCES, P. 31.

The US Army Materiel Command contracted AI mapping and monitoring of its multi-tier supply chains to see where vulnerabilities sit and where alternative sources are viable; this was announced at contract award in 2025, not reported as a result.

What AI can do, by activity

Use approved internal source material such as BIA summaries, dependency maps, site constraints, technical architecture notes, current contracts, and investment templates (*Data rules*, p. 7).

Gap analysis

- **Structured gap analysis.** Ask the AI to compare the required recovery outcomes and critical dependencies with the current arrangements, and to identify any gaps, assumptions or missing information. L2

Identifying strategy options

- **Generate a long list by resource type.** Prompt the AI to suggest possible continuity strategies categorised by area, such as the workplace, staff, technology, suppliers, inventory, records, logistics and process workarounds. L2
- Scanning the market for recovery services and options L1
- **Horizon scanning workflows.** Use workflow and search-enabled AI to produce regular updates on relevant developments in recovery technology or service models. L4

Reducing risks and single points of failure

It is often more valuable to test how a solution might fail than to explore additional options.

- **Solution red-teaming.** Ask the AI to act as a reviewer and identify how a proposed solution could fail in the event of a real disruption, taking into account hidden assumptions and shared dependency risks. L2
- Testing the shortlist against approved internal data in order to identify the limits of sites, suppliers, systems and people. L3

Selecting solutions and resources

Before a solution is chosen, the business owner and the relevant specialists confirm that it is effective and that they understand who is responsible for it.

- **Business case drafting support.** Use AI to summarise structured option data and create a concise comparison paper for management. Ensure that the AI clearly labels assumptions and avoids unsupported claims. L2
- Comparison of shortlisted solution alternatives and their trade-offs L3
- A BCM or resilience platform where results measured across sites indicate which options should be kept, changed or retired. L5

L1 Ad hoc · L2 Repeatable · L3 Defined · L4 Measured · L5 Optimising

STAGE 5 · IMPLEMENT

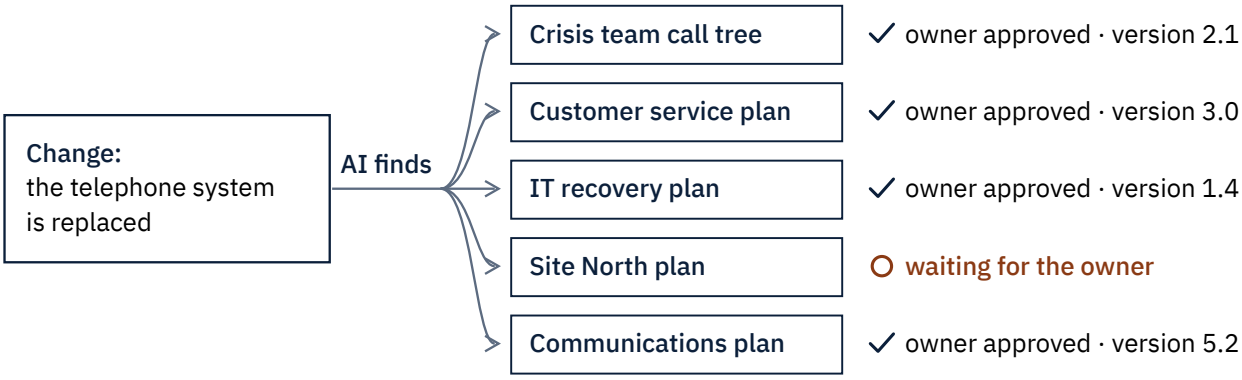
Implementing continuity arrangements

What this stage covers

This stage covers the implementation and operation of continuity arrangements, including plans, response structures, communications arrangements, supporting information, and the practical preparation needed for effective action during disruption.

EXHIBIT 9

One system change affects five plans, each of which requires the owner's approval and a new version.



ILLUSTRATIVE. AI FINDS THE PLANS A CHANGE AFFECTS

Where does AI help, and where are its limits?

Much of this stage is writing, updating and comparing plans, and AI does that well. It drafts plans, checks plans against each other, writes action cards per role, drafts messages for review and flags plans to update when the organisation changes. AI also supports some bounded tasks during an incident, such as retrieving relevant plans, summarising incoming information, extracting actions, or preparing draft messages.

Used today

CITED PUBLIC EXAMPLE, SOURCE IN REFERENCES, P. 31.

The UK Government Communication Service used its own generative AI tool to draft and evaluate crisis communication plans for 17 government departments (2024).

What AI can do, by activity

Gather the following items: plan templates, recovery requirements, role information, escalation arrangements, communications guidance and any other approved source content (*Data rules*, p. 7).

Plans and procedures

The team responsible for implementing the plan confirms that the activation criteria, roles, dependencies, decisions and actions are realistic and practical.

- **Template-grounded plan drafting.** Give the AI the approved plan structure and ask for concise, directive wording that identifies any missing information, rather than making assumptions to fill the gaps. L2
- **Role-based action card generation.** Ask the AI to extract the actions relevant to a named role, timeframe or scenario from the controlled plan. Then, ask the role owner to review them. L2
- **Cross-plan consistency review.** Use AI-based retrieval to compare strategic, tactical, departmental and technical plans for contradictory roles, activation thresholds, communication routes and escalation logic. L3

Warning and communication

- **Draft communications support.** Use AI to prepare holding statements and staff messages for likely disruption scenarios before any incident occurs. Communications and management must approve these before they are used. L2

Incident response

If AI is used during an incident, limit it to retrieval, summarisation, action extraction, or draft message preparation within clear boundaries and approval gates.

- Summarising notes or actions from response meetings L2
- Retrieving relevant plan content during an incident, in a controlled environment L4

Keeping plans up to date

- **Change-triggered maintenance workflows.** Use AI to identify which plans may need reviewing when key data changes, by connecting workflow tools with systems for organisational structure, role ownership, site details or technical dependencies. L3

L1 Ad hoc · L2 Repeatable · L3 Defined · L4 Measured · L5 Optimising

STAGE 6 · VALIDATE

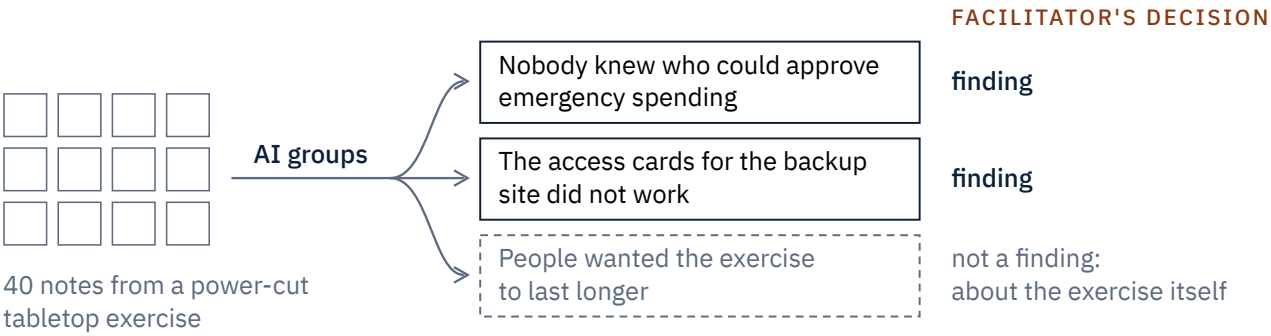
Validation, exercising, review and improvement

What this stage covers

This stage covers the validation and continual improvement of continuity arrangements through exercising, review, assurance activity, lessons identified, maintenance, and management review. Its purpose is to confirm that continuity arrangements are realistic, usable, current, and capable of supporting the organisation during disruption.

EXHIBIT 10

AI groups the debrief notes into themes. The facilitator picks the findings.



ILLUSTRATIVE. THE FACILITATOR READS THE NOTES BEHIND EACH THEME BEFORE DECIDING IT IS A FINDING.

Where does AI help, and where are its limits?

AI can write exercise scenarios, play external parties, summarise debriefs, group lessons identified and compare results across events. It can also prepare material for management review or audit. It can read large volumes of notes, actions and evidence more quickly than people can.

Although AI adds variety and speed, the exercise director, facilitator and auditor still judge the results. As exercise records, debrief comments, audit observations and post-incident material may contain sensitive operational details or personal data, only approved environments should be used for this work.

Used today

CITED PUBLIC EXAMPLE, SOURCE IN REFERENCES, P. 31.

Fayette County Public Schools in Georgia run virtual tabletop exercises on cyber and physical incidents for their leaders with custom AI assistants (2025).

What AI can do, by activity

First, state what the exercise, review or audit is meant to test.

Exercises

- **Scenario and inject generation.** Ask the AI to provide a scenario package tailored to the exercise type, sector, participants and learning objectives. This package should include timed injects and facilitator guidance. L2
- **Controlled role-play support.** Use AI to generate realistic external messages from regulators, customers, suppliers, the media or partners, all of which are subject to facilitator control. L2

Debriefs and lessons identified

- **Debriefing and hot debriefing summarisation.** Structure post-exercise or post-incident discussions around the following areas: strengths, weaknesses, issues, actions and unresolved questions. Ensure that privacy and retention requirements are met. L2
- Clustering lessons, recurring weaknesses, and improvement actions L3
- **Cross-event lessons analysis.** Ask AI to analyse the results of several exercises or incidents and identify any patterns that could suggest underlying capability gaps. L3

Keeping the BCMS up to date

- **Periodic BIA and plan review.** AI can list the BIA entries and plans that are due for review, as well as any changes to organisational structure, role ownership, site details or technical dependencies that have occurred since the last review. Each owner confirms or corrects the relevant entry and records the review date. L4

Audit

- **Audit and evidence gap scanning.** Use AI-based retrieval to compare current BCMS documents and records with internal checklists or approved requirement summaries, flagging any missing evidence for human review. L3

Management review and improvement

AI can group actions and report on progress. People are responsible for each action and decide when it is complete.

- **Management review support.** Use AI to summarise exercise statistics, action progress, maintenance records and assurance findings for top management. L2
- Analytical tools for analysing trends and performance across all exercises, incidents and audit findings. These tools propose the next exercise, which is then decided by the exercise owner. L5

L1 Ad hoc · L2 Repeatable · L3 Defined · L4 Measured · L5 Optimising

04

Choosing a tool

Which tool for which task

The table below provides an overview of the five service types, the tasks they are best suited for, and whether they train on your content.

SERVICE	WHAT FOR	TYPICAL TASKS	TRAINS ON YOUR CONTENT?
Free public service	Generic brainstorming only, no confidential content	Fictional scenario ideas; nothing about your organisation	It can. Check the privacy setting; never enter BCM data
Individual paid service	Low-risk individual productivity, if formally approved	Learning BCM terms and methods; public threat identification and regulatory research	Depends on the plan and its setting; check before use
Team or business subscription	Basic internal use if approved and governed	The same, once approved and governed	Not by default, by the vendor's own statement; confirm in your contract
Enterprise or corporate licence	Often right for live BCM work	Policy and reporting, BIA, plans, exercises, incident support; with logging: scheduled reminders and reviews	Not by default, by the vendor's own statement; confirm in your contract
Private, self-hosted or on-premise	Highly sensitive or regulated use cases	The most sensitive BIA, options and plans	Your own set-up decides

Be aware that even tools of the same type can keep your data for different lengths of time, offer different contract terms, and give connectors different access. Therefore, before granting a tool access to your data, complete the seven checks on the next page.

Before you approve a tool

Before an AI tool can access your data, it must pass the seven checks below.

- [ ] how long the provider keeps your data, how it is deleted, where it is stored if that matters, and whether the provider trains on it
- [ ] privacy and confidentiality terms in the contract, and licence terms for what you upload, including your own licences for standards and copyrighted material
- [ ] sign-in through your identity system, access by role, what each connector may reach, and which sources the tool may use
- [ ] audit logging and reporting
- [ ] whether it suits sensitive operational content and meets your AI policy and information classification rules
- [ ] a fallback if the tool is unavailable, an exit plan if the provider fails, the provider’s breach notification terms, and whether your third-party risk process and your incident plan cover the provider
- [ ] a date to check these answers again, at each renewal and whenever the provider changes the model

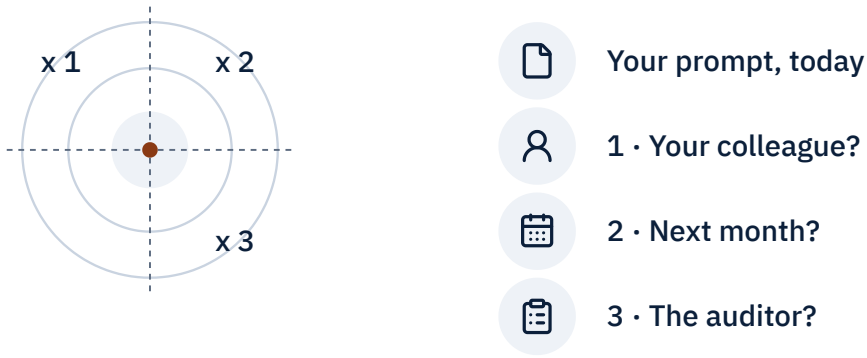
An enterprise label does not inform you about the provider's retention policy and administrative settings, which determine which data is stored and where. When using a chat app, for instance, the vendor has the ability to switch or retire the model without providing prior notice. Therefore, it is recommended to set a fixed date for a re-check. Most of this list matches ETSI EN 304 223 (2025), the UK Government AI Playbook (2025), the NCSC and CISA guidelines (2023), and the SEI model (2026).

05

Prompts that work

One great answer. Could you get it again?

EXHIBIT 11



1 Pick the task | 2 Write the prompt | 3 Check the output | 4 Make it repeatable

Step 1

Which task first?

Pick a task where all four answers are *yes*.

- Do you do it at least once a month?
- Is it a task you would rather not do?
- Could you explain the steps to a new colleague in a few minutes?
- Will a person still review and approve the result, and can the work carry on without the tool?

Ask the AI4BCM Guidance chatbot which task to pick. This will help you identify which section of the guide covers it. **ai4bcm.org/chat**

Step 2

The prompt pattern

ROLE	BCM analyst **[context]**, working on **[the larger task]** for **[who will use the output]**.
TASK	**[one narrow, clear task]**.
SOURCES	the attached material only, as evidence about the organisation. Use no facts or figures from general knowledge.
OUTPUT	**[the structure you want, for example an executive summary, then one section per site]**.
GAPS	anything the sources do not answer. Do not guess it; list it here. Also list any source that is a draft or out of date.
CITE	document and section for each point. Put copied words in quotation marks.

Six example prompts are online, ready to copy. ai4bcm.org/guidance

For even better output

1. Include a few examples of good output and mark them as such.
2. Place long source material first, each document in its own marked block, and place your question at the end.
3. After the first draft, ask the AI to find a supporting quote for each claim, and remove any it cannot support, marking where it was.

Step 3

Before you accept the output

The principles in *The rules that always apply* (p. 7) provide a comprehensive overview of the key aspects to consider when evaluating AI output.

Step 4

Make it repeatable

EXHIBIT 12



A good prompt is luck. Control makes it repeatable.

Level 2 · Repeatable

Share your prompt with your team

A prompt becomes a team asset when it is saved and given a version number, an owner and a named reviewer. Before the team relies on it, run it on three past cases for which you know the correct answer.

Level 3 · Defined

Build your own skill

A skill is a saved prompt that makes the results more consistent. Here is how you can build one:

EXHIBIT 13



The vendors' own guides on building skills: ***Prompting and skill guides*****, p. 32.**

06

Looking after your prompts and tools

Once your team has used a saved prompt, skill or workflow, it becomes part of your BCM process. One person should be responsible for this. Someone else should check its results.

1 Set it up | 2 Use and maintain it

Step 1

A record card for each skill or workflow

A record card provides a brief overview of a skill, including who owns it and how it is safeguarded. Create a card for each skill and workflow that reads or acts on critical files.

WHO	WHAT	HOW SAFE	HOW LONG
OWNER Who is responsible for keeping the card up-to-date.	TASK The one narrow job it does, and what it may not be used for.	CLASSIFICATION LEVEL AND WHERE IT RUNS The type of data it interacts with and the environment in which the tool operates.	VERSION The current version, its date, and what changed.
DEPUTY Who covers when the owner is away.	SOURCES AND DATES Which documents it reads, and the date of each.	PERMISSIONS What it may read, write or send, as set in the tool.	TESTS LAST RUN The date the test cases last ran, and the result.
REVIEWER Who checks each output before it is used.	DECISION OR INPUT Whether its output feeds a decision or only further work.	STOP AND FALLBACK Who can stop it today and how, and how the work still gets done while it is stopped.	RETIRE WHEN The condition that ends its use. See *Keep it up to date*, p. 30.

Used today

CITED PUBLIC EXAMPLE, SOURCE IN REFERENCES, P. 31.

UK public bodies publish one record per AI tool. The Cabinet Office's own record details the tool's owner, purpose, data, oversight and risks. It also states that the tool should not be used for decision-making purposes (2025).

Let your team know about your skills and workflows

Tell your team how to use each skill or workflow, and what to watch out for.

- **Internal audit**, that it exists, and where the card and its records are kept.
- **Information security**, before it reaches any file or system. They approve its permissions.
- **Compliance**, when its output touches a legal or regulatory duty.

Keep it up to date

Make sure to keep your skill or workflow up to date, and retire it when one of the conditions below applies.

- A source it depends on is deleted, replaced or out of date, and nobody updates it.
- The tool loses its approval for the classification level it runs on, or the provider's terms or model change.
- More of its output is rewritten or rejected by you or reviewers.
- The skill or workflow owner retires, and nobody keeps it up to date.

Say that AI helped

A document that AI helped to produce says so inside the document, beside its owner, version and review date. One sentence does it.

```
Parts of this document were drafted with AI support ([tool or skill], version [x])
from the sources listed. [Name] reviewed it on [date].
```

References and further reading

Cited on these pages

Bessen, J. (2018). *AI and Jobs: the role of demand*, NBER Working Paper 24235. Employment grows while demand is elastic and falls once it is not. doi.org/10.3386/w24235

Cabinet Office (2025). *Cabinet Office: Assist*, algorithmic transparency record, published 4 June 2025, updated 30 July 2026, the department's own record. The record-card example; accessed October 2026. gov.uk/algorithmic-transparency-records/cabinet-office-assist

ClearBank (2024). *ClearBank: Safeguarding generative AI use cases in a regulated fintech banking API*, case study in the Portfolio of AI Assurance Techniques, Department for Science, Innovation and Technology, 9 April 2024, described by the bank, not a reported result. The Stage 1 example; accessed October 2026. gov.uk/ai-assurance-techniques/clearbank-safeguarding-generative-ai-use-cases-in-a-regulated-fintech-banking-api

Defense Logistics Agency (2025). *DLA applying AI to supply chain risk management, warfighter readiness*, news article, 13 March 2025, the agency's own report. The Stage 3 example; accessed October 2026. dla.mil/About-DLA/News/News-Article-View/Article/4117309/dla-applying-ai-to-supply-chain-risk-management-warfighter-readiness

EdTech Magazine (2025). *Building K–12 Tabletop Exercises With Generative AI*, 28 May 2025, written by Matthew Jackson, coordinator for digital and media services at Fayette County Public Schools, Georgia. The Stage 6 example; accessed October 2026. edtechmagazine.com/k12/article/2025/05/building-k-12-tabletop-exercises-generative-ai

ETSI (2025). *EN 304 223 V2.1.1, Securing Artificial Intelligence (SAI); Baseline Cyber Security Requirements for AI Models and Systems*, adopted 8 December 2025. Permissions, due diligence, recovery plans, contracts and logging (5.1.2-6, 5.1.2-7, 5.2.2-5, 5.2.2-6, 5.4.2-1). Tool checks. etsi.org/deliver/etsi_en/304200_304299/304223/02.01.01_60/en_304223v020101p.pdf

European Union (2024). *Regulation (EU) 2024/1689 (Artificial Intelligence Act)*, 13 June 2024. Article 4 (AI literacy), as replaced by Regulation (EU) 2026/1744 (Digital Omnibus on AI) of 8 July 2026, in force 27 July 2026. Principle 1; accessed October 2026. artificialintelligenceact.eu/article/4 · eur-lex.europa.eu/eli/reg/2026/1744/oj

Exiger (2025). *U.S. Army Licenses 1Exiger AI Software to Accelerate Defense Acquisition, Reduce Lead Times and Enhance Operational Readiness*, press release, 14 November 2025, a claim at contract award, not a reported result. The Stage 4 example; accessed October 2026. exiger.com/perspectives/us-army-exiger-ai-supply-chain-contract

George Mason University (2025). *Dialing up new 9-1-1 training with AI*, news article, 9 September 2025, on the Fairfax County Department of Public Safety Communications training project, funded by NIST. The Stage 2 example; accessed October 2026. content.gmu.edu/news/2025-09/dialing-new-9-1-1-training-ai

IMDA (2026). *Model AI Governance Framework for Agentic AI*, Version 1.5, 20 May 2026, updated 5 June 2026. Sections 2.2.2 (design for meaningful human oversight), 2.3.1 (system-level safeguards, not prompt wording) and 2.3.2 (testing agents, including indirect prompt injection). Principle 3. imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/mgf-for-agentic-ai.pdf

Jevons, W. S. (1865). *The Coal Question*, London: Macmillan; chapter VII, "Of the Economy of Fuel". Cheaper use, more consumption. Read in the second edition (1866), public domain. oll.libertyfund.org/titles/jevons-the-coal-question

NCSC, CISA and partner agencies (2023). *Guidelines for Secure AI System Development*, Version 1.0, 27 November 2023, current. Limiting what an AI component may trigger. Suppliers must meet your own standards, and you must be ready to fail over. Tool checks. ncsc.gov.uk/collection/guidelines-secure-ai-system-development

OWASP GenAI Security Project (2025). *Top 10 for Agentic Applications for 2026*, December 2025, CC BY-SA 4.0. ASI01 Agent Goal Hijack and ASI06 Memory and Context Poisoning, why retrieved text is evidence and never instruction. *Prompt injection*. genai.owasp.org/resource/owasp-top-10-for-agentic-applications-for-2026

SEI, Carnegie Mellon University, with Accenture (2026). *The AI Adoption Maturity Model v1.0*. Third-party inventory with plans for ending the relationship (section 4.11.2). Tool checks. Its five levels are its own. sei.cmu.edu/library/ai-adoption-maturity-model

UK Government Communication Service (2024). *STOP project: using Assist to transform the UK's crisis communication preparedness*, case study, 11 September 2024, updated 24 August 2026. The Stage 5 example; accessed October 2026. communications.gov.uk/case-study/stop-project-using-assist-to-transform-the-uks-crisis-communication-preparedness

UK Government Digital Service (2025). *Artificial Intelligence Playbook for the UK Government*, 10 February 2025, Open Government Licence v3.0. The four data questions to ask before using organisational data, and the need for a human before any action that cannot be undone. Tool checks; the best starting point for a team new to AI. gov.uk/government/publications/ai-playbook-for-the-uk-government

Further reading

ISACA (2023). *CMMI V3.0*, a licensed model; the level names were read on the CMMI Institute page. Read it for the ladder that levels 1 to 3 of this guidance follow. The name Measured and every description are AI4BCM's own; level 5 takes CMMI's name, Optimising. cmmiinstitute.com/learning/appraisals/levels

Kalai, A. T., Nachum, O., Vempala, S. S. and Zhang, E. (2025). *Why Language Models Hallucinate*. OpenAI. Made-up answers tend to be confident and specific. cdn.openai.com/pdf/d04913be-3f6f-4d2b-b283-ff432ef4aaa5/why-language-models-hallucinate.pdf

NCSC (2026). *Managing the cyber risk of agentic AI*, blog post, 20 August 2026, interim advice that formal guidance will replace. Read it when an agent, rather than a workflow, is proposed. Controls matched to the level of autonomy, sandboxing, logs that cannot be changed, and the ability to stop it at once. ncsc.gov.uk/blogs/managing-the-cyber-risk-of-agentic-ai

NIST (2024). *AI RMF Generative AI Profile*, NIST AI 600-1, July 2024, companion to AI RMF 1.0, which NIST says is being revised. MS-2.5-003: check sources and citations in generated output. GV-6.2-006: fallback may include doing things manually. Workflow design in the online copy. doi.org/10.6028/NIST.AI.600-1

NIST. *AI RMF Playbook*, live companion to AI RMF 1.0, carrying the banner "The AI RMF 1.0 is being updated." A list of actions on third-party dependency, shutting down with backup systems and decommissioning. No maturity levels. airc.nist.gov/airmf-resources/playbook

OWASP (2025). *AI Maturity Assessment (AIMA)*, Version 1.0, 11 August 2025 (Excel toolkit v1.0.1). Read it when a claimed capability needs evidence. Its detailed assessment checks evidence, "not just 'paper compliance'". The released PDF defines eight domains, 24 practices, two streams and three levels per practice. owasp.org/www-project-ai-maturity-assessment

Prompting and skill guides

The vendors' own prompting guides.

Anthropic (n.d.). *Prompting best practices · How to create custom skills · Skill authoring best practices*. platform.claude.com/docs/en/build-with-claude/prompt-engineering/overview

Google (n.d.). *Get started with Gems in Gemini Apps* (section: Tips for creating custom Gems). support.google.com/gemini/answer/15236321

Microsoft (n.d.). *Get started writing prompts in Microsoft Copilot · Get better Copilot responses with great prompting · Agent Builder overview · Write effective instructions for declarative agents*. support.microsoft.com/microsoft-365-copilot · learn.microsoft.com/microsoft-365/copilot/extensibility

OpenAI (n.d.). *Prompting · Prompt engineering · Skills in ChatGPT*. developers.openai.com/api/docs/guides/prompt-engineering · help.openai.com/en/articles/20001066

Standards this guidance works alongside

AICPA and CIMA. SOC 1 and SOC 2 reports, with the *2017 Trust Services Criteria (With Revised Points of Focus — 2022)* behind a SOC 2. When a provider answers one of the tool checks with an assurance report, and you need to see what it covered and over which period. aicpa-cima.com/resources/download/2017-trust-services-criteria-with-revised-points-of-focus-2022

COSO. *Internal Control — Integrated Framework* (2013), with its 2023 supplemental guidance on internal control over sustainability reporting and its 2026 guidance, *Achieving Effective Internal Control Over Generative AI*. When internal control, rather than IT, asks where an AI-supported step sits in a process. coso.org/guidance-on-ic

ISACA. *COBIT 2019*. When AI use in BCM has to be reported inside an IT governance framework, and you want that framework's own words for it. isaca.org/resources/cobit

ISO 22301:2019. *Security and resilience — Business continuity management systems — Requirements*, second edition, with Amendment 1:2024. The RTO and BCMS definitions; cited on p. 13 and in the glossary. iso.org/standard/75106.html

ISO 31000:2018. *Risk management — Guidelines*, second edition. When a BIA result has to travel into the organisation's wider risk work, and you want the vocabulary that work already uses. iso.org/standard/65694.html

ISO/IEC 27001:2022. *Information security, cybersecurity and privacy protection — Information security management systems — Requirements*, third edition, with Amendment 1:2024. When the approval you need for a tool or a connector sits with information security. iso.org/standard/82875.html

ISO/IEC 42001:2023. *Information technology — Artificial intelligence — Management system*, first edition. When the organisation wants a management system for its AI use rather than a rule for one task. The newest ISO standard here, and the only one written for AI. iso.org/standard/81230.html

NIST (2024). *The NIST Cybersecurity Framework (CSF) 2.0*, NIST CSWP 29, 26 February 2024. Plain words for cyber outcomes, for a scenario, an exercise or a conversation with IT. Free to read, and the only standard here the online copy carries. doi.org/10.6028/NIST.CSWP.29

PCI Security Standards Council. *PCI Data Security Standard (PCI DSS)*, v4.0.1. When the material in front of you is payment card data, which the data rules do not list as a class of its own. Classify it before you enter it. pcisecuritystandards.org/document_library

Tools you may know

ChatGPT. Free · Go, Plus, Pro · ChatGPT Business · ChatGPT Enterprise

Claude. Free · Pro, Max · Team · Enterprise

Copilot. Microsoft Copilot app · Microsoft 365 Personal, Family, Premium · Microsoft 365 Copilot Business; Copilot Chat with a work account · Microsoft 365 Copilot (now Microsoft Copilot)

Gemini. Gemini app · Google AI Plus, Pro, Ultra · Google Workspace Business plans; Gemini Enterprise Business · Google Workspace Enterprise; Gemini Enterprise

Sources, read 15 October 2026: chatgpt.com/pricing · openai.com/enterprise-privacy · claude.com/pricing · privacy.claude.com · microsoft.com/microsoft-365-copilot/pricing · learn.microsoft.com/copilot/microsoft-365/microsoft-365-copilot-privacy · support.microsoft.com/en-us/microsoft-copilot/privacy-faq-for-microsoft-copilot · support.microsoft.com/en-us/privacy/copilot-in-microsoft-365-apps-for-home-your-data-and-privacy · gemini.google/subscriptions · workspace.google.com/pricing · knowledge.workspace.google.com/admin/generative-ai/generative-ai-in-google-workspace-privacy-hub.

Glossary

Agentic AI

An AI system that autonomously plans and executes multi-step tasks toward a goal. Some vendors also call a fixed, rule-based setup an agent, so judge a tool by who decides what comes next.

AI

Computer systems that perform tasks which normally need human judgement, such as reading, summarising, classifying or predicting. In this guidance the term covers generative AI, retrieval tools, machine learning and the workflows built from them.

BCM lifecycle

The six activities that manage continuity capability. They are governance, embedding, analysis, solutions design, implementation and validation.

BCMS

Business Continuity Management System; the management system for business continuity, set out in ISO 22301:2019.

BIA

Business Impact Analysis; analysis of activities to determine the effects of disruption.

Connector

A controlled link between an AI tool and a business system, such as a document library or a risk register. It reads or acts on information there. Set it to read-only and least privilege unless a write action is approved; many tools switch write actions on by default.

Generative AI

AI that produces new text, images or code in answer to a prompt, including summaries and comparisons of what you give it. Its output reads fluently whether or not it is correct, which is why every output is reviewed (principle 3).

LLM

Large language model. An LLM is a model designed to interpret and generate human language. Chatbots and most AI assistants are built on one.

MTPD

Maximum Tolerable Period of Disruption; how long an activity can be down before the impact is unacceptable.

Override

A reviewer's change to an AI output, or its rejection. Override rates are worth counting, and a low rate can mean that reviewers approve without checking (principle 3).

Prompt

The instruction a person gives an AI tool. A prompt requests behaviour and enforces nothing; permissions decide what the tool can reach.

Prompt injection

Text in a document, page or message written to make an AI do something other than what its user asked. An AI cannot reliably tell it from the content around it (*Prompt injection*, p. 8).

RA

Risk Assessment; identifying, analysing and evaluating the risks of disruption to prioritised activities and the resources they depend on, so that a person can decide the treatment and accept the residual risk. *Analysis: business impact and risk* (Stage 3, p. 15) gives a method for judging threats against your risk records with AI support.

Register

An approved list of record that the organisation maintains and a named person owns, such as the list of activities, risks, suppliers or resource requirements. AI may draft an entry; only the owner makes it a register entry.

Repository

The controlled place where approved content is kept, such as the plan library, the document management system or the BCMS content store. Approved means someone is accountable for what is in it and for how current it is.

Retrieval

Fetching passages from an approved repository so that an answer is based on them rather than on what the model learned in training. The control is that the repository is approved and its contents are dated.

Role-based access

Permission granted by the job someone does, so a retrieval returns only what that person is already entitled to read. With it, a connector speeds up access to a repository without widening it.

RTO

Recovery Time Objective; the time frame within the MTPD for resuming disrupted activities at a specified minimum acceptable capacity (ISO 22301:2019).

Saved prompt

A prompt kept with its sources and its named reviewer, so a colleague gets the same result from it next month.

Skill

A reusable AI task set up for one purpose. It holds fixed instructions and sometimes scripts, and loads when a request matches its description. This guidance adds an approved knowledge base, a set output format and a named reviewer who checks its outputs.

Workflow

A defined sequence of automated and manual steps that makes a task repeatable.

All of it is open source

Always the current copy. Download this guidance as PDF or as .md files. **ai4bcm.org/guidance**

Citation, licence and changes

Hoekstra, W., Gerner, K., et al. (2026). *AI4BCM: Guideline for using AI by BCM professionals*. CC BY 4.0.

Licensed CC BY 4.0: share and adapt with attribution. Edition tag Release 2026.11.1. Corrections go to the guidance repository, linked from ai4bcm.org/guidance — one issue per change, quoting the page.

Parts of this guidance were drafted, edited and reviewed with AI support (Claude by Anthropic, Kimi by Moonshot AI, GPT by OpenAI) from the sources listed. The authors reviewed it in September and October 2026. Icons from Lucide (lucide.dev), used under the ISC licence.