

AI Maturity Playbook for BCM

1 / 3

Four axes, five levels. Score where you actually stand.

THE FIVE LEVELS

1 · AD HOC

Everyone tries something.
Nobody knows what
produced which BIA.

2 · REPEATABLE

It works because one
person knows how. Almost
none of it is written down.

3 · DEFINED

Method, validation points,
sources and limits are
written down and
versioned.

4 · QUANTITATIVELY MANAGED

Overrides, register age
and boundary hits are
counted.

5 · OPTIMISING

The numbers show what
to fix next. People set the
boundary and check it
afterwards.

THE FOUR AXES

A · METHOD & IMPACT CRITERIA

What the AI is told to do: the BIA
stages, the impact criteria and scales,
the question set.

B · VALIDATION & ACCOUNTABILITY

Who is answerable for what the AI
produced, and at which point in the
BIA they sign for it.

C · RECORD CURRENCY & TRACEABILITY

Whether the BCMS record is still true,
and where each fact in it came from.

D · LIMITS & AUTONOMY

What an agent may do to a live
BCMS record without a person in the
loop.

How to read this · no organisation is “at level 3”. An honest self-assessment names four numbers, for example A4 · B2 · C1 · D1.

Where this sits · ISO 22301:2019 and ISO 31000:2018 carry no maturity concept, so nothing here collides with the standard you already run.

The register: four axes × five levels

2 / 3

One row per axis. Score each on its own.

	AD HOC	REPEATABLE	DEFINED	QUANTITATIVELY MANAGED	OPTIMISING
A METHOD & IMPACT CRITERIA	Analysts prompt freehand. Two people, one interview, two different BIAs	One analyst's prompts get reused. The method lives in their chat history	Stages, impact criteria and question set written down, versioned, agreed with management before elicitation	Rework per stage and answers with no source are counted against that version	The method changes where rework and failed retrievals cluster
B VALIDATION & ACCOUNTABILITY	AI output enters the BCMS unsigned. Nobody can say who agreed it	The analyst checks it. No record that they did, or of what they changed	The process owner validates impact, MTPD and dependencies before the record is saved. Not the analyst, not the tool	Validation rate, override rate and what owners changed are counted per stage	Stages with no overrides for a year drop to spot-check
C RECORD CURRENCY & TRACEABILITY	Nobody knows when the dependency register was last true	A manual refresh once a year, if someone remembers	Every entry carries its source and date. Change triggers are named	Register age and the delay before a change shows up are measured; entries with no source are counted	Operational change proposes the BCMS update; people confirm or reject it
D LIMITS & AUTONOMY	The agent may do whatever the tool allows, including overwriting a validated record	The limits live in the operator's head	Actions tiered by reversibility: draft freely, never write a validated record, never derive recovery plans from a BIA	Boundary hits and refusals are counted	Boundary-hit and refusal rates decide where the limits move next

Evidence, and how you leave each level

Both rows read across the levels, not down an axis.

	AD HOC	REPEATABLE	DEFINED	QUANTITATIVELY MANAGED	OPTIMISING
Evidence WHAT YOU MUST BE ABLE TO SHOW	Nothing. That is the finding	Nothing you could hand an auditor	The method version, the named validators, and the source of each fact	Override rate, register age and boundary hits	Which BCMS change followed from which measurement
Next step HOW YOU LEAVE THIS LEVEL	Run one BIA stage twice the same way	Write the method down, name who validates each stage	Start counting overrides and register age	Define a boundary that holds without step-by- step approval	Pick the next change from the numbers

GLOSSARY

BOUNDARY

The written line around what the AI may do on its own. A boundary hit is one occasion where it tried to cross that line and was stopped.

REWORK

A BIA stage that had to be done again.

SPOT-CHECK

Reviewing a sample instead of every case.

OVERRIDE

A person changed what the AI produced before accepting it. The override rate is the share that gets changed.

REVERSIBILITY

How hard an action is to undo. It decides which actions the agent may take alone.

AUTOMATION BIAS

Trusting the machine's answer over your own judgement.

REGISTER AGE

How long since anyone last confirmed the dependency register is true.

ELICITATION

The interview or workshop where you get the information from people.

END-USER DISCLOSURE

Telling the people affected that AI was involved.

Deliberately out of scope in v0.3 · security as its own axis, third-party agents, automation bias, and end-user disclosure.