

AI GOVERNANCE READINESS CHECKLIST

12 questions your governance committee should be able to answer right now

A working checklist across the five stages of the **Bushey AssureChange™**, Early Design, Development, Deployment, Operation, and Ongoing Use.

Each question below is written to be answerable in committee, on the record. If you can't answer "yes" with evidence, treat it as a "no." Page 2 covers what to do about each "no."

1 Early Design

before you build or buy

- 1 Have we written down, in one place, the business outcome this AI use case must deliver and who signs off if it doesn't? Yes ☐ No ☐
- 2 Have we classified this use case's risk tier using a documented method, rather than a judgement call on the day? Yes ☐ No ☐
- 3 Have we named a single accountable owner for this system who sits outside the vendor and the build team? Yes ☐ No ☐

2 Development

while it's being built or configured

- 1 Did we define "acceptable" performance and failure rate before training or configuration began? Yes ☐ No ☐
- 2 Is there a documented process for reviewing training or input data for bias, provenance, and licensing issues? Yes ☐ No ☐
- 3 Has someone outside the build team formally reviewed this system's design against its risk tier before it progressed? Yes ☐ No ☐

3 Deployment

going live

- 1 Is there a written go/no-go checklist for this system that someone outside the build team must sign off on? Yes ☐ No ☐
- 2 Have we documented what human oversight looks like in production, who can pause or override the system, and how? Yes ☐ No ☐

4 Operation

day-to-day, in production

- 1 Do we monitor this system's outputs against the thresholds we set? Yes ☐ No ☐
- 2 Is there a tested process for flagging a harmful or wrong output, that reaches someone with authority to act? Yes ☐ No ☐

5 Ongoing Use

over the system's lifetime

- 1 Do we re-review this system's risk classification on a set schedule, rather than only after something goes wrong? Yes ☐ No ☐
- 2 If this system were retired tomorrow, could we show, in writing, every material decision made about it since design? Yes ☐ No ☐

What to do about a "No"

A single "no" isn't a crisis it's a gap that's now named, and named gaps are fixable. Here's where to start, by stage.

1 Early Design

Write the outcome and ownership down before the next sprint or vendor call. A 'two-paragraph' charter naming the accountable owner and risk tier is enough to start, traceability is the bar, not perfection.

2 Development

Pause for half a day and get one reviewer outside the build team to check the design and data sourcing against your risk tier. This is the cheapest point in the lifecycle to catch a structural problem.

3 Deployment

Don't go live without a signed-off go/no-go checklist and a named override path. If no one outside the build team can stop this system once it's live, you don't have a safety net, you have a hope.

4 Operation

Stand up a lightweight monitoring routine this week, even a weekly spot-check against your thresholds beats none. Make sure flagged issues reach a named owner with authority to act, not a shared inbox.

5 Ongoing Use

Put a re-review date on the calendar now, **quarterly** is a reasonable default for medium/high-risk systems. Start a single running decision log today; it's far easier to keep one than to reconstruct one later.

Where does your programme stand overall?

This checklist flags individual gaps. For a full picture of your programme's governance maturity, go deeper.

Book a Discovery Call

Talk through your gaps directly with us

bushey.com.au/contact-us

This **Governance Readiness Assessment** is a governance diagnostic tool by Bushey - for more information, contact us at: contactus@bushey.com.au