gwit

MATT JAMES · GREAT WESTERN IT LTD

AI go-live readiness checklist

The checks I'd want agreed before an AI pilot starts or an AI agent goes into live service. Use it alongside your own Service Acceptance Criteria and add anything your organisation needs.

Before the pilot starts

- [ ] What value looks like, using a measure the business already tracks, with where it stands today.
- [ ] When you'll decide, as a date or a point in the work such as the end of a stage, written into the plan.
- [ ] What result means stop, agreed by the board, so that missing it ends the pilot unless the board formally changes the rule.
- [ ] Who makes the call: one named person, usually the sponsor, and who advises them.

The service checks that still apply

- [ ] A named service owner, who is a person with a role rather than a team or a platform.
- [ ] A support group that takes incidents raised against the agent, and a clear view of what it's allowed to change.
- [ ] Monitoring and alerting that reach the support group, tested end to end.
- [ ] An owner and a review cycle for the knowledge the agent relies on.
- [ ] An escalation route to a person, tried with real cases.
- [ ] A rollback plan that someone has rehearsed.
- [ ] A period of early life support after go-live, with extra cover and agreed criteria for handing over to normal support.
- [ ] The agent recorded as a configuration item in the service catalogue or CMDB, with its owner and the systems it depends on.
- [ ] The service desk briefed on what the agent does, and incident categories that let agent problems show up as a trend.
- [ ] Service levels for the agent agreed with the business and reported, such as the share of conversations handed to a person.

What's new for an agent

- [ ] A written scope: what the agent may do, on which records, and what it must hand to a person.
- [ ] Acceptance tests that include awkward cases, such as an ambiguous request, a frustrated user or a case that should go straight to a person. Keep the results as a baseline for spotting drift.
- [ ] An audit trail for every action, back to the agent, the request and the data it used.
- [ ] A way to stop the agent quickly without taking down the service around it, and a named person with the authority to use it.
- [ ] Change control for instructions, models and knowledge sources, with the same approval and testing as configuration.
- [ ] An owner for usage costs and an agreed threshold, with a plan for what happens when it's passed.
- [ ] A date for the first review after go-live, using the same measures as acceptance, and a regular review after that.

Data, security and suppliers

- [ ] A data protection impact assessment, signed off before go-live if the agent handles personal data.
- [ ] Security testing that covers misuse such as prompt injection, and access limited to what the agent's scope needs.
- [ ] Supplier terms checked: how the model is updated, how much notice you get of changes and whether your data is used for training.
- [ ] A fallback for users if the AI service or its supplier is unavailable, covered in your continuity plans.

Telling people it's an AI EU AI Act, Article 50

- [ ] The agent says it's an AI at the start of a conversation, unless that's already obvious.
- [ ] It answers honestly when someone asks if they're talking to a person.
- [ ] Longer or more sensitive conversations, such as HR ones, include a reminder where needed.
- [ ] People can see how to reach a person when they want one.
- [ ] All of this has been tested in user acceptance testing, not only read in the script.

I'm not a lawyer, and this checklist isn't legal advice. Ask your legal or compliance team to add their own checks, so they're part of every go-live whoever is running it.

More on each section at gwit.ltd/blog, including 'Write the stopping rule before the pilot starts' and 'Does your service desk chatbot tell people it's a bot?'

Matt James provides interim ServiceNow programme, delivery and engagement management through Great Western IT Ltd.

gwit.ltd/contact · hello@gwit.ltd