KICKSTANDAI x MY-CC.IO

# AI Agent Readiness Checklist

*A practical checklist for wellness organizations evaluating their first AI agent deployment.*

Wellness organizations are moving fast on AI agents. Scheduling, intake, coaching follow-up, member support. The tools work in a demo. The real test comes after go-live, when an agent makes a call nobody reviewed.

This checklist covers the four areas that determine whether you are ready: data, permissions, integrations, governance. Work through it before your first deployment. Not after your first incident.

## 01 DATA

- ☐ You know exactly what data the agent can see.
  *Not "roughly." A specific, written list of fields and sources.*
- ☐ Health and personal data are separated from general operational data.
  *The agent should not need one access path to reach both.*
- ☐ You have mapped what happens to data the agent generates, not just what it reads.
  *Notes, summaries, and inferred fields count as data too.*
- ☐ Retention rules exist for agent-generated records.
  *If nobody set an expiration, the default is forever.*
- ☐ You know where the data lives, and who else can access it.
  *Including subprocessors and any model provider in the chain.*

## 02 PERMISSIONS

- ☐ The agent's action list is written down, not assumed.
  *If you cannot list what it is allowed to do, you cannot audit it.*
- ☐ Every action on that list has a named owner who approved it.
  *Approval by silence is not approval.*
- ☐ High-stakes actions require human sign-off before execution.
  *Billing changes, care recommendations, anything that touches money or health.*
- ☐ You have defined what the agent does when it hits a scenario outside its permissions.
  *Fail closed, not fail open.*
- ☐ Permission scope gets reviewed on a schedule, not just at launch.
  *Agents get new integrations over time. Permissions should not drift with them unnoticed.*

## 03 INTEGRATIONS

- ☐ You know every system the agent connects to.
  *CRM, EHR, scheduling, payment, messaging. All of them, listed.*
- ☐ Each integration has a documented failure mode.
  *What the agent does, and what the user sees, when a connection drops.*
- ☐ API keys and credentials the agent uses are scoped to the minimum needed.
  *One compromised key should not open every system at once.*
- ☐ You can shut off one integration without shutting down the agent entirely.
  *A single point of failure should not be a single point of full shutdown.*
- ☐ Someone has tested what happens when an integration goes down mid-conversation.
  *Not in theory. Actually tested.*

---

## 04 GOVERNANCE

- ☐ Every agent action is logged, timestamped, and attributable.
  *To a specific decision, not a general session.*
- ☐ Logs are tamper-evident, not just stored.
  *A log that can be quietly edited is not an audit trail.*
- ☐ A named person is accountable for agent behavior.
  *Not a team. A person.*
- ☐ An incident response plan exists for agent-caused errors.
  *Written before you need it, not drafted the day of.*
- ☐ You can produce a complete audit trail on request, without a scramble.
  *If a regulator or a client asked today, you would have an answer by end of day.*

---

**Most wellness organizations pass the data and integrations sections. Fewer pass governance.**

Governance is the part regulators and clients ask about first, and the part most teams build last. If you cannot answer these questions today, that is the deployment to fix before you add another one.

---

**KickstandAi x My-CC.io** | Real-time compliance enforcement for AI agents and automations | Travis@gokickstandai.com