


WHATSAPP COMPLIANCE?



HUMANITY LINK IS COMPLIANT

July 3rd, 2026

Purpose-built, not general-purpose: why WhatsApp's new AI rules don't change anything for us

Earlier this year, WhatsApp updated its Business Solution Terms to draw a clear line around AI on its platform. From January 2026, general-purpose AI assistants, the kind you can ask anything, about anything, are no longer welcome there. What's still allowed, and actively encouraged, are structured, purpose-specific tools: chatbots built to do one job well, with AI used to make that job faster or more responsive, never as an open-ended replacement for the service itself.

We've spent years building chat-based tools for the humanitarian sector, and we've watched a lot of technology hype cycles come and go. This one is different, because for once, a platform policy change lines up exactly with how we've been building from day one.

We never set out to build chatbots. We set out to build services.

When we built one of our WhatsApp assistants for displaced populations in XXX country, the goal was never "let people chat with an AI." The goal was: can someone reach a trusted source of information about services and referrals, in their language, on the app already on their phone, without needing to know what question to ask first. AI helps that conversation flow naturally. It doesn't define what the conversation is about.

Another of our services works the same way. It's a WhatsApp-based assistant, not a general chat companion. Every interaction has a defined shape: intake, guidance, handoff to a human when the situation calls for it. The AI's job is to make that shape feel less like filling out a form and more like talking to someone who understands your situation. It is never invited to wander outside that shape.

Our outbreak-reporting programmes follow the same discipline. When a health worker in XXX country reports a possible case through one of our chat services, the flow is structured on purpose: a defined intake, an AI-assisted segment where the health worker can ask follow-up questions specific to what they just reported, and a structured close. The AI never becomes a general assistant mid-conversation. It stays scoped to the report in front of it.

What "compliant" actually means to us

WhatsApp's policy isn't really about banning AI. It's about banning ambiguity, tools that could be anything to anyone, with no defined purpose and no boundary on what data goes where. So for us, staying aligned with it has meant a few concrete commitments, the same ones we'd hold ourselves to even if no policy existed:

- Every AI-assisted step in our services has a defined job. If a conversation drifts outside that job, the flow hands off, it doesn't improvise.
- None of the data flowing through our chat services trains or improves anyone else's model. It serves the person in front of us, and nothing else.
- Every AI interaction is logged and auditable, so we can show, not just claim, what happened in any given conversation.
- Human escalation is built in, not bolted on. When someone discloses something that needs a person, not a bot, that handoff is immediate.

None of this required us to rebuild anything. It required us to look hard at what we'd already built and confirm the discipline was there. It was, because it's the same discipline good humanitarian programming has always demanded: know your purpose, protect your data, know when to bring in a human.

Why this matters beyond compliance

The humanitarian sector doesn't get to treat platform policy as background noise. We work with people in the most vulnerable moments of their lives, reporting outbreaks, fleeing displacement, disclosing things they've never told anyone else. A chatbot that doesn't know its own boundaries is a real risk to those people, policy or no policy.

So when WhatsApp says AI on their platform needs a defined purpose, a clear boundary, and respect for where user data goes, that's not a new bar for us to clear. It's the bar we were already building to. We'd rather our sector get ahead of that standard than scramble to meet it after the fact.

If you're building chat-based tools for people in crisis, we'd genuinely encourage you to run your own service through the same three questions: What is this AI actually here to do? Where does the data go? What happens when it needs to stop being an AI and become a person? If you can answer all three clearly, you're already most of the way there.