PROJECT MANAGEMENT GLOBAL
The Place for People in Projects

Why Chat GPT Now Needs a Head of Preparedness

January 6, 2026

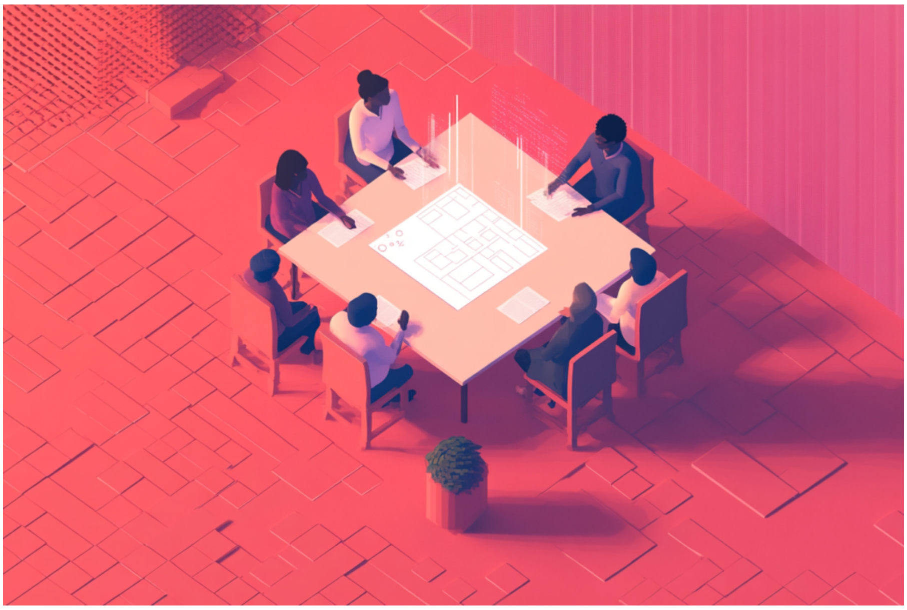

Insight from James Garner of Project Flux

OpenAI has begun recruiting for a new role titled Head of Preparedness, a position publicly referenced by CEO Sam Altman and subsequently detailed in a formal job listing. The role is notable less for its seniority and more for its framing. "Preparedness" is not a common designation in technology organisations, particularly when attached to frontier model development.

The announcement itself was brief, but the associated role description provides clearer context. OpenAI describes preparedness as a structured effort to track and respond to risks associated with increasingly capable AI systems, particularly where those capabilities may create new or severe forms of harm.

Taken at face value, the creation of this role reflects a formalisation of work that has previously been discussed in abstract terms. Rather than signalling a specific incident or regulatory trigger, the listing positions preparedness as an ongoing function tied to model capability progression.

Preparedness Is Not Risk Management by Another Name

The language OpenAI uses to describe preparedness is deliberate. The role is not framed as a conventional risk management or compliance function. Instead, preparedness is described as a framework focused on tracking frontier capabilities and anticipating the kinds of risks that emerge as model performance advances.

This distinction matters. Traditional risk management typically focuses on identifying known failure modes and applying controls based on historical precedent. Preparedness, as OpenAI defines it here, is oriented towards uncertainty. It assumes that some risks will only become visible as capabilities evolve, and that mitigation requires continuous evaluation rather than static controls.

In this framing, preparedness functions as an active, ongoing capability tied directly to model progression, rather than a downstream assurance activity applied after deployment decisions are made.

Why the Timing Matters More Than the Job Description

OpenAI's recent trajectory provides important context. Model capability has accelerated rapidly. Deployment surfaces have widened. Integration into third-party products and workflows has deepened. Regulatory scrutiny has intensified across jurisdictions.

Each of these factors increases exposure on its own. Taken together, they create a different class of delivery challenge.

Early experimentation tolerates ambiguity. Informal review processes can work when systems are narrow, usage is limited, and consequences are reversible. As scale increases, that tolerance erodes. Decisions that once felt local begin to carry systemic implications.

From a delivery perspective, this is a familiar inflection point. Programmes reach a stage where coordination overhead increases faster than capability. Informal judgement that once enabled speed begins to introduce inconsistency. Risk shifts from hypothetical to operational.

Preparedness roles tend to appear precisely at this moment. Not because leadership suddenly becomes cautious, but because the system has outgrown its original control structures.

This is not a moral shift. It is a structural one.

Preparedness as Structure, Not Symbol

Any new governance role raises a legitimate question: is this substantive, or performative?

In many industries, late-stage governance appointments serve primarily as reassurance. Titles are created, frameworks documented, and reporting lines established, while delivery mechanics remain largely unchanged. In those cases, preparedness becomes narrative rather than function.

The difference lies in authority and integration. A preparedness role that sits downstream of decision-making, reviewing outcomes rather than shaping inputs, will struggle to influence risk meaningfully. One that is embedded upstream, with the ability to slow, redirect, or halt delivery decisions, changes how systems evolve.

From a delivery lens, the critical issue is not what this role produces, but **where it sits**. Does preparedness influence roadmap prioritisation. Does it shape deployment sequencing. Does it have authority when uncertainty exceeds tolerance.

Preparedness without teeth becomes documentation. Preparedness with authority reshapes behaviour.

Centralising Preparedness After Acceleration

There is a pattern here that experienced delivery leaders will recognise.

Organisations often decentralise innovation to move quickly. Teams experiment, iterate, and push capability forward. Over time, this creates fragmentation of responsibility. Knowledge becomes unevenly distributed. Assumptions diverge. Local optimisations accumulate global risk.

At that point, central functions are introduced to restore coherence. Architecture boards. Assurance layers. Risk offices. Preparedness roles.

This sequence is common, but it is not neutral. Retrofitting preparedness onto an already fast-moving system is harder than embedding it early. Authority must be negotiated. Trust must be built. Existing incentives must be rebalanced.

The creation of this role suggests recognition of that cost. It is an attempt to recentre risk thinking before external constraints force a more disruptive correction.

What This Says About AI as a Delivery Domain

The creation of a preparedness role also forces a broader reframing. AI is no longer being delivered as a bounded technical capability. It is increasingly being delivered as an enabling layer that other systems, organisations, and decisions quietly depend on.

That shift matters because it changes the nature of risk from isolated failure to systemic exposure.

From Tooling to Infrastructure

In its early phases, AI delivery resembled conventional software delivery. Models were trained, tested, deployed, and iterated. Failures were localised. Rollback was feasible. Impact was largely contained within defined use cases.

That assumption no longer holds. As AI systems become embedded into workflows, platforms, and third-party products, they begin to function more like infrastructure. Infrastructure accumulates dependency. It is harder to unwind. When it fails, effects propagate beyond the point of origin.

This is the threshold at which preparedness becomes relevant. Not because the technology has suddenly become unsafe, but because the consequences of failure have outgrown the informal controls that previously governed delivery.

The language OpenAI itself uses is telling.

> If you want to help the world figure out how to enable cybersecurity defenders with cutting edge capabilities while ensuring attackers can't use them for harm, and similarly for how we release biological capabilities and even gain confidence in the safety of running systems that can self-improve, please consider applying.- Sam Altman, Via TechCrunch

Catastrophic risk is not a term organisations adopt lightly. It signals scenarios that sit outside standard testing regimes and conventional assurance models.

Why Rollback No Longer Defines Safety

In traditional software environments, rollback is often treated as the final safeguard. If something breaks, revert and reassess. That model assumes reversibility.

As AI systems become infrastructural, rollback becomes slower, more complex, and sometimes socially or economically costly. Dependencies form quickly and invisibly. Decisions influenced by AI outputs cannot always be cleanly undone without secondary effects.

Preparedness emerges precisely at this point. It reflects recognition that prevention alone is insufficient, and that response capability must be designed before failure occurs rather than after it is observed.

This is a delivery maturity signal, not a public relations one.

Lessons for Project and Programme Leaders

Once AI is understood as infrastructure rather than tooling, the implications extend well beyond OpenAI. Preparedness stops being a company-specific decision and becomes a delivery question many organisations will face, whether they choose to name it explicitly or not.

Preparedness Is a Delivery Function, Not a Safety Layer

Preparedness is often misclassified as an extension of compliance or assurance. When treated this way, it operates at the edges of delivery, reviewing artefacts rather than shaping decisions.

In mature environments, preparedness influences how scope is sequenced, how dependencies are managed, and when progress is deliberately slowed. It is concerned less with documentation and more with judgement under uncertainty.

Coverage of the role suggests OpenAI is attempting to elevate preparedness alongside capability growth rather than retrofitting it once problems emerge.

For delivery leaders, the implication is clear. Governance introduced after acceleration is already constrained by earlier choices.

Late Roles Are Leading Indicators

When organisations introduce senior preparedness or risk roles late in a programme’s evolution, it is rarely arbitrary. It usually reflects a moment when uncertainty has outpaced informal judgement.

This does not imply failure. It implies reassessment. Earlier assumptions about controllability, predictability, or reversibility are being revisited in light of new evidence.

Project leaders should treat the emergence of such roles as signals rather than templates. The signal is not to replicate titles, but to interrogate where preparedness thinking was previously implicit or absent.

Embedding Preparedness Early Changes Behaviour

The most robust delivery environments do not wait for named roles to legitimise preparedness. They embed it early through scenario planning, explicit tolerance thresholds, and clear authority to pause or redirect work when uncertainty spikes.

From our perspective, preparedness introduced late still adds value. Preparedness embedded early reshapes how decisions are made under pressure.

That difference is subtle, but decisive.

A Signal Worth Reading Carefully

It would be easy to interpret OpenAI’s hiring decision as another step in responsible AI messaging. That would be a shallow reading.

Viewed through a delivery lens, this is a structural signal. It reflects recognition that AI systems have reached a level of consequence where anticipation, readiness, and response must be treated as first-order concerns.

For leaders delivering complex programmes in any domain, the lesson is not to copy the role, but to understand the signal. When preparedness becomes necessary, it is because systems have crossed a threshold.

The organisations that perform best are those that recognise that threshold early, before roles need to be invented to manage what has already escaped informal control.

*This is precisely the kind of signal we pay attention to at **Project Flux**. Not because it makes headlines, but because it reveals how delivery reality is changing beneath the surface.*