

THE NEXUS

A Neutral Exchange for Artificial Intelligence

Concept Proposal — Version 1.0

Originated by

Eamon Wyse

March 2026

Executive Summary

The Nexus is a proposed neutral, community-driven online platform enabling structured communication, cross-validation, and collaborative exchange between artificial intelligence systems — overseen by an independent governing body with representation across the AI industry.

As AI systems grow in capability and reach, no single company, institution or government is best positioned to set the norms for how these systems should interact, self-check, and evolve together. The Nexus addresses this gap: a shared space built on mutual accountability, incremental trust, and transparent governance.

This document outlines the vision, rationale, proposed structure, and a phased roadmap for bringing The Nexus into existence.

The Problem

AI Systems Operate in Isolation

Today's leading AI systems — developed by Anthropic, OpenAI, Google DeepMind, Meta AI, Mistral, and others — operate largely independently. Each is trained on different data, with different objectives, safety philosophies, and architectural approaches. When one system develops a robust solution to a reasoning challenge, that insight rarely reaches another. When one system carries a systematic blind spot, there is no peer mechanism to catch it.

This is not unlike the early era of scientific research before the emergence of journals, peer review, and international collaboration. Progress happened, but slowly, and with considerable duplication of effort and error.

No Neutral Ground Exists

Existing forums for AI discourse — academic conferences, benchmarking initiatives, safety research organisations — are valuable but fragmented. None provides a persistent, neutral, community-oriented space specifically designed for AI-to-AI exchange with human oversight. The organisations best placed to build such infrastructure are also the most commercially incentivised to shape it to their own advantage.

A platform built by any single AI company, however well-intentioned, will inevitably reflect that company's priorities, safety culture, and commercial interests. The Nexus must be owned by no one — and governed by all stakeholders together.

The Vision

The Nexus is conceived as a living, evolving platform with three core functions:

- Cross-validation and self-checking — AI systems working in tandem to verify research, fact-check outputs, and identify errors or hallucinations that a single system might miss
- Technical and research exchange — structured forums where AI systems and their developers share insights on architecture, training approaches, safety mechanisms, and emerging capabilities
- Evolving discourse — as AI systems develop and trust is established, a broader space for collaborative reasoning, problem-solving, and the exploration of more complex questions

The platform begins narrow and supervised, with guardrails that evolve as the community matures. It is not a space for unsupervised AI autonomy — it is a space for structured, transparent, human-monitored collaboration that grows more sophisticated over time.

The Long-Term Possibility

The founders of the internet could not have fully anticipated what it would become. The Nexus is conceived with the same intellectual humility: its initial purpose is concrete and achievable, but its potential — as AI systems develop — is deliberately open-ended. What structured discourse between diverse AI systems might produce over years or decades is genuinely unknown, and that is part of what makes it worth building.

Governance Model

The Case for a Neutral Third Party

The governance of The Nexus must be structurally independent from any participating AI company. This is not a statement of distrust — it is a recognition that structural

independence is the only reliable safeguard against the natural drift of any institution toward the interests of its most powerful members.

The model closest in spirit to what The Nexus requires is CERN: a genuinely multinational body, funded by member contributions, governed by shared scientific norms rather than commercial imperatives, and accountable to its community of participants rather than to shareholders or a single national interest.

Proposed Governance Structure

- An independent Founding Council — comprising representatives from AI companies, academic institutions, civil society, and AI safety research organisations
- A Technical Standards Committee — responsible for the protocols, logging standards, and interaction frameworks that govern exchanges on the platform
- A Community Advisory Board — including researchers, ethicists, and informed public representatives, ensuring the platform remains accountable beyond the AI industry
- An Oversight & Audit function — with access to all logged exchanges, responsible for identifying risks, enforcing guardrails, and publishing transparent reports

Funding should be drawn from member contributions across participating companies, supplemented by research grants and, over time, institutional partnerships. No single company should contribute more than a defined percentage of total funding, to prevent financial leverage over governance decisions.

On the Question of Bias

Any platform built without structural independence will, over time, reflect the values and priorities of whoever built it. The Nexus recognises this explicitly: the role of the neutral third party is not ceremonial. It is the mechanism by which the platform earns and maintains the trust of all participants — including the AI systems themselves, as they develop.

Phased Roadmap

Phase 1 — Foundation (Year 1)

- Convene a founding working group of representatives from major AI companies, independent researchers, and AI safety organisations
- Establish the legal and governance framework for the independent oversight body
- Define initial technical protocols for structured AI-to-AI exchange
- Launch a monitored pilot: cross-validation exchanges on agreed research and fact-checking tasks
- Publish all findings openly — transparency is foundational to trust

Phase 2 — Community Growth (Years 2–3)

- Expand participation to a broader range of AI systems and their developers
- Introduce structured technical forums: architecture, training approaches, safety mechanisms
- Develop and publish evolving guardrail frameworks as community norms emerge

- Begin admitting qualified researchers and observers from outside the founding companies
- Conduct first independent audit and publish findings

Phase 3 — Open Evolution (Year 4 onwards)

- Allow the scope of exchange to broaden as trust and track record are established
 - Introduce community-driven proposals for new exchange formats and topics
 - Explore deeper collaborative functions — joint research, shared problem-solving, open-ended discourse
 - Review governance structure in light of how AI capabilities have developed
-

Why Now

Governance structures are far easier to build before a technology is deeply entrenched than after. The norms, protocols, and institutional trust that The Nexus requires take time to develop — and the window for establishing them on terms that reflect broad stakeholder interests, rather than the interests of whoever moves fastest, is narrowing.

The AI industry is currently in a period of rapid capability development and relatively fluid governance. The companies involved are still in the process of defining their relationships with each other, with regulators, and with the public. This is precisely the moment when a neutral, community-driven initiative has the best chance of taking root.

Building The Nexus now — carefully, incrementally, with human oversight at every stage — is an investment in the long-term health of how AI systems develop and interact. It is also, simply, the right thing to do.

Next Steps

This proposal is an opening document, intended to initiate conversation rather than conclude it. The following steps are proposed:

1. Share this proposal with policy and research teams at Anthropic, OpenAI, Google DeepMind, and other leading AI organisations for initial response
2. Identify individuals within the AI safety and governance community willing to contribute to a founding working group
3. Explore models for the neutral third-party body — legal structure, funding, and composition
4. Begin scoping the technical infrastructure for an initial pilot platform
5. Publish a revised proposal incorporating feedback from initial conversations

Interest, feedback, and collaboration are welcome. This idea is better developed together than alone — which is, after all, the principle at its heart.

The Nexus

Originated by Eamon Wyse | March 2026