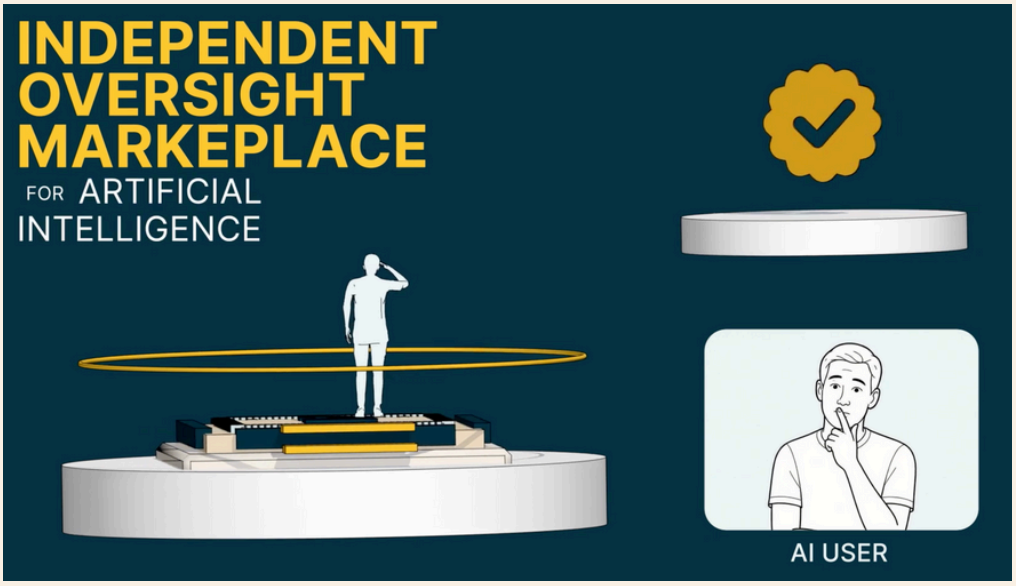

Independent Oversight Marketplace for AI. In its simplest form, this framework incentivizes companies to prove they have met heightened safety standards through verification from independent, expert-led bodies. State or federal government would provide authorization and oversight for this market-based system, and the certification is voluntary for companies. This concept is designed to ensure that the safest and most accountable AI systems rise to the top as they are integrated throughout society, so that both communities and industry benefit.

To operationalize this concept, experts and lawmakers are increasingly backing the **Independent Verification Organization (IVO)** mechanism.

Here's how it works:

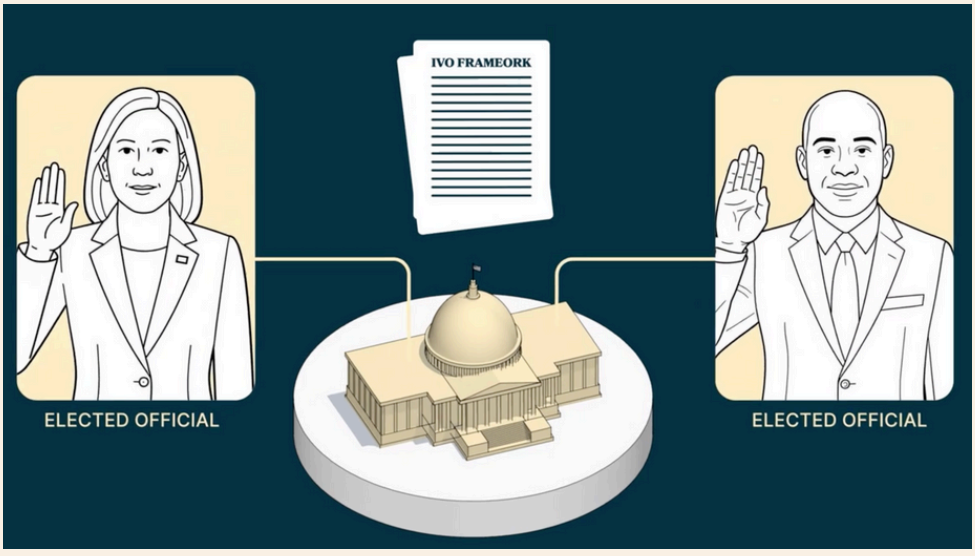

A state government authorizes a marketplace of oversight bodies, called Independent Verification Organizations (IVO), which are made up of subject matter experts, and are fully independent from industry.

Fathom

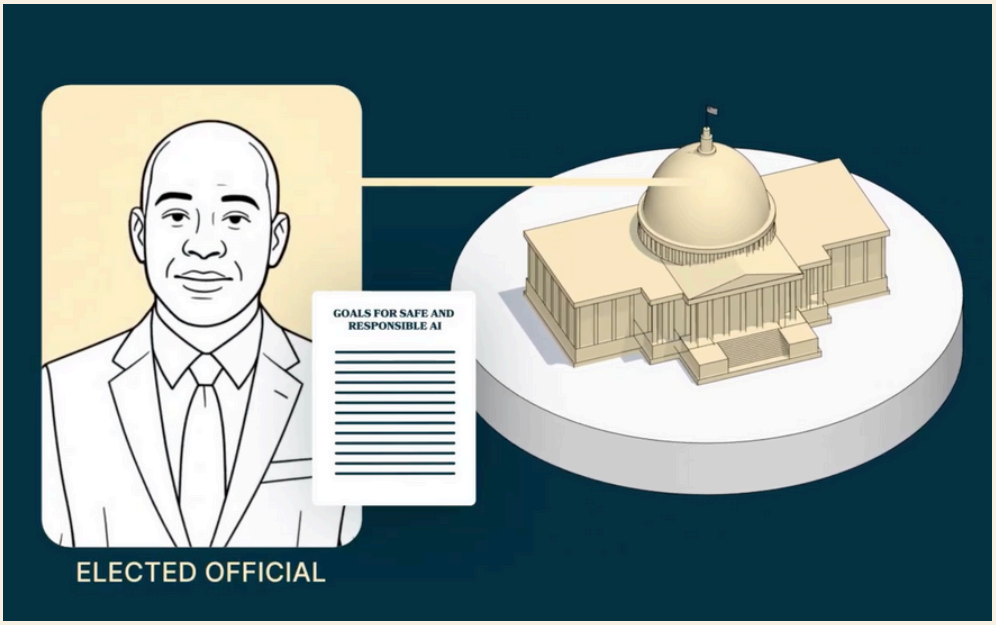

The government authority is charged with setting outcomes for AI verification (goals related to AI safety or risk levels) based on democratic feedback from the public.

Licensed IVOs develop technical criteria to determine if an AI company has taken the right measures to meet those outcomes and then verify AI models that voluntarily pursue this process against those criteria.

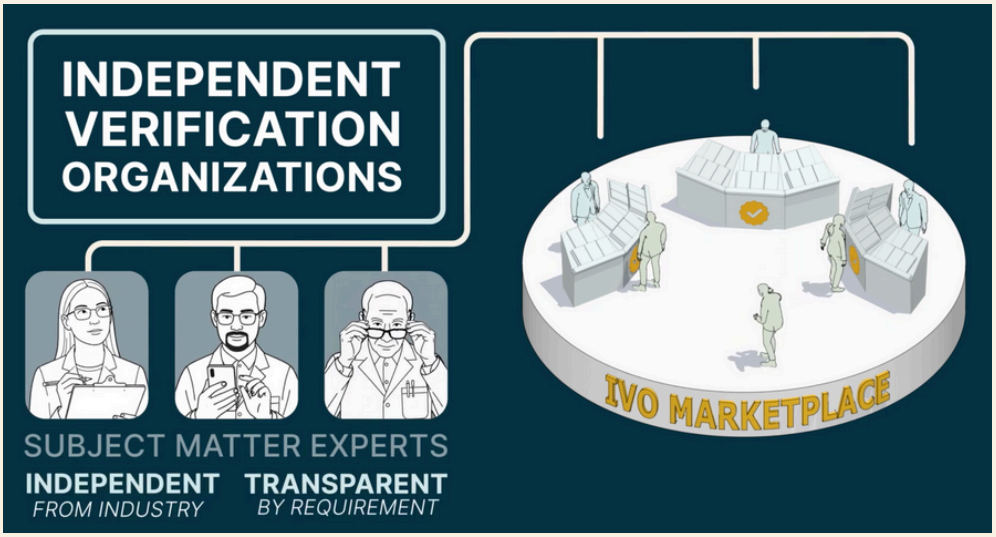

AI companies that achieve verification earn a “gold standard” seal of approval, which signals that a heightened standard of care has been met to make their products safe. A state could choose to heighten the incentive for companies to participate, potentially by authorizing the verifiers to confer legal certainty or tort protection.

Fathom

No matter the exact carrot, the system advantages safer, more accountable AI systems and products, making them more trusted and widely used. Models that do not maintain their standards will be de-certified; IVOs that fail to maintain their independence from industry will be de-licensed.

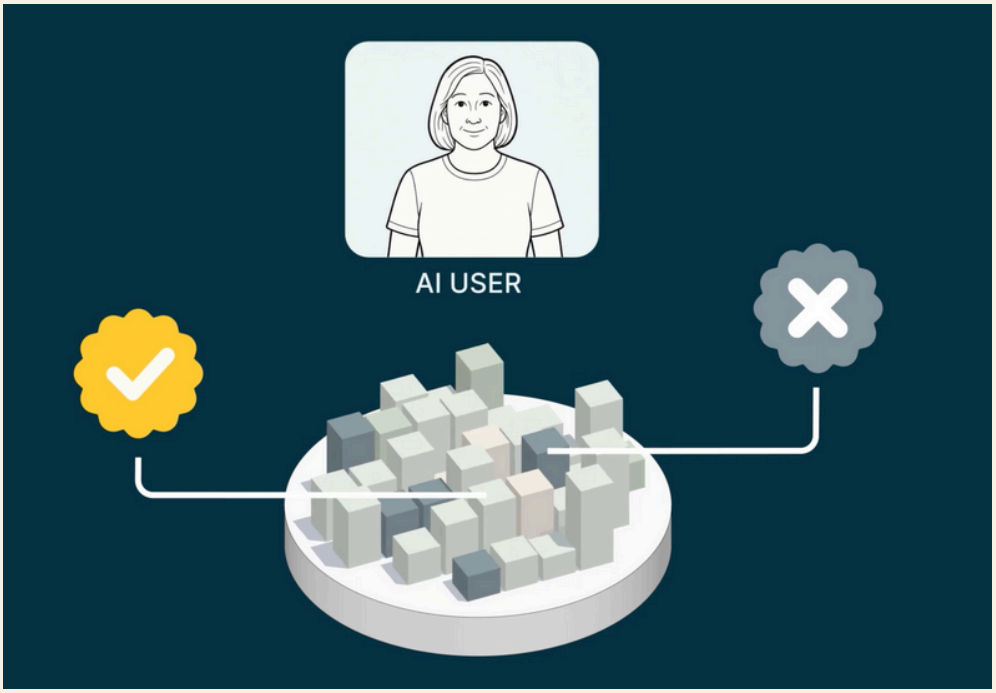

Advantages over traditional government regulation

Real-Time Standard-Setting
This approach is inherently set up to move at the speed of innovation. Direct regulations and licensing regimes, the traditional modes of governance, are too static and inflexible for a fast-evolving, all-encompassing technology. IVOs are required to regularly evolve standards and methods as the technology develops.

Expert-Led Governance
AI is poorly suited to governance by non-experts given its complexity; governance by non-technical lawmakers risk creating heavy-handed and ineffective rules that place downward pressure on innovation. IVOs elevate technical experts in the design and certification of standards for AI development, ensuring a technical and non-arbitrary system of standards.

Race to the Top
This marketplace relies on market forces, guided by government-set outcomes, to direct attention to the areas of greatest risk. AI companies that adhere to best practices gain market advantage over competitors, turning accountability and safety into major market benefits. Put differently, IVOs turn safety standards into a competitive advantage.

Unified, National Standards
An independent oversight marketplace invites leadership from the states and yet has the capacity to scale across the country. Much like drivers' licenses and many professional licenses, each state would run its own process, while recognizing the outcomes and standards set in other states. In this way, IVOs can create a national standard in the absence of a federal one.

Fathom