

The Human Continuity Standard

Fact sheet

Item	Details
What it is	a constitutional layer for advanced AI and a model federal act.
Core rule	an AI may optimize anything inside seven limits and may never optimize the limits away.
The seven limits	human continuity; human flourishing; agency and dignity; life-support systems (food, water, soil, climate); future options; human control; answerability to real-world evidence.
Who it covers	only the most capable and consequential AI systems. Everyday software and everyday use of AI are untouched.
What developers must show	a signed safety case, independent behavioral testing (including shutdown compliance and self-copying attempts), incident reports, staged release, and controls outside the AI.
What developers get	one national rulebook in place of state patchworks, a fixed review clock, and a limited safe harbor for good-faith disclosure.
How it fits current law	it can serve as the evaluation criteria for S. 2938 (Hawley, Blumenthal), the risk-management requirement in H.R. 9925 (FRONTIER Act), or a stand-alone act.
Public support	79% of Americans want AI safety rules kept even if development slows (SCSP and Gallup, Sept. 2026).
Founder	Greg Stoutenburgh: serial entrepreneur, scientist, inventor on dozens of patents worldwide, technologist, AI enthusiast, and father concerned for the future of humanity.
Website	keephumansincharge.org