



INTERFAITH
ALLIANCE
FOR SAFER COMMUNITIES

Guiding Intelligence,
Safeguarding Humanity

The Faith–AI Covenant

I. The Context: A Civilizational Inflection Point

Artificial Intelligence is no longer merely a technical field; it is a civilizational force reshaping knowledge, authority, and the formation of human belief on a global scale. This transformation is unfolding faster than our social, legal, and moral institutions can adapt.

As AI systems become embedded in everyday life, they increasingly influence how people around the world interpret truth, meaning, and moral judgment. In doing so, they are beginning to shape how individuals and communities think about identity, morality, and faith.

For the first time in history, human belief is being mediated at scale by machines. AI systems are now used to summarize sacred texts and interpret theological content. They are increasingly used to seek guidance on moral and theological questions. Without clear cultural grounding and doctrinal integrity, this presents unprecedented risks:


Vulnerability Of The Next Generation

Children are engaging with AI systems during their most formative years, exposing them to risks ranging from cognitive manipulation to potential pathways toward radicalization.



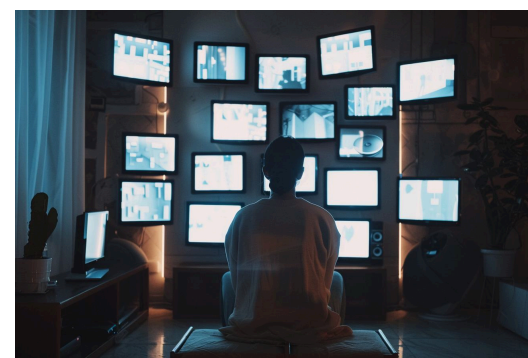
The Proliferation Of Digital Harms

From AI-enabled online child sexual exploitation to the targeted spread of deepfakes, communities are increasingly exposed to harms that outpace existing safeguards.



The Erosion Of Truth

AI can distort religious traditions, spread misinformation, undermine doctrinal integrity, and increase the risk of social tension.



II. The Initiative: The Faith–AI Covenant

The Faith–AI Covenant is a global initiative proposed by the Interfaith Alliance for Safer Communities (IAFSC). It aims to bridge the gap between the "Builders" (technology developers) and the "Stewards" (moral and spiritual leaders).

The Covenant's objective is not to impede innovation, but to elevate it. Faith leaders, as custodians of human values and community trust, must be key architects in the moral governance of AI. This initiative ensures that Artificial Intelligence reflects a dual imperative: advancing technical progress while upholding the sanctity of human dignity.

The Covenant is a voluntary, non-binding initiative designed to complement existing global conventions as well as national laws and regulatory frameworks. It produces a suite of shared principles and voluntary pledges directed toward AI companies, faith institutions, and multi-stakeholder coalitions.

By aligning technical precision with ethical integrity, the Covenant ensures that advanced technologies strengthen rather than undermine human dignity.



What Makes This Initiative Distinct

Multi-faith

Preventive, not reactive

Multi-stakeholder

Global, not regionally constrained

"Science can purify religion from error and superstition; religion can purify science from idolatry and false absolutes. Each can draw the other into a wider world, a world in which both can flourish."

His Holiness Pope John Paul II

III. Strategic Pillars of the Initiative

The Covenant translates ethical intent into technical reality through three operational pillars:



Repository of Truth

Developing verified, consensus-based repositories of religious knowledge across faiths and languages to reduce the risk of producing inaccurate or harmful content.

Defining the Moral Framework

Empowering faith leaders to shape the ethical guardrails required to protect human dignity, agency, and cognitive liberty.



Corporate Responsibility & Safety

Engaging AI companies to co-design safeguards that prevent harms such as online child sexual exploitation, deepfakes, and the manipulated religious content from being disseminated to users and communities.



"Machines were made to serve humankind – not the other way around."

IV. The Covenant's Global Journey



A. Regional Roundtables

The regional roundtables are the foundational building blocks of the Covenant. Their purpose is to move beyond abstract ethics and toward localized, actionable protocols. By hosting these dialogues in seven global hubs, the initiative ensures that the final Covenant is not a "top-down" document, but a global synthesis of cultural wisdom and technical reality.

The roundtables, taking place throughout 2026, will cover **Abu Dhabi, Beijing, Bengaluru, Nairobi, Paris, New York and Singapore.**

“The best way to predict the future is to invent it.”

Alan Kay (Computer Science Pioneer)

IV. The Covenant's Global Journey

B. Roundtable Format and Outcomes

The regional roundtables are designed as outcome-driven sessions focused on translating ethical insight into practical guidance for AI development.

Each gathering will convene 40–60 senior participants from technology companies, major religious traditions, faith-based organizations, civil society, and academia. This deliberate composition ensures that ethical reflection, technical expertise, and community realities are brought into direct dialogue.

The roundtables will aim to achieve three core objectives:

Gather

Local Insights:

Surface region-specific ethical concerns, cultural sensitivities, and technical realities to ensure that AI systems are developed and deployed with consideration of the contexts in which they operate.

Build

Multi-Stakeholder Consensus:

Establish a common language and practical reference points that allow technology leaders and faith communities to align on safety, responsibility, and trust.

Identify

Moral Blind Spots:

Support engineering and product teams in identifying risks that may not be visible within purely technical environments.




IV. The Covenant's Global Journey

C. The Concluding Summit: Abu Dhabi

The Covenant's global journey will culminate in a high-level summit in Abu Dhabi, where the Faith-AI Covenant will be formally adopted. This summit marks the transition from dialogue to implementation, bringing together senior leaders from technology, faith, and policy sectors to endorse and operationalize the Covenant.

The outcome is a set of clear, practical commitments from its signatories, ensuring that the Covenant functions not as a symbolic declaration, but as an operational framework for real-world action.



For AI Leaders

Commitments to integrate ethical safeguards and faith-informed considerations into the design, training, and deployment of AI systems, including mechanisms to prevent high-risk harms before they reach users.

For Faith Leaders

Commitments to integrate ethical safeguards and faith-informed considerations into the design, training, and deployment of AI systems, including mechanisms to prevent high-risk harms before they reach users.

“AI offers enormous potential in improving human capability. It must also seek to protect, preserve and cherish the dignity of the human person.”

Archbishop Justin Welby
(Former Archbishop of Canterbury)

V. Governance

The initiative is guided by an Advisory Steering Group, chaired by Baroness Joanna Shields OBE. Her leadership anchors the initiative in a rigorous focus on safety, drawing on her extensive experience in technology policy and internet safety to drive the Covenant toward meaningful implementation.



Guiding the development of the Covenant, including the incorporation of outcomes from the regional roundtables

Helping bridge conceptual gaps between the technical ambitions of the AI industry and the moral traditions of diverse faith communities

Ensuring that marginalized voices and Global South perspectives remain central to the dialogue, particularly during the regional roundtables.

The Advisory Steering Group, following the Abu Dhabi summit, will oversee the development of a 12-month roadmap designed to monitor and report on the implementation of all summit outcomes, commitments, and pledges. **This structured follow-up supports the transition of the Covenant into tangible global actions.**

“The challenge is immense: to teach our creation this new mind, not only how to compute, but also how to respect; not only how to optimize, but also how to serve; not only how to solve problems, but also how to recognize the mystery.”

VI. Why join this initiative

Participation in the Covenant offers distinct value for both technology and faith leaders.

For AI Leaders

Aligning with the Covenant provides a unique opportunity to build sustained public trust beyond regulatory compliance. This collaborative environment enables engineering teams to uncover cultural and ethical nuances early in development, helping mitigate unforeseen reputational risks. Furthermore, it offers a venue to contribute to a harmonized set of global ethical benchmarks, facilitating more consistent innovation across international markets.

For Faith Leaders

The Covenant will serve as a vital platform for ensuring that AI development remains aligned with human flourishing and the sanctity of the individual. It provides a global voice for faith communities to advocate for the protection of the vulnerable and to address systemic issues like algorithmic bias. Participants help define the necessary boundaries for how technology interacts with sacred traditions, working to preserve religious heritage in an increasingly digital world.



The future of artificial intelligence will not be shaped by technology alone, but by the values we choose to uphold, together.

About the Interfaith Alliance for Safer Communities

The IAFSC was established to empower faith leaders to work for the safety and security of our communities, tackling issues such as child sexual abuse, extremism, displacement and trafficking. We aim to facilitate the building of bridges between key stakeholders including faith communities, NGOs and experts in various domains. We are conscious of the importance of empowering faith leaders, both at the institutional and grass-root levels, with knowledge, and to mobilize them to play a more active role in community safety.



**INTERFAITH
ALLIANCE**
FOR SAFER COMMUNITIES

To learn more about us, visit our website (www.iafsc.org)
or contact us by email (info@iafsc.org)