

Religious Institutions on AI — A companion guide for the moral layer of AI governance

Secular frameworks (NIST, ISO, EU AI Act, OECD, UNESCO) define **what AI must do**. Religious and civilizational traditions answer a different question: **what AI is for**. This guide synthesizes how the world's major religious institutions are responding to AI — starting with Pope Leo XIV's brand-new encyclical *Magnifica Humanitas*, and extending to Catholic Vatican doctrine, Orthodox, Protestant, Islamic, Jewish, and Buddhist thought. Pair this document with my other resources for full stack.

CONTENTS

01	Why This Matters Now	02	Magnifica Humanitas — The New Encyclical
03	Antiqua et Nova — The Doctrinal Foundation	04	Orthodox Christian Thought
05	Protestant & Evangelical Perspectives	06	Islamic Scholarship on AI
07	Jewish Ethical Thought on AI	08	Buddhist Perspectives
09	The Convergence — What All Traditions Agree On	10	How This Complements Secular AI Governance
11	Closing — The Moral Layer Is Not Optional	12	Resources & Sources

1 · Why this matters now

Religious institutions are speaking on AI with unprecedented speed, depth, and clarity. The moral layer is closing the gap that secular frameworks deliberately leave open.

The major AI governance frameworks — NIST AI RMF, ISO/IEC 42001, EU AI Act, OECD AI Principles, UNESCO AI Ethics — share an intentional limitation. They define **what AI systems must do** to be safe, fair, transparent, and accountable. They do not, by design, attempt to define **what AI is ultimately for**, what the human person is, or what civilizational ends technology should serve.

That question lives in the moral, theological, and philosophical traditions of the world's religions. And in the past 16 months, those traditions have begun answering it with extraordinary speed:

- **January 2025** — The Vatican publishes **Antiqua et Nova**, a 117-paragraph doctrinal note from the Dicastery for the Doctrine of the Faith and the Dicastery for Culture and Education.
- **May 2026** — Pope Leo XIV signs his first encyclical, **Magnifica Humanitas**, on the 135th anniversary of Leo XIII's Rerum Novarum. The signature is intentional: he frames AI as the new industrial revolution.
- **2024-2026** — The International Islamic Fiqh Academy, Al-Azhar's Islamic Research Academy, the Muslim World League, and the Assembly of Muslim Jurists of America are all actively developing AI ethics positions grounded in classical jurisprudence.
- **Ongoing** — Jewish, Buddhist, Orthodox, and Protestant institutions contribute distinct philosophical frameworks that address what secular governance cannot.

For consultants, leaders, and policymakers working in AI governance, this is not optional reading. It's the layer that gives the technical frameworks their human meaning.

2 · Magnifica Humanitas — Pope Leo XIV's first encyclical

Signed May 15, 2026 · Published May 25, 2026 · 42,300 words · 5 chapters · The most sweeping religious statement on AI in modern history.

¶ The framing — and why the signing date matters

The full title is "**Magnifica Humanitas: On Safeguarding the Human Person in the Time of Artificial Intelligence.**" Pope Leo XIV signed it on **May 15, 2026 — the 135th anniversary** of the promulgation of Pope Leo XIII's **Rerum Novarum** (1891).

That choice is deliberate. **Rerum Novarum** was the Catholic Church's response to the Industrial Revolution — the document that established Catholic social teaching's positions on labor, capital, dignity of work, and the role of the state. By signing on its anniversary, Leo XIV is explicitly framing **AI as the new industrial revolution** — and Magnifica Humanitas as the church's structural response in the same lineage.

"Technology is not a force antagonistic to humanity, nor is it inherently evil. However, technology is never neutral, because it takes on the characteristics of those who devise, finance, regulate, and use it."

POPE LEO XIV · MAGNIFICA HUMANITAS · PARAPHRASED CORE PREMISE

⚡ The central thesis — "Disarm" AI

The encyclical's most distinctive contribution is the call to **"disarm" AI**. Not to ban it. Not to reject it. To strip it of the mentality of military, economic, and cognitive competition that currently drives its development.

"To disarm means discrediting the assumption that technical power automatically confers the right to govern. To disarm does not mean rejecting technology, but preventing it from dominating humanity."

POPE LEO XIV · MAGNIFICA HUMANITAS · CORE DIRECTIVE

✦ The transhumanism critique

A significant portion of the encyclical addresses **transhumanism and posthumanism** — the philosophical movements that frame human limitations (mortality, embodiment, cognitive bounds) as defects to be overcome through technological enhancement.

Leo XIV's counter-position is direct: **limitations are not defects to be eliminated; they are a constitutive dimension of the human person.** The encyclical's signature line on this point:

"Humanity flourishes not despite limitations, but often through them."

POPE LEO XIV · MAGNIFICA HUMANITAS · ON HUMAN NATURE

▶ The five chapters — what the encyclical actually covers

The full text is divided into five chapters, each developing one dimension of the human-AI relationship:

CH. 01

Safeguarding the Human Person

The core anthropological position: AI must serve, not redefine, what it means to be human.

CH. 02

Promoting Truth

AI's relationship to information, knowledge, and the erosion of shared truth in synthetic-content era.

CH. 03

Dignity of Work

The direct line to Rerum Novarum — labor displacement, the meaning of work, the worker's moral standing.

CH. 04

Social Justice

Access, equity, the global digital divide, AI's role in concentrating or distributing power.

CH. 05

Peace

Autonomous weapons, AI-enabled warfare, the disarmament of AI systems from military logic.

3 · Antiqua et Nova — The doctrinal foundation

January 28, 2025 · 117 paragraphs · Co-issued by the Dicastery for the Doctrine of the Faith and the Dicastery for Culture and Education

¶ What Antiqua et Nova set up

Antiqua et Nova ("Ancient and new") preceded Magnifica Humanitas by 16 months. It's a doctrinal note — less authoritative than an encyclical, but doctrinally foundational. It addresses "the relationship between artificial intelligence and human intelligence" and offers reflections on "the anthropological and ethical challenges raised by AI."

The document covers AI's implications across **education, economy, labour, health, human and international relations, and war** — providing the building blocks Leo XIV later draws into Magnifica Humanitas.

! The central philosophical move

Antiqua et Nova draws a clear philosophical line between what machines do and what the human mind is. The document's crucial distinction:

"AI performs some tasks exceptionally well, but can only mimic human qualities such as empathy and understanding. Artificial intelligence should be a tool to complement and not replace the richness of human intelligence."

ANTIQUA ET NOVA · VATICAN DOCTRINAL NOTE · JANUARY 2025

✦ Why this matters for governance

Antiqua et Nova made the distinction that secular regulation often skips: **simulation is not possession**. An AI system that produces empathy-shaped outputs has not become empathetic; it has become better at producing the outputs. The moral, legal, and relational consequences of this distinction propagate through every downstream policy question — from accountability to consent to the question of who bears responsibility for AI-mediated harm.

Magnifica Humanitas builds directly on this foundation, extending the philosophical claim into a comprehensive social and political vision.

4 · Orthodox Christian thought on AI

Engagement led by the Hellenic College Holy Cross School of Theology and broader Orthodox bioethics scholarship.

Eastern Orthodox

BIOETHICS-ROOTED · THEOLOGICAL ANTHROPOLOGY

LEAD INSTITUTIONS

Huffington Ecumenical Institute (Hellenic College Holy Cross), Public Orthodoxy platform, individual Orthodox bioethics scholars. Engagement formally began with February 2024 theological panels on AI's ethical repercussions for churches.

CORE LENS

Orthodox engagement with AI flows from the church's existing **theological anthropology** — the human as **imago Dei** (image of God), embodied, communal, oriented toward theosis (union with God). AI is evaluated against whether it serves or distorts that anthropology.

PRACTICAL CONCERNS

Bioethical questions around AI in medicine, the impact of AI on liturgical and pastoral practice, the question of whether AI-mediated relationships violate the church's understanding of personhood. Less programmatic than the Catholic engagement; more theological and pastoral.

5 · Protestant & Evangelical perspectives

Decentralized response anchored in scriptural anthropology — Imago Dei, stewardship, and human flourishing.

Protestant & Evangelical

SCRIPTURAL ANTHROPOLOGY · STEWARDSHIP

LEAD VOICES

C.S. Lewis Institute, Wheaton College ethics faculty, evangelical think tanks. Less centralized than Catholic or Orthodox response — characteristic of Protestant institutional structure.

THREE ANCHORING CONCEPTS

Imago Dei (humans bear God's image — non-transferable to machines), **Stewardship** (humans are responsible for creation and the technologies built within it), **Human Flourishing** (the biblical good life as the metric against which AI is evaluated).

DISTINCTIVE CONTRIBUTION

Protestant engagement frequently emphasizes **the mission of the church in the AI age** — how Christian witness, formation, and pastoral care should adapt without surrendering theological core. Less programmatic policy, more formational guidance for individuals and congregations.

6 · Islamic scholarship on AI

Built on classical jurisprudence — *maslaha*, *maqasid al-shari'ah* — applied to AI ethics through institutional *fiqh* councils.

Islamic

FIQH-BASED · VIRTUE ETHICS · PLURALIST BENCHMARKING

LEAD INSTITUTIONS

International Islamic Fiqh Academy (IIFA), Al-Azhar Islamic Research Academy, Muslim World League, Assembly of Muslim Jurists of America (AMJA). AMJA centered its 2026 Imam's Conference on AI ethics.

CORE JURISPRUDENTIAL FRAME

AI is evaluated against **maqasid al-shari'ah** (the higher objectives of Islamic law) and **maslaha** (the common good). The five universals of maqasid — preservation of life, intellect, religion, lineage, property — function as a checklist for whether a given AI system serves or threatens them.

DISTINCTIVE CONTRIBUTION

Islamic AI ethics proposes a **virtue theory** grounded in Islamic sources — cultivating virtues in both individuals and collectivities. The emphasis is on "**from means to meanings**" and "**from minds to hearts**". AI is not just a means to ends; it shapes the inner life of those who use it.

PRACTICAL PROJECTS

Co-designed Islamic AI ethics frameworks (UK Muslim community partnerships), pluralist ethical benchmarking proposals (recognizing AI must be evaluated against multiple ethical traditions, not just secular Western liberalism), debate around AI in religious guidance ("can you fatwah-shop with AI?").

7 · Jewish ethical thought on AI

Anchored in collective responsibility, the concept of Lashon Hara, and the long Jewish tradition of textual reasoning.

Jewish

COLLECTIVE RESPONSIBILITY · TEXTUAL REASONING

CORE LENS

Jewish AI ethics frequently challenges purely rights-based individualism with the framework of **collective responsibility**. The question is not only "what may I do with AI?" but "what does the community owe each other in an AI-mediated world?"

LASHON HARA

The concept of **Lashon Hara** — the prohibition on hateful true speech about a person — has been applied with surprising force to AI ethics. The framework challenges rights-based views of online speech: the question is not "is it true?" but "does sharing it harm a person who cannot meaningfully respond?" Highly relevant to AI-generated content and large-scale automated discourse.

METHODOLOGICAL STRENGTH

Jewish ethical reasoning has 2,000+ years of practice in **applying ancient texts to genuinely novel situations**. This is the closest pre-modern analog to what AI ethics now requires: the structured, communal, textually-grounded extension of moral principles to circumstances the original authors could not have imagined.

8 · Buddhist perspectives on AI

Reducing suffering, the centrality of intention, and the interdependence framework.

Buddhist

DUKKHA · CETANA · INTERDEPENDENCE

REDUCING DUKKHA

Dukkha — usually translated as suffering, more precisely as **dissatisfaction or unease** — is the first noble truth. Buddhist AI ethics asks: does this AI system reduce suffering for those who interact with it, including those affected indirectly? Attention-engineered systems that intentionally amplify dukkha (compulsive use, anxiety, fragmented attention) face particular Buddhist critique.

CETANA – INTENTION

Cetana (intention) is morally weight-bearing in Buddhist ethics. The intention behind a system's design — what the designers actually meant to optimize for — is a primary object of evaluation, often more so than the system's immediate effects. AI systems designed for engagement rather than human flourishing carry the moral weight of that design intention regardless of stated mission.

INTERDEPENDENCE

The Buddhist principle that all phenomena arise in dependence on conditions — **pratityasamutpada** — provides a natural framework for thinking about AI systems as embedded in vast webs of data, labor, energy, supply chains, and downstream effects. No AI is autonomous; every AI is a node in a system of conditions, and ethical responsibility extends across that network.

PRACTICAL APPLICATIONS

Buddhist-values-based applications (meditation tools, attention-awareness apps) explicitly designed to reduce dukkha rather than maximize engagement — offered as a counter-example to mainstream attention-capture design.

9 · The convergence — what all traditions agree on

Despite radically different theological frames, six principles appear across every tradition examined.

Comparative analysis across the traditions surveyed reveals striking convergence around six principles. Each tradition arrives at them through its own theological grammar, but the destinations align. This convergence is what makes inter-tradition AI ethics possible — and what gives religious institutions a credible voice in pluralistic policy debates.

PRINCIPLE	HOW EACH TRADITION FRAMES IT
Human dignity & agency	Catholic: Imago Dei + safeguarding human person · Orthodox: theosis-oriented embodiment · Protestant: Imago Dei + flourishing · Islamic: hifz al-aql (preservation of intellect) · Jewish: Tzelem Elohim · Buddhist: each being's capacity for awakening
Fairness & justice	Catholic: social justice (Magnifica Humanitas Ch. 4) · Islamic: maslaha (common good) · Jewish: tzedek (justice) · Buddhist: extension of compassion across interdependent networks
Accountability & stewardship	Catholic: technology takes characteristics of those who design it · Protestant: stewardship of creation · Islamic: amanah (trust) · Jewish: collective responsibility · Buddhist: cetana (intention bears moral weight)
Non-harm	Catholic: dignity of the person · Islamic: la darar wa la dirar (no harm) · Jewish: Lashon Hara · Buddhist: first precept (non-harming) extended to system design
Compassion	Catholic: caritas · Orthodox: philanthropia · Protestant: love of neighbor · Islamic: rahma · Jewish: chesed · Buddhist: karuna
Truthfulness & transparency	Catholic: promoting truth (Magnifica Humanitas Ch. 2) · Islamic: sidq · Jewish: emet · Buddhist: right speech (samma vaca)

10 · How this complements secular AI governance

The full stack: technical frameworks define obligation; religious traditions define meaning. Both layers are required.

For consultants, leaders, and policy makers, the practical insight is this: secular technical frameworks and religious moral traditions are not in competition. They occupy **different layers of the same stack**. Pair them.

SECULAR · WHAT

NIST AI RMF 1.0 · Risk management functions (govern, map, measure, manage). **Answers:** what controls does my organization need?

SECULAR · WHAT

EU AI Act · Risk-tiered binding regulation. **Answers:** what is legally required?

SECULAR · WHAT

ISO/IEC 42001 + 24027 · Management systems + bias auditing. **Answers:** what is the operational standard?

SECULAR · WHAT

OECD AI Principles + UNESCO Ethics · Internationally agreed values. **Answers:** what is the global norm?

RELIGIOUS · WHY

Magnifica Humanitas + Antiqua et Nova · Catholic anthropology and social teaching applied to AI. **Answers:** what is the human person, and what is AI for in a human life?

RELIGIOUS · WHY

Islamic maqasid framework · The five universals as a checklist. **Answers:** does this AI serve or threaten the goods that constitute a worthwhile human life?

RELIGIOUS · WHY

Jewish collective responsibility + Lashon Hara · Communal ethics applied to AI-mediated speech and decision. **Answers:** what do we owe each other in an AI-saturated world?

RELIGIOUS · WHY

Buddhist dukkha + cetana + interdependence · Suffering-reduction and intention-evaluation framework. **Answers:** does this AI reduce or amplify dissatisfaction, and what was it actually designed for?

11 · Closing — The moral layer is not optional

The position this guide leaves you with.

Technical safety alone is not safety.

You can build an AI system that passes every NIST control, every ISO audit, every EU AI Act compliance test — and still build a system that **fails the human person**.

Technical frameworks define the floor. They do not define the ceiling.

The world's religious traditions, with their 2,000-to-5,000 years of practice in asking what makes a human life worth living, have begun answering the question secular governance leaves open: **what is AI for?** The answers from Magnifica Humanitas, Antiqua et Nova, Islamic maqasid, Jewish collective responsibility, and Buddhist dukkha-reduction are not interchangeable — but they converge on a position that matters:

AI must serve the flourishing of human persons — not the concentration of power, not the elimination of human limitations, not the optimization of engagement, not the maximization of capital efficiency. Anything that fails that test is a technical success and a civilizational failure.

For consultants and leaders working in AI governance, the practical takeaway is direct: **read both layers of the stack**. NIST and the Vatican. ISO and the maqasid. EU AI Act and Lashon Hara. The technical frameworks tell you what to build. The religious traditions tell you what you are building for.

— Jorge Martín Jarrín Žak · Companion to the Unconscious Bias in AI Playbook · June 2026

12 · Resources & sources

→ Catholic / Vatican

Primary sources:

- [Magnifica Humanitas — full encyclical text \(vatican.va\)](#)
- [Antiqua et Nova — full text \(vatican.va, January 2025\)](#)

Secondary analysis:

- [Pope Leo's Magnifica Humanitas — Vatican News overview](#)
- [TIME — Pope Leo Uses First Major Papal Text to Warn About Dangers of AI](#)
- [USCCB — Pope Leo urges world to 'disarm' AI](#)
- [Catholic Outlook — Magnifica Humanitas comprehensive reading](#)
- [Transparency Coalition — Guide to Pope Leo XIV's encyclical on AI](#)
- [National Catholic Bioethics Center — The Ethics of AI in Antiqua et Nova](#)
- [Markkula Center — Antiqua et Nova: Anthropology and Social Ethics](#)
- [Antiqua et Nova — Wikipedia overview](#)

→ Orthodox & Protestant

- [Public Orthodoxy — Artificial Intelligence: Bioethical Considerations and Limitations](#)
- [Juicy Ecumenism — AI: Christian Churches Respond](#)
- [C.S. Lewis Institute — Artificial Intelligence: Imago Dei, Stewardship, and the Mission of the Church](#)

→ Islamic

- [Artificial Intelligence and the Islamic Theology of Technology \(MDPI\)](#)
- [Co-Designing Islamic AI Ethics: Insights from the UK Muslim Community \(CHI 2026\)](#)
- [Artificial Intelligence in Islamic Ethics: Towards Pluralist Ethical Benchmarking for AI \(Springer\)](#)
- [MuslimMatters — Faith and Algorithms: From an Ethical Framework for Islamic AI to Practical Application](#)

→ Comparative & multi-tradition

- [AI and Faith — Religious Ethics in the Age of AI and Robotics](#)
- [Religious ethics in the age of AI: comparative faith-based approaches to AI governance](#) (Springer, 2026).
- [Artificial Intelligence in Religious Education: Ethical, Pedagogical, and Theological Perspectives](#) (MDPI).

Religious Institutions on AI — A Companion Guide · v1.0 · June 2026 · Jorge Martín Jarrín Žak

Pairs with: [unconscious-bias-in-ai.html](#) · Both live in [astrylis-memoirs/docs/playbooks/](#)

Print to PDF: ⌘P → Save as PDF.