Anglican Diocese of
BENDIGO

ARTIFICIAL INTELLIGENCE (AI) POLICY

Version: 1.0	*Date of Issue:* September 2026
Contact: Registrar	*Review Date:* September 2027

Introduction

As Artificial Intelligence (AI) becomes more prevalent in our society, workplaces, and places of ministry, the Anglican Diocese of Bendigo is committed to using this technology ethically, safely, and in a manner that upholds Christian values of compassion, justice, dignity, and inclusion.

This policy provides a framework for the ethical use, development, and management of AI within the Diocese. It supports the responsible use of AI tools and systems in ways that align with our mission to uphold the teachings of Christ and serve the community.

Theological Framework

Temple's measure (1942)

Christianity and Social Order — three principles for any social question:

Freedom	Fellowship	Service
The God-given room to become who God calls you to be. The labeller in Logos is a person, not "a human in the loop."	We are made for one another. An intelligence that gathers all power to two centres and leaves the family as raw material sins against fellowship — however efficient.	Power is for the good of the whole. The test of any AI is not what it can do, but whom it serves — and whether the last are served first.

Source: AI, the Gig Economy, and the Anglican Word: The Rt Rev'd Andrew Doyle, Anglican Communion Science Commission

Our policy aims to:

The aims of this policy are to:

- Promote fair, transparent, and human-centred AI use (as understood through Temple's measure of Freedom, Fellowship and Service)
- Safeguard the dignity, privacy, and rights of individuals
- Provide guidance to staff and leadership in using AI tools responsibly
- Support innovation while ensuring AI aligns with the Diocese's values and governance structures
- Reduce harm and prevent misuse of AI technologies, especially regarding vulnerable populations

Definitions

Artificial Intelligence (AI):

The simulation of human intelligence processes by machines, especially computer systems. These processes include learning, reasoning, and self-correction.

Personally Identifiable Information (PII):

Any data that can be used to identify a specific individual, either on its own or when combined with other information. Common examples include names, email addresses, phone numbers, dates of birth and government-issued ID numbers.

Policy Scope & Coverage

This policy applies to all employees, clergy, licenced lay ministers and volunteers within the Anglican Diocese of Bendigo who use, manage, or make decisions related to AI systems.

1. Guiding Principles:

The Diocese adopts the *Australian Government's AI Ethics Principles*[1], underpinned by Christian ethical values. These principles guide the Diocese in ensuring that all AI use is safe, secure and reliable:

- **Human, societal and environmental wellbeing:** AI systems should benefit individuals, society and the environment.
- **Human-centred values:** AI systems should respect human rights, diversity, and the autonomy of individuals.
- **Fairness:** AI systems should be inclusive and accessible, and should not involve or result in unfair discrimination against individuals, communities or groups.
- **Privacy protection and security:** AI systems should respect and uphold privacy rights and data protection, and ensure the security of data.
- **Reliability and safety:** AI systems should reliably operate in accordance with their intended purpose.
- **Transparency and explainability:** There should be transparency and responsible disclosure so people can understand when they are being significantly impacted by AI, and can find out when an AI system is engaging with them.
- **Contestability:** When an AI system significantly impacts a person, community, group or environment, there should be a timely process to allow people to challenge the use or outcomes of the AI system.
- **Accountability:** People responsible for the different phases of the AI system lifecycle should be identifiable and accountable for the outcomes of the AI systems, and human oversight of AI systems should be enabled.

[1] Australia's AI Ethics Principles | Department of Industry Science and Resources *accessed 18/3/2026*

2. *CARE Framework for Ethical Use of AI:*

The Diocese encourages the use of the *CARE Framework*[2] *(Appendix 1)* to support ethical decision-making when considering AI use:

- **Consequences:** anticipate the positive and negative outcomes of deploying AI in your work
- **Accountability:** consider who is accountable if the AI makes a mistake
- **Responsibility:** consider how AI aligns with our mission to serve our diocese, its parishes and its people
- **Explainability:** people should be able to understand why and how AI systems make decisions

This structured approach helps users evaluate AI's impact prior to use.

Consequences: Understanding the Ripple Effects

Every action has consequences, and AI is no exception. The first step in ethical decision-making is to anticipate the positive and negative outcomes of deploying AI in your work.

Questions to ask:

- What benefits will this AI system deliver (e.g., efficiency, improved services)?
- Could it inadvertently cause harm, such as amplifying biases or excluding vulnerable groups?

How might these impacts evolve over time?

- By weighing these consequences, you can make informed choices that maximise benefits while minimising harm.

Accountability: Who Takes Ownership?

AI systems often operate in a space where responsibility can seem unclear. Accountability ensures someone is ready to step in when things go wrong.

Questions to ask:

- Who is accountable if the AI makes a mistake—the developer, the organisation, or both?
- Are there safeguards to review and correct errors?

Why it matters:

- Establishing clear accountability fosters trust and ensures swift action when issues arise.

[2] CARE'ing about AI: an ethical AI framework|Digital Transformation Hub *accessed 23/6/2026*

Responsibility: Ethical Obligations to Stakeholders

Responsibility is about your organisation's duty of care towards those affected by AI. NFPs, in particular, must consider how AI aligns with their mission to serve and protect their stakeholders.

Questions to ask:

- Does this use of AI uphold fairness, transparency, and equity?
- How might it impact vulnerable groups, and what safeguards are in place to prevent harm?

Why it matters:

- Taking responsibility for ethical considerations strengthens public confidence in your organisation.

Explainability: Building Transparency and Trust

For AI to be trusted, its decisions must be explainable. Stakeholders should be able to understand why and how AI systems make decisions.

Questions to ask:

- Can the reasoning behind AI decisions be easily communicated?
- How transparent are the data and processes used by the AI system?

Why it matters:

- Transparent systems foster trust and reduce misunderstandings, particularly in high-stakes areas like resource allocation or donor engagement.

3. *Acceptable Uses of AI*

Examples of where AI may be used, having first applied the CARE Framework:

- Assist in drafting material (with human review and theological oversight).
- Enhance administrative efficiency and personal productivity (e.g., summarising reports, drafting routine correspondence, formatting documents, organizing personal task lists)
- Summarisation (e.g. summarising complex documents on non-sensitive topics, recording and summarising transcripts of meetings).
- Improve accessibility (e.g., generating alternative formats for communication, language translation).
- Support creative design for events, newsletters, and digital communication.
- Aid in research or background preparation for training, teaching, and strategic initiatives from publicly available data.

4. *Prohibited Uses of AI*

AI must not be used to:

- Handle confidential data (e.g. personally identifiable information, sensitive pastoral details, financial data, personnel records, or confidential meeting minutes). Third-party AI platforms (e.g. free versions of ChatGPT, Claude, Copilot) must **never** be used to process personally identifiable information.
- Make pastoral decisions (e.g., counselling advice, providing spiritual advice, theological declarations) without human discernment.
- Generate confidential reports involving personal data (e.g., investigations, personnel matters).
- Replace human relationship and discernment in sacramental, spiritual, or pastoral contexts.
- Draft legal documents without professional legal oversight.
- Create or distribute content that may mislead, offend, or misrepresent Christian teaching or Diocesan policy.
- Draft sensitive communications (e.g., responses to complaints, disciplinary matters, major announcements) without explicit approval and rigorous human oversight.
- Create sacred art, liturgy, or prayers without explicit theological and aesthetic review and approval.
- Facilitate automated decision-making affecting individuals (e.g. volunteer screening, resource allocation) without robust human review and appeal processes.

5. Approved AI Tools

Approved AI Tools		
User Group	*Tool*	*Parameters*
Registry Office	Microsoft Co-Pilot	Within Microsoft 365 Tenant – Bendigo Anglican network
Active Clergy	Not prescribed	

6. Reporting Concerns

Suspected breaches of this policy, data privacy incidents related to AI, or concerns about inappropriate AI use should be reported to the Registrar.

This policy was created using resources from the Anglican Diocese of Bunbury and the Anglican Diocese of Melbourne.

Policy Responsibility

- Bishop in Council

Related Documents: Internal

- Cyber Security Policy
- Privacy Policy
- Safe Communication Policy

Related Documents: External

- Australian Government: Australia's AI Ethics Principles
- Digital Transformations: CARE Framework for Ethical Use of AI

Approved by Bishop in Council:

8 September 2026

Policy History	
Policy created	September 2026
Policy reviewed	
Policy amended	

Anglican Diocese of
BENDIGO

Appendix 1: CARE'ing about AI: an ethical AI framework

Developing a framework for ethical use of AI in the not-for-profit sector

Source: Digital Transformation: CARE'ing about AI: an ethical AI framework|Digital Transformation Hub

To support ethical decision-making in this complex landscape, we've developed the CARE framework: Consequences, Accountability, Responsibility, and Explainability. This structured approach helps leaders evaluate AI's impact and ensure its use aligns with their organisational values.

Why CARE matters

The CARE framework is more than an ethical checklist; it's a practical tool for navigating the complexities of AI. By embedding these principles into your decision-making processes, your organisation can harness the potential of AI while upholding the values that define the NFP sector.

AI is neither inherently good nor bad—it's a tool. The real question is how we, as human stewards, choose to wield it. Through frameworks like CARE, NFPs can make informed, ethical decisions that ensure AI serves the greater good.

The CARE Framework: Four pillars of ethical AI

1. Consequences: Understanding the Ripple Effects

Every action has consequences, and AI is no exception. The first step in ethical decision-making is to anticipate the positive and negative outcomes of deploying AI in your work.

Questions to ask:

- What benefits will this AI system deliver (e.g., efficiency, improved services)?
- Could it inadvertently cause harm, such as amplifying biases or excluding vulnerable groups?

How might these impacts evolve over time?

- By weighing these consequences, you can make informed choices that maximise benefits while minimising harm.

2. Accountability: Who Takes Ownership?

AI systems often operate in a space where responsibility can seem unclear. Accountability ensures someone is ready to step in when things go wrong.

Questions to ask:

- Who is accountable if the AI makes a mistake—the developer, the organisation, or both?
- Are there safeguards to review and correct errors?

Why it matters:

- Establishing clear accountability fosters trust and ensures swift action when issues arise.

3. Responsibility: Ethical Obligations to Stakeholders

Responsibility is about your organisation's duty of care towards those affected by AI. NFPs, in particular, must consider how AI aligns with their mission to serve and protect their stakeholders.

Questions to ask:

- Does this use of AI uphold fairness, transparency, and equity?
- How might it impact vulnerable groups, and what safeguards are in place to prevent harm?

Why it matters:

- Taking responsibility for ethical considerations strengthens public confidence in your organisation.

4. Explainability: Building Transparency and Trust

For AI to be trusted, its decisions must be explainable. Stakeholders should be able to understand why and how AI systems make decisions.

Questions to ask:

- Can the reasoning behind AI decisions be easily communicated?
- How transparent are the data and processes used by the AI system?

Why it matters:

- Transparent systems foster trust and reduce misunderstandings, particularly in high-stakes areas like resource allocation or donor engagement.

Applying the CARE framework in real life

In a parish context: AI might be used to draft a weekly pew sheet, newsletter item, or social media post about upcoming services and community events. The CARE framework would prompt the parish to consider the consequences of inaccurate information, ensure a nominated staff member or volunteer is accountable for checking details before publication, take responsibility for protecting pastoral sensitivities and avoiding personal information, and make clear to readers that the material has been reviewed by a person before being shared.

In a parish context: AI could be used to help prepare a draft roster for Sunday ministries, such as readers, intercessors, welcomers, and morning tea volunteers. Applying CARE would mean checking the consequences for volunteer workload and inclusion, assigning accountability to the parish administrator or ministry coordinator rather than the AI tool, taking responsibility for ensuring pastoral care needs and availability are respected, and explaining to volunteers that the roster is only a draft and can be adjusted through human conversation.

In a diocesan administration context: AI might be used to summarise non-confidential meeting papers or draft routine correspondence. Applying CARE would mean considering whether the summary could omit important nuance, ensuring a diocesan officer remains accountable for accuracy, taking responsibility for excluding confidential or personally identifiable information, and explaining that AI has assisted with drafting only, not decision-making.

In a worship planning context: AI could be used to suggest hymn themes, intercessory prayer prompts, or seasonal service resources. Applying CARE would mean checking the theological and liturgical consequences of the suggestions, assigning accountability to the incumbent or authorised worship leader, taking responsibility for ensuring resources are pastorally appropriate and consistent with Anglican worship, and explaining that AI-generated material must be reviewed before inclusion in public worship.

In a parish council or governance context: AI could be used to identify themes in non-confidential feedback from parishioners about strategic priorities. Applying CARE would mean considering whether minority perspectives might be overlooked, ensuring council members remain accountable for interpretation and decisions, taking responsibility for privacy and respectful consultation, and explaining that AI can support reflection but cannot determine the direction of parish life or ministry.

Disclosure: AI was used to develop these examples