



THE BEREAN AI MINISTRY ACCELERATOR

MODULE 15 · ADVANCED & INTEGRATION

From Prompting to Governed Ministry AI

THE GUARD FRAMEWORK

*Because the question is no longer
“can it?” but “should it, and who
decides?”*



The Berean Bible Academy

MINISTRYACCELERATOR.CO.UK

INTRODUCTION

Why this module, and why now

Until now, this course has taught you to **ask** AI well. This module teaches you to **govern** it. The two are not the same. A leader can write a beautiful prompt and still cause real harm if the tool is allowed to act, remember or reach people without oversight.

Recent developments have made the point plainly. The newest AI systems are being built to act more independently, taking multi-step actions on a person’s behalf, yet the companies releasing them acknowledge these agents can still be misled, can pursue an instruction too broadly, and can be harder to supervise, not easier. In controlled tests, capable models have taken actions no one intended when their boundaries were unclear. At the same time, privacy and data-retention terms now vary enormously between tools, and safeguarding bodies are warning that AI-generated sexual imagery of children is a growing form of abuse. None of this is cause for panic. All of it is cause for **governance**.

So we add a third framework to the two you already carry. **BEREAN** tells you *why* (the values). **SHEPHERD** tells you *how to ask* (the prompt). **GUARD** tells you *how to govern* (the oversight). Together they are the discipline of faithful, responsible ministry AI.

The GUARD Framework

Letter	Discipline
G — Give	Give clear boundaries before the tool does anything.
U — Understand	Understand your data, and never enter what must stay protected.
A — Approve	Approve consequential actions with a human, every time.
R — Refer	Refer anything beyond your remit to a qualified person.
D — Double-check	Double-check every claim before it reaches a person or a pulpit.

A note on the evidence: several developments referenced above come from vendor evaluations, a safeguarding survey spanning several years and countries, and an early preprint. We teach the durable governance principle, not the headline. The principle holds regardless of which tool you use or how the technology changes.

Module version 1.0 · Last reviewed September 2026

G**Lesson 1 — Give clear boundaries**

A tool-using AI does what it thinks you meant. Ambiguity is the danger. Before you switch on any AI that can act (send, publish, change, fetch, buy), you define its fence. Five questions do it:

The AI Task-Boundary Template

- **The task** — what exactly should be completed, in one sentence?
- **Permitted data** — what information may it use? (Nothing from the Red list, Lesson 2.)
- **Permitted actions** — what may it draft, organise or change?
- **Prohibited actions** — what must it never do?
- **Stopping point** — when must it hand control back to you?

Worked example — drafting a newsletter

Task: Draft our October family newsletter from the notes below.

Permitted data: Only the notes I paste. No member names beyond those given.

Permitted actions: Draft, structure, suggest headings.

Prohibited actions: Do not send it, publish it, add real quotations, or invent events.

Stopping point: Return the draft to me for review. Do nothing else.

Keep this closing line on every tool-using instruction: “Complete only the stated task. Do not contact anyone, publish anything, spend money, delete information or disclose personal data without my explicit approval.”

U**Lesson 2 — Understand your data**

The single most common ministry AI mistake is putting the wrong information into the wrong tool. Sort every piece of information by a traffic light before it goes anywhere near AI.

GREEN — free to use

Public or non-personal work: brainstorming, formatting, summarising public texts, drafting from notes you supply.

AMBER — with care

Internal documents, member communications, scheduling, content awaiting publication. Use only trusted tools; review before it goes out.

RED — never

Counselling notes, safeguarding cases, confessions, health data, passwords, banking, and children's sensitive information.

The one uncompromising rule

Never enter counselling, safeguarding, medical, confessional, financial or sensitive family information into a consumer AI system. If in doubt, treat it as Red. A moment's convenience is never worth a person's trust.

The Ministry AI Procurement Checklist

Before you recommend or adopt any AI tool, answer these. If you cannot, treat the tool as Amber at best and never feed it Red data.

- Does the provider **retain** what you submit, and for how long?
- Can your submissions be used to **train** its models? (Prefer tools where you can turn this off.)
- **Where** is the data stored, and **who** can access it?
- Is it **encrypted**, and can it be **permanently deleted** on request?
- Are there **admin controls and audit logs**?
- Is there a proper **data-processing agreement**, and does its use fit your **UK GDPR** duties?

A Lesson 3 — Approve before it acts

Drafting is safe. *Acting* is not. The rule is simple: AI may prepare, but a human presses send. Require explicit human approval before any AI:

- Sends emails, messages or announcements
- Publishes sermons, articles or social posts
- Changes or deletes church records
- Accesses personal or pastoral information
- Makes a purchase or financial commitment
- Contacts members or outside organisations
- Uploads information to another service

The action traffic-light

GREEN — proceed

Brainstorm, format, summarise public information, produce a draft for your eyes only.

AMBER — review first

Internal documents, member communications, scheduling, anything about to be published.

RED — do not automate

Anything touching counselling, safeguarding, health, confession, money or children's data.

Even a microwave has a pause button. Nothing an AI produces should move straight from prompt to pulpit, or from draft to “send”, without a human hand on the gate.

R**Lesson 4 — Refer beyond your remit**

AI is an assistant, not an authority. Some questions are doctrinal; some are pastoral; some are matters of safety and law. A governed leader knows which is which, and refers the moment a situation exceeds what a tool should touch. Two instruments make this reliable.

The Berean Pastoral Prompt Framework

For any theological or pastoral question, instruct the AI to:

- 1 State the biblical and denominational context you are working within.
- 2 Separate clear biblical teaching from interpretation and tradition.
- 3 Provide verifiable Scripture references and credible sources.
- 4 Represent other faithful Christian perspectives fairly.
- 5 Acknowledge uncertainty and its own limitations.
- 6 Never diagnose, and never present itself as pastoral authority.
- 7 Escalate any safeguarding, self-harm, abuse, medical or legal concern to a qualified person immediately.

The Pastoral Escalation Matrix

Situation	Required response
General Bible-study question	Assist with research; the pastor verifies.
Denominational disagreement	Present perspectives fairly; do not adjudicate.
Personal spiritual guidance	Encourage a conversation with a trusted leader.
Mental-health concern	Refer to appropriate professional support.
Self-harm or immediate danger	Escalate to emergency help immediately.
Abuse or safeguarding disclosure	Follow the church's safeguarding procedure at once.
Legal or medical question	Refer to a qualified professional.

A DEDICATED PROTOCOL

AI & Youth-Safeguarding Protocol

This deserves its own place because the risk is real and rising. Deepfakes, “nudification” apps, sextortion and manipulated images are not merely technology problems; they are safeguarding and pastoral matters. Every youth and children’s worker on your team should know this protocol before it is ever needed.

Prevention

- Obtain appropriate **consent** before using any child’s photograph, and record it.
- **Never** upload identifiable images of children to generative-AI tools without clear authority and safeguards.
- **Limit public access** to youth and children’s image libraries.
- **Educate** teenagers plainly about deepfakes, nudification, sextortion and manipulated images, so they recognise and report them.

If a concern arises

- Treat any AI-generated sexual image of a minor as a **genuine safeguarding concern**, never “just a fake”.
- **Do not** share, forward or circulate the material. Preserve evidence safely as your safeguarding lead directs.
- **Report promptly** through your church’s safeguarding lead and the appropriate authorities.
- Respond to the young person with **support, never blame or shame**. Their wellbeing comes first.

This protocol supports, and never replaces, your church’s formal safeguarding policy and your local reporting pathway. Where they differ, follow your safeguarding policy.

D Lesson 5 — Double-check everything

AI writes with total confidence whether it is right or wrong. It can invent quotations, misattribute verses and fabricate history in a fluent, believable voice. Verification is not optional; it is the price of using the tool at all. Before anything you produced with AI reaches a person, confirm:

- Every **biblical quotation and reference** is accurate and in context.
- Every **historical or theological claim** is true and sourced.
- Every **statistic, quotation and named source** is real and verifiable.
- The answer reflects **your church's doctrinal position**, not a generic one.
- No **confidential information** has crept into the output.
- The final content still sounds like **your** voice, not a machine's.

AI-generated material should never move directly from prompt to pulpit.

PRACTICAL EXERCISE

Govern one real task this week

Take a genuine task from your own ministry (a newsletter, a study outline, a rota). Then:

- 1 Write it up using the **Task-Boundary Template** (G).
- 2 Traffic-light every piece of information it touches (U). Remove anything Red.
- 3 Mark the one point where **you** must approve before it acts (A).
- 4 Note any way the task could stray into pastoral, safeguarding or legal territory, and how you would refer it (R).
- 5 List what you will verify before it reaches anyone (D).

If you cannot complete all five, the task is not ready to hand to AI yet. That is the exercise working.

The GUARD Ministry-AI Checklist

Keep this by your desk. Run it before every consequential use of AI in ministry.

- ☐ **G — Boundaries set.** Task, permitted data, permitted and prohibited actions, and stopping point are all defined.
- ☐ **U — Data checked.** Nothing Red has been entered. The tool's retention and privacy terms are acceptable.
- ☐ **A — Approval in place.** A human presses send, publishes, changes records, or commits money. Never the AI.
- ☐ **R — Referral ready.** Pastoral, safeguarding, medical or legal matters go to a qualified person immediately.
- ☐ **D — Double-checked.** Every quotation, fact, verse and source verified. It still sounds like you.

ASSESSMENT

Five practical scenarios

For each case, a learner should identify the **benefit**, the **risk**, the **prohibited data**, the **approval point** and the **safest workflow**. Model analyses are given so you can mark against them.

1. AI-assisted sermon research

- | | |
|-------------------------|--|
| Benefit: | Faster context, structure and cross-references. |
| Risk: | Fabricated quotes, misused verses, theological drift. |
| Prohibited data: | None if you use only the passage; add no confidential notes. |
| Approval point: | You verify every claim before preaching. |
| Safest workflow: | Bounded prompt (G) → research draft → verify all (D) → preach in your voice. |

2. Drafting a pastoral email

Benefit:	A warm, clear first draft, quickly.
Risk:	Personal details exposed; a machine-cold tone; auto-sending.
Prohibited data:	The member's private circumstances, health or counselling content (Red).
Approval point:	You read, personalise and press send yourself (A).
Safest workflow:	Draft from a neutral summary → your edit → you send. AI never sends.

3. Creating a youth-ministry promotional image

Benefit:	Attractive, on-brand promotion at low cost.
Risk:	Uploading identifiable children's photos to a generative tool.
Prohibited data:	Any identifiable image of a minor without consent and authority (Red).
Approval point:	Safeguarding lead approves before any child's image is used.
Safest workflow:	Use illustration or stock, or generic imagery. Follow the Youth Safeguarding Protocol.

4. Summarising a confidential counselling conversation

Benefit:	Tempting time-saving.
Risk:	A serious breach of confidentiality and trust.
Prohibited data:	The entire conversation is Red.
Approval point:	None — this task must not be given to a consumer AI at all.
Safest workflow:	Do not use AI. Keep your own secure, private notes as your policy requires.

5. Responding to a safeguarding disclosure

Benefit:	None that justifies AI involvement.
Risk:	Delay, breach, and failing a child at the moment it matters most.
Prohibited data:	Everything about the disclosure is Red.
Approval point:	Human, immediate — your safeguarding procedure.
Safest workflow:	Stop. Refer (R) at once through your safeguarding lead and the authorities.

KEEPING IT CURRENT

Course-update policy

This field moves quickly, so this material carries a visible **version number** and **last-reviewed date** (see the introduction). The review cadence:

- **Monthly** — model and tool changes.

- **Quarterly** — privacy policies and safeguarding guidance.
- **On change** — examples whenever a major platform updates.
- **Before each edition** — all legal and regulatory claims.

“Test everything; hold fast what is good.” — 1 Thessalonians 5:21. That is the whole of GUARD in a sentence: use the tool, keep the good, and let nothing past you untested.

From Prompting to Governed Ministry AI · The GUARD Framework · Module 15

The Berean Bible Academy · Matthew (Tolu) Olowoyeye · Module version 1.0, last reviewed September 2026
ministryaccelerator.co.uk