Delbarton School
Artificial Intelligence Policy Framework

A Delbarton Benedictine–DELTA Framework

FINAL DRAFT-SUITABLE FOR DISTRIBUTION
August 28, 2026

I. Delbarton's Philosophy and Core Values Regarding Artificial Intelligence (AI)
Our Approach Flows from Delbarton's Identity and Mission

Delbarton School approaches artificial intelligence (AI) as first a matter of human formation. Thus, our Mission Statement, Belief Statements, Profile of a Graduate, and Hallmarks of a Benedictine School are the lenses through which the School evaluates AI.

The *Rule* of St. Benedict and the Catholic Church's social and educational teachings also inform this discernment. The University of Notre Dame's DELTA framework (**D**ignity, **E**mbodiment, **L**ove, **T**ranscendence, **A**gency) serves as a complementary practical tool for applying these commitments to emerging technologies.

This hierarchy matters. Delbarton does not adopt technology simply because it is available, efficient, or widely used. We ask whether a particular use helps a student glorify God, love learning, listen with the ear of his heart, demonstrate humility and patience, live honorably in community, regard creation as sacred, practice moderation, and grow in service to God and neighbor. In other words, is it a tool that will help us achieve our mission as a Benedictine Catholic school?

Benedictine Principles for AI Use

- *Listening and Obedience*: Technology should support attentive listening, critical reflection, responsible freedom, and responsiveness to truth, teachers, community, and conscience, not replace them.
- *Work and Discipline*: Intellectual labor has formative value. AI should not remove the skills of reading, writing, reasoning, practice, perseverance, or craftsmanship that an assignment is designed to foster in a student.
- *Humility*: Students and adults must recognize the limits of their own knowledge and the limits of AI systems, including their capacity to generate errors, bias, or fabricated information.
- *Community and Stability*: Technology should strengthen, not displace, the human relationships through which learning and formation occur.
- *Hospitality and Love of Neighbor*: AI should be used in ways that honor the dignity of others and serve rather than manipulate, exploit, impersonate, or exclude them.
- *Stewardship and Moderation*: Data, time, attention, intellectual property, material resources, and creation are entrusted to our care. AI should be used prudently and proportionately, not compulsively or wastefully.

DELTA (University of Notre Dame) as a Complementary Discernment Tool

- *Dignity*: Does this AI use honor each person as made in the image of God, or reduce persons to data, outputs, rankings, or problems to be optimized?
- *Embodiment*: Does this AI use strengthen embodied learning, attention, and face-to-face relationships, or displace them?
- *Love*: Does this AI use foster authentic care and responsibility toward real persons? AI may simulate the language or appearance of care, but it cannot love.
- *Transcendence*: Does this AI use preserve openness to truth, beauty, goodness, wisdom, worship, wonder, and meaning beyond utility and optimization?
- *Agency*: Does this AI use strengthen the human capacity to attend, reason, listen, choose what is true and good, accept responsibility, and act with integrity, or does it surrender those tasks to a machine?

Cross-cutting Benedictine question

- *Stewardship and the Common Good*: Is this use moderate, equitable, protective of privacy and creation, respectful of human work and intellectual property, and ordered toward the good of the community, especially those who may be vulnerable or excluded?

Our Commitment
Delbarton will use AI to help form young men as brothers and leaders who think critically, work diligently, act honorably, serve the common good, and use technology wisely in lifelong service to God and neighbor. We will ask not only what AI can do, but what habits, relationships, judgments, and persons its use helps us form in our students.

II. A Four-Tier AI Policy Framework for Delbarton
The four-tier framework below identifies where AI may touch the learning process at Delbarton. The teacher determines the tier based on the assignment's learning objective or a particular component of it.

Students do not select their own tier. If a teacher or course policy has not designated an AI tier, students must ask before using AI on assessed work.

Two school-wide rules apply at every permitted tier:

1. Students remain responsible for the accuracy, integrity, appropriateness, and explainability of everything submitted.
2. Any permitted AI use must be disclosed according to the teacher’s directions. Disclosure never converts an otherwise prohibited use into a permitted one.

Tier 1, RED: Prohibited Uses

RED means that AI may not be used for the designated assignment, assessment, task, or component. The following uses are prohibited school-wide unless expressly authorized for a legitimate educational purpose:

- Submitting AI-generated content as one's own work or concealing AI involvement when disclosure is required.
- Using AI on tests, quizzes, examinations, or other RED assessments.
- Using AI to perform the intellectual work that an assignment is specifically designed to assess, including writing, reasoning, translation, problem-solving, analysis, reflection, or coding when those skills are the learning objective.
- Inventing, altering, or fabricating sources, quotations, data, images, recordings, or evidence.
- Creating or distributing deceptive impersonations, deepfakes, synthetic voice or image content, or other media that misrepresent a real person without legitimate authorization.
- Using AI to harass, humiliate, sexualize, exploit, manipulate, discriminate against, or otherwise harm another person.
- Entering confidential, sensitive, personally identifiable, health, counseling, disciplinary, accommodation, personnel, safeguarding, or other protected information into an unapproved AI system.
- Using AI to evade school safety, security, filtering, monitoring, or academic-integrity controls.
- Using AI as a substitute for personal prayer, sacramental or pastoral care, spiritual direction, counseling, authentic friendship, or an assigned personal reflection whose purpose is the student's own interior examination.

Academic integrity response

The School will address violations through its academic-integrity and disciplinary procedures. Responses should be proportionate to the seriousness, intent, and repetition of the violation and should include a formative component when appropriate; for example, a conversation with the teacher, demonstration of the student's actual learning, revision or resubmission, or reflection on how integrity and trust can be restored.

Tier 2, ORANGE: Restricted Uses

ORANGE permits AI for pre-writing and inquiry support when the teacher has authorized it. The student must still perform the substantive research, reading, thinking, outlining, and drafting required by the assignment. In the STEM subjects, ORANGE permits the use of technology and AI to develop conceptual understanding. ORANGE does not permit the use of technology or AI for the generation of STEM problem set work and solutions, lab results or lab analysis in submitted work.

- AI may suggest search terms, research questions, concepts to investigate, possible counterarguments, or potential source leads.
- AI may be used for brainstorming or to test a student-created thesis, claim, or plan.
- AI is not a scholarly database and should not be treated as an authoritative source finder. Students must independently locate, read, evaluate, and verify all sources through appropriate primary texts, library databases, or other authoritative resources.
- No AI-generated prose, analysis, solution, translation, or other substantive content may appear in the final submission unless a different tier has been expressly designated for that component.
- Students must provide the required AI Use Statement or process log.

Tier 3, YELLOW: Restricted Revision Support
YELLOW includes all authorized ORANGE uses and permits AI to provide feedback on a student-created draft. The student, not the AI, must perform the revision. In the STEM disciplines, YELLOW permits the student to check answers and revise work based on errors found in the work.

- AI may flag grammatical errors, unclear sentences, organizational problems, repetition, possible gaps, or places where an argument is difficult to follow.
- AI may ask questions or identify areas for reconsideration; it may not supply replacement prose or new substantive content that the student inserts as his own.
- The student must compose all revised language himself and decide which suggestions to accept or reject.
- The student must retain and submit, when required, the pre-AI draft, version history, or other evidence of original authorship.
- STEM solutions and lab analysis/results may only be reviewed in AI when the teacher has explicitly permitted it by stating that the assignment is coded YELLOW

Tier 4, GREEN: Teacher-Directed Generative Use
GREEN permits AI-generated content only when the teacher explicitly designates the assignment or a component as GREEN because the educational objective includes learning to use, critique, direct, compare, transform, or evaluate AI-generated material.

- Students may use AI to generate designated portions of the work, up to and including a full work product when the assignment is expressly designed for that purpose.
- Students must critically evaluate AI output for accuracy, bias, omissions, fabricated information, ethical implications, and suitability for the task.
- Students must clearly identify AI-generated content and provide the required AI Use Statement, process log, prompt record, code comments, or other documentation specified by the teacher.
- Students remain responsible for every claim, citation, image, line of code, calculation, and other elements submitted and must be able to explain or defend the work without relying on the AI tool.

- GREEN should not be used when the primary learning objective is the very skill the AI would perform for the student.

Standard AI Use Statement
Unless a teacher specifies another format, a student who uses AI on an assessed assignment should include a brief disclosure such as:

> *AI Use: Flint was used to generate search terms and test my working thesis. I independently located and read the sources. No AI-generated text appears in this submission.*

The AI Use Statement documents *process*; it does not replace normal citation of the actual sources, authors, data, or primary texts on which the work relies.

Four-Tier Summary

Tier	Permitted Role of AI	Student Obligation	Key Boundary
RED	No AI on the designated work.	Complete the work independently.	AI may not touch the assessed process or product.
ORANGE	Inquiry, research leads, brainstorming, testing ideas.	Verify sources; do the reading, thinking, outlining, and drafting.	No AI-generated substantive content in the final submission.
YELLOW	Feedback on student-created work.	Student performs every revision and preserves original draft evidence when required.	AI may diagnose or question; the student writes the revision.
GREEN	Teacher-directed generation, transformation, or task completion.	Disclose, verify, critique, explain, and take responsibility for all output.	Only when explicitly assigned; not when AI would replace the skill being assessed.

Examples: Applying the Framework

Example 1 RED: AI writes an assessed essay
A student enters an essay prompt into an AI tool, receives a five-page essay, and submits it for credit as his own work.

- *Dignity*: The student fails to honor and exercise the intellectual capacities entrusted to him and treats the teacher as an output processor rather than a mentor in learning.
- *Embodiment*: The use bypasses the reading, writing, discussion, feedback, and vulnerability that support learning in the student-teacher relationship.
- *Love*: The concealed use damages trust and substitutes simulated communication for the student's own encounter with the teacher and subject.

- *Transcendence*: It bypasses opportunities for discovery, wonder, and meaning that often arise through the difficult process of forming an argument.
- *Agency*: It surrenders reasoning, authorship, and moral responsibility to a tool and constitutes academic dishonesty.

Summary: RED. AI has replaced the intellectual work and relational process that the assignment is designed to form.

Example 2 ORANGE: AI supports research planning
A junior writing a U.S. History paper asks Flint to suggest search terms and possible scholarly sources on the 1919 Red Scare and to identify weaknesses in his working thesis. He then independently locates and reads the sources through library databases, rejects irrelevant or unreliable suggestions, revises the thesis in his own words, outlines the paper himself, and drafts with no AI-generated prose. He submits the required AI Use Statement or process log.

- The student remains the author and exercises judgment rather than outsourcing it.
- AI accelerates discovery but does not replace reading, annotation, analysis, drafting, or teacher conference.
- Every source suggestion is independently verified; a model-generated citation is never accepted merely because it looks plausible.

Summary: ORANGE. This tier is explicitly for pre-writing investigation and pre-problem solving skill development. No AI-generated prose, problem solving, coding, or any other assessed work shall be part of the submitted work.

Example 3 YELLOW: AI flags revision needs
A senior writes a complete history essay (thesis, evidence, analysis, organization, and conclusion) without AI. He then asks Flint to flag unclear sentences, comma splices, repetition, and passages where the argument is difficult to follow. He decides which observations are useful and rewrites each sentence himself. He submits the pre-AI draft or version history if required.

- AI acts as a diagnostic reader, not a co-author.
- The student preserves his own voice by writing the revised language himself rather than installing AI-generated replacements.
- Peer review and teacher feedback remain preferred relational forms of revision; AI supplements rather than displaces them.

Summary: YELLOW. The governing line is: AI may help identify where revision is needed for clarity; the student must decide and write what the revision will be and make the revisions himself.

Example 4 GREEN: AI output becomes an object of study
Students research an ethical dilemma using authoritative sources. The teacher then directs them to use AI to generate a case study based on their research. Students compare the generated scenario with their sources, identify bias, omissions, and oversimplifications, revise it, and debate the resulting case in class.

- AI generation is explicit and teacher-directed.
- The intellectual objective is not merely to produce the case study but to evaluate, refine, and morally reason about AI-generated material.
- The activity returns to embodied discussion, shared inquiry, and accountable human judgment.

Summary: GREEN. AI-generated content is permitted because generation itself is part of the learning design and the students remain responsible for critical evaluation.

III. Governance and Implementation

AI Implementation Committee

We will establish an AI Implementation Committee that will apply Benedictine principles of listening, mutual obedience, humility, communal discernment, and shared responsibility. Membership should include Mission and Ministry, Academic Affairs, Student Affairs, Technology, faculty, students, and appropriate parent input. The Headmaster or his designee approves major policy changes and high-impact institutional uses.

Responsibilities include:

- Evaluate significant AI implementations through Delbarton's Mission, Profile of a Graduate, Benedictine Hallmarks, Catholic Social Teaching, and the DELTA questions.
- Maintain the four-tier framework, disclosure practices, approved-tool standards, and discipline-specific guidance.
- Oversee professional development and age-appropriate student formation.
- Monitor equity, accessibility, privacy, security, and unintended consequences.
- Review incidents, complaints, and requests for reconsideration or recourse related to AI-supported decisions.
- Consult parents and other constituencies when policies materially affect student formation, privacy, or family responsibilities.

Data Privacy and Stewardship
Student and employee data are a sacred trust. Delbarton will practice data minimization, appropriate security, transparency, and human accountability. Enter personally identifiable or confidential information only into systems the School has approved for that category of data and use.

Unless specifically authorized by the School, users may not enter into public or unapproved AI systems:

- Student or employee names together with grades, discipline records, health information, counseling information, accommodations, safeguarding information, financial information, or other sensitive records.
- Confidential personnel, legal, governance, donor, financial, security, or institutional information.
- Identifiable student work when the system's terms allow retention or secondary use beyond the School's authorized educational purpose.

Vendor review should consider data retention and deletion, model-training practices, security controls, age requirements, accessibility, bias, contractual ownership of data and output, incident notification, and the ability to export or delete institutional data.

Human Review and Due Process
AI output may inform but may not by itself determine consequential decisions about students or employees. Individuals affected by significant AI-supported recommendations must have access to meaningful human review and an avenue for correction or appeal.

Automated AI-detection scores or similar probabilistic indicators may be considered as one piece of information, but they may not establish an academic-integrity violation on their own. Faculty must evaluate the student's process, drafts, version history, demonstrated understanding, and other relevant evidence.

Approved AI Tools
The School shall maintain a separate, current list of approved AI tools and the purposes for which each is authorized. Placement on the approved-tools list does not override assignment-level RED/ORANGE/YELLOW/GREEN restrictions.

Because products, features, data practices, and age requirements change rapidly, the approved-tools list should be maintained administratively rather than embedded permanently in the core policy. Appendix 4 includes a current August 2026 list for review.

IV. Faculty and Staff Guidelines

Course-Level Policies

Faculty have appropriate discretion to set course- and assignment-specific AI tiers according to learning objectives, subject to school-wide prohibitions and departmental or School standards. Each course should communicate clearly:

- Which assignments or components are RED, ORANGE, YELLOW, or GREEN.
- What disclosure, prompt logs, draft history, code comments, or other documentation is required.
- How AI use affects assessment, authorship, and academic integrity.
- That students should ask before using AI whenever an assignment's tier is unclear. Students may not exploit an initial lack of clarity in the assignment. Faculty bear a responsibility to mitigate any ambiguity prior to the commencement of the assignment. Students shall always be responsible for working with the teacher to determine the appropriate level/tier of AI usage, if any.

Appropriate Faculty and Staff Uses

- Lesson planning, curriculum development, generating practice materials, administrative drafting, brainstorming, formatting, and other low-stakes support tasks, subject to privacy rules and professional judgment.
- Differentiating materials and supporting accessibility when the tool and data use are School-approved, and the teacher remains responsible for educational suitability.
- Using AI to generate examples or counterexamples that students will critique, when that use advances the learning objective.

Uses Requiring Heightened Discernment and Human Judgment

- Evaluation of student work: AI may assist with limited analysis or organization only when consistent with School guidance, but the teacher must personally read the work and exercise final evaluative judgment.
- Personalized learning recommendations: AI may suggest options, but the teacher or appropriate professional must assess the student's actual needs and context.
- Routine communications: AI may assist with drafting, but the employee is responsible for accuracy, tone, confidentiality, and authentic authorship.
- Pastoral, disciplinary, evaluative, recommendation, condolence, counseling, or other highly personal communications should remain substantially human-authored because the relationship and witness of the writer are themselves part of the communication.

Prohibited or Non-Delegable Faculty and Staff Uses

- Delegating final grading or evaluative judgment to AI tools.
- Allowing AI to make or effectively determine decisions about student discipline, formation, admissions, financial aid, accommodations, counseling, safeguarding, college recommendations, hiring, employment evaluation, or other consequential matters.
- Using AI to simulate a personal relationship, pastoral presence, spiritual direction, or counseling relationship that the School owes to a student, family, colleague, or community member.
- Uploading confidential or protected information to a system not approved for that use.

Professional Development
All faculty and relevant staff should receive ongoing professional development in: Delbarton's Benedictine–Catholic framework for AI; technical and information literacy; privacy and data stewardship; bias and misinformation; discipline-specific pedagogy; academic integrity and authorship; age-appropriate student formation; and practical use of the four-tier system.

Professional learning should include opportunities for peer exchange and revision based on lived classroom experience.

V. Student Education and Formation

AI Literacy as Moral Formation
AI literacy at Delbarton is not merely technical fluency. Students should learn how AI systems work, where they fail, how they influence attention and judgment, and how to use them responsibly. They should also learn that *a machine is not a person*: it can generate convincing language, images, recommendations, or simulations of care without possessing conscience, wisdom, friendship, faith, love, or moral responsibility.

Student formation should therefore include both technological competence and practices that AI cannot replace: sustained attention, silence, reading, writing, memorization where appropriate, conversation, prayer, worship, service, friendship, physical activity, craftsmanship, and face-to-face community.

Grade-Specific AI Curriculum

Grades 7–8: Foundations

- What AI is and is not; recognizing AI-generated content in daily life.
- Digital citizenship, privacy, consent, impersonation, manipulation, and safe use.
- Foundational concepts of dignity, agency, listening, honesty, moderation, and human relationships.

- Why AI is a tool rather than a person, friend, counselor, spiritual guide, or conscience.

Grades 9–10: Ethics and Responsibility

- Full introduction to the Delbarton Benedictine–DELTA framework and four-tier system.
- Academic integrity, authorship, disclosure, research verification, misinformation, fabricated citations, and deepfakes.
- Bias, fairness, privacy, data protection, intellectual property, and the dignity of human work.
- Case studies in ethical decision-making and the common good.

Grades 11–12: Advanced Discernment and Leadership

- Advanced engagement with Catholic social teaching and AI: dignity, common good, solidarity, subsidiarity, justice, work, universal destination of goods, and care for creation.
- AI and the future of work, leadership, civic life, culture, truth, art, science, and human relationships.
- Governance, platform power, surveillance, privacy, equitable access, and environmental/resource costs.
- Technical deep dives for interested students and critical evaluation of model limitations.
- Leadership in using technology for service, justice, human flourishing, and the common good.

Parents and Families
Delbarton recognizes parents as primary partners in the education and formation of their sons. The School should provide parents with clear explanations of the four-tier system, approved tools, privacy protections, age-appropriate expectations, and ways families can reinforce responsible use, moderation, truthfulness, and healthy human relationships at home.

VI. Assessment, Accountability, and Continuous Improvement

Regular Assessment
At least annually, the AI Implementation Committee should review the effect of AI policy and practice on student learning and formation, faculty experience, academic integrity, equity and accessibility, privacy and security, the teacher-student relationship, and alignment with Delbarton's Foundation Documents and the DELTA framework.

Feedback and Recourse

- Student feedback, with anonymous options where appropriate.
- Faculty and staff discussion forums and discipline-specific review.
- Parent communication and opportunities for input.

- A clear channel for reporting concerns about bias, privacy, safety, academic integrity, or inappropriate AI-supported decisions.
- Meaningful human review and correction when an individual believes an AI-supported process has produced an inaccurate, unfair, or harmful result.

Incident Response
When an AI system causes or materially contributes to harm, the School should take proportionate immediate action, which may include suspending the AI tool's use; investigating the cause and scope; protecting and remediating affected individuals; communicating appropriately and transparently consistent with privacy and legal obligations; and revising policy, training, or technical controls to reduce recurrence.

Continuous Improvement
This policy is a living document. Delbarton will review the core policy at least annually, and revise it when educational experience, Catholic teaching, legal requirements, vendor practices, or technological developments materially warrant change. Operational materials, such as approved tools, examples, and technical guidance, may be updated more frequently. Benedictine humility requires acknowledging when an approach is not working and having the courage to change course.

VII. Conclusion: Formation for an AI Age
Delbarton is not simply managing technology. We are forming young men who must learn to live wisely and faithfully in a world increasingly shaped by artificial intelligence. Our goal is not to produce students who are merely competent with AI, but graduates who love learning, seek truth, work diligently, listen with the ear of the heart, exercise freedom responsibly, honor the dignity of others, practice moderation, care for creation, and serve the common good.

AI can assist learning when it remains a tool within the human relationships and purposes of education. It must not become a substitute for thinking, authorship, conscience, friendship, community, pastoral care, prayer, or the teacher's personal responsibility for students. Efficiency is a legitimate good, but it is not the highest good; some forms of struggle, practice, patience, presence, and human work are themselves part of education and formation.

The Delbarton Benedictine–DELTA framework therefore rests on four enduring commitments:

- *Mission*: AI use must serve Delbarton's Benedictine Catholic purpose and the formation described in the Profile of a Graduate.
- *Transparency and Truth*: We disclose permitted AI use, verify information, distinguish generated content from authoritative sources, and refuse deception.
- *Human Responsibility*: Consequential judgments about persons remain human; students and adults remain accountable for what they submit, communicate, or decide.
- *Discernment for the Common Good*: We judge technology not simply by capability or efficiency but by whether it honors dignity, strengthens relationships and agency, protects

privacy and creation, and contributes to human flourishing and lifelong service to God and neighbor.

Appendix 1: Delbarton–DELTA Discernment Questions
Use these questions when evaluating a new AI tool, practice, assignment, or institutional use. They do not replace the four-tier rules; they guide discernment about whether and how a use should be authorized.

Principle	Key Question
Dignity	Does this honor persons as made in the image of God, or reduce them to data, rankings, outputs, or problems to be optimized?
Embodiment	Does this strengthen embodied learning, attention, presence, and face-to-face relationships, or displace them?
Love	Does this foster authentic care, responsibility, and community? Is simulated relationship being mistaken for genuine human love or presence?
Transcendence	Does this preserve openness to truth, beauty, goodness, wisdom, worship, wonder, gratitude, and meaning beyond utility?
Agency	Does this strengthen the capacity to attend, reason, listen, choose what is true and good, accept responsibility, and act with integrity?
Stewardship / Common Good	Is the use moderate, equitable, protective of privacy and creation, respectful of human work and intellectual property, and ordered toward the good of others, especially those who may be vulnerable or excluded?

Appendix 2: Sample Course-Level Policies
The following illustrate how the same school-wide framework can produce different rules across disciplines. These are only examples and each faculty member is responsible to work with co-teachers of the same course to adapt a consistent course-specific policy rather than simply adopt these examples.

English 10
The default for analytical writing is RED for drafting. A teacher may designate YELLOW for final revision after a complete student-written draft, with the pre-AI draft or version history available. A research paper may be ORANGE for search terms, source leads, and thesis-testing.

Any AI-suggested source must be independently located and verified. Undisclosed or unauthorized use is an academic-integrity issue. When in doubt, ask before using AI.

Algebra II
Tests and quizzes are RED. Nightly problem sets may be ORANGE when the teacher permits AI to explain a concept or walk through a different example problem; the student may not enter the assigned problem and copy the result. The test is whether the student can reproduce and explain the reasoning with the tool closed.

Long-term modeling projects may be GREEN for computation or graphing when those are not the skills being assessed, and RED for mathematical reasoning and written interpretation.

Spanish III
Assessed writing, speaking, and translation are RED when language production or translation is the skill being assessed. A teacher may permit ORANGE or YELLOW support for specified learning activities that do not replace the target skill.

Students may ask for an explanation of a grammatical structure they have already encountered; they may not ask AI to produce assessed language that they submit as their own.

Chemistry
Tiers may vary by section of a lab report. Background research may be ORANGE. Procedure, data analysis, and discussion may be RED when those sections assess the student's scientific reasoning. Graphing, unit conversion, or spreadsheet formulas may be GREEN when the technical production is not the learning objective.

The student remains responsible for verifying all calculations and interpreting what the data mean.

AP Computer Science
Foundational units and assessments may be RED while later projects may include GREEN components. Students are responsible for every line of code submitted, regardless of who or what generated it, and must be able to explain arbitrary sections of their code.

AI-generated blocks must be marked as directed by the teacher, and students must document what they changed and why.

Sacraments
Personal reflections, prayer reflections, and service reflections are RED when the purpose is honest engagement with the student's own experience, conscience, faith, or interior life.

Research on Church teaching and historical context may be ORANGE. AI-generated references to Scripture, official Church documents or Church history must be verified against authoritative primary sources before use.

Appendix 3: Standard Disclosure and Documentation

Student AI Use Statement (Short Form)
AI Use: [Tool] was used to [purpose]. I used the output to [how it affected the process]. I independently verified [sources/facts/calculations as applicable]. [No AI-generated content appears in the final submission / The following AI-generated components were expressly permitted: ____].

Teacher-Directed Process Log (Optional Elements)

- Tool and account used.
- Date(s) of use.
- Prompts or representative prompts.
- Outputs retained or rejected.
- Sources independently located and verified.
- Changes the student made after AI feedback.
- Identification of any AI-generated content explicitly permitted in the final submission.

Appendix 4: Delbarton Approved AI Tools (August 2026)

This appendix will be updated administratively as products, contracts, age restrictions, features, and data practices change. Inclusion here does not authorize use on an assignment whose tier does not permit it.

- *Flint K-12*: School-licensed primary AI tool for faculty, staff, and students, subject to assignment-level permissions and School configuration.
- *Google Gemini*: Available to faculty and staff through Delbarton's Google Workspace, subject to School configuration and data-governance rules. Student access through School accounts remains restricted unless School policy changes.
- *Google Notebook LLM*: Available to faculty and staff through Delbarton's Google Workspace, subject to School configuration and data-governance rules. Student access through School accounts remains restricted unless School policy changes. Notebook LLM is primarily used in the sciences for lab data analysis.
- *Speechify*: Approved for reading assistance as authorized by the School or an appropriate educational accommodation. School-wide licensing status should be reviewed periodically.
- *Kurzweil 3000*: Approved reading and accessibility support where provided through School systems, subject to assessment-specific directions and accommodations.

Any additional AI tool should undergo School review before being required for students or used with protected School data.

Appendix 5: Reference Documents Informing this Revision

- Delbarton School, Mission and Beliefs; Profile of a Delbarton Graduate; Hallmarks of a Delbarton Benedictine Education (current School Foundation Documents).
- *Rule* of St. Benedict, especially the Prologue and themes of listening, obedience, humility, work, stability, community, hospitality, stewardship, and moderation.
- Dicastery for the Doctrine of the Faith and Dicastery for Culture and Education, Antiqua et Nova: Note on the Relationship Between Artificial Intelligence and Human Intelligence (28 January 2025), especially nn. 77–84 on education and nn. 90–94 on privacy.
- Pope Leo XIV, *Magnifica Humanitas*: On Safeguarding the Human Person in the Time of Artificial Intelligence (15 May 2026), especially nn. 59–75 on Catholic social doctrine and nn. 143–147 on schools and education.
- University of Notre Dame, DELTA: Human Formation in the Age of AI (Dignity, Embodiment, Love, Transcendence, Agency).

Appendix 6: Summary of Major Proposed Changes and Rationale

Area	Proposed Change	Rationale
Identity hierarchy	Delbarton foundation documents now govern; DELTA is complementary.	Makes the policy institutionally authentic and rooted first in Delbarton's Benedictine Catholic mission.
Benedictine content	Adds work, discipline, listening/obedience, humility, community/stability, hospitality, stewardship, and moderation throughout.	Prevents Benedictine identity from being confined to the introduction.
Catholic social teaching	Adds common good, solidarity, subsidiarity, dignity of work, care for creation, privacy, justice, and equitable access.	Broadens the ethical lens beyond academic integrity.
Agency	Reframed as freedom and responsibility ordered toward truth and the good.	Avoids interpreting agency as mere autonomy or self-direction.
Tier authority	Teachers assign tiers; students do not self-select.	Closes a practical enforcement gap.
ORANGE	AI suggests search terms/leads; sources must be independently located and verified.	Corrects the misleading comparison of generative AI with scholarly databases.
YELLOW	AI diagnoses or questions; the student writes the revision.	Creates a clear line between feedback and co-authorship.
GREEN	Explicit teacher designation, disclosure, verification, critique, explainability.	Prevents "encouraged" from being read as blanket permission.
Disclosure	Standard AI Use Statement distinguishes process disclosure from source citation.	Improves academic practice and clarity.
RED misuse	Adds deepfakes, harassment, fabricated evidence, unsafe content, privacy breaches, safety circumvention, and relational/pastoral substitution.	Recognizes school-wide harms beyond cheating.
Privacy	Adds concrete categories of protected data and vendor-review criteria.	Turns a strong principle into an operational safeguard.

Area	Proposed Change	Rationale
High-stakes decisions	Expands non-delegable human judgment to grading, discipline, admissions, aid, counseling, accommodations, recommendations, safeguarding, and personnel matters.	Protects dignity, due process, and professional responsibility.
AI detection	Prohibits treating detector scores as sufficient proof of misconduct.	Preserves human judgment and fair process.
Student curriculum	Clarifies grade bands and adds AI anthropomorphism, misinformation, labor, IP, creation, platform power, and CST.	Strengthens formation for an AI-saturated culture.
Parents	Adds parent partnership in formation and policy communication.	Reflects Catholic understanding of parents as primary educators and strengthens consistency between School and home.
Approved tools	Moves vendor list to administrative appendix.	Allows rapid operational updates without destabilizing the core policy.
Review cadence	Uses “at least annually” for the core policy and more frequent operational updates as needed.	Removes ambiguity in “biannual” and supports continuous improvement.
Theological language	Revises statements about “degrading” dignity or “destroying” wonder.	Preserves the Catholic conviction that dignity is inherent while maintaining moral seriousness.