MAINE'S PUBLIC UNIVERSITIES
UNIVERSITY OF MAINE SYSTEM

REQUEST FOR PROPOSALS #2026-065
Enterprise Generative AI Services
RESPONSE ADDENDUM #1
March 11, 2026

Table of Contents

RFP General Instruction Questions

1. Can you please give us an extension of 1-2 weeks to submit our proposal?
   ANSWER: No. The deadline of March 27, 2026, 11:59 p.m. EST remains firm as stated in the RFP.

2. What is the University's timeline for making a decision and proceeding with the software purchase after the RFP process is complete?
   ANSWER: Timeline is identified in RFP Section 2.0. In addition, the final decisions is dependent on Board of Trustees (BOT) approval

3. The RFP references respondent presentations scheduled for April 8–9. Will UMS provide vendors with a standardized presentation format, agenda or list of topics to be addressed or are vendors expected to structure their presentations independently? Additionally, will vendors be permitted to conduct a live demonstration of their solution during this session and if so, what technical environment or access will UMS provide to facilitate this?
   ANSWER: Refer to RFP Section 1.3.7.

4. Is this contract intended to be awarded to a single vendor or to multiple vendors?
   ANSWER: Refer to Section 2.2 of the RFP.

5. What was the annual spend for the previous year on this Project?
   ANSWER: This is a new initiative. There is no previous annual spend for an enterprise AI agreement.N/A we have had an enterprise AI agreement in the past.

6. Are you open to a hybrid delivery model with a mix of offshore and onshore resources?
   ANSWER: No

7. Work will be onsite or remote?
   ANSWER: We are requesting licenses for a cloud solution.

8. Could the University expand upon the discrepancy between two items? In Appendix A, Item 1, it says "This pricing structure contained herein will remain firm for a period of 90 days from the date and time of the quote deadline date." However in Appendix C, item 8, it says that "Pricing will be guaranteed by the vendor for the term of the Agreement." which could be as many as 12 years from now.
   ANSWER: Appendix A is for the evaluation team's consideration of the submission material only and is unrelated to the term of the Agreement and pricing requirements.

9. What are the current limitations and challenges that you are facing?
   ANSWER: The University does not currently have an enterprise AI solution in place. Individual departments may have experimented with various AI tools independently, but there is no unified, institution-wide platform. The primary challenges include the lack of a secure, FERPA-compliant enterprise AI service, inconsistent access across campuses, and the need for a centralized platform that supports teaching, learning, research, and administrative functions for all UMS users.

10. What duration of post-launch support is expected (e.g., 90 days), and should vendors propose ongoing maintenance options beyond the initial launch?

ANSWER: Just initial implementation support, if any. Ongoing maintenance and support should be included in the subscription for the Agreement term. Just initial implementation support, if any.

11. Regarding the Accessibility and Information Security pass/fail gates, if a solution is 'Partial' on a specific requirement but has a documented roadmap for compliance, will that result in a 'Fail' for the entire section?
    ANSWER: The evaluation team will determine the pass/fail outcome based on the overall response. A "Partial" response with a documented roadmap does not automatically result in a fail; however, the decision will depend on the completeness, feasibility, and timing of the proposed remediation as evaluated by the team relative to University requirements.

12. Can you please explain the pass/fail criteria for Accessibility and IT Security, and what evidence is required to pass?
    ANSWER: The evaluation team will determine the pass/fail outcome based on the overall response using the results of the RFP Submission Appendix I information and using a consensus scoring approach. Please note that RFP Appendix H reflects the "Required" elements for the category Accessibility and Data Privacy & Security (FERPA).

13. What success metrics will the University use to evaluate solution effectiveness (e.g., adoption rate, cost savings, student engagement) and over what timeframe will initial performance be assessed?
    ANSWER: Refer to RFP Section 2.0 for the RFP evaluation and process.

14. Please provide top 5 key agenda on AI adoption and modernization
    ANSWER: The University prioritizes broad enterprise AI adoption across the following areas: (1) Teaching and learning enhancement, including writing assistance, content creation, and formative assessment support; (2) Research support, including literature review, data analysis, and research assistance tools; (3) Administrative efficiency through document analysis, workflow automation, and operational support; (4) Student services and accessibility, including translation, transcription, and text-to-speech capabilities; (5) Institutional data analysis and decision support through secure, grounded AI interactions with University data.

15. Could you provide 2–3 high-priority 'use case' scenarios that represent a successful outcome for this requirement? Are you solving for a specific challenge, such as reducing faculty workload, improving student writing outcomes?
    ANSWER: The University envisions several high-priority use cases: (1) Faculty use: AI-assisted course content development, writing feedback tools, and formative assessment support integrated with Brightspace via LTI 1.3; (2) Student use: Research assistance with citation-grounded responses (RAG), writing improvement tools, code generation sandbox for CS courses, and multilingual support for international students; (3) Administrative use: Document analysis and summarization, policy research, and operational workflow support across all campuses.

16. Do you anticipate usage caps or governance controls (e.g., limiting token consumption per user, restricting certain models to specific departments, or requiring approval workflows for sensitive use cases)?
    ANSWER: The University does not anticipate imposing hard usage caps that would impede educational use, consistent with Appendix H Requirement #7 (unlimited access or predictable flat-rate license). Internal governance controls regarding model access, sensitive use cases, or departmental policies will be managed by UMS and are not expected to be vendor-imposed.

17. To ensure we provide the most cost-effective solution, do you prefer the performance of proprietary models or the flexibility and transparency of open-source alternatives?
ANSWER: The University is model-agnostic and values capability over licensing model. Both proprietary and open-source models are acceptable. Appendix H Requirement #3 lists multi-model support as Preferred to avoid vendor lock-in. Respondents should propose whichever approach delivers the best combination of capability, security, and cost-effectiveness.

18. Do you expect this to be a product offering (COTS), or can it be a solution that we build specifically for the UMS's requirements?
ANSWER: The University prefers a commercial off-the-shelf (COTS) product that meets the requirements outlined in Appendix H. Custom-built solutions will be considered only if they demonstrably meet or exceed the stated requirements. Respondents should clearly identify which capabilities are COTS and which would require custom development.

19. Does the University have an existing AI governance or responsible AI framework that vendors must support?
ANSWER: Yes. The University of Maine System has an existing AI governance framework and responsible AI policies in place. Vendors are not expected to develop governance policies but should ensure their platform supports the enforcement of institutional AI governance controls, such as usage policies, data handling rules, and content guardrails.

20. Are vendors expected to provide mechanisms to detect bias, hallucinations, or unsafe AI outputs?
ANSWER: Yes. Appendix H Requirements #46 (Bias Mitigation and Transparency) and #47 (AI Safety Guardrails) are both marked as Required. Vendors must provide mechanisms to detect and mitigate bias in model outputs, prevent generation of harmful content, and provide transparency documentation on their efforts in these areas.

21. Are collaborative AI environments required for research teams or academic departments?
ANSWER: Yes. Appendix H Requirement #42 lists Team dashboards/workspaces as Preferred. The University envisions collaborative AI environments for research teams, academic departments, and class-based workspaces where data can be shared within approved groups while maintaining isolation from other teams.

22. Does UMS have existing landing zones, or cloud governance practices the solution must align with?
ANSWER: The University does not have specific cloud governance landing zones that the solution must align with. The University is seeking a vendor-hosted SaaS solution. The vendor is responsible for hosting and infrastructure. There is no preferred cloud provider.

23. Are there current genAI tools already licensed/approved at any campus (and if so, what are they and what gaps are driving this procurement)?
ANSWER: No. There are no enterprise AI tools currently licensed or approved at any UMS campus. This is a new procurement to establish the first enterprise-wide AI service.

24. For given RFP scope, have you done any analysis on technical platform and architecture needed? If not, would you consider this in scope for this RFP?
ANSWER: The University has defined its requirements in Appendix H (Solution Requirements Matrix) and Appendix J (Information Technology). Respondents are expected to propose a technical

architecture that meets these requirements. Platform and architecture recommendations from Respondents are welcome as part of their submission.

25. Will UMS accept a joint response from a technology platform provider and a systems integrator for implementation and training?
ANSWER: No. The University prefers a single vendor response. Joint responses from a technology platform provider and a systems integrator are not being accepted for this procurement.

26. The University intends to handle bias, fairness, and ethics as required by law and policy.
    - Does the University have a specific AI Ethics Committee or governing body that will provide the "human oversight", and what is their expected interface with the vendor?
    - Since the University seeks a "long-term partnership" but the specifics are "unknown," what are the key Key Performance Indicators (KPIs) the University will use after the first 12 months to determine if the service is successful

ANSWER: The University has AI governance in place. KPIs for evaluating service effectiveness will be defined during implementation in collaboration with the awarded vendor. Initial assessment will focus on platform availability, adoption rates across user populations, and alignment with institutional requirements.

INCUMBENT ANSWER

ANSWER: N/A we have had an enterprise AI agreement in the past.

QUESTION(S)

1. Who are previous incumbents on this project?

ANSWER: There are no previous incumbents. This is a new initiative; the University of Maine System has not previously had an enterprise AI agreement.

2. Who is the incumbent vendor currently providing similar services?

ANSWER: There is no incumbent vendor. This is a new procurement.

3. Is the incumbent vendor allowed to participate in this RFP?

ANSWER: N/A. There is no incumbent vendor. This is an open competitive procurement.

IMPLEMENTATION

Answer

The RFP does not prescribe a single deployment pattern or a fixed go-live date. Respondents should propose a **realistic implementation plan**—including dates for project start, pilot(s), go-live, and stabilization—that fits their solution and UMS's multi-institution context. Pricing for one-time implementation tasks should be placed under **Appendix C, Exhibit 1 (Table 1): "Initial Cost 'One-Time' Implementation,"** with any role-based services shown in **Table 2: Professional Services Rate Schedule** and any scoped fixed deliverables in **Table 3**.

UMS **is open to phased or full system-wide launches**; vendors should justify the approach in their **implementation roadmap**. The RFP specifically asks vendors to provide a **Technology Implementation Roadmap** (see **Appendix J – Question 13**) that looks ahead **18 months** and covers key areas (e.g., agentic workflows, multimodal features, deeper integrations). Use that roadmap to indicate **pilot timing, enterprise rollout timing, and major milestones**

UMS's evaluation will look for a plan that delivers **production availability** aligned to the roadmap and that supports **adoption across key user groups** (faculty, students, staff/end-users, administrators, IT support) with appropriate training. As this is an enterprise service, success should be framed as:

- A secure, production-ready service **available to intended user populations** as scoped for the phase (pilot or system-wide), and
- A **training/onboarding program** mapped to user roles and responsibilities, with one-time training shown in **Table 1** and any role-based services in **Table 2**, plus optional campus-specific add-ons in **Table 4** (see training pricing placement guidance in Appendix C).

Respondents should specify **assumptions about UMS/campus participation** (e.g., project management, communications, local champions, testing, data connections). Any **vendor-led** support required should be priced as **professional services** (Table 2) or as **fixed deliverables** (Table 3), while the **baseline implementation** and **one-time training** belong in **Table 1**.

Question(s)

1. What is your desired timeline for deployment — are you looking for a phased rollout (e.g., pilot group first) or a full institution-wide launch, and by what date?

ANSWER: The RFP does not prescribe a single deployment pattern or a fixed go-live date. Respondents should propose a realistic implementation plan including dates for project start, pilot(s), go-live, and stabilization. UMS is open to phased or full system-wide launches. See Appendix J Question 13 for the 18-month roadmap requirement.

2. Does the University anticipate phased implementation or pilot deployments before enterprise-wide rollout across all campuses?

ANSWER: Yes, UMS is open to phased implementation or pilot deployments before enterprise-wide rollout. Vendors should justify the approach in their implementation roadmap.

3. What is the University's target go-live date, and is a phased rollout (pilot campus followed by full system deployment) preferred over a simultaneous system-wide launch?

ANSWER: No fixed go-live date is prescribed. Respondents should propose a realistic timeline. UMS is open to either phased rollout or simultaneous system-wide launch.

4. What is the anticipated project start and go-live date?

ANSWER: The Agreement is expected to begin after Board of Trustees approval. The specific start and go-live dates should be proposed by the Respondent in their implementation roadmap.

5. What are the expectations and timeline for ongoing support and maintenance after the system is implemented?

ANSWER: Ongoing maintenance and support should be included in the subscription for the Agreement term. Implementation support is expected only for the initial deployment period.

6. What does UMS imagine a successful implementation to be? For example, is it that all students, faculty, and staff have access to the AI service within a certain period of time? Or is it that some subset of these end-users have been trained on and can use the AI service, tailored to their use case and context within a certain timeframe?

ANSWER: A successful implementation means a secure, production-ready service available to intended user populations (faculty, students, staff, administrators) with appropriate training and onboarding mapped to user roles.

7. What level of roll-out support are they assuming from their internal resources at the campus-level?

ANSWER: UMS anticipates a joint implementation model in which the vendor leads deployment and enablement, while UMS provides campus-level participation (e.g., project management, communications, local champions, testing). Respondents should specify assumptions about UMS/campus participation.

8. Is the University open to a phased or multi-wave implementation approach (for example, piloting on select campuses or use cases first, then scaling over time), potentially with different timing by campus?

ANSWER: Yes. UMS is open to phased or full system-wide launches; vendors should justify the approach in their implementation roadmap.

9. The submission form package requests an 18-month implementation roadmap. Are there any milestones within that timeline by when you would like end users to have the solution in hand, of the AI service? Are you open to a phased / multi-wave implementation approach (for example, piloting on select campuses or use cases first, then scaling over time), potentially with different timing by campus? Can the implementation period go beyond 18 months?

ANSWER: The RFP does not prescribe specific milestones within the 18 months. Respondents should propose milestones in their implementation roadmap. UMS is open to phased/multi-wave implementation.

10. What internal resources and personnel will the University commit to supporting implementation and what level of change management, training, and adoption support is expected from the vendor?

ANSWER: UMS anticipates a joint implementation model where the vendor leads deployment and UMS provides campus-level participation. Respondents should state and price their assumptions about UMS resource commitments in their proposal.

11. Is the University open to a pilot or phased contract structure prior to committing to the full five-year term, and if so, what would the minimum pilot duration and evaluation criteria be?

ANSWER: The RFP contemplates a 5-year initial term with optional renewal up to 10 years as specified in Appendix D. Respondents may propose phased implementation approaches within the Agreement structure.

12. What is the expected implementation timeline and phasing?
    - Simultaneous deployment to all campuses or phased rollout?
    - If phased, what is the expected timeline to reach all campuses?

ANSWER: UMS is open to either simultaneous deployment to all campuses or phased rollout. Vendors should propose and justify their approach. If phased, include the expected timeline to reach all campuses in the 18-month implementation roadmap.

IMPLEMENTATION SUPPORT

Answer

Yes. UMS anticipates a joint implementation model in which the vendor leads the deployment and enablement work, and UMS provides campus-level participation for project success. Respondents should state and price their assumptions clearly in the cost exhibits and implementation plan. One-time implementation and onboarding belong in Appendix C, Exhibit 1 (Table 1) — "Initial Cost 'One-Time' Implementation/Training," with any role-based services in Table 2 (Professional Services Rate Schedule) and any fixed-scope deliverables in Table 3.

UMS envisions a shared-responsibility operating model after go-live:

- Vendor Responsibilities (post-implementation):
  - Operate and support the service per the proposed SLA, including model/service availability, updates, and reliability. (Address availability, rate limits, and failover in Appendix J – Availability & Limits / Failover Strategy.)
- Maintain secure operations of the SaaS or hosted service, including data protection controls, and provide administrative tools and usage dashboards for visibility into tokens/costs/active users per Appendix J – Monitoring.
- Provide ongoing enablement for new features (e.g., release briefings, update training). Price these either within the annual subscription lines in Table 1 (Years 1–5; optional 6–10) or as optional services in Tables 2/4.

Question(s)

1. Is UMS assuming support from internal resources at the campus level to support roll-out? If yes, please describe (level of support, number / types of resources, etc.)

ANSWER: Yes. UMS anticipates a joint implementation model in which the vendor leads deployment and enablement work, and UMS provides campus-level participation for project success. Respondents should state and price their assumptions clearly.

2. How does UMS envision operational ownership of the solution post-implementation?

ANSWER: UMS envisions a shared-responsibility model: the vendor operates and supports the service per the proposed SLA (availability, updates, security), while UMS handles internal adoption, policy, and campus-level administration.

SLA & SUPPORT

Question(s) & Answer(s)

1. Have SLA's for ongoing support already been defined, and if so, what are they?
   ANSWER: SLAs have not been pre-defined. Respondents should propose their standard SaaS SLAs. The University expects standard enterprise SaaS service levels including 99.9% minimum uptime, business-hours technical support, and documented incident escalation procedures.

2. What level of ongoing support are you envisioning? Are there existing teams in place who can provide the ongoing support needed? Would UMS be open to leveraging digital agents for ongoing support to reduce the effort needed by resouces.
   ANSWER: The University has internal IT teams that can provide first-level support. Respondents should propose a support model that includes vendor-provided technical support for platform issues, escalation procedures, and access to product expertise. The University is open to leveraging digital agents or self-service tools as part of the support model.

3. Can you please share more about the Administrative Tech Support and if there would be a need to train internal IT teams to effectively maintain and enhance the AI service?
   ANSWER: Respondents should propose training and documentation for internal IT administrators. This should include admin console training, troubleshooting guides, and ongoing access to product documentation and release notes. The scope of internal IT training should be detailed in the Respondent's implementation plan.

4. Where do you want vendor support to begin/end (tiering), and what SLAs matter most (incident response, uptime, support response time)?
   ANSWER: Respondents should propose support tiering based on their standard model. Key priorities for UMS are platform uptime/availability, incident response times for critical issues, and support responsiveness for administrative and technical questions.

5. What minimum uptime SLA is required, what are the mandated support hours and is 24/7 coverage expected for critical production incidents?

ANSWER: The University expects a minimum 99.9% uptime SLA. Business-hours support (Monday-Friday, 8 AM - 6 PM EST) is acceptable for standard issues. 24/7 coverage is not required, though critical platform outages should have an after-hours escalation path.

6. What are the expected incident response and resolution times for critical outages and does the University require a dedicated Technical Account Manager?
   ANSWER: Respondents should propose incident response and resolution times appropriate for an enterprise SaaS platform. A dedicated Technical Account Manager is not required but may be proposed as a value-add.

TRAINING

Question(s) & Answer(s)

1. Can you please elaborate on the scale of training, or learning support, desired (e.g., type of audience, number of individuals, duration) and any preferences around training formats (e.g., facilitated in person, hybrid or virtual, recorded on-demand, quick reference guides)? How should student use be guided differently from faculty or staff use, if at all?
   ANSWER: We expect Respondents to understand their solution and propose options for training, learning support, etc.

RFP Appendix C

Budget

Budget Answer

Submissions will inform the Board to set a budget. The University of Maine System fiscal year starts July 1st.

Budget Question(s)

1. If this is a new Contract, What is the annual Budget for this?

ANSWER: Submissions will inform the Board to set a budget. The University of Maine System fiscal year starts July 1st.

2. Please outline the anticipated budget range for the proposed solution and identify the key cost drivers, assumptions, and dependencies associated with the budget range.

3. Do you have an estimated monthly or annual budget range for this initiative, and is pricing preferred on a per-user, per-token/usage, or flat-rate basis?

ANSWER: Submissions will inform the Board to set a budget. Respondents should provide their best-value pricing per Appendix C instructions.

4. What is the approved budget or range allocated for this project?

ANSWER: No budget range has been pre-defined. Submissions will inform the Board to set a budget.

5. Is this project budgeted for in the current fiscal year, or would you be able to share any details about the financial backing?

ANSWER: Submissions will inform the Board to set a budget. The University of Maine System fiscal year starts July 1st.

AUTHORIZED PARTICIPATING ENTITIES

Answer

The RFP's Master Agreement states that the Maine Community College System and Maine Maritime Academy "may elect to utilize this Agreement under the same terms," and that the Contractor must extend all products, services, terms, conditions, and pricing to these institutions. The Respondent should not include submission information for them.

Question(s)

1. The RFP notes that Maine Community College System and Maine Maritime Academy may also use the solution. Should vendors include pricing and licensing assumptions for these institutions in the initial response, or treat them as optional extensions?

ANSWER: Pricing for the University of Maine System must be included. MCCS and MMA may elect to utilize the Agreement under the same terms. Respondents should price for UMS and may include optional pricing for participating entities.

2. Will the Maine Community College System (MCCS) be included in this award? If so, please provide their total FTE count, as this will materially affect the pricing structure and per-unit cost modeling.

ANSWER: MCCS and Maine Maritime Academy may elect to participate under the Agreement terms but are not guaranteed participants. Their FTE counts are not included in the UMS totals in Appendix C.

3. The RFP mentions the Maine Community College System and Maine Maritime Academy as potential participating entities (Section 27 of the Agreement). Should respondents price and scope for these entities from the outset, or treat them as optional add-ons?

ANSWER: Respondents should price for the University of Maine System. MCCS and MMA participation is optional under the Agreement terms. Respondents may include optional pricing for these entities but it is not required.

FTE / LICENSING

Answer

Pricing for the University of Maine System must be included. Appendix C specifies the combined FTE (Students 25,262; Employees 5,650; Total 30,912) for pricing where FTE-based tiers apply.

University of Maine System is unaware of any existing enterprise agreements (e.g., Microsoft E5/Copilot, Google Workspace) that could constrain or influence how this solution is licensed, however the answer may be influenced by the Respondents' submission.

Question(s)

1. What is the estimated number of users (students, faculty, staff) expected to access the platform in the first year?

ANSWER: Appendix C specifies the combined FTE: Students 25,262; Employees 5,650; Total 30,912. Proposals should be based on these figures.

2. Should licensing proposals be based on all 30,000 plus FTE or only projected active users?

ANSWER: Pricing should be based on the full FTE count of 30,912 as specified in Appendix C.

3. Of the 30,000 plus FTE, how many are expected to be active monthly users, and how should the vendor segment the user base (e.g., students, faculty/power users, administrative staff) for licensing purposes?

ANSWER: Active monthly user projections are not available. Pricing should be based on the full FTE count of 30,912 as specified in Appendix C.

4. The RFP covers 30,000 students and 5,000 staff. Does the University want a single enterprise-wide tenant, or should the solution allow for sub-tenants managed by individual campuses (e.g., Maine Law or the Graduate Center)?

ANSWER: Per Appendix J Question 14, the solution should support multiple institutions within a single tenant with logical data separation. Campus-level workspaces are preferred.

5. What is the total number of anticipated users across students, faculty, and staff, and how are they distributed across each category?

ANSWER: Appendix C specifies: Students 25,262; Employees 5,650; Total 30,912. Distribution across individual campuses is available in the RFP documents.

6. Are there any existing enterprise agreements (e.g., Microsoft E5/Copilot, Google Workspace) that could constrain or influence how this solution is licensed?

ANSWER: The University of Maine System is unaware of any existing enterprise agreements (e.g., Microsoft E5/Copilot, Google Workspace) that would constrain or influence how this solution is licensed.

PRICING

Answer

Respondents should provide University of Maine System enterprise pricing which will include the following:

- University of Maine System Office
- University of Maine
- University of Maine at Machias (regional campus of University of Maine
- University of Maine at Augusta
- University of Maine at Farmington
- University of Maine at Fort Kent
- University of Maine at Presque Isle
- University of Maine School of Law
- University of Southern Maine

Pricing for the University of Maine System must be included. Appendix C specifies the combined FTE (Students 25,262; Employees 5,650; Total 30,912) for pricing where FTE-based tiers apply.

The University encourages price incentives for enterprise, multi-year, or multi-institution awards; include such incentives in the Appendix C tables (show list, extended costs, discounts, totals).

Use the Appendix C format and Exhibit 1 (Tables 1–4); add rows/columns as needed and leave non-applicable tables blank. We do require Table 1 completed. Excel versions are required.

Show Years 1–5 and optional Years 6–10, including extended costs, discounts, and totals.

Include all costs necessary for full compliance (solution + services), reflect total cost of ownership, and submit best-value pricing at initial response (the University will not seek a BAFO). Pricing must be guaranteed for the Agreement term.

Any alternative or optional scenarios are segregated (e.g., Table 4 for growth/enhancements by institution). Tiered pricing (essential vs. full capabilities) may be presented as options with precise descriptions and unit costs. All optional items must be clearly labeled as optional.

SaaS subscription proposals: If any hourly professional services (e.g., implementation, training beyond baseline, configuration) are billable, include Table 2 – Professional Services Rate Schedule. If you anticipate custom features/deliverables, use Table 3 – Pricing for Custom Features Deliverables (role, hours, rate, cost). If no hourly services or custom deliverables are proposed, Tables 2 and 3 may be left blank

Question(s)

1. For Appendix C (Cost Template), should vendors estimate AI consumption costs (tokens/compute units) based on defined usage scenarios, or propose fixed-fee tiers to cover reasonable usage?

ANSWER: Respondents should provide their best-value pricing using the Appendix C format. Both consumption-based and fixed-fee models are acceptable. The University prefers predictable pricing per Appendix H Requirement #7.

2. About how many people on your team would regularly use the AI solution?

ANSWER: The combined FTE is 30,912 (Students: 25,262; Employees: 5,650). All users should have access per Appendix H Requirement #7.

3. We have a widely adopted product across US Government. Our pricing model doesn't appear to fit into the model of this RFP, however. Is there a way for us to respond with product pricing and professional services as opposed to the current structure outlined in Appendix C.

4. What training and onboarding support is expected from the vendor? The Cost Template includes "Initial Cost' One-Time' Training" but doesn't specify scope.
   Please clarify:
   - Who needs training (students, faculty, staff, administrators, IT support)?
   - What training modalities are required (in-person, virtual live, self-paced online)?
   - Should training be provided at each campus separately or system-wide?
   - Is train-the-trainer support needed for UMS staff to provide ongoing user support?
   - What ongoing training is expected for new features and updates?

ANSWER: Respondents should propose training appropriate to their solution. Training should cover all user types (students, faculty, staff, administrators, IT support). Modalities should be proposed by the vendor. Training may be system-wide or campus-specific. Train-the-trainer options are welcome. Ongoing training for new features should be included.

5. Should vendors provide separate pricing breakdowns by individual campus, and is the University open to multi-year or enterprise flat-rate pricing incentives as outlined in Appendix C?

ANSWER: Enterprise-wide pricing for the full University of Maine System is required. The University encourages price incentives for enterprise, multi-year, or multi-institution awards.

6. Which pricing models in Appendix C (Cost Template) should proposals include? Appendix C includes multiple pricing tables.

   Please clarify:

   ● Should all tables be completed, or only

ANSWER: Use the Appendix C format and Exhibit 1 (Tables 1-4). Complete all tables applicable to the proposed solution. Table 1 is required.

   those applicable to the proposed solution?

   ● If proposing a SaaS subscription model, should Tables 2 and 3 (hourly rates and custom features) still be completed?

   ● Is Table 4 intended for alternative pricing scenarios or optional add-ons?

ANSWER: If proposing a SaaS subscription model, Tables 2 and 3 should be completed if any hourly professional services are billable. Table 4 is intended for alternative pricing scenarios, optional add-ons, or campus-specific enhancements.

7. How does the University recommend structured "firm pricing" given that usage costs for models like GPT-4 or Gemini typically vary based on consumption?

ANSWER: Respondents should include all costs necessary for full compliance, reflect total cost of ownership, and submit best-value pricing at initial response. Pricing must be guaranteed for the Agreement term. The University will not seek a BAFO.

8. Would the University consider a fixed-cost model specifically for hosted open-source models (e.g., Llama) if they are hosted directly within the University's cloud platform?

ANSWER: The University is model-agnostic and open to various pricing models. Respondents should propose their best-value approach using the Appendix C format.

9. Will pricing need to be broken out separately per institution (e.g., UMaine vs. USM vs. UMA), or is a single enterprise-wide price acceptable?

ANSWER: Enterprise-wide pricing for the University of Maine System is acceptable. Pricing does not need to be broken out per individual institution unless the Respondent chooses to offer campus-specific options in Table 4.

10. Does the University require the ability to scale seat counts up or down mid-term to accommodate seasonal enrollment fluctuations across the different institutions?

ANSWER: Respondents should describe their licensing model flexibility. The University values predictable pricing per Appendix H Requirement #7.

11. Should proposals include multiple pricing scenarios? Should vendors provide:
    - Single pricing model for the assumed scope?
    - Multiple pricing options (e.g., Orono-only, phased rollout, full UMS)?
    - Tiered pricing based on different feature sets (essential vs. full capabilities)?

ANSWER: Respondents should provide system-wide UMS enterprise pricing as the primary proposal. Alternative or optional scenarios (e.g., tiered pricing by feature set) may be segregated in Table 4.

12. What is the intended licensing scope and user population for Year 1 pricing?
    - Should pricing be for University of Maine (Orono) only (~12,000 users) with expansion pricing for other campuses?
    - Or should it be system-wide for all UMS campuses (~35,000 users) from Year 1

ANSWER: Pricing should be system-wide for all UMS campuses (30,912 FTE) from Year 1. See Appendix C for campus breakdown.

13. Does the University prefer unlimited enterprise licensing, consumption-based pricing or a hybrid model - and are token usage estimates or sample workload scenarios available to inform consumption-based cost projections?

ANSWER: The University prefers predictable pricing without usage caps per Appendix H Requirement #7. Respondents should propose their best-value model.

14. Will the University purchase licenses for underlying models separately, or should the cost of model access be bundled into the Respondent's proposal?

ANSWER: Yes. All costs including model licensing should be bundled into the Respondent's pricing proposal per Appendix C instructions.

15. How should "unlimited access" be defined in the pricing model?
    - Unlimited prompts/tokens per user with per-user pricing?
    - Flat institutional fee regardless of usage volume?
    - Are there any anticipated usage caps that would trigger overage charges?

ANSWER: The University does not anticipate imposing hard usage caps per Appendix H Requirement #7. Respondents should propose a predictable pricing model (flat-rate or per-user) that provides broad access without overage charges.

16. Appendix H Item 8 (Users): Can you clarify how many total users each year would be required for no caps?

ANSWER: The combined FTE is 30,912 as specified in Appendix C. All users should have access without caps per Appendix H Requirement #7.

17. Could the University provide an estimate of the number of users or departments expected to use the Generative AI services in the initial phase?

ANSWER: The University envisions enterprise-wide deployment. Specific phasing details should be proposed by the Respondent in their implementation roadmap.

18. Is there an expected usage volume (e.g., number of queries, interactions, or AI workloads) vendors should consider when designing the solution architecture?

ANSWER: Detailed usage volume estimates are not available. The platform should support the full FTE count (30,912) without usage caps per Appendix H Requirement #7.

19. Could you provide an estimate of how many end users (teachers, students, and researchers) are expected to utilize the platform across the University of Maine system?

ANSWER: Appendix C specifies: Students 25,262; Employees 5,650; Total 30,912.

RFP Appendix D&D1

MAINE'S FREEDOM OF ACCESS ACT (FOAA)

Answer

The University must adhere to the provisions of the Maine Freedom of Access Act (FOAA), 1 MRSA §401 et seq. As a condition of agreement, a respondent must accept that, to the extent required by the Maine FOAA, any ensuing contractual documents and RFP submissions, are considered public records and therefore are subject to freedom of access requests.

Question(s)

1. Are there any portions of the RFP response or resulting contract that the University anticipates redacting under Maine's Freedom of Access Act (FOAA), particularly with respect to vendor pricing and proprietary solution architecture?

ANSWER: The University must adhere to the Maine Freedom of Access Act (FOAA). Respondents should clearly mark any proprietary or confidential information. Specific FOAA redaction decisions are made on a case-by-case basis.

2. How should respondents mark proprietary content under Maine FOAA, and what may still be disclosed?

ANSWER: Respondents should clearly mark proprietary and confidential content in their submissions. However, the University cannot guarantee that any information will be exempt from disclosure under FOAA.

3. For FOAA handling, do you expect the vendor to provide search/export tooling for prompts/outputs, or will UMS handle that outside the platform?

ANSWER: UMS will handle FOAA requests outside the platform. Vendors are not expected to provide specific search/export tooling for FOAA compliance.

MAINE'S PRIVACY LAWS

Answer

The University will consider a vendor's AI-specific Data Processing Addendum (DPA) as a supplement to the Master Agreement, provided it does not conflict with or supersede the terms of the Master Agreement. The DPA should address FERPA obligations, applicable state privacy laws, and data handling practices specific to AI services. The Master Agreement remains the governing document in all cases.

Question(s)

1. Will the University consider a vendor's AI-specific Data Processing Addendum (DPA) alongside the Master Agreement, particularly to address obligations under FERPA, HIPAA and applicable state privacy laws?

ANSWER: Yes. The University will consider a vendor's AI-specific Data Processing Addendum (DPA) as a supplement to the Master Agreement, provided it does not conflict with or supersede the terms of the Master Agreement.

EULA

Question(s) & Answer(s)

1. Under Section 1.2.1.3 of the RFP it states that by submitting a response, we agree that UMS is not subject to any EULA, terms of service, or otherwise. Could the University elaborate on the reason behind this request? Some OEMs do not have the ability to remove the EULA on any license sold.

   ANSWER: The requirement in Section 1.2.1.3 exists to ensure that all purchases made by the University of Maine System are governed solely by the negotiated Agreement between UMS and the awarded vendor. As a public institution, UMS must ensure that its contractual obligations are contained within a single, consistent agreement that has been reviewed and approved in accordance with University policies, state requirements, and legal standards.

   Click-through agreements, EULAs, online terms of service, or similar end-user terms often contain provisions that conflict with state procurement rules, public record obligations, data-security expectations, venue/jurisdiction requirements, or indemnification limits. Because these terms are typically non-negotiable and can be presented directly to end users, UMS must expressly disallow them to avoid unintentionally accepting conflicting or unauthorized contractual conditions.

   UMS understands that some OEM solutions may technically require the display of a standard EULA that cannot be removed. In such cases, the intent of the requirement is that UMS is not legally bound by those click-through, shrink-wrap, or embedded terms, and that the UMS Agreement takes precedence in all instances. The presence or display of a non-removable EULA is acceptable, provided that it does not supersede, modify, or conflict with the negotiated Agreement.

CLAUSE REVIEW: APPENDIX D1

Answer

The RFP Appendix D1 provides a section for each Respondent to respond to specific language clauses. Please provide your response and/or concern in the Respondent Exception area for any clause where you are requiring or requesting a modification, for the evaluation team's consideration. Scoring of Appendix D1 and D is noted in RFP Section 2.1.2.2.

In addition, the vendor can red line the Master Agreement directly. A response to Appendix D1 is required.

RFP Filename: 06 - 2026-065-RFP-IT-Appendix D-Agreement provides the language, coverage types and minimum limits in Rider B.

The term end date on Appendix D, Section 2, will be modified based on the decision of the overall solution provided at the time of the award notification. You should follow instructions in Appendix C and provide the costs as outlined in the instructions and pricing tables provided.

Question(s)

1. Which Master Agreement clauses are negotiable vs. non-negotiable, and how should exceptions be submitted (e.g., red line in Word)?

ANSWER: The RFP Appendix D1 provides a section for each Respondent to note exceptions. Please provide your response and/or concern in the Respondent Exception area for any clause requiring modification. The vendor may also red-line the Master Agreement directly.

2. What cybersecurity insurance coverage types and minimum limits will be required of the vendor under Rider B?

ANSWER: RFP Filename: 06 - 2026-065-RFP-IT-Appendix D-Agreement provides the language, coverage types, and minimum limits in Rider B.

3. The Master Agreement's "Entire Agreement" clause would supersede standard vendor platform terms of service. How does the University expect this to function operationally for cloud-based AI platforms that rely on standard API terms from underlying model providers such as OpenAI or Anthropic?

ANSWER: The Master Agreement governs the contractual relationship. Standard vendor platform terms of service, EULAs, or clickthrough agreements do not bind UMS. See the EULA section answer for additional detail.

4. Could the University please elaborate on the length of the contract, Appendix D, Pg 2, Line 1 says the contract shall terminate on June 30, 2038, however we were only asked to price between 5 and 10 years.

INTELLECTUAL PROPERTY RIGHTS

Answer

The University's Master Agreement governs all intellectual property (IP) created under the contract. As stated in the Agreement: "Any information and/or materials, finished or unfinished, produced in performance of this Agreement, and all of the rights pertaining thereto, are the property of the University and shall be turned over to the University upon request."

1. AI-Generated Outputs

   All AI-generated outputs produced through the contracted solution during the term of the Agreement—whether text, images, summaries, datasets, workflows, or other generated content—are considered materials produced in performance of the Agreement.
   Therefore, the University of Maine System holds all intellectual property rights to these outputs and may use them without restriction.

2. Fine-Tuned Models Developed Using University Data

   If any fine-tuned, adapted, or customized model is created using University data or developed specifically in performance of the Agreement, then:

   - The resulting fine-tuned model,
   - All derivative model artifacts, and
   - Any associated training datasets, embeddings, prompts, or configuration files are considered materials produced in performance of the Agreement.

   Accordingly, UMS retains ownership of all rights in such models, and the vendor must make them available (including weights, configuration, or exportable artifacts where applicable) upon request.

This ownership applies regardless of hosting location (vendor-hosted, cloud-hosted, or UMS-hosted).

3. Vendor Foundation Models

   This requirement does not transfer ownership of the vendor's underlying foundation models, proprietary architectures, or generalized platform capabilities that existed prior to the Agreement. However:

   - Any UMS-specific fine-tuning,
   - Any derivative model created using UMS data, and
   - Any custom artifacts produced during the engagement fall under University ownership as stated above.

4. Return and Transfer Requirement

   Per the Master Agreement, the vendor must turn over all such materials to the University upon request, including at contract termination.

Question(s)

1. Who holds intellectual property rights to AI-generated outputs and any fine-tuned models developed using University data during the term of the agreement?

ANSWER: Per the Master Agreement, all AI-generated outputs produced through the contracted solution during the Agreement term are considered materials produced in performance of the Agreement and are the property of the University. See the Intellectual Property Rights section for full details.

2. Will the University retain full ownership of AI-generated outputs and derivative works created by users?

ANSWER: Yes. Per the Master Agreement, the University retains ownership of all AI-generated outputs and derivative works. This includes text, images, summaries, datasets, workflows, and other generated content.

RFP Appendix H

General Questions

Answer

Appendix H is represented in filename: 05 - 2026-065-RFP-IT-AppendixH-SolutionRequirementsMatrix

Appendix H, provided in the file "05 – 2026-065-RFP-IT-AppendixH-SolutionRequirementsMatrix," requires Respondents to indicate whether their proposed solution meets each listed requirement. For every item in the matrix, the Respondent must select "Yes," "Partial," or "No" based on the capabilities of the solution being proposed.

Each requirement in Appendix H is labeled as either "Required" or "Preferred," providing guidance on its relative importance:

- Required items represent capabilities the University identifies as essential to the solution.
- Preferred items represent desirable features that may strengthen a proposal but are not mandatory for compliance.

If a Respondent marks a requirement as "Partial" or "No," Appendix H instructions require the Respondent to provide an explanation describing what is needed to meet the requirement, whether additional cost is involved, and where that cost appears within the Appendix C pricing exhibits, along with an estimated timeline for meeting the requirement.

Respondents must submit Appendix H in the provided Excel format as part of their final RFP submission.

This is a new solution for the University of Maine System. As such, the University expects Respondents to demonstrate a thorough understanding of both their proposed solution and the University's stated requirements. Respondents must clearly specify:

- whether their solution meets each requirement,
- if it partially meets the requirement,
- if it does not meet the requirement, and
- where applicable, provide options, enhancements, or alternatives that would enable the solution to meet UMS's needs.

This level of clarity is essential to ensure accurate evaluation and alignment with the University's enterprise strategy and total cost-of-ownership expectations.

Question(s)

1. Appendix H — Solution Requirements Matrix: Appendix H, which carries 45 of 100 evaluation points, does not appear to have been included in the publicly released RFP package. Given its significant weight in the overall scoring and the direct impact it will have on how vendors structure their responses, we respectfully request that UMS release the detailed version of Appendix H at

the earliest opportunity. Timely access to this document is critical to ensure vendors can submit fully informed and competitive responses prior to the March 27 deadline.

2. Beyond the capabilities referenced in Appendix H, has the University developed a prioritized list of required versus desired features - and if so, can it be shared to help BizAcuity understand which functional gaps would be disqualifying versus acceptable? Specifically, we would benefit from understanding whether the University has defined must-have features (e.g., RAG, SSO, multi-campus administration, role-based access), preferred features (e.g., voice interaction, agentic workflows, LMS integration), and features that are aspirational or roadmap-dependent. This will allow Bizacutity to accurately scope the proposed solutions, identify any gaps transparently and avoid both over-engineering and under-delivering on the proposed solution.

ANSWER: Appendix H (Solution Requirements Matrix) provides the prioritized list. Each requirement is labeled as Required or Preferred. Respondents must indicate compliance with each item per the Appendix H instructions.

3. Is there a list of primary enterprise systems (e.g., Brightspace, PeopleSoft, or Salesforce) that the generative AI service is expected to integrate with on day one?

ANSWER: Brightspace (LMS) and Shibboleth (IdP) are the primary systems for integration. Other enterprise system integrations are preferred but not required at go-live.

4. Requirement #28 lists MCP (Model Context Protocol) support as Required. Can you provide examples of the specific MCP server connections or academic databases you envision connecting to?

ANSWER: No specific MCP server connections or academic databases are identified at this time. Respondents should describe their MCP capabilities and available integrations with common higher education systems.

5. Reference: AppendixH-SolutionRequirementsMatrix, Tech/Admin, #41- Please provide a list of desired standard formats? (.doc, .xls., .cvs, .pdf)

ANSWER: Common document formats including PDF, DOCX, XLSX, CSV, TXT, PPTX, and similar standard formats should be supported. Respondents should list all supported formats in their Appendix H response.

6. Reference: AppendixH-SolutionRequirementsMatrix, Core Functionality, #16 Please clarify this requirement, including which tools or products would be needed to enable writing assistance in Microsoft Office or similar applications.

7. Reference: AppendixH-SolutionRequirementsMatrix, Base, #5- Please elaborate on this requirement. Does the solution require support training of SLM model?

ANSWER: Appendix H Requirement #5 (Required) refers to the ability for users or the University to provide custom datasets to work with smaller or private models. Respondents should describe whether their platform supports fine-tuning, RAG, or other approaches for working with University-provided data.

8. What are the primary use cases for multimodal content generation capabilities? Appendix H requires several content generation capabilities:
   Text to Image Generation - Required Multimodal Capabilities – Required Code Generation and Execution Sandbox - Required
   Please clarify the priority use cases:
   - Educational content creation by faculty for course materials?
   - Student projects and assignments?
   - Research visualization and data analysis?
   - Administrative communications?
   - Are there discipline-specific requirements (e.g., scientific diagrams, architectural renderings, chemical structures)?

ANSWER: All listed use cases are relevant: educational content creation by faculty, student projects and assignments, research visualization, and administrative communications. There are no specific discipline requirements at this time. Respondents should describe their multimodal capabilities.

9. What are the requirements for "End user provided links to Google Drive or Microsoft OneDrive and SharePoint folders"? This is listed as Preferred in Appendix H but lacks implementation details.
   Please clarify:
   - Should the integration support

ANSWER: This is a Preferred requirement (Appendix H #6). Respondents should describe their approach to cloud storage integration, including authentication methods, folder-level access, and any limitations.

   real-time synchronization or one-time import?
   - What authentication methods are required (OAuth 2.0, service accounts)?
   - Should users be able to grant access to specific folders or entire drives?
   - Should the system maintain active links to source documents for citation purposes?
   - Are there file size or total volume limitations per user?

ANSWER: Respondents should describe their integration approach. Real-time synchronization is not required; on-demand access is acceptable. Authentication should align with standard OAuth 2.0. Users should be able to grant access to specific folders. File size and volume limitations should be described in the proposal.

10. What system logs and audit data can be retained while maintaining "Zero Data Retention"?For security, compliance, and support purposes, certain logging is typically necessary even with zero data retention for user content. Please clarify what can be retained:
    - Access logs (who logged in, when)?
    - Usage statistics (number of prompts, features used, but not content)?
    - Error logs for troubleshooting?
    - Security audit logs for compliance?

- What is the required retention period for these system logs?

ANSWER: Zero Data Retention (Appendix H #9) applies to user prompts and AI outputs. System-level operational logs (access logs, usage statistics, error logs, security audit logs) may be retained for operational and compliance purposes. Specific retention periods will be configured during implementation.

11. How should "Data supplied by one user would not be used for or shared with any other user without explicit permission" work with team workspaces? Appendix H requires strict data isolation (row 13) but also requires "Team dashboards/workspaces" (row 42).
    Please clarify:
    - How should data sharing work within approved teams or research groups?
    - Should team workspace data remain isolated from other teams?
    - What consent mechanism is required for data sharing within teams?
    - Can instructors access student work submitted through the platform

ANSWER: Data sharing within approved teams should be permission-based with explicit user consent. Team workspace data should remain isolated from other teams by default. Instructors may access student work submitted through the platform within the context of their course. Respondents should describe their data isolation and sharing model.

12. What is the intended use case and oversight model for "Analyze and assign a grade for content based on a set of given criteria"? This is listed as Preferred in Appendix H but raises questions about academic integrity and faculty oversight.
    Please clarify:
    - Is this intended for formative assessment (practice quizzes, self-study) or summative assessment (final grades)?
    - Should the system produce suggested grades or final grades?
    - What level of faculty oversight and approval is expected?
    - Should grading rubrics be customizable by individual instructors?
    - Are there specific assessment types that should NOT use AI grading?

ANSWER: This Preferred requirement envisions AI-assisted formative assessment (practice quizzes, self-study, draft feedback). The system should produce suggested grades, not final grades. Faculty oversight and approval is expected. Grading rubrics should be customizable by individual instructors. High-stakes summative assessment should not rely solely on AI grading.

13. What are the requirements for "User or University provided data sets for SLM (Small Language Models) / Private instances"? Appendix H lists this as Required but does not specify volume, formats, or use cases.
    Please provide guidance on:
    - Expected total volume of University-provided datasets (in GB or TB)?
    - Expected number of discrete datasets (by department, campus, research group)?

- Anticipated file formats beyond those listed in row 38 (PDF, DOCX, CSV, TXT, PPTX)?
- How frequently will datasets need to be updated (real-time, daily, weekly, monthly)?
- Should datasets be shareable across organizational units or remain siloed

ANSWER: Specific volume, format, and frequency details are not available at this time. Respondents should describe their platform capabilities for dataset ingestion, supported formats, volume limits, and update mechanisms. Datasets may be shared across organizational units or siloed depending on administrative configuration.

14. Could the University provide the initial priority use cases for the Generative AI services (e.g., student services, administrative automation, knowledge search, research support, etc.)?

ANSWER: Priority use cases include: teaching and learning (writing assistance, research support, code generation), administration (document analysis, workflow support), accessibility (transcription, translation), and research (data analysis, literature review). See the General Questions section answer for additional detail.

Skills / Multi-Agent Support

Answer

The requirement in RFP Appendix H Lists the Skills Support: Example: having the ability to load and chain on-demand skills for more complex workflows...this is not possible in a contained CustomGPT/Gem/Project environment, as Preferred. The Respondent should specify what is supported in the Respondent's Comments section of the spreadsheet for the Skills Support row.

Question(s)

1. What is the expected scope for "Skills support," "Multi-agent support," and "Local-file support"? These are listed as Preferred in Appendix H with minimal description:
   - Skills support - "having the ability to load and chain on-demand skills for more complex workflows"
   - Multi-agent support - "using coordinated agents for holistic instructional design and verification"
   - Local-file support - "being able to provide knowledge in and have an LLM output to a local folder on a device"

   Please clarify:
   - Are these capabilities needed for Year 1 or future years?
   - What specific use cases would these capabilities enable?
   - Should vendors propose specific implementation approaches or just indicate capability?

ANSWER: These are listed as Preferred in Appendix H. Respondents should describe whether their platform supports these capabilities. Specific use cases and implementation approaches are welcome. Year 1 availability is preferred but not required.

2. Is the University seeking a commercial off-the-shelf (COTS) solution or a custom-built platform? The RFP requirements in Appendix H include both standard enterprise AI capabilities and specialized features (MCP support, Skills support, Multi-agent support, Local-file support). Additionally, Section 1.1.3 states "we are looking for you to provide a higher-education solution based on your expertise."
   Please clarify:
   - Is the University seeking an existing commercial platform with standard features?
   - Is custom development acceptable or required to meet specific UMS requirements?
   - Should proposals include both COTS capabilities and custom development options?
   - If custom development is proposed, what is the expected maintenance and update model for custom features?

ANSWER: The University prefers a commercial off-the-shelf (COTS) product. Custom development is acceptable only if it demonstrably meets or exceeds the requirements. Respondents should include both COTS capabilities and any custom options. See Appendix H for the full requirements list.

LLM/SLM

Answer

The answers provided above and below address the individual questions. Multi-model support is Preferred per Appendix H Requirement #3. The University is model-agnostic and values capability, security, and cost-effectiveness.

Question(s)

1. What is the scope and timeline for Learning Management System (LMS) integration? Appendix H lists "LTI 1.3 Support" as Required with the description "Seamless integration with Learning Management Systems (e.g., Canvas, Blackboard, Brightspace)."
   Please clarify:
   - Which LMS platform(s) does UMS currently use?
   - Should LTI 1.3 integration be functional at go-live (July 1, 2026)?
   - What specific LTI integration features are required (deep linking, grade passback, roster sync)?
   - If integration is not needed immediately, what is the expected timeline?

   ANSWER: UMS currently uses Brightspace as its LMS. LTI 1.3 integration is required per Appendix H Requirement #34, but is not required to be functional at go-live. Priority LTI features include deep linking and, if applicable, grade passback. The specific integration timeline will be determined during implementation planning.

2. In the “Statement of Work” section, the utilization and creation of “Private LLMs” is mentioned. However, this requirement is not listed in the Solutions Requirements Matrix. Could you clarify the

scope of this capability, specifically with respect to fine-tuning models? Additionally, what personas within the University system would be the primary users of this capability?
ANSWER: Requirement #5 in Appendix H refers to the ability for users or the University to provide custom datasets to work with smaller or private models within the platform. This could include fine-tuning, RAG with institutional data, or using university-provided knowledge bases. Research teams, IT staff, and select faculty would be the primary users. The scope of fine-tuning vs. RAG-based customization should be described by the Respondent.

3. Should the solution support multiple LLM providers (multi-model architecture) or will a single LLM provider be acceptable?
ANSWER: Yes. Multi-model support is Preferred per Appendix H Requirement #3. The University prefers a platform that supports multiple LLM providers to avoid vendor lock-in.

4. Will the resulting LLMs be required to be available as an API to the University, similar to standard commercial LLM offerings?
ANSWER: Yes. API access is listed as Preferred in Appendix H Requirement #36. The University desires API availability for developers, researchers, and IT teams to integrate AI capabilities into custom applications.

5. Item 5 in Appendix H states: User or University provided data sets for SLM (Small Language Models) / Private instances (Required). We would appreciate clarification regarding your intended definition of "Private instances."
This term can be interpreted in several ways in the context of generative AI platforms. For example:
    1. Private tenant/workspace isolation – the University operates within a logically isolated tenant environment where customer data is segregated from other customers, and workspaces or projects are isolated within that tenant
    2. Dedicated, single-tenant infrastructure – The provider hosts a dedicated environment for the University using infrastructure reserved for a single tenant
    3. Customer-hosted deployment – the University deploys and manages the instance within its own cloud environment or preferred region

    Can you clarify which of the above options, if any, aligns with the intent of this requirement?
    ANSWER: Option 1 (private tenant/workspace isolation with logical data segregation) is acceptable and aligned with the University's requirements. A dedicated single-tenant infrastructure is not required, provided the vendor can demonstrate adequate data isolation, security controls, and compliance with the data privacy requirements in the RFP.

6. Ability to select from or switch between multiple underlying LLMs (e.g., GPT-4, Claude, Gemini, Llama) to avoid vendor lock-in. In order to achieve the goal stated in Item 3, a single platform will have to be developed (by UMS or by a third-party) that would allow UMS users to select one or more of multiple available models to use in any given situation, and the prompt to be delivered and received back from the model. This is called a "router". For that to work, someone (UMS or a third-party) would have to develop that platform, and UMS would also need to obtain licenses to each of those models.
    - Does UMS intend to develop this router, or are they asking for this to be provided by the respondent?

ANSWER: The Respondent should provide the platform including model routing/switching capabilities. The University does not intend to develop a separate model router.

- To support this, UMS will need to license the multiple models that they want available to them. Does UMS intend to license each of these models separately, or does UMS want the respondent to quote each of the models in its response. If the latter, can UMS provide an appropriate number of users for each model to be licensed?

ANSWER: The Respondent should provide the platform including model routing/switching capabilities. Model licensing costs should be bundled into the Respondent's pricing proposal per Appendix C instructions. The University does not intend to license individual models separately.

7. Could you please clarify whether the State requires vendors to build and host a dedicated/private Large Language Model (LLM) environment as part of the solution, or if solutions leveraging existing AI platforms with multiple pre-integrated LLMs (hosted in a secure environment and configured for the State's use) would also be considered acceptable?
   ANSWER: Solutions leveraging existing AI platforms with multiple pre-integrated LLMs hosted in a secure, compliant environment are acceptable and preferred. The University is not requiring vendors to build and host a dedicated private LLM from scratch. The solution must meet the data privacy, security, and compliance requirements specified in the RFP.

Translation Services

The requirement in RFP Appendix H Lists the Translation Services of English, French, Spanish, Chinese, and other major languages, as Preferred. The Respondent should specify which ones are supported in the Respondent's Comments section of the spreadsheet for the Translation Services row.

Question(s)

1. What languages and translation capabilities are required vs. preferred?
   The RFP mentions international students and Appendix H lists "Translation services - Preferred: English, French, Spanish, Chinese, and other major languages."
   Please clarify:
   ● Which languages are Required vs. Preferred for the platform interface?
   ● Which languages are Required vs. Preferred for AI interaction (prompts and responses)?
   ● Should translation capabilities include technical/academic terminology?
   ● Are there specific use cases prioritized (translating research articles, supporting multilingual classrooms)?

ANSWER: Translation Services (Appendix H #26) is Preferred. English is the primary language. French, Spanish, and Chinese are listed. Respondents should describe all supported languages and translation capabilities, including technical/academic terminology support.

2. While four languages are explicitly specified in the Solution Requirements Matrix(#26), what is the total number of languages the LLM is expected to support overall? Which languages are required vs. preferred?

ANSWER: English, French, Spanish, and Chinese are explicitly listed. Support for additional major languages is preferred. Respondents should list all supported languages in their Appendix H response.

SaaS, Hosting and Resources

Answer

Please note that RFP Appendix H reflects the "Required" elements for the category Base as 'Cloud-based SaaS" and well as other items listed under the Base category.

All licenses for a cloud solution, hosted by the Respondent and all data remains strictly within U.S. borders.

Question(s)

1. Does the University prefer a vendor-hosted SaaS platform, or will private cloud / on-premise deployments also be considered?

ANSWER: The University requires a vendor-hosted, cloud-based SaaS solution per Appendix H Requirement #1 (Required).

2. Does UMS require vendor-hosted SaaS, UMS-hosted cloud, on-prem, or are multiple options acceptable?

ANSWER: Vendor-hosted SaaS is required per Appendix H Requirement #1.

3. The RFP specifies U.S. data residency. Is there flexibility for global redundancy (e.g., disaster recovery backups outside the U.S.), or must all data remain strictly within U.S. borders?

ANSWER: All data must remain strictly within U.S. borders, including disaster recovery and backups. There is no flexibility for global redundancy outside the U.S.

4. Who will be responsible for hosting?

ANSWER: The vendor is responsible for hosting the SaaS solution.

5. If the SaaS solution will be hosted on your cloud platform, is this requirement still applicable?

ANSWER: Yes. Regardless of hosting arrangement, all data must remain within U.S. borders and the vendor is responsible for hosting and operating the service.

6. Is the expectation for the solution to be hosted on the vendor's infrastructure, or will it be hosted within the University's specific cloud platform environment?

ANSWER: The solution should be hosted on the vendor's infrastructure as a SaaS offering.

7. Please clarify if the SaaS solution will be hosted on your cloud platform, or ours (Impact Makers)?

ANSWER: The vendor is responsible for hosting. The University expects a vendor-hosted SaaS solution.

8. Can the University clarify if this bid is intended for professional services (consulting/integration) or a licensed SaaS platform?

ANSWER: The University is seeking a licensed SaaS platform, not professional services (consulting/integration) alone. Implementation support services may be included as part of the proposal.

9. Does the University have a preferred deployment model for the Generative AI solutions (e.g., vendor-hosted SaaS, University-hosted, public cloud, or hybrid)?

ANSWER: Vendor-hosted SaaS is the preferred and required deployment model per Appendix H Requirement #1.

10. Could you please clarify whether respondents are expected to provide a fully hosted Generative AI platform, implementation services, or a combination of both?

ANSWER: The University expects a fully hosted Generative AI platform. Implementation support services should be proposed as needed.

11. Do you accept offshore resources?

ANSWER: No. All resources must be U.S.-based.

12. Do resources need to be US-based or can we explore using nearshore or offshore resources? Is this true for both the implementation and the ongoing maintenance and support?

ANSWER: All resources must be U.S.-based for both implementation and ongoing maintenance/support. Offshore and nearshore resources are not permitted.

13. Do you expect the vendor to perform any tasks on-site, or can all work be performed remotely?

ANSWER: All work may be performed remotely. The University is requesting licenses for a cloud solution.

14. Does UMS require that all development, implementation, support and maintenance activities be performed exclusively by personnel located within the United States or are offshore and nearshore resources permissible provided that appropriate data security controls, access restrictions and compliance safeguards are in place?

ANSWER: Yes. All development, implementation, support, and maintenance activities must be performed by personnel located within the United States.

15. If a cloud-based solution is preferred, does the University have a preferred or approved cloud provider (e.g., AWS, Microsoft Azure, Google Cloud)?

ANSWER: The University does not have a preferred or approved cloud provider. The vendor is responsible for selecting and managing cloud infrastructure.

16. What is your preferred hosting service provider (AWS or Azure)?

ANSWER: No preference. The vendor should propose the cloud infrastructure appropriate for their solution.

SSO

Question(s) & Answer(s)

1. Do all six universities within the System utilize a single, unified platform for Single Sign-On (SSO)?
   ANSWER: Yes

2. Appendix H mentions LTI 1.3 and SSO. Are there specific LMS platforms (Canvas, Blackboard, Brightspace) or identity providers (Azure AD, Okta, Shibboleth) currently in use that vendors should prioritize?

ANSWER: UMS uses Brightspace as its LMS and Shibboleth as its identity provider. Vendors should prioritize integration with these platforms.

3. Does the University require integration with Single Sign-On (SSO) platforms such as Azure AD, Okta, or Shibboleth?
   ANSWER: Yes. The University requires SSO integration with Shibboleth. All six universities use a single, unified Shibboleth-based SSO platform.

4. Will the University require integration with its existing Single Sign-On (SSO) and Multi-Factor Authentication (MFA) systems for all 10 campuses?
   ANSWER: Yes. All campuses use a single unified SSO (Shibboleth). MFA is managed at the identity provider level and does not need to be implemented separately by the vendor.

5. Which identity provider is currently used for Single Sign-On (e.g., Shibboleth, Azure AD, Okta), and will the University require role-based restrictions on model or feature access by user type?
   ANSWER: Shibboleth is the identity provider. Role-based access restrictions on model or feature access by user type are preferred per Appendix H Requirement #39.

6. What are the specific identity provider and authentication requirements? The RFP requires "Single Sign-On (SSO) integration" but does not specify protocols or providers. Additionally, Section 1.1.2 mentions the University of Maine System consists of multiple universities and centers.
   Please clarify:
   - Which SSO protocol(s) must be supported (SAML 2.0, OAuth 2.0, OpenID Connect, CAS, Shibboleth)?
   - Is there a single identity provider for all UMS institutions or multiple IdPs?
   - What user attributes must be passed through SSO (username, email, role, department, campus affiliation)?
   - Must the solution support multi-factor authentication (MFA)?

   ANSWER: SAML 2.0 via Shibboleth is the required SSO protocol. There is a single identity provider for all UMS institutions. User attributes passed through SSO should include at minimum: username, email, role, and campus affiliation. MFA is managed at the IdP level and does not require separate vendor implementation.

RFP Appendix I

Question(s) & Answer(s)

1. Is it mandatory to submit the HECVAT and VPAT along with our proposal, or can they be submitted once the contract award is made?
   ANSWER: Yes

2. Under Appendix I, Item 3-A, it says " The HECVAT FULL must be completed when ... the solution includes providing consulting services" Can the University please elaborate on it's definition of "consulting services.
   ANSWER: Consulting services includes any professional services including but not limited to, support, implementation, training, professional services offered.
   Based on the University's guidance, we understand that "consulting services" refers to any professional services associated with the proposed solution, including but not limited to:
   - Implementation services
   - Configuration or deployment support
   - Training or enablement services
   - Ongoing support or advisory services
   - Any additional professional services offered as part of delivering or maintaining the solution
     We acknowledge that when our offering includes any of the above service types, the HECVAT Full assessment is required.

   In addition, we recognize that completion of the HECVAT Full is also required when:
   a. The data transmitted, stored, processed, or accessed includes protected health information (PHI) or any data subject to the Health Insurance Portability and Accountability Act (HIPAA).
   b. The solution involves processing credit or debit card payment transactions, including any workflows subject to PCI-DSS requirements.

   We will ensure the appropriate HECVAT version is completed in accordance with these criteria.

3. Beyond the HECVAT, are there additional compliance frameworks the solution must satisfy (e.g., HIPAA, FedRAMP), and does the University require customer-managed encryption keys?
   ANSWER: FERPA compliance is required. HIPAA is not anticipated for the initial deployment of this solution. The University does not require customer-managed encryption keys at this time. See Appendix I for the full set of compliance requirements including HECVAT and VPAT.

4. The RFP requires SOC 2 Type II or ISO 27001. Will equivalent certifications (e.g., FedRAMP, HIPAA, or other recognized standards) be considered acceptable?
   ANSWER: Yes

5. Regarding Requirement #12 (SOC 2 Type II or ISO 27001): Will the University accept evidence of an in-progress SOC 2 Type II audit with a documented completion timeline within 30-40 days and compensating controls documentation, or is a completed certification required at the time of proposal submission?
6. ANSWER: Submit your proposal and note when you expect to obtain the certification.

7. Does the University require SOC 2, ISO 27001, or other compliance certifications for the platform?

ANSWER: Yes. SOC 2 Type II or ISO 27001 certification is required per Appendix H Requirement #12. Equivalent certifications such as FedRAMP are also acceptable.

8. What data classifications will the AI system be permitted to access (e.g., public data, FERPA-protected student records, research data, PHI), and does the University require Zero Data Retention (ZDR) across all model providers?
   ANSWER: The AI system will primarily handle FERPA-protected student data and general institutional/research data. PHI (Protected Health Information) is not anticipated for the initial deployment. Zero Data Retention should be available as a configurable option per Appendix H Requirement #9.

9. What types of data (e.g., student records, research data, financial data, PII, FERPA-protected data) will be processed through the platform, and what are your data classification and compliance requirements?
   ANSWER: Primary data types include FERPA-protected student records, research data, general administrative data, and user-generated content (prompts, uploads). PHI is not anticipated. Compliance requirements include FERPA, U.S. data residency, SOC 2 Type II or equivalent, and the data privacy provisions outlined in the Master Agreement and Appendix I.

10. Do you require that all data remain within a specific geographic boundary (e.g., U.S.-based data centers), and are there specific regulatory frameworks (FERPA, HIPAA, ITAR, etc.) that the solution must comply with?
    ANSWER: Yes. All data must remain within U.S. borders, including disaster recovery and backups. FERPA compliance is required. HIPAA is not currently anticipated. See Appendix H Requirements #10-13 and Appendix I for full compliance requirements.

11. What FERPA-specific controls must the AI solution support when handling student data and academic records?
    ANSWER: The solution must be capable of acting as a School Official under FERPA. Key controls include: strict no-training policy on University data (Requirement #8), zero data retention option (Requirement #9), data isolation between users (Requirement #13), U.S. data residency (Requirement #11), and automated data retention/deletion policies (Requirement #40). See Appendix H for the full list.

12. Does the University require centralized administrative controls for managing AI usage across campuses and departments? Must the solution comply with WCAG 2.1 AA accessibility standards, and is VPAT documentation required for submission?
    ANSWER: Yes

13. CAG 2.1 AA compliance is listed as required. Should vendors also anticipate WCAG 2.2 compliance in the near term?
    ANSWER: Yes

14. Does the WCAG 2.1 AA requirement apply only to the user interface (UI), or does it also extend to AI-generated outputs such as exported PDF reports, data visualizations, and code snippets?
    ANSWER: Yes

15. Must the solution comply with WCAG 2.1 AA accessibility standards, and is VPAT documentation required for submission?
    ANSWER: Yes

16. Should the AI interface support screen readers and other assistive technologies across both web and mobile environments?
    ANSWER: Yes. Per WCAG 2.1 AA (and upcoming 2.2) requirements, the AI interface must support screen readers and assistive technologies across both web and mobile environments.

MAINE'S PUBLIC UNIVERSITIES
UNIVERSITY OF MAINE SYSTEM

RFP Appendix J

Question(s) & Answer(s)

1. What are the primary content repositories you expect users to connect (e.g., document management, cloud storage, shared drives), and are those standardized across campuses?
   ANSWER: The University does not prescribe specific content repositories at this time. Respondents should describe their available integration capabilities with common platforms (e.g., SharePoint, Google Drive). Specific integrations will be determined during implementation.

2. What is the Student Information System (SIS) (and any key admin systems like HR/finance/CRM) you expect AI to integrate with, and is integration required at go-live or later?
   ANSWER: SIS and administrative system integration is not required at go-live. Respondents should describe available integration capabilities. Specific system integrations will be planned during implementation.

3. Any constraints for managed vs unmanaged devices (BYOD), browser support, VDI, or mobile app requirements?
   ANSWER: The solution must be browser-based and accessible from any modern browser on any OS (Windows, macOS, iOS, Android, ChromeOS) per Appendix H Requirements #2 and #37. BYOD is expected; no special device management is required.

4. Any network/security controls that will affect access (e.g., IP allowlisting, VPN requirements, campus network segmentation)?
   ANSWER: No special network constraints. Standard HTTPS access is expected. No IP allowlisting, VPN, or campus network segmentation requirements for accessing the solution.

5. Do you have an existing enterprise search / Knowledge Base platform you want AI to integrate with (and via what mechanism: API, connector, search index)?
   ANSWER: No existing enterprise search or Knowledge Base platform is identified for AI integration at this time.

6. Should the solution provide APIs or SDKs for integration with custom University applications?
   ANSWER: Yes. API access is Preferred per Appendix H Requirement #36. Respondents should describe available APIs and SDKs.

7. Do you require API access for developers, researchers, or IT teams to build custom applications or workflows on top of the generative AI platform?
   ANSWER: Yes. Per Appendix H Requirement #36, API access for developers, researchers, and IT teams is Preferred.

8. Does the University expect vendors to propose specific AI applications or solutions, or is the focus primarily on providing AI strategy and advisory services?
   ANSWER: The University is seeking a fully functional AI platform/solution, not AI strategy or advisory services. Respondents should propose specific AI applications and capabilities as outlined in Appendix H.

9. Are there any restrictions on using third-party AI APIs or external AI services for model inference?

ANSWER: No restrictions beyond the data security, privacy, and compliance requirements specified in the RFP (Appendix H Requirements #8-13 and Appendix I).

10. Is there an incumbent enterprise AI solution currently in place and will Bizacuity be expected to manage a formal transition or data migration from that system?

ANSWER: No. There is no incumbent enterprise AI solution currently in place. This is a new procurement. No transition or data migration from a previous system is expected.

11. Who at the University of Maine is requesting this information and leading the project for the Enterprise Generative AI Services initiative?
    ANSWER: All inquiries should be directed to the contact information provided on the RFP cover page: UMSResponses@maine.edu.

12. Is this RFP intended to cover the first year of a larger, multi-year project, or is it designed to address the entire enterprise AI initiative?
    ANSWER: The RFP covers the full Agreement term, which includes an initial 5-year term with optional renewals up to 10 years. See Appendix C and Appendix D for details.

13. Have you already set up AI governance for responsible use of AI? If you have not yet defined, would you consider this is scope for this RFP?

ANSWER: Yes. The University of Maine System has an existing AI governance framework and responsible AI policies in place. This is not in scope for this RFP. Vendors should ensure their solution supports the University's existing governance framework.

14. Does the University expect a single enterprise AI platform for all campuses or multiple approved solutions that departments may adopt independently?
    ANSWER: The University prefers a single enterprise AI platform for all campuses, as stated in RFP Section 2.2.

15. Which enterprise systems should the AI solution integrate with (e.g., Learning Management Systems, identity providers, research repositories)?
    ANSWER: Brightspace (LMS) and Shibboleth (IdP) are priorities for integration. Other enterprise system integrations are preferred but not required at go-live. Respondents should describe their integration capabilities in Appendix J.

16. What payment processing options are in scope for vendor vs your current cloud provider?
    ANSWER: Payment processing is not in scope for this procurement.

17. What is UMS's expectation regarding audit logging — who needs access to logs, at what granularity, and for how long must they be retained?
    ANSWER: Respondents should propose audit logging capabilities. The University expects access logs, usage statistics, and security audit logs. Content logs should align with configurable data retention/deletion policies per Appendix H Requirement #40. Specific retention periods will be defined during implementation.

18. Should authentication be systemwide or per-campus, and do you anticipate multiple IdPs in scope at go-live?
    ANSWER: System-wide authentication using a single IdP (Shibboleth). One IdP at go-live.

19. Do you want authorization driven entirely by IdP groups/roles, or do you expect in-app role management too?
    ANSWER: Both. Authentication via the IdP (Shibboleth); in-app role management for granular access control per Appendix J Question 11b (RBAC).

20. Do you need separate administrators per campus/department with scoped visibility, and should admins be able to set distinct policies per unit?
    ANSWER: Yes. Per Appendix J Question 14, multi-institution support with logical data separation is required. Campus-level administrators with scoped visibility and the ability to set distinct policies per unit are Preferred.

21. Which LMS is primary, and which LTI workflows are "day 1" vs "later" (e.g., assignment support, feedback, content creation)?
    ANSWER: Brightspace is the primary LMS. Specific LTI workflows and timeline (day 1 vs. later) will be determined during implementation planning. LTI 1.3 support is Required per Appendix H #34.

22. Do you anticipate different feature sets or controls by population (student vs faculty vs staff), or one consistent experience?
    ANSWER: Respondents should propose role-based access options. Different feature sets or controls by population (student vs. faculty vs. staff) are preferred per Appendix H Requirement #39 (Granular access control).

23. Are there specific categories of University data (e.g., student records under FERPA, research data under sponsored grants, PHI) that must be explicitly prohibited from being uploaded to the platform, or is the expectation that guardrails prevent misuse?
    ANSWER: Primarily FERPA-protected student records and general institutional/research data. PHI is not anticipated. See Appendix H Requirements #8-13 for data handling requirements.

24. Is residency a constraint on storage only, or also on processing/inference and support operations?
    ANSWER: U.S. residency is a constraint on all data at rest, in transit, and during processing/inference. All data must remain strictly within U.S. borders.

25. Is it acceptable to retain system information logs for identifying performance issues or future enhancements. This ability could be turned off.
    ANSWER: Yes. System-level logs (access logs, usage statistics, error logs, security audit logs) may be retained for operational and compliance purposes, provided they do not contain user prompt or response content. Content retention should be configurable per Appendix H Requirement #40.

26. Are vendors required to ensure that University prompts and data are not used to train public models?
    ANSWER: Yes. Appendix H Requirement #8 is Required: Vendor MUST guarantee that University data (prompts, uploads, outputs) is NOT used to train, fine-tune, or improve foundation models or third-party models.

27. Are there defined data retention and deletion policies for prompts, uploaded documents, and generated outputs?
ANSWER: Yes. Appendix H Requirement #40 (Required) specifies admin ability to set automated deletion schedules for user data. Specific retention periods will be configured during implementation.

28. Are there any restrictions or policies governing AI-generated research content or academic publications?
ANSWER: The University has AI governance policies in place. Vendors should ensure the platform supports content guardrails and safety features per Appendix H Requirement #47. Specific policies regarding AI-generated research content will be managed by UMS.

29. Should administrators be able to define usage policies, quotas, or restrictions by user role (students, faculty, staff)?
ANSWER: Yes. Per Appendix H Requirement #39 (Granular access control, Preferred) and Requirement #43 (Enterprise management capabilities, Required), administrators should have the ability to define usage policies and restrictions by department or user group.

30. Are usage analytics and monitoring dashboards required for administrators to track adoption and system performance?
ANSWER: Yes. Per Appendix H Requirement #44 (Usage Analytics, Preferred) and Appendix J Question 12 (Monitoring), dashboards showing adoption rates, active users, token usage, and costs are expected.

31. Are there requirements for private connectivity (e.g., VPN/private link) versus public internet access?
ANSWER: No private connectivity (VPN/private link) is required. Standard HTTPS access over the public internet is acceptable, provided data is encrypted in transit using TLS 1.2 or higher per the Technical Interface requirements.

32. Can you provide specific examples of the teaching, learning, and research use cases you envision? (e.g., curriculum development, research paper analysis, tutoring assistance, lab data analysis, grant writing, etc.)
ANSWER: The University envisions broad use across teaching, learning, research, and administration. Specific use cases include: writing assistance and feedback, research article summarization and citation, course content creation, code generation for CS education, document analysis, transcription/translation for accessibility, and administrative workflow support.

33. What administrative functions do you anticipate leveraging generative AI for? (e.g., policy drafting, budget analysis, HR communications, student services, compliance reporting, etc.)
ANSWER: Administrative functions include document analysis and summarization, policy research, report generation, data analysis, and operational workflow support. The platform should serve as a general-purpose AI tool, not specialized for a single administrative function.

34. Will all users require the same level of access and capabilities, or do you envision tiered access levels (e.g., basic access for students vs. advanced/administrative access for faculty and researchers)?

ANSWER: The University anticipates differentiated access by role (student, faculty, staff, administrator) per Appendix H Requirement #39. Respondents should propose role-based access tiers that align with the RBAC capabilities described in Appendix J Question 11b.

35. What existing systems (e.g., LMS such as Canvas/Blackboard, SIS, ERP, identity providers like Active Directory/SSO/Shibboleth) would need to integrate with the generative AI platform?
ANSWER: Brightspace (LMS) and Shibboleth (IdP) are the primary systems. Other system integrations will be planned during implementation. Respondents should describe available connectors and integration capabilities.

36. Regarding the 'analysis' of Google Drive or Microsoft OneDrive and SharePoint folders, could you clarify the expected output? For example, are you looking for metadata extraction, or full-text semantic search across these data sources?
ANSWER: This Preferred requirement (Appendix H #6) envisions users linking to their own cloud storage folders so the AI can analyze documents stored there. Respondents should describe their approach, including authentication method, whether it supports one-time import or ongoing sync, and any file size or volume limitations.

37. Does your requirement for zero data retention imply a strictly 'Ephemeral Mode' for all user interactions? Specifically, do you require that all document indexes and session contexts be purged from memory immediately upon session termination, even if this requires re-indexing the data for future sessions?
ANSWER: Zero Data Retention (Appendix H Requirement #9) means user prompts and AI outputs are not retained by the vendor after the session ends. System-level operational logs (access logs, usage statistics, error logs) may be retained. The option should be configurable at the administrative level, allowing UMS to choose between ephemeral mode and standard retention with automated deletion schedules per Requirement #40.

38. To better understand the scope of 'Writing assistance,' are you looking for a specialized tool tailored to your specific domain (SLMs and private data), or do you intend to use a general-purpose assistant like Microsoft Copilot for these tasks?
ANSWER: Writing assistance (Appendix H Requirement #16) refers to AI-powered tools within the platform for refining drafts, improving tone, checking grammar, and structuring content. This is a capability of the AI platform itself, not a plugin or extension for Microsoft Office.

39. Regarding the 'assign a grade/feedback' requirement, do you envision the AI providing a final grade directly to the student, or is it intended to provide a 'draft grade' and feedback for a faculty member to review and approve before release?
ANSWER: The grading/feedback requirement (Appendix H #18, Preferred) envisions AI-assisted formative assessment where the AI provides suggested feedback and scoring based on instructor-defined criteria/rubrics. Faculty oversight and final approval are expected. This is not intended for final summative grading without human review.

40. Regarding 'Transcription,' do you require real-time, live speech-to-text (e.g., for live lectures or meetings), or is the primary use case the asynchronous processing of pre-recorded audio and video files?

ANSWER: Transcription (Appendix H Requirement #20, Required) includes both real-time speech-to-text and pre-recorded audio/video transcription capabilities. Respondents should describe which modes are supported.

41. Does your requirement include transcription for non-English languages or technical terminology (e.g., medical, legal, or STEM-specific jargon)? This will help us determine if a specialized SLM is needed to handle your department-specific vocabulary.
ANSWER: Yes. Non-English language transcription is desired, consistent with the Translation Services requirement (Appendix H #26). Technical/academic terminology support is preferred.

42. When it comes to audio transcription, is "good enough" accuracy fine, or would you prefer higher precision for more technical conversations? No consideration has been given for fine-tuning the model as of now
ANSWER: High accuracy is preferred, especially for academic content. Respondents should describe their transcription accuracy rates and any options for domain-specific tuning.

43. Do you have requirements for multi-modal capabilities beyond text generation, such as image generation, audio processing, video analysis, or data visualization?
ANSWER: Yes. Appendix H Requirement #24 (Multimodal Capabilities, Required) includes the ability to process and reason across mixed inputs such as images, charts, and documents. Respondents should describe their full multimodal capabilities.

44. When you reference "private LLMs," can you elaborate on what this means for your institution? Are you looking to fine-tune models on proprietary university data, host self-managed models, or create isolated datasets/knowledge bases within a shared platform?
ANSWER: Private LLMs (Appendix H Requirement #5) refers to the ability for users or the University to provide custom datasets that are used with smaller or private model instances within the platform. This includes RAG-based knowledge grounding with institutional data and, where supported, fine-tuning capabilities. The data used remains University property and must be isolated per Requirement #13.

45. Do you have existing datasets, knowledge bases, or document repositories that you would like ingested into the platform for retrieval-augmented generation (RAG)?
ANSWER: The University has various institutional datasets, knowledge bases, and document repositories across campuses. Specific sources for initial RAG ingestion will be identified during implementation. Respondents should describe their data ingestion capabilities, supported formats, and connector options.

46. Does the University require document ingestion and retrieval-augmented generation (RAG) capabilities using institutional data sources?
ANSWER: Yes. RAG is Required per Appendix H Requirement #17 (Citation and Grounding). The solution should support ingestion of University documents and provide citations/references to source materials in AI responses.

47. What level of training and change management support do you require for onboarding students, faculty, and staff onto the platform?

ANSWER: Respondents should propose training and change management appropriate to their solution. See the Training section answer and Appendix C pricing instructions for how to structure training costs.

48. Does your team prefer to manage the deployment infrastructure directly, or would you prefer a serverless approach to minimize infrastructure management?
ANSWER: The University expects a vendor-hosted SaaS solution. The vendor is responsible for deployment infrastructure. See the SaaS/Hosting section answer.

49. Could you provide more detail on the expected workflow for 'User or University provided data sets'? Specifically, do you intend to host the SLM on your own private instances, and what is the expected volume of the data sets you will be providing?
ANSWER: Appendix H Requirement #5 envisions users or administrators uploading datasets (PDF, DOCX, CSV, etc.) to create private knowledge bases or fine-tuned model instances within the platform. Respondents should describe the upload workflow, supported formats, size limits, and how data isolation is maintained.

50. Text to Speech (Accessibility): Could you please provide a high-priority 'use case' scenarios that repsent a successful outcome for this requirement?
ANSWER: Text-to-Speech (Appendix H Requirement #21, Required) is an accessibility feature enabling AI-generated text to be read aloud. Priority use cases include accessibility compliance for visually impaired users, language learning support, and content consumption across devices.

51. Image to Text (OCR / Description): Could you please provide a high-priority 'use case' scenarios that represent a successful outcome for this requirement?
ANSWER: Image-to-Text (Appendix H Requirement #22, Required) includes OCR and image description capabilities. Priority use cases include extracting text from scanned documents, describing images for accessibility, and analyzing charts/diagrams for research and coursework.

52. Text to Image Generation: Could you please provide a high-priority 'use case' scenarios that repsent a successful outcome for this requirement?
ANSWER: Text-to-Image (Appendix H Requirement #23, Required) supports content creation by faculty and students. Use cases include educational materials, visual aids for presentations, research visualization, and creative projects. No specific discipline-specific rendering requirements (e.g., chemical structures) are mandated.

53. Translation services: Could you please provide a high-priority 'use case' scenarios that repsent a successful outcome for this requirement
ANSWER: Translation Services (Appendix H Requirement #26, Preferred) support UMS's international student population and multilingual research needs. Priority use cases include translating research articles, supporting multilingual classrooms, and enabling non-English-speaking students to interact with the platform in their preferred language.

54. Does the solution perform research assistance against a pre-integrated source of academic journals, or does it require the user to manually upload PDFs to a local library for analysis?
ANSWER: Research assistance (Appendix H Requirement #27, Preferred) may integrate with external academic databases as described in MCP support (Requirement #28). The specific sources for integration will depend on the Respondent's capabilities and available connectors.

55. When you say 'chaining skills,' are you looking for the system to automatically move from Task A to Task B (like an autonomous agent), or do you want a drag-and-drop workflow builder?
ANSWER: Skills support (Appendix H #29, Preferred) envisions the system allowing users to load and chain on-demand capabilities for complex multi-step workflows. Respondents should describe whether their platform supports this type of composable workflow automation.

56. Does this requirement imply a locally-installed application (thick client), or a web-based tool that uses a local agent/bridge to access the file system?
ANSWER: Local-file support (Appendix H #31, Preferred) refers to the ability for the AI to read from and write to a local folder on a user's device. This could be implemented through a browser-based interface, desktop application, or API integration. Respondents should describe their approach.

57. Will there be any other systems besides Canvas, Blackboard, Brightspace that need to be integrated with the AI solution? And what kind of data each of these systems hold?
ANSWER: Brightspace is the primary LMS for UMS. LTI 1.3 is the required integration standard. No other LMS platforms are identified for day-one integration.

58. What is the expected behavior for user provisioning? Should the system automatically create a new user account upon the first LTI launch ('Just-in-Time' provisioning), or will accounts be pre-synced?
ANSWER: User provisioning should leverage the existing Shibboleth SSO integration. The system should support automatic provisioning/deprovisioning based on SSO authentication with appropriate role attributes.

59. Can you provide the full list of applications which need to be integrated with the AI solution and a brief details of each systems in terms of functionality and data availability?
ANSWER: Brightspace (LMS) and Shibboleth (IdP) are the primary integration requirements. Respondents should list all available pre-built connectors and API capabilities. A comprehensive integration list is not prescribed; Respondents should propose based on their solution's capabilities.

60. For document processing, can you give a sense of how much internal content it would need to handle (e.g., approximate size, number of files, and the types of documents)?
ANSWER: The University does not have a specific estimate for internal content volume at this time. Respondents should describe their platform's content ingestion capacity, supported formats, and any volume limitations.

61. How often do you expect people to interact with the Chatbot? For example, per person per day, and do you see any patterns when it's used?
ANSWER: Detailed per-user interaction estimates are not available at this time. The platform should support the full FTE count (30,912) as described in Appendix C without usage caps per Requirement #7.

62. Who would be using the Chatbot most often? Could you provide a ballpark estimate of the user numbers, along with the types of roles or access levels?
ANSWER: All user populations (students, faculty, staff) are expected to use the platform. The combined FTE is 30,912 (Students: 25,262; Employees: 5,650) per Appendix C.

63. What kind of data sensitivity or compliance requirements should we keep in mind (e.g., encryption, GDPR, other standards)?
ANSWER: FERPA compliance is Required. PHI/HIPAA is not anticipated. See Appendix H Requirements #8-13 for data sensitivity and compliance requirements, and Appendix I for security assessment requirements.

64. What are the user interfaces required for the Chatbot? Or will it get integrated with the existing business applications?
ANSWER: The primary interface is a web-based application accessible via standard browsers. LTI 1.3 integration with Brightspace is also required. API access is Preferred. Respondents should describe all available interfaces (web, mobile, API, LMS embed).

65. Additional details on the chatbot use cases – from a Business and Technical perspective.
ANSWER: The University envisions broad enterprise use cases. Respondents should describe how their platform addresses the specific requirements in Appendix H, including conversational AI, document analysis, code execution, and multimodal capabilities.

66. Will the selected vendor be responsible for provisioning and managing infrastructure, or will the solution be deployed within the University's existing cloud environment?
ANSWER: No. The University expects a vendor-hosted SaaS solution. The vendor is responsible for provisioning and managing infrastructure.

67. Does the University have a preferred Generative AI technology stack or model providers (e.g., OpenAI, Azure OpenAI, Anthropic, or open-source models)?
ANSWER: No. The University is model-agnostic and does not prescribe a specific technology stack or model provider. Multi-model support is Preferred per Appendix H Requirement #3.

68. Are vendors allowed to propose open-source large language models hosted in a secure environment if they meet the University's security requirements?
ANSWER: Yes. Open-source LLMs hosted in a secure, compliant environment are acceptable, provided they meet the data privacy, security, and performance requirements of the RFP.

69. Is the University planning to begin with a pilot or proof-of-concept phase before broader deployment?
ANSWER: The University is open to phased implementation approaches. Respondents should propose a realistic implementation plan as described in the Implementation section answer.

70. What are the primary use cases the University expects to support in the first year (e.g., student support chatbot, faculty research assistant, administrative automation) and is the goal a single centralized AI platform or multiple specialized tools?
ANSWER: Year 1 use cases span teaching/learning (writing assistance, research support, code generation), administration (document analysis, workflow support), and accessibility (transcription, translation, text-to-speech). See Appendix H for the full requirements matrix.

71. Is guest access or a public-facing AI interface (e.g., a website chatbot) within scope and what peak concurrent user loads should Bizacuity plan for?
ANSWER: Guest access or a public-facing AI chatbot is not in scope for this procurement. The solution is for authenticated UMS users only.

72. What are the University's expectations regarding prompt and response log retention periods, and is there a requirement for bulk export of chat logs for e-discovery or archival purposes?
ANSWER: Prompt and response log retention should be configurable by administrators per Appendix H Requirement #40. Zero Data Retention should be available as an option per Requirement #9. Specific retention policies will be determined during implementation.

73. Does the University prefer a vendor-neutral, multi-model platform supporting providers such as OpenAI, Anthropic and Google, or is a single-model solution acceptable? Are there restrictions on using non-US cloud infrastructure?
ANSWER: Yes. The University prefers a vendor-neutral, multi-model platform per Appendix H Requirement #3. Respondents should describe which models are available and how model switching is supported.

74. Is there a requirement for open-source models or private/on-premises model hosting, and what maximum response latency is acceptable for model inference at scale?
ANSWER: Open-source models are acceptable. Private/on-premises model hosting is not required but may be proposed as an option. The primary requirement is a vendor-hosted SaaS solution per Appendix H Requirement #1.

75. Which data sources should be prioritized for initial RAG ingestion (e.g., Canvas LMS, SharePoint, University websites, research repositories), and what is the estimated volume of documents to be indexed at launch?
ANSWER: Specific data sources for initial RAG ingestion will be identified during implementation. Respondents should describe available connectors and ingestion capabilities.

76. Does the University currently operate a vector database or knowledge management system the solution should leverage, and is document-level access control within the RAG system a requirement?
ANSWER: No. The University does not currently operate a vector database or knowledge management system for AI purposes.

77. Which core platforms must the AI solution integrate with at launch (e.g., Canvas LMS, Student Information System, Microsoft 365, Google Workspace), and does the University require real-time API integrations or will scheduled batch integrations be sufficient?
ANSWER: Brightspace (LMS) and Shibboleth (IdP) are the required integrations at launch. Other integrations are preferred but not required at go-live.

78. Should the solution support multiple institutions operating within a single tenant with logical data separation by campus, and should billing and token usage reporting be separated by institution?
ANSWER: Yes. Per Appendix J Question 14, the solution should support multiple institutions within a single tenant with logical data separation (workspaces). Billing separation by campus is also required.

79. Should administrators at individual campuses have the ability to independently create and manage departmental workspaces, or will administration be centralized at the system level?
ANSWER: Yes. Campus-level administrators should have the ability to configure policies, manage workspaces, and control access within their institution, per Appendix J Questions 14 and Appendix H Requirement #39.

80. Is the University seeking a vendor that builds proprietary models, or a provider that integrates existing models (e.g., GPT-4, Gemini) via API?
ANSWER: The University is open to either proprietary or open-source approaches, or a platform that provides access to multiple model providers. The key requirement is meeting the functional, security, and compliance requirements in the RFP.

81. Regarding the use of external platforms, is the expectation to use links from other platforms only for limited, specific use cases?
ANSWER: Appendix H Requirement #6 (Preferred) envisions users providing links to their cloud storage (Google Drive, OneDrive, SharePoint) for the AI to analyze documents stored there. Respondents should describe their integration approach.

82. What is the scope of MCP (Model Context Protocol) integration requirements?
    - Should vendors propose custom MCP server development for specific University databases?
    - Or only provide platform capability for University IT to deploy their own MCP servers?
    - Are there specific systems identified for integration (library systems, research databases, etc.)?

ANSWER: No specific systems beyond Brightspace and Shibboleth are identified for mandatory integration. Respondents should describe available integrations with common higher education systems (library systems, research databases, etc.).

83. Does UMS envision using institutional content (e.g., course materials, research repositories, policy documents) primarily to ground responses, or does UMS also want to fine-tune/train custom private models on University data? If either or both, which campuses or departments have content ready and what is the approximate volume and format?
ANSWER: Yes. The University envisions using institutional content for RAG-based knowledge grounding per Appendix H Requirement #17. Specific content sources will be determined during implementation.

84. Are there known peak periods or concurrency expectations we should design around (term start/finals), and do you have historical indicators?
ANSWER: The academic calendar drives usage patterns, with higher demand during fall/spring semesters and peak periods around midterms and finals. Specific concurrency estimates are not available. The platform should support the full FTE count of 30,912.

85. Does UMS have preferences regarding the use of commercial foundation models vs. private or fine-tuned models?
ANSWER: The University is model-agnostic. Both commercial and open-source models are acceptable per Appendix H Requirement #3.

86. What is your existing footprint of AWS, Google, Azure, and any other signficant installed base of cloud services?
ANSWER: The University does not have a preferred cloud provider. The vendor is responsible for hosting the SaaS solution and should describe their cloud infrastructure in their response.

87. How does UMS anticipate communicating / raising awareness of the tool? (associated with launch and regular student on-boarding)
ANSWER: Internal communications and awareness campaigns will be managed by UMS. The vendor should propose training and adoption resources as described in the Training section.

88. What metrics have been establish to evaluate the effectiveness and impact of this initiative on the University's operations?
ANSWER: Effectiveness metrics will be defined during implementation in collaboration with the awarded vendor. Initial focus areas include platform adoption rates, uptime/availability, and user satisfaction.

89. How is IT governance structured across the seven universities — is there a central IT authority, or does each campus maintain independent control? Who will own the admin console?
ANSWER: IT governance is centralized at the UMS System Office level with campus-level IT teams for local operations. The AI solution should support centralized administration with campus-level delegated management per Appendix J Questions 14 and Appendix H Requirements #39 and #43.

90. What are the priority use cases the University intends to support initially (e.g., teaching assistance, research analysis, administrative productivity)?
ANSWER: Priority use cases include writing assistance, research support, code generation, document analysis, transcription/translation, and administrative workflow support. See the response to Questions 14 and 15 above and Appendix H for the full requirements.

91. Should the platform support fine-tuning or training of private models using University datasets
ANSWER: Fine-tuning and private model training using University datasets is addressed by Appendix H Requirement #5 (Required). Respondents should describe whether their platform supports fine-tuning, RAG-based customization, or both, and how University data ownership and isolation are maintained.