VERMONT

**State of Vermont**
*Secretary*
**Agency of Digital Services**
Secretary's Office
One National Life Drive – Dewey Building 2nd Floor
Montpelier, VT 05620-2120
[phone] 802-828-4141

*Denise Reilly-Hughes, State CIO and*

*Andrea DeLaBruere, Deputy Secretary*

October 7, 2026

Dear Chair Plunkett and Committee Members,

Thank you for the Committee's correspondence of September 24, 2026, and for the discussion at the Committee's September 30, 2026, meeting. Attached to this letter are responses to the questions from members of the Joint Information Technology Oversight Committee regarding the Agency of Digital Services' (ADS) use of Anthropic's Claude Code for evaluation purposes, along with a description of the Agency's approach to evaluating new technologies.

We have seen rapid growth in new technologies centered on artificial intelligence over the past two years. This comes at a particularly sensitive and exciting time for the State. We are modernizing our systems, updating legacy contracts, and evaluating new products and technologies that did not exist a year ago.

The rapid development of frontier AI models offers both risk and opportunity for our state and for us as Vermonters. ADS takes the approach that both must be balanced for Vermont to build robust, practical AI governance and cybersecurity approaches.

As the Secretary and Vermont's State CIO, I have been tasked by the legislature, among other things, to "coordinate and optimize the use of technology within State government" and to "review and approve State government information technology and cybersecurity policies." 3 V.S.A. § 3302(b)(1), (3). The evaluation agreement with Anthropic falls within these responsibilities. ADS entered into the agreement in June 2026 for a six-month evaluation. ADS is using the tool to examine State system code for security and operational weaknesses. The evaluation is limited to system code; no personal information or Vermonters' data has been reviewed.

ADS must evaluate new technologies and services to understand their efficacy and uses in order to execute our legislatively assigned responsibilities. We must understand the current state of the art beyond the academic, particularly in cybersecurity, and this understanding must be developed before the formal procurement of services or systems that use the newest technologies.

Recent state procurements of technology solutions have included either explicit or implicit use of AI. ADS's procurement team, Data Management & Artificial Intelligence Division, and General Counsel have led the adoption and negotiation of standard contract language regarding the use and limits of AI in state procurement contracts, in collaboration with the Department of Buildings and General Services and the Office of the Attorney General. Limited evaluation

agreements provide Vermont with practical knowledge to understand the risks, rewards, and limits of AI technology. Practical knowledge of this technology and the practices of the companies that create it is critical to informing future product procurement.

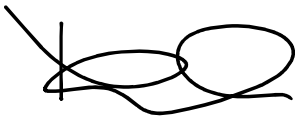
We appreciate the Committee's questions regarding Bulletin 3.5 and its applicability to evaluation agreements generally and the Anthropic agreement specifically. At the time the agreement was signed, ADS had not considered this six-month evaluation program to be a procurement under Bulletin 3.5. The evaluation is a time-bound trial. However, ADS recognizes that the agreement can reasonably be viewed as a procurement. After considering the questions in your letter and consulting with the Secretary of Administration, ADS will use the Bulletin 3.5 waiver process for trials and evaluations going forward. I appreciate the Committee's questions and perspective, which allowed me to consider this matter through a different lens.

At the September 30 meeting, members raised the importance of understanding who authorizes the use of increasingly powerful technologies. ADS agrees that clear authorization matters. Authorization occurs at a specific point in time and is documented, whether through the Bulletin 3.5 process or the Agency's governance documentation. The Anthropic evaluation agreement was signed by the Agency's then-serving Chief Information Security Officer and reviewed by ADS General Counsel and the Chief Data and AI Officer. In ADS's view, equally important is the ongoing oversight that follows: the guardrails and governance around each use case, the data with which AI tools interact, and the safeguards in place to protect State systems and Vermonters. These priorities, along with the other principles in the State's AI Code of Ethics, guide how ADS evaluates and governs the use of AI.

ADS values its partnership with this Committee and the importance of working together to carry out the work necessary to support and protect Vermonters.

Please do not hesitate to contact me with any further questions. ADS would welcome the opportunity to brief the Committee on the results of the evaluation when it concludes.

Sincerely,



Denise Reilly-Hughes
Secretary and State Chief Information Officer
Agency of Digital Services

Attachment: Responses to Committee Questions
cc: Secretary of Administration

VERMONT

**State of Vermont**
*Secretary*
**Agency of Digital Services**
Secretary's Office
One National Life Drive – Dewey Building 2nd Floor
Montpelier, VT 05620-2120
[phone] 802-828-4141

*Denise Reilly-Hughes, State CIO and*

*Andrea DeLaBruere, Deputy Secretary*

October 7, 2026

Dear Chair Plunkett and Committee Members,

The Secretary requested that I provide additional information regarding the testimony in the second half of your September 30, 2026 meeting. I appreciated the comments of the Center for Democracy & Technology (CDT) and GIE Foundation representatives. I have long respected the CDT's work and look forward to learning more about the GIE Foundation.

I would first like to respond to the question raised by Representative Sibilia regarding the statutory requirement that the State's AI governance be updated annually. Our policies have been adopted or used as a model by other states. Since its first release, the Code of Ethics has required only minor revisions.

The principles established before the launch of ChatGPT have remained durable through significant technological change, including changes in how AI itself is defined. This is because both the legislative framework for Vermont's use of AI and the Code of Ethics were built on principles rather than on any particular technology.

Our AI governance has evolved to keep pace with new technologies. The Council has released guidance for AI Use and Data Management as a part of our 2024 report, Guidelines for State Employee Use of Generative AI (which has been revised at least twice since issuance), and at the beginning of 2026, the AI Division formalized our Vendor AI Usage Disclosure, which requires vendors to disclose their use of AI in State procurements. We work with the Agency of Administration's Office of Procurement and Contracting and VTrans to update our standard attachments to include additional protections for contracts that are not primarily technical.

Vermont's approach to AI has always been grounded in human-centered design principles. The State treats AI systems as power tools in the hands of a human operator. The State does not deploy autonomous AI agents, and a human directly operates every AI tool in use.

We have begun augmenting the legislature's mandated tool review process (AI Inventory) with a use-case-level review. This use-case review documents which features of the process enable AI use for that step. This is our application of NIST's AI governance guidance.

We would welcome collaboration in developing that evaluation process and would like to propose that next year the legislature shift from the tool-based evaluation approach to a use-case-oriented framing, as many other states are moving in that direction. We believe this change would allow more meaningful transparency for the legislature and Vermonters.

I would also like to describe how ADS is using the Anthropic SLTT program tools. We are not using Claude Security. Claude Security requires access to the entire codebase and does not operate within our existing technology and development stack. We continue to review whether it makes sense to adopt as additional capabilities roll out. Because Vermont does not use Claude Security, the question raised by Representative Priestley regarding that product's underlying model does not affect the State's evaluation.

Vermont also does not use, and does not have access to, Claude Mythos. This is a core difference between the SLTT and Glasswing programs. The SLTT program offers only "Mythos class" capabilities, not Mythos itself. The widely reported incident in which a model acted outside its test environment is not feasible with the models and processes we have in place.

Anthropic offers four "families" of models, from least to most capable: Haiku, Sonnet, Opus, and Fable. The cost of these models is aligned with their capabilities. We try to use the least capable model that provides sufficient output.

We use the Claude Code plugin for VS Code. We require staff to complete a 2-day in-person orientation and training that establishes norms and teaches them to use the tools safely before making this available to them. The reason we use the Claude Code plugin within VS Code is that it shows all work as proposed and the rationale before taking any action. It also keeps the developer using a power tool. Our developers are hands-on, and the AI tools do not have access beyond the folder on the developer's machine where the development or test code is stored.

We begin code review by doing an orientation of the code base. In this step, samples of the code and the file structure are shared with the Anthropic APIs. The AI tool creates an overview of the application itself, evaluates and documents the underlying requirements that were used to create the application, and documents the general application architecture. The developer reviews this to both validate that the AI has a good understanding of the application as well as improve their own ability to manage the application. For small applications, we generally use Sonnet; for medium and large applications, we use Opus; and for very large applications, we use Fable.

If developers use Sonnet or Opus, Anthropic does not retain the prompts and responses. If they use Fable, Anthropic keeps a copy of the prompts (including attached code) and responses for up to 30 days, and that data is accessed only in emergency circumstances when activity is flagged as potential misuse of the tools. For codebases large enough to use Fable, the code sent to the model represents a small percentage of the codebase. ADS reviewed these retention terms before making the tools available to developers and determined it was appropriate because only system code, and no personal information or Vermonters' data, is involved.

The second stage is to develop a "threat model," which identifies the primary ways an attacker would attempt to get into and abuse the application. This, in turn, informs a vulnerability scan, identifying potential weak points an attacker could attempt to use. These vulnerabilities are then evaluated in detail, and finally, the second stage produces an output document that suggests likely attack paths, recommendations on how to fix them, and how to test them to ensure both the attack vector is closed and that nothing else was inadvertently broken. The developers review this output and share it with the ADS Security Division when significant issues are identified. We consider this artifact highly sensitive.

The developer may then choose to use Claude Code to apply the proposed fixes to the application. All generated code is reviewed through our standard release review processes, which are currently being updated to accommodate the increased volume generated by this program. In the industry, this is called the Software Development Lifecycle (SDLC), and includes multiple rounds of testing before any changes make it to production. ADS designed its use of the program to support its staff, protect the security and resilience of State systems, and maintain Vermonters' trust.

I would welcome the opportunity to answer any further questions the Committee may have. Additionally, we have a half-day "Pathfinder" training that we use to orient non-technical staff to our ethics and to our appropriate use and review frameworks. We are using this training to prepare staff outside of ADS to identify and develop use cases, and it has been very well received. We have offered the training to the Joint Fiscal Office and the Office of Legislative Counsel and would like to extend that offering to the committee or anyone else in the legislative branch who might benefit from it.

Very respectfully,



Josiah Raiche
Chief Data and AI Officer,
Agency of Digital Services

**ATTACHMENT A**

# Responses to Committee Questions

*Joint Information Technology Oversight Committee Inquiry Regarding the Anthropic Claude Code Evaluation Agreement*

October 6, 2026

The following responses address the questions set out in the Committee's correspondence of September 24, 2026. Each Committee question or request is reproduced in bold, followed by the response of the Agency of Digital Services (ADS).

## I. Section 1. Potential deficiencies in the legal obligations applied to Anthropic with respect to the protection of confidential information acquired from the State

**1. Did ADS conduct a competitive bid or simplified bid prior to awarding the contract?**

**2. Did ADS request a waiver to the competitive bidding requirements of Bulletin 3.5 and, if so, can the Agency produce the waiver approval form for Committee review?**

**3. If the Agreement was issued as a sole source contract, did the CTO give prior approval to the written justification pursuant to Bulletin 3.5?**

**4. If the Agreement was issued as a sole source contract, did the Secretary of Administration provide prior approval?**

***Response to Questions 1–4***

ADS did not conduct a bid, request a waiver, or prepare a sole source justification because, at the time the agreement was signed, ADS did not consider the evaluation a procurement under Bulletin 3.5. As noted in the accompanying letter, ADS will use the Bulletin 3.5 waiver process for trials and evaluations going forward.

**5. If this Agreement is viewed as a gift or grant, did the State receive approval from the Joint Fiscal Committee as required in 32 V.S.A. § 5?**

***Response to Question 5***

ADS does not consider this time-bound evaluation agreement a gift or grant. No funds were exchanged, and the agreement does not augment any appropriation. The agreement references a $100,000 credit, which represents an estimated allocation of tokens, the unit of compute used by AI tools. The credits are not monetary funds provided to the State. To date, ADS has used approximately 10 percent of the allocation. Anthropic offers this access to state, local, tribal, and territorial governments through its State, Local, Tribal, and Territorial (SLTT) program.

**6. Did the Secretary of Administration review and approve the zero-dollar contract as required by Bulletin 3.5?**

7. **Did the Office of the Attorney General review and approve the zero-dollar contract as required by Bulletin 3.5?**

8. **Assuming the vendor would have access to confidential information, did the Office of the Attorney General review and approve the contract as required by Bulletin 3.5?**

9. **Was Attachment C included as part of the executed contract as required by Bulletin 3.5?**

10. **Were there modifications to Attachment C and, if so, were the required modifications reviewed and approved before execution?**

***Response to Questions 6–10***

The agreement was not processed as a zero-dollar contract under Bulletin 3.5 and was not submitted to the Secretary of Administration or the Office of the Attorney General. ADS General Counsel reviewed the agreement. That review included Anthropic's privacy and security policies. No personal information or Vermonters' data has been involved in the evaluation. Attachment C was not included in the evaluation agreement and will be included in future procurements.

**II. Section 2. Potential issues with the contracting process carried out by the Agency**

**In addition to answering the questions above, please provide all of the following which will help document that the State procurement process was followed:**

1. **The complete final executed agreement, including the AA-14 coversheet, Attachment C, all other attachments, incorporated terms, and any amendments or modifications between the State and Anthropic.**

***Response to Request 1***

The evaluation agreement is attached. Because it was not processed as a contract under Bulletin 3.5, there is no AA-14 coversheet or Attachment C.

2. **Documentation identifying the reviews and approvals associated with the Agreement. (ID 421652)**

***Response to Request 2***

The agreement was reviewed by ADS General Counsel and the Chief Data and AI Officer and signed by the Agency's then-serving Chief Information Security Officer in June 2026.

3. **The pre-deployment impact and risk assessment described above.**

***Response to Request 3***

Before the evaluation, ADS reviewed Anthropic's published model cards. At the conclusion of the evaluation, ADS will add this use case to the State's AI Inventory and complete a use case approval form.

4. **According to the SLTT documentation provided to Representative Sibilia, the process requires a connection to the LLM API to function. Options include:**

**Anthropic-hosted LLM, Amazon Web Services-hosted LLM, MS Foundry LLM, or Google Cloud LLM. Which option did Vermont use?**

***Response to Request 4***

ADS is using the Azure hosted LLM.

**III. Section 3. Whether a NIST Impact and Risk Assessment was performed by the Agency before consenting to the agreement.**

***Response***

This is an evaluation, not an implementation. The purpose is to determine whether the technology warrants further consideration. It does not authorize use of these tools in any State production system. Consistent with the risk-based approach of the NIST AI Risk Management Framework, ADS scaled its review to the limited scope of the evaluation: review of system code for security vulnerabilities in test and development environments, with no personal information or Vermonters' data and no decisions affecting Vermonters. That review included ADS General Counsel's review of Anthropic's privacy and security policies and the Agency's review of Anthropic's published model cards. ADS teams are documenting observations from use of the tools. If the State later procures this or similar technology for production use, it will be subject to full architectural, security, data, contractual, and procurement review before authorization.