Can Federal Contractors Use ChatGPT and Other Generative AI Tools?

Tom Dunlap is a founding partner at Dunlap Bennett & Ludwig, a veteran-owned law firm with over 75 attorneys nationwide. As a former prosecutor and Army officer with extensive litigation experience, Tom has secured notable victories, including a $31 million patent jury verdict and multiple Federal Circuit appeals. His expertise spans government contracts, trademarks, patents, and general intellectual property law, currently serving as an outside counsel for US HHS and the Commonwealth of Virginia. A true renaissance professional, Tom holds a JD, MBA, and an MS in Biotechnology and is an inventor with multiple patents, including technology used in COVID testing. Recognized in the WTR 1000 as one of the world's leading trademark lawyers and a Virginia Super Lawyer for over a decade, Tom combines practical experience as a former CEO and founder of Ceres Nanosciences with academic contributions as an adjunct professor at Washington & Lee University and Lafayette College, where he teaches intellectual property law. Outside of his legal practice, Tom is a PADI Master Diver, instrument pilot, saber fencer, and former Burmese National Kickboxing Silver Medalist.

Tom Dunlap, Partner



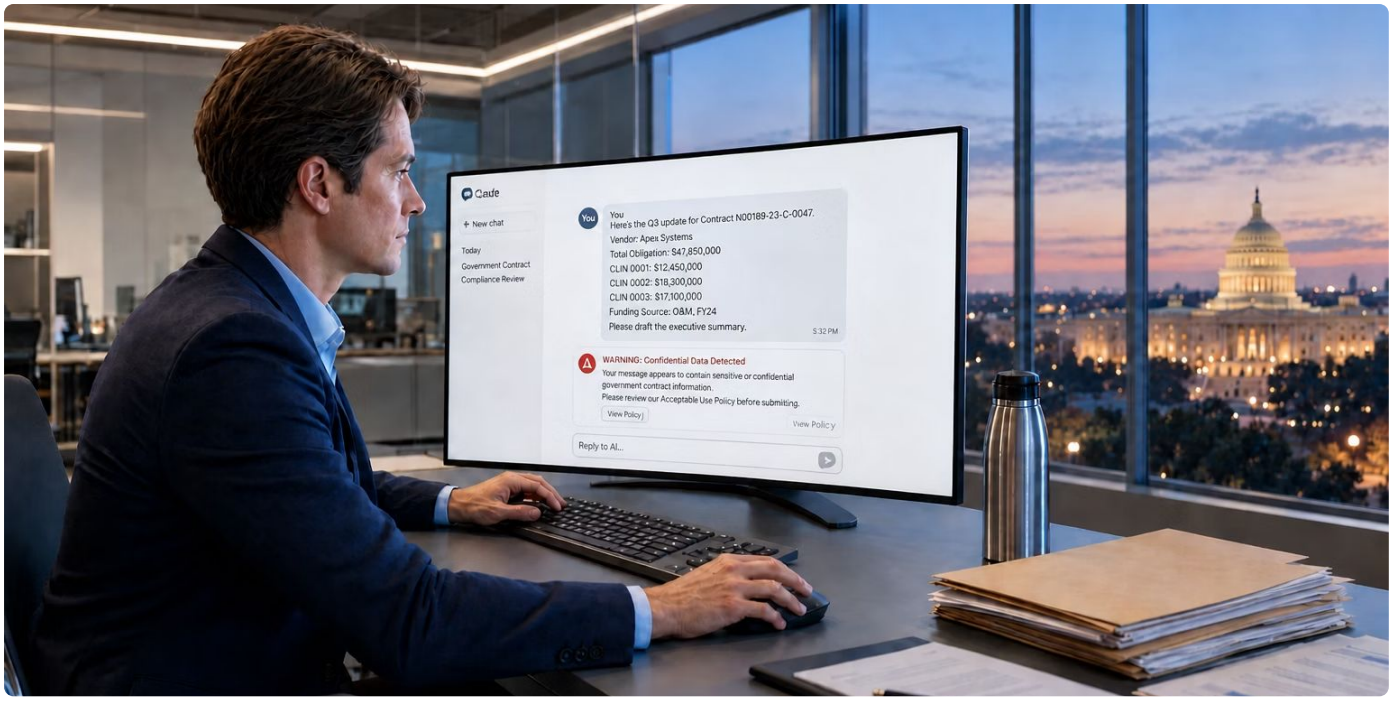

Yes. Federal contractors can use ChatGPT, Claude, Gemini, Copilot, and the rest of them, and most already do, whether or not anyone approved it. Nothing in federal law bans the tools. The one categorical prohibition on the books is narrow. Section 1532 of the FY2026 National Defense Authorization Act (P.L. 119-60, signed December 18, 2025) bars AI developed by DeepSeek, its parent High Flyer, and their affiliates from Defense Department systems and from contractor performance of DoD contracts, subcontractors included. Everything else lives in your contracts, your solicitations, and the data you handle.

That last part is the whole game. The question I get is almost never "can we use AI." It is which information can go into which tool under which contract. A marketing draft can go almost anywhere. A drawing marked CUI cannot go into the free version of anything. The interesting cases sit in between, and that is where companies get hurt.

The ground moved several times this year. GSA proposed an AI clause in March and rewrote it in June. The FAR Council re-proposed the government-wide CUI rule in June. The Pentagon suspended the planned CMMC Phase 2 rollout in July 2026, halting the expansion of mandatory third-party assessments. A September DFARS class deviation continued that suspension as binding direction for contracting officers. This piece describes the rules as they stand on September 29, 2026 and says plainly which ones are proposed rather than final. It reflects how we work through these questions with contractors at Dunlap Bennett & Ludwig. How they apply to your company depends on your contracts and your facts.

Which Data Can Go Into an AI Tool?

Think of your information as a ladder. Every rung down, fewer tools qualify.

- **Public and ordinary business information.** Marketing copy, research, internal memos with no government or third-party strings. Commercial AI carries the least legal risk here. Your own confidentiality is still your problem, and so are the vendor's terms of service, but nobody from the government is going to call.
- **Federal contract information (FCI).** Anything the government gave you, or you produced for it, that is not public. The basic safeguarding clause, FAR 52.204-21 (which appears as 52.240-93 in

- solicitations issued under the FAR overhaul deviations since February 1, 2026), imposes 15 security requirements on any system that handles FCI. A consumer chatbot is not that system.
- **Controlled unclassified information (CUI).** The hard line for most contractors. It gets its own section below.
- **Export-controlled technical data.** Harder still. Putting it into a commercial AI service can be the violation itself, not just a policy breach.
- **Classified information.** Never touches a commercial AI tool. That one needs no paragraph.

Two categories work across every rung. Competitors' bid and proposal information and source selection information are restricted in any system under the Procurement Integrity Act (41 U.S.C. §§ 2101–2107 and FAR 3.104). And anything you received under an NDA carries its own restrictions. Your AI vendor is a third party. Read the NDA before you paste.

Can Contractors Put CUI Into ChatGPT, Claude, Gemini, or Copilot?



Not in the consumer or standard business versions. Not with an enterprise license either, unless that license comes with a specific deployment your security assessment actually covers.

The mechanism is a contract clause. DFARS 252.204-7012 requires defense contractors to protect covered defense information under NIST SP 800-171, and paragraph (b)(2)(ii)(D) requires any external cloud service that stores, processes, or transmits that information to meet the FedRAMP Moderate baseline or its equivalent. (FedRAMP relabeled "Authorized" as "Certified" in May 2026 and replaced the impact levels with certification classes. The DFARS clause still says Moderate, so that is the word to use with your contracting officer.) A consumer chatbot meets none of it.

The trap is the enterprise plan. Vendors sell "we don't train on your data" as though it settled the question. It settles a privacy question. It does nothing for the security question. An enterprise plan

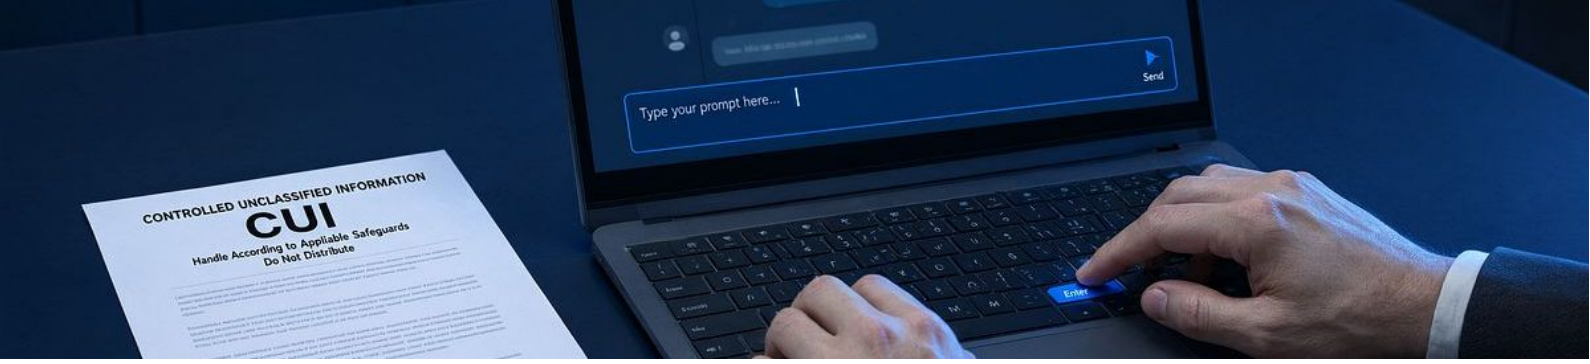

that runs in the vendor's ordinary commercial cloud is still outside the boundary your System Security Plan describes and still not a FedRAMP Moderate environment.

Compliant setups exist, and they are getting easier to buy. Copilot inside a Microsoft GCC High tenant. Azure OpenAI, Amazon Bedrock, or Google's Vertex AI in their government regions, deployed inside your assessed boundary. A model you host yourself on hardware you control. The vendor's name does not decide the question. The deployment does. Check the service's certification status on the FedRAMP Marketplace, then confirm that the specific configuration you are buying is the one that was assessed. A service certified in one region or configuration is not certified in every configuration, and the sales rep is not the one who signs your SPRS affirmation.

Civilian-agency contracts handle CUI through agency-specific terms today. That is changing. On June 23, 2026 the FAR Council re-proposed the government-wide CUI rule (FAR Case 2017-016) as part of the Revolutionary FAR Overhaul. It moves CUI requirements into a new FAR Part 40, adds a solicitation provision at 52.240-6 and a contract clause at 52.240-7, sets a 72-hour clock for reporting CUI incidents (down from the eight hours in the January 2025 draft), and points civilian contractors at NIST SP 800-171 Revision 3 while DoD still requires Revision 2 under DFARS 7012 and 32 C.F.R. Part 170. Comments closed July 23. Nothing is final. But a company that sells to both DoD and a civilian agency should plan on carrying two baselines for a while, and any AI tool it approves for CUI has to satisfy both.

What About Export-Controlled Data?

This is the strictest rung, and the one where the tool itself can commit the violation. Typing ITAR technical data into a commercial AI service can be an unauthorized export under 22 C.F.R. § 120.50 in two ways at once. The provider's engineers and support staff may be foreign persons, which makes the disclosure a deemed export. And the data may be processed or stored on servers outside the United States.

People point to the 2020 encryption carve-out at 22 C.F.R. § 120.54 and its EAR cousin at 15 C.F.R. § 734.18. Those provisions were written for storage and transmission. They require end-to-end encryption to a FIPS 140-2 standard with the decryption keys withheld from every third party, including the cloud provider. A language model has to read your text in the clear to answer you. The carve-out does not survive that. Penalties run per violation, and the civil maximum under ITAR is well into seven figures.

Our advice here is the same for every client, and it is not subtle. No export-controlled technical data goes into any outside AI service without a specific written authorization and a legal review



first. The bright line costs you some convenience. It is cheaper than a voluntary disclosure to DDTC.

Does the CMMC Pause Change Anything?

Less than the headlines said, and in one respect it cuts the other way.

The facts first. On July 13, 2026 the DoD Chief Information Officer, Kirsten Davies, suspended CMMC Phase 2, the point at which third-party C3PAO certification at Level 2 would have become a condition of award for CUI contracts starting November 10, 2026. Phases 3 and 4 were frozen with it. A CMMC Reform Task Force got 60 days to review the program, and a request for information drew more than 1,100 responses. On September 3, DoD issued a class deviation (2026-O0025, Revision 3) directing contracting officers to strip third-party assessment requirements from solicitations and contracts, which turns a policy memo into an enforceable rule. The task force's 60-day clock ran out on September 11. Its recommendations go to the CIO first, and whether she publishes them, and when, is her decision. Industry expects something public in late September or October. I would not bet a compliance budget on that date.

What did not change is everything that actually binds you today. DFARS 252.204-7012 and NIST SP 800-171 Revision 2. The self-assessment score in SPRS. The annual affirmation under 252.204-7021. The 72-hour incident reporting clock. The safeguarding terms already sitting in your prime contract, which your prime did not suspend because the Pentagon issued a memo. (The February 1, 2026 DFARS class deviations also eliminated 252.204-7019 and renumbered 7020 as 252.240-7997. The substance survived the renumbering.)

So why do I say the pause raises the AI stakes? Because with no auditor coming, your own filed self-assessment is the compliance record, and it is a representation to the government. A score that describes a controlled CUI environment becomes false the day an engineer pastes covered data into a consumer chatbot, not at some future audit. The Justice Department has pursued False Claims Act cases against contractors over cybersecurity misrepresentations of exactly this kind. MORSECORP paid $4.6 million in March 2025 after reporting a SPRS score of 104 that a later assessment put at negative 142. Raytheon and its affiliates paid $8.4 million in May 2025. LOGZONE, an Alabama defense contractor, paid $507,144 in June 2026 over a near-perfect self-assessment that a DoD audit found inaccurate. These cases start with whistleblowers. Your employees know which tools their teams really use, and some of them know what a relator's share is.

Where nothing requires disclosure, accuracy still does. Representations about human involvement and how deliverables are produced carry weight in evaluation, and they need to match the actual process.



Do Contractors Have to Tell the Government They're Using AI?

Sometimes, and the list is getting longer.

There is no across-the-board disclosure rule today. The obligation comes from a solicitation or contract term, from agency policy, or from the nature of the work, most obviously when AI is part of what you are selling or it processes government information. Where nothing requires disclosure, accuracy still does. If your proposal describes a review process with a human in it, there had better be a human in it.

GSA is the agency to watch. Its first attempt, a draft clause posted March 6, 2026, would have applied whenever AI was used in performing a Schedule contract, and it drew enough industry fire that GSA extended comments and pulled it from the spring refresh. The rewrite, published June 17, 2026 as proposed GSAR 552.239-7001, "Basic Safeguarding of Data within Large Language Model Artificial Intelligence Systems," is narrower and broader at the same time. Narrower, because it applies only when government data, meaning prompts, source material, and outputs, is processed by a large language model, and it expressly carves out LLMs embedded in ordinary commercial products like word processors and LLMs incidental to the contract's purpose. Broader, because it amends the GSAR generally and would reach every GSA-administered contract, not just the Schedules. Under the clause you may not train or improve a model on government data, you owe CISA and the contracting officer a report within 72 hours of an incident, and you must give 30 days'

notice before adding or swapping an LLM or provider. Comments closed August 3. GSA expects to roll the clause into existing Schedule contracts by mass modification at Refresh 32.

A services contractor that uses a third-party LLM to draft deliverables on a GSA task order is inside that trigger even though the contract has nothing to do with AI. This is why I tell clients to build the inventory now. Which tools, which contracts, which data. The clause will ask for it, and so will the next one.

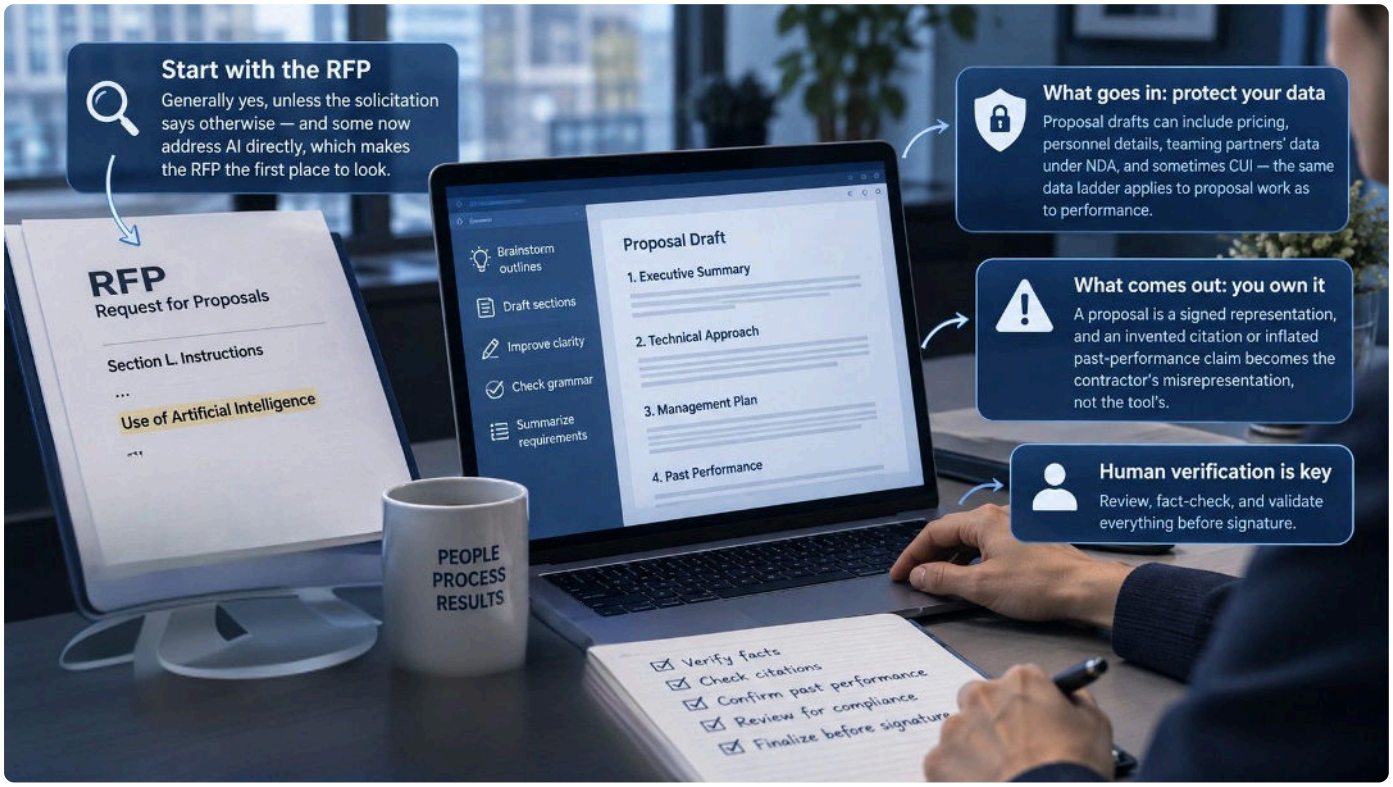

Can Contractors Use AI to Write Proposals?

Generally yes, unless the solicitation says otherwise, and some now do. Read the RFP first. A few agencies have started asking offerors to disclose AI use in proposal preparation, and a handful restrict it.

Two risks. The first is what goes in. Proposal drafts carry pricing, key personnel resumes, teaming partners' data under NDA, and sometimes CUI. The same ladder applies to proposal work as to performance, and a capture team at 2 a.m. the night before a deadline is the group most likely to forget it. We had a client whose proposal writer fed a teaming partner's past-performance narrative into a consumer chatbot to "tighten it up." The partner's NDA had a no-third-party-systems clause. Feeding the proposal to the AI was fine. The teaming relationship was not.

The second risk is what comes out. A proposal is a signed representation. An invented citation, a past-performance reference the tool embellished, or a compliance matrix that says "fully compliant" where the model was guessing becomes your misrepresentation, not the software's. Nobody has successfully blamed the chatbot yet, and I would not want to be the test case. Human verification before signature is the control point. Budget the hours for it.

What Does a Workable Contractor AI Policy Cover?

I have read a heck of a lot of contractor AI policies this year. The ones that work share a few features. The ones that fail share one.

The working ones decide in advance which rungs of the ladder may touch which tools, so no engineer is making a judgment call at 11 p.m. They map the contracts' AI-relevant terms, including what flows down to subcontractors, because a prime's promise is only as good as its subs' behavior. They state, without ambiguity, whether each approved tool sits inside or outside the CUI boundary. They block consumer AI services on systems that handle FCI or CUI, technically, not just on paper. They train people on the scenarios that actually come up, meaning proposal drafting, code help, deliverable review, and summarizing a customer's email. And they fold AI tools into the spillage procedure and assign one person to track the three open rulemakings.

The failing ones only prohibit. Employees reach for consumer tools when no approved tool can do the work, and a policy that bans without providing gets violated inside a month. The fix costs money. A GCC High Copilot seat or a government-region model deployment is not free, and neither is the assessment work to bring it inside your boundary. It is still cheaper than an FCA settlement, and it is the only version of a policy I have seen hold up.

What If Someone Already Pasted CUI Into ChatGPT?

Preserve first. The conversation is usually the best record of what went in, which account was used, and what the settings were. Deletion has its place, after the facts are documented and the reporting obligations are assessed, not before. Our companion guide, **An Employee Pasted Confidential Information Into ChatGPT. What Should the Company Do Next?**, walks through the general assessment.

Federal contracts add a layer. On defense work you need a prompt answer to whether this is a reportable cyber incident under 7012, because the clock is 72 hours and it started when you discovered the event. The contract or agency procedures may require notice to the contracting officer, and primes and subs often owe each other notice under the flowdowns. The one move that closes every option without answering any question is quietly deleting the chat and hoping. I have seen that tried. It does not work, and the deletion becomes its own problem when the whistleblower complaint lands.

Where This Leaves You

Use the tools. Match them to the data. Keep the paperwork honest. The contractors who get in trouble this year will be the ones whose SPRS score, proposal language, and actual desktop behavior tell three different stories. Using AI is fine. Telling the government one story about your

environment while your engineers live another is the problem, and it is a problem that predates AI by a couple of decades.

If you want a second set of eyes on your AI policy or your contract flowdowns before the GSA clause and the FAR CUI rule land, contact us.

The Bottom Line

Federal contractors can use AI, and most already do. The data draws the line: ordinary business information travels freely, FCI needs safeguarded systems, CUI confines AI to authorized environments, export-controlled data stays out of outside tools entirely, and classified information is categorically off the table. The contractors best positioned to capture what AI offers are the ones whose certifications stay accurate, whose tool approvals map to their data, and whose employees have a compliant tool that works — so that contract eligibility never rides on a single paste.

This article is provided for general informational purposes only and does not constitute legal advice. Reading it does not create an attorney-client relationship. Laws, regulations, and provider terms change, and readers should evaluate current requirements and the circumstances of their own situation.

Dunlap Bennett & Ludwig works with federal contractors and other organizations on AI governance, government contracts, cybersecurity compliance, intellectual property, commercial agreements, internal investigations, and disputes involving emerging technologies. How these rules apply depends on the specific contracts, agencies, and facts involved.

To learn more about Dunlap Bennett & Ludwig, call **888-306-4030** or email **clientservices@dbllawyers.com**.

Attorney Advertising. Prior results do not guarantee a similar outcome.