August 13, 2026

Ms. Leanne Condren
Office of the Department of War Chief Information Officer

Subject: Reforming CMMC and Reducing Compliance Burden for the DIB

Dear Ms. Condren:

Deep Fathom, Inc. submits the attached response to the Department of War Request for Information on reforming CMMC and reducing compliance burden for the Defense Industrial Base.

Our central point is simple:

- Protect CUI and strengthen operational resilience without lowering the near-term security baseline.
- Keep more suppliers in the DIB and get new entrants ready faster by clarifying scope and reducing duplicated evidence work.
- Make verification trustworthy and scalable by matching its method and frequency to risk and material change, using structured evidence and carefully evaluated AI assistance to focus expert judgment.

Recent advances in AI make source-grounded evidence review for defined tasks practical enough to evaluate. When evidence is appropriately scoped and linked to its source, machine-assisted review may reduce repetitive work, identify gaps and contradictions, and help reviewers focus on consequential questions. The Department should validate those benefits before relying on them and leave assessment, contracting, and Government decisions with authorized people.

Deep Fathom builds security and compliance software for the DIB. Our working platform represents all 110 NIST SP 800-171 Revision 2 requirements and all 320 CMMC Level 2 assessment objectives in one software graph connecting scope, controls, evidence, responsibility, findings, and remediation. It provides a practical basis for the recommendations in this response.

Within 60 days, the Department can publish a graduated verification model, open a vendor-neutral CMMC evidence-profile process, clarify commercial-service and inherited-control treatment, and charter a no-reliance evaluation of AI-assisted verification. These actions support the Department's stated goals of stronger cybersecurity, lower barriers for small and non-traditional businesses, access to commercial capability, and greater speed to capability.

Thank you for considering this input.

Respectfully,

Steven Hess
Co-Founder and Chief Executive Officer
Deep Fathom, Inc.

Kevin Hunt
Co-Founder and Chief Technology Officer
Deep Fathom, Inc.

Deep Fathom, Inc.
UEI WCJSF1ZPQH55 | CAGE 12LN3 | DUNS 119472153
405 Anglers Drive, Unit D, Steamboat Springs, CO 80487
Kevin Hunt | 650-597-0413 | kevin@deepfathom.ai | Fax: none
Steven Hess | 404-316-2412 | steven@deepfathom.ai | Fax: none

No portion of this submission is proprietary. Deep Fathom consents to public release in whole or in part.

Subject: Reforming CMMC and Reducing Compliance Burden for the DIB

Notice ID: DoDCIOReformingCMMCforDIB001

Respondent: Deep Fathom, Inc. | UEI WCJSF1ZPQH55 | CAGE 12LN3 | DUNS 119472153

Point of contact: Kevin Hunt, Founder and Chief Technology Officer | kevin@deepfathom.ai | 650-597-0413

Address: 405 Anglers Drive, Unit D, Steamboat Springs, CO 80487

Proprietary information: None

BLUF (Bottom Line Up Front)

The Department should test whether CMMC cost can be reduced, supplier readiness accelerated, and confidence in implementation improved. In the near term, preserve the current Level 2 security baseline, but replace the binary choice between self-assessment and C3PAO assessment with graduated verification tied primarily to data sensitivity, mission consequence, threat exposure, supplier criticality, and material change. Evidence completeness and quality should influence sampling and escalation, not determine the supplier's risk tier. A common CMMC evidence profile could allow one source-linked record to support self-assessment, machine-assisted review with accountable human validation, independent assessment, and targeted Government review. AI should organize and test evidence; it should not make award, certification, affirmation, or enforcement decisions. The Department should evaluate this model through a vendor-neutral, no-reliance test before assigning legal effect.

Basis for Our Response

Deep Fathom builds supplier-facing security and compliance technology for the DIB and works with DIB contractors, CMMC advisors, MSPs, security practitioners, and assessment organizations. That work consistently exposes three sources of burden:

- Scope and responsibility are often unclear until late in the readiness process.
- The same implementation facts are rewritten across disconnected artifacts.
- Verification is primarily point-in-time even though systems, threats, and evidence change continuously.

Our platform models all 110 NIST SP 800-171 Revision 2 requirements and all 320 CMMC Level 2 assessment objectives in an OSCAL-native compliance graph. The graph connects systems, boundaries, controls, implementation statements, evidence, inherited responsibilities, findings, and remediation work. The same source information can produce human-readable SSP and assessment views and machine-readable exchange packages.

Deep Fathom's AI-assisted assessment workflow collects and evaluates evidence for each assessment objective, identifies missing or contradictory information, and drafts findings for human review. A human reviewer confirms, changes, or rejects each conclusion. Source evidence, machine analysis, human judgment, and authorized decisions remain distinguishable and traceable.

This working implementation and our engagement with DIB participants provide the practical basis for this response.

Where Burden Accumulates

The recommendations below follow four stages of supplier work:

1. **Scope correctly.** Identify CUI, the system boundary, applicable requirements, service providers, and inherited responsibilities before suppliers invest against the wrong scope.
2. **Build readiness.** Implement safeguards and capture implementation statements and evidence as the work is performed, instead of reconstructing a package later.
3. **Verify in proportion to risk.** Reuse evidence with known sources, dates, scope, and review history, and direct independent human effort toward uncertainty, exceptions, and high-consequence decisions.
4. **Update after material changes.** Refresh evidence and trigger review when systems, scope, providers, responsibilities, or risk materially change.

Correct scope prevents unnecessary investment. Evidence captured during implementation is easier to reuse. Updating affected evidence after material changes is more current than reconstructing an entire package on a fixed cycle. Structured data and deterministic checks can support this workflow today. Recent AI capabilities warrant

controlled evaluation for the narrower tasks of mapping source evidence, comparing statements with artifacts, identifying missing or conflicting support, and prioritizing human review.

Q1. Top Five Cost Drivers, Administrative Burdens, or Operational Challenges

Official question: "Identify the top five most prohibitive cost drivers, administrative burdens, or operational challenges your organization has experienced, or anticipates to experience, when attempting to comply with the CMMC framework and NIST SP 800-171 Rev 2."

Answer: The following anticipated burdens are based on building supplier assessment workflows, working with DIB participants, and analyzing the program's requirements and published cost estimates. CMMC cost includes the cost to become secure, produce reliable evidence, and verify and repeatedly repackage that evidence.

Rank	Driver	Why it creates cost or delay	Who bears it
1	Scope and responsibility uncertainty	Ambiguous CUI boundaries, contract-level assignment, service-provider classification, and inherited responsibility cause over-scoping, late redesign, and rework	Suppliers and primes
2	The gap between current maturity and required security	Suppliers must fund and sustain people, architecture, technology, and remediation before credible evidence exists. These are necessary security investments	Suppliers, with incomplete or delayed recovery through contract prices
3	Repeated narrative and evidence production	The same implementation fact is rewritten in SSPs, policies, procedures, objective narratives, evidence indexes, and remediation records	Suppliers
4	Late disagreement about evidence sufficiency	Without a common evidence profile, suppliers, assessors, primes, and customers can disagree about coverage, freshness, provenance, and sufficiency after substantial work is complete	Primarily suppliers
5	Point-in-time reconstruction and limited reuse	Evidence is recollected for self-assessment, C3PAO review, Government review, primes, customers, and insurers even when the underlying fact has not changed	Suppliers and reviewers

Available public estimates measure different parts of this problem. SBA currently estimates that total small-firm compliance costs can reach approximately $388,600 for self-assessment and $593,800 for third-party certification. Those figures establish the scale of current small-business concern, but SBA's public release does not decompose the estimate or identify a representative average.[1]

The Department's 2024 rulemaking model is narrower. It assumes the contractor has already implemented NIST SP 800-171 and estimates $34,277 for Level 2 self-assessment and initial affirmation and $101,752 for certification assessment and initial affirmation. Within that $67,475 difference, $31,234 is modeled C3PAO cost and approximately $36,241 is additional contractor preparation, participation, reporting, and affirmation labor.[3]

The figures should not be treated as competing estimates of the same thing. Reform should measure three cost layers separately:

Cost layer	What it includes	Primary reform opportunity
Implement and sustain security	Controls, technology, architecture, personnel, remediation, and operations	Targeted assistance, commercial capability, correct scope, and risk-based requirements
Produce and maintain evidence	SSP content, implementation statements, mappings, provenance, updates, and reconciliation	Common evidence profile, source-linked records, and AI-assisted assembly
Verify and reuse evidence	Self-assessment, independent assessment, Government review, reporting, and downstream requests	Graduated verification, automated checks, sampling, and controlled reuse

Structured evidence, deterministic checks, and carefully evaluated AI assistance are most relevant to the second and third layers. These methods do not eliminate the cost of necessary safeguards, technology, or security operations. Better scope and earlier feedback may reduce avoidable implementation rework; the more direct savings arise from maintaining evidence, reconciling records, preparing for review, and reusing verified information.

Q2. Security Controls Delivering the Most Tangible Uplift

Official question: "Which specific security controls has your organization found to deliver the most tangible uplift of cybersecurity and actual risk reduction?"

Answer: Based on our platform work and engagement with DIB contractors, CMMC advisors, MSPs, security practitioners, and assessment organizations, the highest-value areas are foundational access control, secure configuration, flaw and vulnerability management, boundary and data protection, system integrity and active monitoring, and logging, incident response, and recovery. We tested that practitioner experience against three sources of support.

First, CMMC Level 1 incorporates the 15 basic safeguarding requirements in FAR 52.204-21. These are the Department's foundation for protecting Federal Contract Information and include access restriction, identity and authentication, physical and media protection, boundary protection, flaw remediation, malicious-code protection, and scanning.[5]

Second, DIBCAC's analysis of 117 High NIST SP 800-171 assessments conducted from 2019 through 2022 identifies the requirements most often assessed as Other Than Satisfied: FIPS-validated cryptography, MFA, flaw remediation, periodic risk assessment, vulnerability scanning, three logging requirements, incident-response testing, and baseline configuration.[4] This is valuable evidence of implementation and verification difficulty. It is not, by itself, a ranking of security impact.

Third, we consider whether a practice interrupts a common attack path, limits exposure or blast radius, improves detection and response, or supports recovery. Together, these sources support the following priorities:

Practice area and representative requirements	Security uplift	Refreshable evidence	Human judgment that remains
Access control, identity, and MFA (3.1.1, 3.1.2, 3.5.1, 3.5.2, 3.5.3, 3.5.4)	Limits unauthorized access and reduces the value of stolen credentials	Account and factor coverage, privileged-role assignments, exceptions, access changes	Whether access paths, privileges, factors, and exceptions fit the environment and risk
Asset knowledge and secure configuration (3.4.1, 3.4.2, 3.4.6, 3.4.7)	Establishes what must be protected and reduces insecure defaults, unnecessary functionality, and configuration drift	Inventory coverage, approved baseline, enabled functions and services, configuration drift, unsupported assets, exceptions	Whether the inventory and boundary are complete and the baseline provides only necessary capabilities across IT, OT, and inherited services
Flaw and vulnerability management (3.11.2, 3.11.3, 3.14.1)	Reduces exposure to known weaknesses and connects identification to risk-informed remediation	Scan coverage, finding age, patch latency, remediation status, accepted exceptions	Exploitability, operational constraints, remediation priority, compensating safeguards, and risk acceptance
Boundary and data protection (3.1.3, 3.13.1, 3.13.5, 3.13.6, 3.13.8, 3.13.11)	Limits unauthorized paths, protects CUI flows and transmission, and reduces incident spread	Approved connections, network zones, exposed services, allow rules, rule changes, cryptographic mode and scope	Whether actual data flows and implementation provide adequate protection
System integrity and active monitoring (3.14.2, 3.14.3, 3.14.4, 3.14.5, 3.14.6, 3.14.7)	Blocks malicious code, keeps protections current, and detects attacks and unauthorized use	Protection and update coverage, scan results, alert disposition, monitoring coverage, detected unauthorized activity	Whether coverage, update cadence, alert handling, and monitoring depth fit the environment and threat
Logging, incident response, and recovery (3.3.1, 3.3.2, 3.3.3, 3.3.4, 3.3.5, 3.6.1, 3.6.3)	Improves detection, attribution, investigation, containment, reporting, and practiced response	Logging-source health, user traceability, review records, alert failures, exercise and restore-test results	Whether coverage, attribution, escalation, response quality, and recovery capability fit mission consequence

In the near term, Deep Fathom recommends preserving the Level 2 baseline while reforming verification. Over the longer term, the Department and NIST can revisit the control set using incident data, DIBCAC findings, threat intelligence, pilot results, and mission consequence. Neither Deep Fathom nor the available DIBCAC analysis establishes that the present 110 requirements are the only valid long-term baseline.

Where appropriately scoped source data is available, structured evidence and automated review can compare current results with the maintained record and flag material changes for human review. How readily a practice can be observed affects verification and sustainment; it should not determine whether the practice belongs in the security baseline.

Q3. Highest Overhead With Least Measurable Security Improvement

Official question: "Conversely, which specific regulatory requirements or security controls create the highest administrative overhead and financial burden with the least measurable improvement to your actual cybersecurity posture?"

Answer: The available evidence does not support removing any current NIST SP 800-171 safeguard. The clearest low-value overhead instead arises from CMMC documentation and assessment practices that repeatedly restate and reconstruct the same implementation facts.

Administrative practice	Why incremental value is low	Better treatment
Repeated free-form narratives for one fact	Multiple versions become inconsistent and stale without making the source fact more reliable	Maintain one authoritative implementation statement and derive required views
Separate prose for every objective	One source fact can support several objectives	Map the fact explicitly to every supported objective
Evidence without provenance	Screenshots and exports without source, time, scope, or limitations are difficult to validate or refresh	Require provenance, relevant time, integrity, coverage, and limitations
Reconstructing inherited controls for each customer	Repeating a provider's common responsibility does not strengthen it	Use scoped provider records with explicit customer responsibilities and unresolved conditions
Fixed-cycle package reconstruction	Rebuilding an unchanged package costs labor and can still miss changes between cycles	Use type-specific freshness and material-change triggers

The profile should give each security assertion a stable identifier and record which CMMC requirement and assessment objective it addresses; the implementing system or service; the evidence source, collection method, time, and scope; limitations and inherited responsibilities; reviewer disposition; integrity and conflict checks; freshness rules; and the changes that require renewed review. NIST OSCAL provides a standards-based foundation; the Department should define the CMMC-specific profile and its acceptable uses.[6]

Structured tooling can validate the profile and reuse its contents across required views. Structured review can replace much of today's avoidable subjectivity with consistent, traceable checks. Deterministic rules should decide objective questions; AI can organize and compare interpretive evidence and direct reviewers to exceptions. Qualified reviewers remain responsible for scope, sufficiency, disputed facts, and findings.

The profile should reduce existing duplication rather than add another mandatory artifact. Human-readable SSP and assessment views and machine-readable exchange packages should render from the same source record. The SSP remains a useful explanation of scope, architecture, implementation, responsibility, and risk.

Q4. Commercial Cybersecurity Capabilities and Department Recognition

Official question: "Describe how your organization utilizes existing commercial cybersecurity capabilities, platforms, managed services, or any other additional strategies or initiatives to safeguard data, improve operational resiliency, and reduce cybersecurity risk, and how the DoW might better recognize or accept these commercial solutions within a compliance or risk framework."

Answer: Commercial technology already performs much of the security work. The barrier is the cost and uncertainty involved in proving that each component is acceptable inside a CMMC environment.

Deep Fathom uses commercial identity, hosting, monitoring, development, backup, and resilience capabilities across Azure commercial and Azure Government environments. Our architecture is also designed for customer-cloud and

future on-premises deployment. This gives us direct experience with how deployment location, data handling, inherited services, and customer responsibility change the compliance boundary.

Across the DIB, commercial services provide identity and MFA, endpoint protection, vulnerability management, configuration management, logging, backup, managed security, and evidence workflow. Small suppliers often consume these functions through an MSP rather than operate them directly.

A commercial component can solve one security problem while creating a much larger compliance surface. Adding a service may require the supplier to define and prove its access control, authentication, logging, configuration, vulnerability management, data handling, incident obligations, and inherited responsibilities. In cloud deployments, whether a service is a CSP, another type of ESP, handles CUI, or handles security protection data can determine its scope and whether a FedRAMP Moderate authorization or equivalency path applies.[8][9]

When Department guidance leaves acceptance conditions open to assessor judgment, suppliers face a high-stakes assessment without being able to use the assessment team as an implementation advisor. They therefore design to the most conservative plausible interpretation, even when a narrower commercial solution may satisfy the requirement. In Deep Fathom's experience, variation in what assessment teams accept turns that uncertainty into cost and pushes the market toward high-water-mark architectures.

The Department should:

1. **Publish worked examples for CSPs, other ESPs, security protection assets, CUI, and common non-CUI security data.** The examples should show how service function, data handling, incident obligations, and customer responsibility determine scope.
2. **Define evidence requirements for commercial products with limited, clearly defined functions and data boundaries.** Do not require every useful component to become a stand-alone compliance environment when its function can be safely bounded, its contribution to a control can be verified, and remaining responsibilities can be explicitly assigned.
3. **Formalize shared responsibility and inheritance.** Allow providers to publish scoped verification information once, including customer responsibilities, limitations, and unresolved conditions, and permit authorized downstream reuse.
4. **Accept source evidence where practical.** A signed API response showing MFA coverage, patch latency, log-ingestion health, encryption state, or restore-test outcome can provide stronger and fresher support than a manually transcribed screenshot.
5. **Clarify CMMC treatment of FedRAMP 20x Class C.** FedRAMP is a government-wide program rather than a DoW-owned program, but the Department controls how its outputs are treated under CMMC. The Department should state whether and under what conditions a Class C result satisfies current authorization or equivalency expectations.[7][9]

Structured provider evidence can be translated into the common CMMC profile, with inherited responsibilities and open customer obligations stated explicitly. This can lower integration and assessment cost while requiring evidence of what a commercial service actually provides.

Q5. Phase I Self-Assessments and a More Dynamic Posture

Official question: "Regarding Phase I self-assessments, what specific administrative or technical challenges does your organization face in maintaining, verifying, and reporting compliance, and how could this process be fundamentally streamlined? Have your self-assessments led to a more dynamic cyber posture approach, or are they performed only for compliance purposes?"

Answer: Self-assessment becomes dynamic when the reported score is generated from a maintained, source-linked record and affected evidence is refreshed after material change. The score alone does not create a dynamic posture. The Department's suspension of Phase II leaves Phase I self-assessment requirements in place, making the quality and maintainability of those assessments the immediate issue.[2] Our perspective comes from building and operating Deep Fathom's assessment platform and supporting supplier workflows with DIB contractors, CMMC advisors, MSPs, security practitioners, and assessment organizations.

The recurring challenges are incomplete scope and inheritance, inconsistent evidence sufficiency, separation of an SPRS score from its supporting facts, delayed feedback, and point-in-time results that outlive the systems and conditions that produced them.

Deep Fathom connects scope, policies, procedures, architecture and data-flow diagrams, implementation statements, evidence, findings, and remediation in one operating record that generates linked SSP, assessment, and POA&M views. When a system, provider, responsibility, configuration, or evidence source changes, the platform can identify affected requirements and surface the resulting gap. The dynamic posture comes from maintaining the basis for the score, not merely resubmitting a score.

Self-assessment can be streamlined and strengthened by:

- linking every material claim to source, scope, collection method, relevant time, and limitations;
- applying deterministic completeness, integrity, freshness, and conflict checks before submission;
- using AI to map evidence, identify unsupported claims, and prioritize review;
- keeping machine analysis separate from the accountable human conclusion;
- applying material-change and type-specific refresh rules;
- returning structured feedback so suppliers can correct interpretation or source-data errors; and
- using risk-based sampling and independent review to check whether evidence-backed self-assessments remain accurate.

A score without its evidentiary basis is a periodic compliance assertion. A maintained, source-linked record can support continuous improvement and risk-appropriate verification.

Q6. Specific Actions for the Next 60 Days

Official question: "What specific, actionable policy changes or regulatory reforms should the CMMC Reform Task Force recommend over the next 60 days to drastically reduce costs and barriers to entry for small, medium, and non-traditional businesses without degrading the protection of federal data?"

Answer: The Department should publish a common evidence format, define graduated verification, remove recurring sources of scoping and inheritance uncertainty, and test AI-assisted verification before relying on it.

1. **Publish a graduated Level 2 verification model and risk criteria.** Define when evidence-backed self-assessment, AI-assisted evidence review with accountable human validation, C3PAO assessment, and targeted Government review are appropriate. Use data sensitivity, mission consequence, threat exposure, supplier criticality, and material change as the primary criteria. Use evidence completeness and quality to trigger sampling, correction, or escalation within the applicable tier. For each method, define the evidence producer, reviewer, decision owner, permissible use, correction process, challenge rights, and escalation path.
2. **Open a vendor-neutral CMMC evidence-profile process.** Publish an initial draft grounded in OSCAL and the existing CMMC assessment methodology. Require human-readable and machine-readable representations from the same source record. Include small manufacturers, OT operators, MSPs, RPOs, C3PAOs, primes, Government assessors, NIST, software vendors, and security researchers in the process. Publish conformance tests so no vendor owns the standard.
3. **Publish a supplier readiness, CUI, and scoping playbook.** Provide worked examples for CUI boundaries, contract-level assignment, provider classification, inheritance, customer responsibility, common non-CUI security data, and commercial-cloud conditions. Show an immature supplier how to establish scope, implement foundational safeguards, capture acceptable evidence, and obtain nonbinding feedback before an assessment with program consequences.
4. **Formalize commercial-service evidence and inherited controls.** Allow a provider to publish scoped verification information once, with explicit customer responsibilities, limitations, unresolved conditions, and validity periods. Permit downstream reuse where scope and policy alignment remain intact.
5. **Launch a vendor-neutral, no-reliance evaluation of AI-assisted verification.** Invite multiple qualified implementations to review the same synthetic evidence packages under a Government-defined protocol. Include representative small manufacturers, legacy IT, OT, MSP-supported environments, and inherited cloud services. Compare outputs with reference findings established or adjudicated by qualified reviewers, and publish performance, error, disagreement, correction, and failure results. Experimental output must not affect CMMC status, SPRS, award, affirmation, or enforcement.

The evaluation should predeclare measures in three categories:

Measure	Examples
Burden	Time to scope, supplier evidence-production labor, reviewer labor, elapsed assessment time, and cost by layer
Quality	Evidence completeness, agreement and disagreement with authorized findings, false positives, false negatives, unresolved results, and reuse without loss of context
Safety	Performance by supplier maturity and technical environment, security and privacy failures, model drift, adversarial inputs, and correction effectiveness

Feasibility, Risks, Challenges, and Innovations

Within its 60-day report, the Task Force can recommend a focused evaluation and include a draft charter, an initial objective set, proposed measures, data-handling conditions, reviewer qualifications, and success and stop criteria. The Department can then launch a time-limited follow-on evaluation under the appropriate authority.

This evaluation can begin with a limited subset of objectives and expand only after the method demonstrates reliable performance. Near-term reform should preserve the current 110 Level 2 requirements while verification methods are tested. Longer-term changes to the control set should be evidence-based and handled separately from the immediate need to improve the credibility of verification and reduce cost.

The evaluation should require source-linked findings, accountable human decisions, measured error and drift, Government-defined metrics, multiple implementations, protected evidence, and representative supplier environments. It should include incomplete, conflicting, and adversarial evidence. Results should remain separate from official status and procurement consequences until the Department establishes acceptable error limits, responsible decision-makers, correction procedures, and appeal rights.

The Department can publish the graduated-verification criteria, scoping and inheritance definitions, evidence-profile draft, and evaluation charter within 60 days. Binding changes should follow the appropriate acquisition, guidance, or rulemaking path after evidence from the evaluation is available. Existing SSPs, evidence, assessments, providers, and supplier investments should have a defined transition path.

Q7. Specific Actions to Improve Operational Resilience

Official question: "What specific, actionable policy changes or regulatory reforms should the CMMC Reform Task Force recommend over the next 60 days to drastically improve operational resilience against cyber attacks at your organization?"

Answer: The Department should improve operational resilience by maintaining strong cybersecurity and credible verification across the DIB, while adding continuity requirements for suppliers whose disruption would have material mission consequences. Modern automation, including AI restricted to defined evidence tasks, can help make verification more current and affordable without replacing independent judgment with unsupported self-reporting.

Cybersecurity reduces the likelihood and impact of compromise. Operational resilience also requires an organization to continue essential work during disruption, restore affected capabilities, and adapt after incidents and exercises. Additional resilience requirements should be reserved for suppliers and dependencies whose interruption would affect the mission.

Through our platform operations and work with DIB participants, Deep Fathom maintains connections among systems, responsibilities, implementation status, evidence, findings, and remediation. This makes it possible to identify which requirements and records are affected when an environment changes. The same principle can support resilience across the DIB: preserve foundational security for every supplier and require additional continuity evidence according to mission consequence.

The Task Force should recommend five actions:

1. **Maintain a strong cybersecurity foundation and credible verification.** Access control, MFA, secure configuration, flaw remediation, boundary protection, logging, incident response, containment, and recovery all contribute to resilience. Evidence-backed self-assessment can have a role, but independent or

Government verification remains appropriate where data sensitivity or mission consequence demands greater confidence.

2. **Tier additional resilience requirements by mission and production consequence.** The Department should consider supplier replaceability, sole-source and concentration risk, production lead time, inventory buffers, operational demand, and the effect of an outage on fielded capability. A supplier whose disruption could halt critical munitions or weapons production warrants different resilience evidence and review frequency than a readily substitutable commodity supplier.
3. **Identify the dependencies that can stop critical production.** Mission-critical suppliers should map the IT systems, operational technology, cloud and managed services, logistics functions, facilities, and upstream suppliers whose failure could interrupt delivery. The Department should distinguish the resilience of a supplier's production environment from the cybersecurity and resilience of IT or OT products delivered for Government use. These involve different systems, evidence, and remedies.
4. **Require practical continuity and recovery evidence.** Where mission consequence warrants it, suppliers should exercise containment, degraded-mode operations, manual workarounds, alternate providers or production paths, and restoration procedures. Evidence should include achieved recovery times, exercise results, unresolved findings, and demonstrated restoration of critical systems and data. This provides a more reliable measure than plans and policy documents alone.
5. **Make resilience threat-informed and adaptive.** Refresh critical evidence after material architecture or provider changes, major incidents, failed exercises, newly exploited vulnerabilities, or changes in mission criticality. Where law and sensitivity permit, return aggregate incident and threat patterns to similarly situated suppliers as warnings, targeted assistance, and temporary verification priorities. Carefully evaluated automation can help identify stale or degraded conditions and direct human attention to the suppliers and dependencies with the greatest mission consequence.

The Department can maintain a strong security baseline and rigorous verification while concentrating additional continuity requirements on critical production. Applying modern automation to defined evidence tasks may reduce the cost and increase the frequency of verification; the proposed evaluation would determine where those benefits are reliable.

Proposed Deep Fathom Contribution

Deep Fathom can contribute a working implementation to the AI-assisted verification evaluation proposed in Q6. Our platform models all 110 CMMC Level 2 requirements and 320 assessment objectives in a source-linked compliance graph and performs AI-assisted, objective-level evidence review with accountable human validation.

We propose a Government-defined, vendor-neutral evaluation in which Deep Fathom and other qualified approaches review the same synthetic Level 2 evidence packages. Government-designated reviewers, potentially including qualified Government assessors, could establish the reference findings and adjudicate disagreements. The evaluation should use Government-defined metrics and should not confer certification, authorization, or procurement preference.

Deep Fathom can contribute:

- **Structured control and objective mapping** across the full Level 2 assessment model.
- **Source-linked evidence analysis** that shows what supports each conclusion and where evidence is missing, stale, contradictory, or out of scope.
- **Deterministic conformance checks** for machine-testable requirements, combined with AI-assisted analysis where evidence requires interpretation.
- **Human review workflows** that preserve reviewer authority, document overrides, and maintain an auditable decision history.
- **Recurring and change-triggered assessment workflows** so material changes can be reviewed without rebuilding an entire assessment package.
- **Secure deployment options** that exchange only the evidence needed for review and can keep sensitive evidence in the supplier's or customer's approved environment.

The Department should assess these capabilities through independent evidence rather than vendor claims. Deep Fathom is prepared to demonstrate them against a synthetic Level 2 package and participate in a Government-defined, vendor-neutral evaluation within the Task Force's proposed follow-on period.

Source Notes

1. U.S. Small Business Administration, "SBA Commends U.S. Department of War's Suspension of CMMC Phase II for Small Defense Contractors," July 13, 2026: https://www.sba.gov/article/2026/07/13/sba-commends-us-department-wars-suspension-cmmc-phase-ii-small-defense-contractors.
2. Department of War, "Forging the Arsenal of Freedom: Department of War Suspends CMMC Phase II Requirements," July 13, 2026: https://www.war.gov/News/Releases/Release/Article/4542329/forging-the-arsenal-of-freedom-department-of-war-suspends-cmmc-phase-ii-require/.
3. CMMC Program final rule, 89 FR 83092, including small-entity cost analysis at 83183-86: https://www.govinfo.gov/content/pkg/FR-2024-10-15/pdf/2024-22905.pdf.
4. DCMA DIBCAC, "Top 10 Other Than Satisfied Requirements," based on 117 High NIST SP 800-171 assessments conducted from 2019 through 2022: https://www.dcma.mil/DIBCAC/ and https://www.dcma.mil/Portals/31/Documents/DIBCAC/DIBCAC_Top_OTS_Reqts.pptx.
5. FAR 52.204-21, "Basic Safeguarding of Covered Contractor Information Systems," and CMMC Level 1 Assessment Guide: https://www.acquisition.gov/far/52.204-21 and https://dowcio.war.gov/Portals/0/Documents/CMMC/AssessmentGuideL1v2.pdf.
6. NIST Open Security Controls Assessment Language: https://csrc.nist.gov/projects/open-security-controls-assessment-language.
7. FedRAMP 20x package guidance: https://www.fedramp.gov/2026/agencies/use/packages/20x/.
8. CMMC Program final rule, 32 CFR 170.19, external service-provider scoping: https://www.ecfr.gov/current/title-32/subtitle-A/chapter-I/subchapter-D/part-170/section-170.19.
9. Department of War CIO, "CMMC Frequently Asked Questions," July 2026 excerpt, External Service Providers: https://dowcio.war.gov/Portals/0/Documents/CMMC/FAQsv6.pdf.