

PERFORMANCE WORK STATEMENT

Federal AI Governance, Compliance, and Responsible Use Program Services

Issuing Agency: [Agency Name]

Contract Type: Firm-Fixed-Price (FFP)

Acquisition Authority: FAR Part 12 (Commercial Services)

Period of Performance: Base Year plus 4 Option Years

Contract Start Date: July 1, 2026

Version: 1.0 | Date: April 2026

Regulatory Framework

Executive Order 14110 (Safe, Secure, and Trustworthy AI) | OMB M-24-10 (Advancing Governance, Innovation, and Risk Management for Agency Use of AI) | FAR Part 12 (Commercial Items) | FAR 37.602 (Performance-Based Acquisition)

1. Introduction

1.1 Purpose

This Performance Work Statement (PWS) defines the performance outcomes, standards, and deliverables required for contractor support of the [Agency Name] Federal AI Governance, Compliance, and Responsible Use Program. The contractor shall establish, operate, and continuously improve an enterprise-wide AI governance capability in accordance with Executive Order 14110, OMB Memorandum M-24-10, and applicable agency policy.

1.2 Background

[Agency Name] operates, procures, or uses artificial intelligence systems across its mission functions. Executive Order 14110 (October 30, 2023) and OMB Memorandum M-24-10 (March 28, 2024) require each covered agency to designate a Chief AI Officer, maintain a comprehensive inventory of AI use cases, conduct impact assessments for rights-impacting and safety-impacting AI, implement risk management practices consistent with the NIST AI Risk Management Framework, and publish annual public AI use case inventories.

The Agency currently manages an estimated 200 to 500 AI systems across approximately 20 organizational components, with an anticipated growth rate exceeding 25 percent per year. Current governance activities are distributed across multiple offices without a consolidated operating model, making enterprise-level inventory, assessment, and compliance reporting resource-intensive and inconsistent. This acquisition establishes a contracted, performance-based AI governance program to stand up the operating model, execute the inventory and assessment workload, and maintain continuous compliance with federal requirements through the full period of performance.

1.3 Scope Summary

The contractor shall deliver an integrated AI governance capability consisting of: (1) maintenance of a complete, accurate inventory of agency AI systems classified by risk tier; (2) execution of AI impact assessments for rights-impacting and safety-impacting systems within defined cycle times; (3) development and continuous maintenance of AI policy, including standards aligned to NIST AI RMF; (4) compliance monitoring and reporting to OMB, agency leadership, and the public; (5) workforce training on responsible AI use and governance processes; (6) operation of the agency Governance, Risk, and Compliance (GRC) platform configured for AI governance workflows; (7) review of AI-related procurement actions for governance compliance; (8) triage and response to AI-related incidents and harms; and (9) stakeholder engagement across agency components.

Work is performed principally at the Agency HQ in the Washington DC metropolitan area, with virtual support extending to all agency components. The contractor shall not perform inherently governmental functions. Decisions requiring final agency determinations (e.g., waiver approvals, final risk acceptances, policy issuance) remain with authorized Government officials.

1.4 Applicable Documents and Standards

- Executive Order 14110, Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (October 30, 2023)

- OMB Memorandum M-24-10, Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence (March 28, 2024)
- NIST AI 100-1, Artificial Intelligence Risk Management Framework (AI RMF 1.0)
- NIST AI 600-1, Artificial Intelligence Risk Management Framework: Generative AI Profile
- FAR Part 12, Acquisition of Commercial Products and Commercial Services
- FAR 37.602, Performance-Based Acquisition
- FAR 52.212-4, Contract Terms and Conditions — Commercial Products and Commercial Services
- Agency AI Policy (as promulgated by the Chief AI Officer during performance)
- Section 508 of the Rehabilitation Act of 1973, as amended
- Federal Information Security Modernization Act (FISMA) of 2014

2. Definitions and Acronyms

2.1 Defined Terms

AI System: An automated system that uses machine learning, statistical models, or similar techniques to generate outputs that influence or inform decisions, recommendations, or actions, as defined in OMB M-24-10.

Rights-Impacting AI: AI whose output serves as a principal basis for a decision or action affecting individual rights, opportunities, or access to services, as defined in OMB M-24-10 Appendix I.

Safety-Impacting AI: AI whose output has the potential to significantly impact the safety of human life or well-being, critical infrastructure, or the environment, as defined in OMB M-24-10 Appendix I.

AI Impact Assessment: The structured evaluation required by OMB M-24-10 for rights-impacting and safety-impacting AI, documenting intended use, risks, mitigations, and stakeholder consultation.

AI Use Case Inventory: The comprehensive list of AI systems in use, development, or procurement at the Agency, maintained pursuant to EO 13960 and OMB M-24-10.

Chief AI Officer: The senior official designated by the Agency pursuant to OMB M-24-10 Section 3, responsible for coordinating AI use, promoting AI innovation, and managing AI risks.

GRC Platform: The Governance, Risk, and Compliance technology platform used by the Agency to manage AI governance workflows, inventory, assessment documentation, and compliance reporting.

AQL: Acceptable Quality Level, the maximum allowable defect rate or minimum performance threshold at which a performance standard is considered met.

2.2 Acronyms

Acronym	Full Form
AI	Artificial Intelligence

AIMS	AI Management System
AQL	Acceptable Quality Level
ATO	Authority to Operate
CAIO	Chief AI Officer
CIO	Chief Information Officer
CLIN	Contract Line Item Number
COR	Contracting Officer's Representative
EO	Executive Order
FAR	Federal Acquisition Regulation
FISMA	Federal Information Security Modernization Act
GRC	Governance, Risk, and Compliance
HQ	Headquarters
KPI	Key Performance Indicator
NIST	National Institute of Standards and Technology
OMB	Office of Management and Budget
POP	Period of Performance
PWS	Performance Work Statement
QASP	Quality Assurance Surveillance Plan
RMF	Risk Management Framework
SME	Subject Matter Expert

3. Requirements — Performance Objectives

This section sets forth the performance objectives the contractor shall achieve. Each objective identifies the required outcome, the performance standard, the Acceptable Quality Level, and the method of Government assessment. The contractor retains the freedom to propose and implement the approach, staffing, and methods necessary to meet or exceed the performance standards, consistent with FAR 37.602.

3.1 AI Use Case Inventory and Classification

Required Outcome: The contractor shall achieve and maintain a complete, current, and accurately classified inventory of all AI systems in development, procurement, or operation at the Agency, in conformance with OMB M-24-10 inventory requirements and EO 13960.

Performance Standard: All rights-impacting and safety-impacting AI systems are inventoried and risk-tier classified within 30 calendar days of the Agency becoming aware of the system. The inventory is published internally in the GRC platform and externally (as applicable) on the Agency's public AI use case inventory page on the statutory reporting cadence.

Acceptable Quality Level (AQL): 100 percent of high-risk (rights-impacting or safety-impacting) AI systems inventoried within 30 days; 95 percent of all AI systems inventoried within 60 days.

Method of Assessment: Monthly government audit of the inventory against known AI procurement records, CIO system catalogs, and randomly sampled component interviews.

3.2 AI Impact Assessments

Required Outcome: The contractor shall complete AI Impact Assessments for all rights-impacting and safety-impacting AI systems in accordance with OMB M-24-10 Section 5 requirements, including risk identification, mitigation planning, stakeholder consultation, and documentation.

Performance Standard: Completed AI Impact Assessments are delivered to the Chief AI Officer for review within 60 calendar days of a system being identified as rights-impacting or safety-impacting. Assessment packages include the elements required by OMB M-24-10 Appendix II and are structured for rapid Chief AI Officer review.

Acceptable Quality Level (AQL): 95 percent of assessments delivered on time; 90 percent approved by the Chief AI Officer on first review without substantive rework.

Method of Assessment: Government review of assessment deliverables for completeness, analytical rigor, and policy alignment. Quarterly Chief AI Officer sign-off on the assessment backlog.

3.3 Risk Management and NIST AI RMF Alignment

Required Outcome: The contractor shall maintain an AI risk management program aligned to the NIST AI Risk Management Framework (AI RMF 1.0 and the Generative AI Profile), including risk identification, measurement, mitigation planning, and continuous monitoring.

Performance Standard: All active rights-impacting and safety-impacting AI systems have a current risk profile with identified risks categorized by NIST AI RMF function (Govern, Map, Measure, Manage), documented mitigations, and a review cycle not to exceed 12 months.

Acceptable Quality Level (AQL): 100 percent of high-risk systems with current risk profiles at any point in time; 95 percent on time for scheduled reviews.

Method of Assessment: Quarterly Government sampling of 10 percent of active high-risk systems for risk profile currency, mitigation closure, and framework alignment.

3.4 Policy Development and Maintenance

Required Outcome: The contractor shall develop, maintain, and periodically revise Agency AI policies, standards, and operating procedures aligned with federal requirements and emerging best practices.

Performance Standard: AI policy library is reviewed and updated within 45 calendar days of any material change to federal AI policy (EO, OMB memo, NIST framework update) or identified internal gap. Each policy has an assigned owner, version history, and next scheduled review date.

Acceptable Quality Level (AQL): 90 percent of federal policy changes triggering updates completed within 45 days; no policies past due for scheduled review by more than 30 days.

Method of Assessment: Monthly policy library currency report; Government spot-checks during Chief AI Officer policy council meetings.

3.5 Compliance Monitoring and Reporting

Required Outcome: The contractor shall produce accurate, timely compliance reporting to OMB, Agency leadership, and the public as required by EO 14110, OMB M-24-10, and Agency policy.

Performance Standard: Annual OMB AI use case inventory report is prepared and submitted on time with zero material errors identified in OMB quality review. Quarterly leadership dashboards are delivered within 10 business days of quarter close. Public AI use case inventory is updated on the Agency website within 30 calendar days of inventory changes.

Acceptable Quality Level (AQL): 100 percent on-time delivery for OMB-required reports; 95 percent on-time for internal leadership reports; zero material reporting errors.

Method of Assessment: Direct Government review of submitted reports; OMB quality feedback; internal leadership feedback on dashboard utility.

3.6 Workforce Training and Awareness

Required Outcome: The contractor shall develop, deliver, and maintain role-based training on responsible AI use, AI governance processes, and compliance obligations for the Agency workforce and contractor population using AI tools.

Performance Standard: Role-based training curricula cover general AI awareness (all staff), AI system owners and developers (intermediate), and AI governance staff and Chief AI Officer council (advanced). 90 percent of targeted audiences complete required training within 60 days of assignment. Training content is updated within 60 days of any material federal policy change.

Acceptable Quality Level (AQL): 90 percent completion rates on all required training; positive feedback (4.0 of 5.0 or higher) on training effectiveness surveys.

Method of Assessment: Learning Management System completion reports; participant feedback surveys; quarterly training program review.

3.7 GRC Platform Configuration and Operations

Required Outcome: The contractor shall stand up, configure, and operate the Agency GRC platform for AI governance workflows, including inventory management, assessment intake and routing, risk register, policy library, and reporting.

Performance Standard: Platform is fully operational for all AI governance workflows within 120 calendar days of contract start. Platform availability is maintained at 99.5 percent or better during business hours (06:00-20:00 ET, weekdays) for the remainder of the period of performance. Workflow configurations are updated within 30 days of process changes.

Acceptable Quality Level (AQL): 99.5 percent platform availability during business hours; 100 percent of scheduled releases executed with zero production defects affecting compliance workflows.

Method of Assessment: Automated platform monitoring; Government review of change records and release notes.

3.8 AI-Related Procurement Review

Required Outcome: The contractor shall review Agency procurement actions that involve AI acquisition or AI-enabled services for compliance with EO 14110, OMB M-24-10 Section 4, and applicable FAR provisions, providing findings to contracting officers and the Chief AI Officer.

Performance Standard: All AI-related procurement actions routed to contractor review receive a documented review finding within 10 business days of intake. Findings include identification of applicable governance requirements, recommended PWS or SOW language, and flagging of pre-award risks.

Acceptable Quality Level (AQL): 95 percent of reviews completed within 10 business days; zero post-award governance findings that should have been caught at pre-award review.

Method of Assessment: Monthly review cycle time report; quarterly audit of post-award findings traceable to pre-award review gaps.

3.9 AI Incident Response and Harms Triage

Required Outcome: The contractor shall operate an AI incident intake and triage function receiving reports from internal stakeholders, the public, and automated monitoring; classifying incidents by severity; and coordinating Agency response in accordance with OMB M-24-10 incident reporting requirements.

Performance Standard: All incident reports are triaged with initial severity classification within one business day of receipt. High-severity incidents (affecting rights-impacting or safety-impacting systems with active harm) are escalated to the Chief AI Officer within four hours of classification. Federal AI incident reporting obligations (to OMB or NIST as applicable) are met within required timeframes.

Acceptable Quality Level (AQL): 100 percent of high-severity incidents escalated within four hours; 95 percent of all incidents triaged within one business day; 100 percent on-time federal reporting.

Method of Assessment: Incident log review; Chief AI Officer confirmation of escalation timeliness; federal reporting confirmation.

3.10 Stakeholder Engagement and Communications

Required Outcome: The contractor shall maintain regular engagement with Agency components, the Chief AI Officer council, external oversight bodies (as directed), and the public (through inventory publication and appropriate disclosures) to ensure the governance program reflects agency operational needs and external accountability expectations.

Performance Standard: Monthly Chief AI Officer council meetings are supported with pre-briefs, agenda materials, and action item tracking. Component engagement sessions are held with each agency component at least semi-annually. Public communications (inventory, statutorily required disclosures) are released on committed schedules.

Acceptable Quality Level (AQL): 100 percent of Chief AI Officer council meetings supported; 100 percent of components engaged on the semi-annual cadence; 100 percent of public releases on schedule.

Method of Assessment: Meeting log and action item tracker; component feedback survey; public release archive audit.

4. Deliverables

The contractor shall produce and submit the deliverables set forth below. Deliverable acceptance is per the acceptance criteria specified. The Contracting Officer's Representative (COR) is the designated acceptance authority.

ID	Deliverable	Description	Format	Frequency	Due	Acceptance
D-01	Program Management Plan	Describes contractor approach to meeting all performance objectives, including staffing model, quality control, risk management, and subcontracting (if any).	PDF, editable source	Once, updated annually	30 days post-award; annual updates within 30 days of option exercise	Complete per template; COR approval

D-02	Monthly Status Report	Progress against each performance objective, AQL attainment, issues, risks, and forward-look.	PDF + Excel metrics	Monthly	10 business days after month end	Complete; correct metrics; narrative addresses all AQL variances
D-03	AI Use Case Inventory	Current inventory of agency AI systems with classification and required OMB M-24-10	GRC platform database + machine-readable export	Continuous; export monthly	Continuous maintenance; monthly export by 5th business day	Inventory completeness verified; OMB metadata complete
D-04	Quarterly Leadership Dashboard	Rollup of governance program health, inventory trends, assessment throughput, compliance posture, and risks requiring leadership attention.	PDF + interactive dashboard	Quarterly	10 business days after quarter close	Content complete; data validated; narrative actionable
D-05	AI Impact Assessment Packages	Individual assessment packages per OMB M-24-10 Appendix II for rights- and safety-impacting	PDF package with supporting artifacts	On demand, per system	60 days from identification trigger	Complete per OMB Appendix II; CAIO approved
D-06	Annual OMB AI Use Case Inventory Report	Statutory annual report to OMB per EO 13960 and M-24-10.	OMB-specified format	Annually	Per OMB schedule	Accepted by OMB; zero material errors

D-07	AI Policy Library	Complete set of Agency AI policies, standards, and SOPs.	PDF + controlled-copy source in GRC platform	Continuous; baseline delivery and semi-annual certification	Baseline: 90 days post-award; semi-annual recertifications	Policy coverage complete; owner-signed currency
D-08	Training Curricula and Completion Reports	Role-based curricula, training materials, and LMS completion reports by audience tier.	SCORM-compatible modules + monthly completion CSV	Curricula: baseline, updated; Reports: monthly	Baseline curricula: 120 days post-award; reports monthly	Curricula reviewed by Chief AI Officer; reports accurate
D-09	AI-Related Procurement Review Findings	Documented review of AI-related contract actions.	Findings memo per action	On demand, per action	10 business days from intake	Complete; recommendations actionable
D-10	AI Incident Reports	Triaged incidents with classification, root cause, mitigation, and federal reporting status.	Incident log + per-incident reports	Continuous; summary monthly	Continuous; summary by 10th business day of following month	Complete log; federal reports timely
D-11	Transition-In Plan	Plan for inheriting incumbent artifacts and stabilizing service operations.	PDF	Once	15 days post-award	COR approved before transition-in execution
D-12	Transition-Out Plan	Plan for knowledge transfer, artifact return, and incoming contractor cooperation.	PDF	Once, updated annually	Draft: 180 days before end of final period; final: 90 days before end	COR approved

D-13	Final Close-Out Report	End-of-contract summary of performance, lessons learned, and outstanding items.	PDF	Once	30 days after contract end	Complete; all artifacts returned
-------------	------------------------	---	-----	------	----------------------------	----------------------------------

5. CLIN Structure

This is a Firm-Fixed-Price contract under FAR Part 12 (Commercial Services). The CLIN structure is by period of performance. All performance objectives in Section 3 are priced within the applicable period CLIN. Travel and ODCs are reimbursable at cost with no fee markup, subject to FAR 31.205-46 (Travel) and the Limitation of Cost clause applicable to the CLIN.

CLIN	Description	Pricing Basis	Scope
0001	Base Year Services	FFP	All Section 3 performance objectives for the Base Year period, inclusive of Transition-In activities.
0002	Base Year Travel	Cost-reimbursable (not to exceed)	Travel in support of Section 3 objectives, subject to FAR 31.205-46.
0003	Base Year ODCs	Cost-reimbursable (not to exceed)	Other direct costs per approved Program Management Plan.
1001	Option Year 1 Services	FFP	All Section 3 performance objectives for OY1.
1002	Option Year 1 Travel	Cost-reimbursable (NTE)	Same basis as CLIN 0002.
1003	Option Year 1 ODCs	Cost-reimbursable (NTE)	Same basis as CLIN 0003.
2001	Option Year 2 Services	FFP	All Section 3 performance objectives for OY2.
2002	Option Year 2 Travel	Cost-reimbursable (NTE)	Same basis as CLIN 0002.
2003	Option Year 2 ODCs	Cost-reimbursable (NTE)	Same basis as CLIN 0003.
3001	Option Year 3 Services	FFP	All Section 3 performance objectives for OY3.
3002	Option Year 3 Travel	Cost-reimbursable (NTE)	Same basis as CLIN 0002.
3003	Option Year 3 ODCs	Cost-reimbursable (NTE)	Same basis as CLIN 0003.

4001	Option Year 4 Services	FFP	All Section 3 performance objectives for OY4, inclusive of Transition-Out activities.
4002	Option Year 4 Travel	Cost-reimbursable (NTE)	Same basis as CLIN 0002.
4003	Option Year 4 ODCs	Cost-reimbursable (NTE)	Same basis as CLIN 0003.

6. Period of Performance

Base Year: July 1, 2026 — June 30, 2027 (twelve months)

Option Year 1: July 1, 2027 — June 30, 2028

Option Year 2: July 1, 2028 — June 30, 2029

Option Year 3: July 1, 2029 — June 30, 2030

Option Year 4: July 1, 2030 — June 30, 2031

Option exercise is at the sole discretion of the Government pursuant to FAR 52.217-9, with the notice periods stated in the resulting contract.

7. Place of Performance

Primary: [Agency Name] Headquarters, Washington DC metropolitan area.

Virtual support: All Agency components nationwide, via secure remote access to the GRC platform and Agency collaboration tools.

Incidental travel: To Agency component sites as required to support inventory completeness, assessment execution, and component engagement per Section 3.10. All travel is subject to GSA Federal Travel Regulation rates and requires advance COR authorization.

8. Government-Furnished Property and Information

- Access to the Agency GRC platform (government-owned or government-furnished under separate license).
- Access to the Agency network and collaboration tools for all contractor personnel requiring such access.
- Agency AI policy baseline documents (as they exist at contract start).
- Agency AI use case inventory baseline (as it exists at contract start).
- Workspace at Agency HQ for designated key personnel (subject to space availability).

The Government does not furnish contractor personnel laptops, productivity software, or non-Agency-specific tools. The contractor is responsible for all contractor-owned equipment and software.

9. Security Requirements

9.1 Personnel Security

All contractor personnel requiring access to Agency systems or data shall have, at minimum, a favorable Tier 2 (Moderate Risk Public Trust) investigation prior to network access. Personnel requiring access to sensitive but unclassified systems or data shall meet the position sensitivity level designated by the Agency for the role. Key Personnel (Section 10) shall meet Tier 3 (High Risk Public Trust) or as designated by the Agency.

9.2 Information Security

Contractor handling of Agency information shall comply with FISMA, NIST SP 800-53 controls appropriate to the information system categorization, and Agency information security policy. All contractor systems storing, processing, or transmitting Agency information must be authorized by the Agency per the Risk Management Framework prior to use.

9.3 Handling of AI Governance Artifacts

AI impact assessment packages, risk registers, and incident reports may contain sensitive agency operational information and personally identifiable information. Contractor shall apply Controlled Unclassified Information (CUI) handling per 32 CFR 2002 to all such artifacts unless otherwise designated by the Agency.

9.4 Incident Reporting

Contractor shall report suspected information security incidents involving Agency information or systems to the Agency Security Operations Center within one hour of discovery, consistent with OMB M-22-05 and Agency policy.

10. Key Personnel

The contractor shall designate the following roles as Key Personnel. Substitution of any Key Personnel during contract performance requires prior written approval of the Contracting Officer per the substance of FAR 52.237-2 (Substitution of Personnel). The Government may require removal and replacement of any contractor personnel whose performance is determined to be unsatisfactory.

10.1 Program Manager

Single point of accountability for contract performance. Minimum qualifications: 10 years of experience managing complex federal services contracts, of which at least three years managing information technology governance, risk, or compliance programs. Current PMP certification or equivalent demonstrated experience. Active Public Trust investigation (Tier 2 or higher) at time of proposal.

10.2 AI Governance Lead

Senior subject matter expert for AI governance program operations. Minimum qualifications: 7 years of experience in technology or AI governance, risk management, or regulatory compliance, of which at least two years with federal AI policy (EO 13960, EO 14110, or OMB M-24-10). Advanced degree in a relevant field or equivalent experience. Demonstrated expertise in NIST AI RMF application.

10.3 Chief AI Ethics Officer

Senior advisor on responsible AI, bias assessment, and rights-impacting AI considerations. Minimum qualifications: 10 years of experience spanning public policy, applied ethics, civil rights, or responsible AI practice. Demonstrated publications, speaking engagements, or prior government service establishing standing in the field.

10.4 Technical Integration Lead

Lead architect for GRC platform configuration and system integrations. Minimum qualifications: 10 years of experience in enterprise systems architecture, including at least three years with GRC platforms (ServiceNow GRC, Archer, or equivalent). Current industry certification in the platform being operated at the Agency.

11. Reporting and Oversight

11.1 Reporting Cadence

Activity	Frequency	Timing
Contract Kickoff Meeting	Once	Within 15 business days of contract start
Weekly Status Stand-up	Weekly	Each week, COR and Program Manager
Monthly Status Report (D-02)	Monthly	By 10th business day of following month
Monthly Chief AI Officer Council Support	Monthly	Per Chief AI Officer schedule
Quarterly Leadership Dashboard (D-04)	Quarterly	10 business days after quarter close
Quarterly Contract Performance Review	Quarterly	30 days after quarter close
Annual Program Management Plan Update	Annual	Within 30 days of option exercise
Annual OMB Inventory Report Submission (D-06)	Annual	Per OMB schedule

11.2 Points of Contact

- Contracting Officer: As designated in the resulting contract.
- Contracting Officer's Representative (COR): As designated in the resulting contract. Primary day-to-day Government technical point of contact.
- Contract Specialist: Secondary acquisition point of contact.

- Chief AI Officer: Senior Agency sponsor and approval authority for AI impact assessment packages, policy issuance, and risk acceptance decisions.

12. Quality Assurance Surveillance Plan (Summary)

This section summarizes the Government surveillance methods and performance thresholds for each performance objective. A detailed QASP may be issued separately and updated during performance. The methods below take precedence.

Objective	Performance Standard	AQL Threshold	Method	Frequency
3.1 Inventory	100% high-risk within 30 days; 95% all within 60 days	95% threshold for CPARS Satisfactory	Monthly audit	Monthly
3.2 Assessments	95% on-time; 90% first-review approval	90% on-time for CPARS Satisfactory	CAIO review + sampling	Quarterly
3.3 Risk Management	100% currency; 95% on-time reviews	90% currency for CPARS Satisfactory	Quarterly sampling	Quarterly
3.4 Policy	90% of updates in 45 days; no past due >30 days	80% for CPARS Satisfactory	Monthly currency report	Monthly
3.5 Compliance Reporting	100% on-time OMB; 95% internal; zero material errors	95% on-time for CPARS Satisfactory	Direct report review	Per report cycle
3.6 Training	90% completion; 4.0/5.0 feedback	85% completion for CPARS Satisfactory	LMS reports + survey	Monthly
3.7 GRC Platform	99.5% availability; zero release defects	99.0% for CPARS Satisfactory	Automated monitoring	Continuous
3.8 Procurement Review	95% within 10 business days	90% for CPARS Satisfactory	Cycle time report + audit	Monthly
3.9 Incident Response	100% high-severity ≤4hr; 95% triage ≤1 business day	95% triage for CPARS Satisfactory	Incident log review	Monthly
3.10 Stakeholder Engagement	100% CAIO council; 100% semi-annual components	95% for CPARS Satisfactory	Meeting log + surveys	Per engagement

13. Transition

13.1 Transition-In

The contractor shall complete transition-in within 90 calendar days of contract start. Transition-in activities include: (a) review and inheritance of incumbent artifacts (inventory baseline, draft policies,

assessment backlog); (b) stand-up of the GRC platform for AI governance workflows; (c) key personnel onboarding and security clearance verification; (d) establishment of recurring Government touchpoints (weekly stand-up, monthly CAIO council support, monthly reporting cadence); and (e) delivery of the initial AI use case inventory export demonstrating inventory completeness. During transition-in, performance against Section 3 AQLs is measured on a ramp-up basis documented in the Program Management Plan.

13.2 Transition-Out

The contractor shall prepare a Transition-Out Plan beginning no later than 180 calendar days prior to the end of the final period of performance, with final Plan delivered 90 calendar days prior to period end. Transition-out includes: (a) complete knowledge transfer to the Agency or incoming contractor; (b) return of all Government-furnished information and artifacts; (c) export of complete GRC platform configuration and data; (d) documentation of all active incidents, in-flight assessments, and pending policy actions; (e) incumbent cooperation with the incoming contractor for a period of up to 60 calendar days after contract end, subject to separate bridge authorization if beyond the base period.

14. Constraints and Assumptions

14.1 Constraints

- Contractor shall not perform inherently governmental functions (FAR 7.503). Final Agency determinations including waiver approvals, risk acceptances, and policy issuance remain with authorized Government officials.
- Contractor personnel shall not authorize Government expenditures or bind the Government contractually.
- Work is performed in accordance with Agency policy and applicable federal law in effect during performance; material changes in federal AI policy may trigger scope adjustment per the Changes clause.

14.2 Assumptions

- The Agency will maintain an active Chief AI Officer designation throughout the period of performance.
- The Agency will provide access to the designated GRC platform, Agency network, and necessary systems within 30 days of contract start.
- Agency AI policy baseline documents and the current AI use case inventory will be provided to the contractor no later than the contract kickoff meeting.
- Chief AI Officer will maintain availability for monthly governance council meetings and for rights-impacting and safety-impacting assessment approvals within the cycle times in Section 3.
- Commercial item determination pending; to be confirmed by the Contracting Officer prior to award. If determined non-commercial, the resulting contract will incorporate FAR Part 15 clauses accordingly.

Appendix A: Current Environment Description

The Agency's current AI governance environment is characterized by distributed ownership across program offices, an incomplete centralized inventory, and ad hoc assessment processes. Estimated 200 to 500 AI systems in use, development, or procurement across approximately 20 organizational components. Current inventory coverage is estimated at 40 to 60 percent of active systems. Impact assessments have been completed for an estimated 15 to 25 percent of identified rights-impacting or safety-impacting systems. The Agency has a GRC platform in operation but is not yet configured for AI governance workflows.

Anticipated growth rate: 25 to 35 percent per year in AI system count, driven by increased adoption of generative AI tools and AI-enabled COTS procurement.

Appendix B: Volume Data and Historical Metrics

Metric	Estimated Volume
Active AI systems (estimated baseline at contract start)	300 systems
Rights-impacting or safety-impacting (estimated baseline)	75 systems
Anticipated new AI systems entering inventory per year	75 to 125
Anticipated AI impact assessments per year	40 to 60
Anticipated AI incident reports per year	20 to 40
Anticipated AI-related procurement reviews per year	80 to 120
Target population for general AI awareness training	15,000 personnel
Target population for intermediate training (AI system owners/developers)	2,000 personnel
Target population for advanced training (governance staff and CAIO council)	200 personnel

Appendix C: System Interface Specifications

The GRC platform shall integrate with the following Agency enterprise systems. Integration specifications are maintained in the Agency Enterprise Architecture repository and shall be made available to the contractor post-award.

- Enterprise Data Catalog (for cross-reference of AI system metadata)
- Identity Management System (for user provisioning on GRC platform)
- Procurement System (for AI-related procurement action feeds)
- Security Operations Center platform (for AI incident cross-reference)

- HR Information System (for training assignment by role and component)
- Public-facing AI use case inventory page (for automated inventory publication)

Appendix D: Acronym List

See Section 2.2 for the full acronym list. Additional acronyms that may be encountered during performance are defined in the Agency AI Policy library as maintained per D-07.