



VENDOR PROPOSAL COMPLIANCE
MANAGEMENT USE CASE

Agentic AI Lab: Proof of Concept Document

Acknowledgements

Alyssa Feola, Seventeen Sierra
Anil 'Neil' Chaudhry, US Department of Transportation
Chris Robinson, GSA
Dani Hill, AWS
Frank Indiviglio, NOAA
Jim St. Clair, C3HIE
KJ Lian, AWS
Matt Craddock, AWS
Scott Hilner, AWS

Disclaimer

This document was prepared by the members of the ATARC Agentic AI Lab in their personal capacity. The opinions expressed do not reflect any specific individual nor any organization or agency they are affiliated with and shall not be used for advertisement or product endorsement purposes.

1. Background and Strategic Fit

The Agentic AI Lab working group, established by ATARC, aims to explore and test various use cases for artificial intelligence in government contexts. Through a member voting process, the group selected Vendor Proposal Compliance Management as its inaugural project. This initiative aligns with ATARC's mission to foster innovation and collaboration between government, academia, and industry in the field of emerging technologies.

The lab's focus on agentic AI, systems that can autonomously act on behalf of users, positions it at the forefront of AI applications in government procurement and compliance management. By leveraging cutting-edge AI technologies, the lab seeks to address critical challenges in federal procurement processes while promoting open standards and accessibility.

2. Business Challenges

The primary challenge addressed by this initiative is the complexity and dynamism of federal procurement regulations. Specifically, Federal agencies struggle to stay current with the ever-changing Federal Acquisition Regulation (FAR) and Defense Federal Acquisition Regulation Supplement (DFARS). New Executive Orders (EOs) from the current administration add another layer of complexity to compliance requirements. Vendors, especially small businesses, face difficulties with ensuring their proposals align with the latest regulations and orders. The review process for proposal compliance is time-consuming and prone to human error, potentially leading to non-compliant submissions or missed opportunities.

These challenges collectively contribute to inefficiencies in the federal procurement process, potentially impacting the quality and competitiveness of vendor proposals and the government's ability to procure optimal solutions. Questions were raised regarding security risks associated with Agentic AI systems and potential mitigation strategies. We previously invited thought leader Ken Huang to present to the ATARC Agentic AI community, where he shared his [MAESTRO framework](#).

3. Proof-of-Concept (PoC) Environment

The Proof of Concept (PoC) environment and solution developed by the ATARC Agentic AI Lab consists of the following components:

3.1 High-Level Architecture

The solution architecture comprises five integrated layers that work together to deliver automated compliance assessment capabilities. The User Interface provides a web-based interface for proposal submission and report output, detailed in Appendix A. The Agentic AI Framework initially implemented [crew.ai](#) but subsequently migrated to [Strands Agents SDK](#) for improved output formatting and enhanced agent coordination. The AI Model Integration connects to Amazon Bedrock to access Claude 4.0 Sonnet for complex regulatory interpretation. The Data Collection Layer implements web crawling capability for gathering the latest regulatory information with updates for FAR, DFARS, and Executive Orders. The Infrastructure Layer employs containerized architecture ensuring deployment flexibility across cloud platforms. Enumerated aforementioned architecture components:

- **User Interface:** Web-based interface for proposal submission and report output. ([Appendix A: User Interface Development](#)).
- **Agentic AI Framework:** Multi-Agent implementation of Strands Agents to integrate with content and display results on [interface](#). (migrated for improved output formatting)
- **AI Model Integration:** Connects to Amazon Bedrock to access the [Claude 4.0 Sonnet](#) model for GovCloud compliance and analysis.
- **Data Collection Layer:** Web crawling capability for gathering latest regulatory information with the latest updates for FAR, DFARS, and Executive Orders.
- **Infrastructure Layer:** Containerized Architecture ensuring deployment flexibility across cloud platforms.

3.2 Multi-Agent Architecture and Amazon Bedrock Implementation

The solution employs specialized agents for different compliance domains, with each agent designed to address specific regulatory requirements and analytical perspectives. The ATARC Agentic AI Lab's vendor proposal compliance assessment solution leverages Amazon Bedrock as its foundational AI infrastructure, providing generative AI capabilities with built-in security, privacy, and responsible AI features.

3.2.1 Foundation Model Selection and Rationale

The solution employs a dual-model strategy that balances sophisticated reasoning capabilities with cost-effective performance optimization. Claude 4.0 Sonnet serves as the primary reasoning engine for complex regulatory interpretation and compliance analysis. This model excels at understanding nuanced legal language found in FAR guidance, and synthesizing information across multiple regulatory documents to form well-rounded compliance assessments. Claude 4.0 Sonnet's sizable context window enables the system to process lengthier vendor proposals alongside regulatory documentation without losing context that might impact model output quality.

Amazon Nova Pro could potentially complement Claude 4.0 Sonnet by providing high-performance natural language processing for analytical tasks. Nova Pro's price performance makes it ideal for high-volume workloads such as document chunking, initial relevance scoring, and metadata extraction. This allocation of model tasks is performant and efficient in that it ensures intensive reasoning tasks leverage Claude's capabilities while routine operations can be addressed using Nova Pro.

3.2.2 Knowledge Base Architecture and Implementation

The system's knowledge bases provide agents with comprehensive access to federal procurement regulations through two primary repositories: 1) the FAR corpus and 2) the Executive Orders index. The FAR knowledge base is organized to preserve the hierarchical structure of regulations and their subsection relationships. This ensures agents retrieve complete regulatory requirements with full context, mitigating model misinterpretation.

The Executive Orders knowledge base captures presidential mandates with additional compliance requirements beyond standard FAR guidelines (i.e., applicability information that determines which orders apply to specific contract types, dollar thresholds, etc.). Both knowledge bases use semantic search capabilities that enable agents to identify relevant regulations, enabling accurate retrieval even when vendor proposals use different terminology than official regulatory language.

To ensure knowledge base currency and compliance accuracy, the system implements an automated update pipeline through AWS Lambda functions that periodically crawl official government sources including the Federal Register and SAM.gov to identify new Executive Orders and FAR amendments while preserving version history.

3.2.3 Multi-Agent Orchestration and Specialization

The solution's migration from crew.ai to Strands Agents was a deliberate architectural decision to improve output formatting, agent coordination, and result synthesis. Strands Agents provides a more flexible framework for defining agent behaviors, managing inter-agent communication, and orchestrating parallel analysis workflows. The multi-agent architecture employs three specialized agents, each with their respective responsibilities, and knowledge and data access patterns.

- The **FAR Agent** serves as the primary compliance evaluator, focusing exclusively on Federal Acquisition Regulation requirements. This agent has been instructed via prompting to analyze proposals systematically against FAR provisions, identify compliance requirements, and assess the completeness of required documents. The FAR Agent queries the FAR knowledge base using semantic search to retrieve relevant regulatory sections based on proposal content, then determines compliance status. Its output includes specific FAR citations, compliance determinations with supporting evidence, and prioritized remediation recommendations.

- The **Executive Order Agent** operates in parallel, focusing on presidential directives that may impose additional requirements beyond standard FAR provisions. This agent maintains awareness of the temporal applicability of Executive Orders, understanding which directives apply to specific contract types, dollar thresholds, and procurement categories. The EO Agent's specialized knowledge enables it to identify compliance requirements that might not be immediately apparent to human reviewers, particularly when Executive Orders reference other regulatory frameworks or impose requirements through indirect mechanisms.
- The **Technical Agent** addresses technical compliance aspects including security frameworks, technical specifications, and performance requirements. This agent evaluates whether vendor proposals adequately address a solution's technical needs (i.e., cybersecurity requirements, data privacy standards, etc.). Ultimately, the Technical Agent ensures vendor proposals meet substantive technical requirements and procedural compliance.

3.3 Multi-Agent Orchestration and Specialization

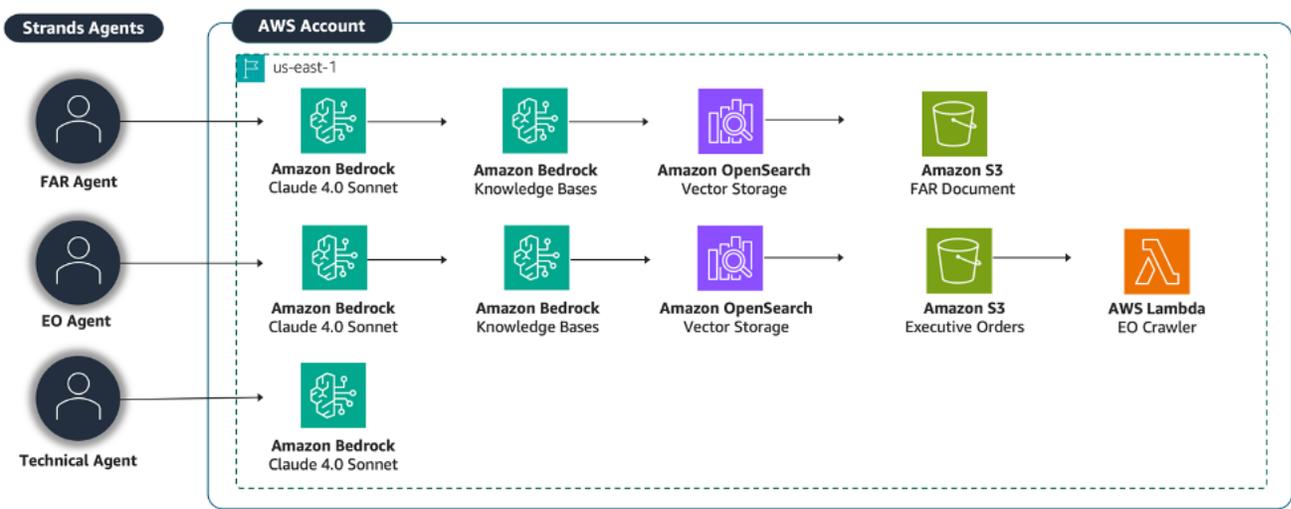


Figure 1: Solution Architecture

3.4 AWS Technical Implementation

The POC leverages AWS services within a structured architecture deployed in the us-east-1 region. Storage is managed through dedicated S3 buckets for both FAR documents and Executive Orders, providing scalable and durable object storage. For vector storage and similarity search capabilities, the architecture utilizes OpenSearch as the vector database. The AI/ML layer is powered by Amazon Bedrock with Nova Pro or Claude Sonnet, enabling advanced natural language processing and generative AI capabilities. Automation is handled through AWS Lambda functions that perform Executive Order crawling and updates, ensuring the system remains current with the latest regulatory information. Further details below:

Storage Setup:

Amazon S3 provides the document storage layer through two dedicated buckets. The FAR documents bucket maintains the repository of Federal Acquisition Regulations that agents query during compliance analysis, while the Executive Orders bucket stores presidential directives that may impose additional procurement requirements.

Vector Database:

Amazon OpenSearch serves as the vector database for semantic search capabilities, storing embeddings generated from regulatory text. When agents analyze vendor proposals, they query OpenSearch to retrieve relevant FAR sections and Executive Orders based on semantic similarity rather than simple keyword matching.

Foundation Model Infrastructure:

Amazon Bedrock serves as the central AI infrastructure, providing managed access to foundation models without requiring infrastructure management. The solution leverages Claude 4.0 Sonnet through Amazon Bedrock (Strands Agents default as a function of agent orchestration).

Agent Orchestration:

The Strands Agents SDK orchestration framework coordinates the three specialized agents (FAR Agent, Executive Order Agent, and Technical Agent) as they analyze vendor proposals against regulatory requirements.

Automation Functions:

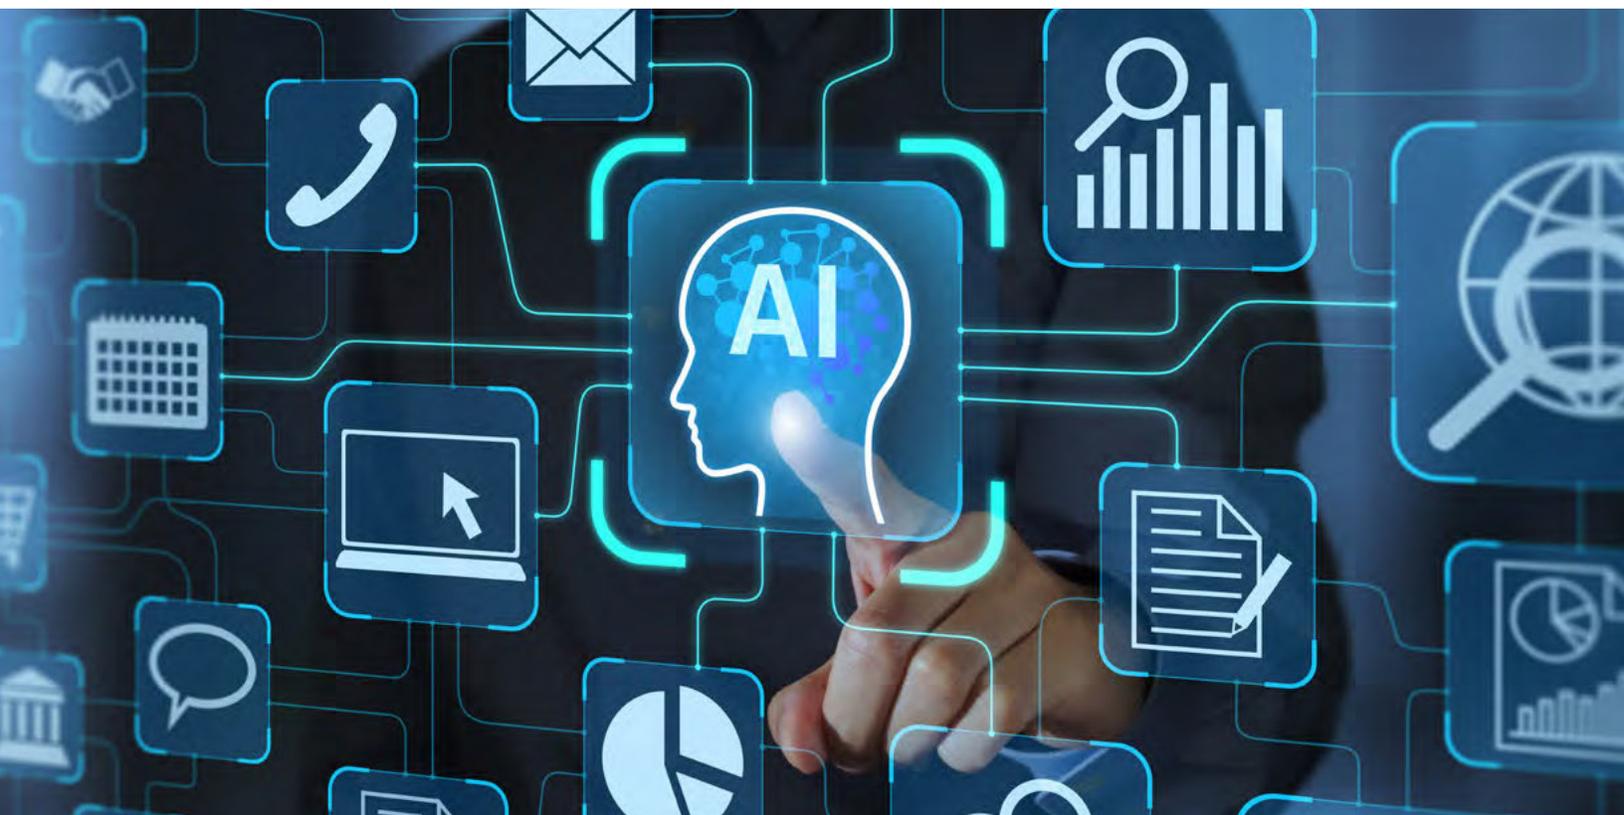
AWS Lambda implements the Executive Order crawler functionality, providing serverless compute for regulatory data maintenance. The Lambda function periodically queries official government sources to identify newly published Executive Orders, automatically updating the knowledge base without manual intervention.

Frontend Build:

The user interface built with Next.js 16 and React 19 provides users with document upload capabilities, analysis tracking, and compliance results display. TypeScript 5.9 ensures type-safe development, reducing runtime errors and improving code maintainability. The web-based interface enables access from any device with a browser.

4. Implementation Output Analysis

The ATARC Agentic AI Vendor Proposal Compliance Management use case demonstrates a sophisticated implementation that leverages AI agents to automatically analyze vendor proposals against current FAR and Executive Orders. The system successfully correlates synthetic proposals with the latest regulatory requirements, identifying both compliance achievements and critical discrepancies that could impact procurement decisions ([Appendix B: UI Implementation and Sample Outputs](#)).



The demonstration analyzed PatriotCloud Hyperscaler Inc.'s \$8.5 million proposal, revealing a nuanced compliance landscape. The system performed dual-analysis from two complementary perspectives:

Analysis 1 focused on tactical FAR compliance with immediate remediation steps, while **Analysis 2** provided strategic regulatory framework context with formal risk classifications. Both analyses converged on critical findings, demonstrating the system's reliability and thoroughness.

4.1 Key Findings

4.1.1 Critical Compliance Gap

The most significant finding across both analyses was the **missing Certificate of Current Cost or Pricing Data (FAR 15.406-2)**. This represents a non-waivable requirement that creates an insurmountable barrier to contract award. The system correctly identified this as creating "immediate rejection risk" despite the proposal's strong technical merit.

4.1.2 Secondary Deficiencies

The analysis uncovered additional compliance gaps including:

- Inadequate evaluation criteria compliance and documentation
- Insufficient requirement traceability
- Missing contract type justification
- Gaps in socioeconomic program requirements (small business utilization)

4.1.3 Positive Compliance Elements

Notably, the system also identified areas of strong compliance, including excellent alignment with Executive Orders on AI policy and robust security framework adherence.

4.1.4 Analytical Strengths

The implementation demonstrates several powerful capabilities:

- a. **Convergent Validation:** Both analytical perspectives independently identified the same critical deficiency, providing confidence in the assessment
- b. **Risk Stratification:** Clear classification of issues by severity (critical vs. secondary)
- c. **Actionable Recommendations:** Prioritized remediation steps with specific timelines (0-10 days for critical actions)
- d. **Regulatory Citation:** Specific FAR references enabling rapid verification and correction

4.1.5 Potential Gaps and Limitations

Areas of Improvement	Gap	Mitigation
Temporal Challenge	The system's ability to stay current with rapidly evolving regulations (FAR updates, new Executive Orders) requires continuous knowledge base updates. There's no indication of the last regulatory update date in the output.	Implement automated regulatory monitoring with version control, displaying the "as-of" date for all regulatory references. Establish quarterly validation cycles against official FAR sources.
Context-Specific Interpretation	FAR compliance often involves nuanced interpretation based on contract type, dollar thresholds, and agency-specific requirements. The analysis may not capture all contextual exceptions or waivers.	Incorporate agency-specific FAR supplements and deviation databases. Add confidence scores to compliance assessments indicating interpretation certainty levels.
False Positive/Negative Risk	Automated analysis may flag compliant sections as deficient (false positive) or miss subtle non-compliance (false negative), particularly in complex technical proposals.	Implement human-in-the-loop validation for high-risk findings. Maintain audit trails showing the specific proposal text triggering each compliance flag for expert review.

5. Forward Looking

5.1 Recommendations for Future Enhancement

5.1.1 Real-Time Regulatory Feeds & Temporal Compliance

While the current PoC implementation employs automated crawling through AWS Lambda functions to maintain currency, future iterations should include real-time regulatory feed evaluations that enable the system to reference most up-to-date compliance requirements. This enhancement would involve direct connections to government sources (i.e., SAM.gov, the Federal Register, etc.) and be event-driven, replacing the current batch-based setup.

5.1.2 Confidence Scoring & Uncertainty Management

Future enhancements will incorporate confidence scoring mechanisms that quantify the certainty of compliance evaluations. Confidence scores might include the clarity of a proposal's regulatory language, proposal completeness, consistency of findings between agents, and more. Low confidence scores would trigger human-in-the-loop validation, ensuring that higher-risk decisions receive appropriate human intervention.

5.1.3 Interactive Clarification & Adaptive Analysis

Current compliance assessments operate in a single-pass mode where agents analyze proposals without the ability to request additional information or clarification from vendors. Future implementations should include interactive clarification capabilities that enable the system to request additional context to address proposal ambiguities, as well as refine assessments based on vendor responses before generating final compliance assessment outputs.

5.1.4 Responsible AI with Amazon Bedrock Guardrails

Amazon Bedrock's native safeguarding feature provides a safeguard policy feature, ensuring that AI-generated analyses maintain appropriate boundaries and avoid unethical and irresponsible outputs. The implementation configures guardrails to filter sensitive procurement information and prevent the generation of inappropriate compliance advice that could mislead users. Future iterations should also consider content filtering prevents the system from processing or generating information that could compromise procurement integrity.

5.2 Expanded Use Cases Across Federal Agencies

The architectural patterns and capabilities demonstrated in this compliance assessment solution have potential to translate to other US federal compliance use cases – specifically for agencies looking to improve operational efficiency and reduce risk and waste:

Grant Proposal Evaluations presents benefits for agencies such as the National Science Foundation (NSF), National Institutes of Health (NIH), and Department of Energy (DOE). These agencies process thousands of grant applications annually, each requiring evaluation against program-specific guidelines, regulatory requirements, and merit-based criteria.

Contract Modification Analysis can assist teams needing to evaluate proposed contract changes against original terms, applicable regulations, and agency policies. Federal contracts frequently require modifications due to scope changes and schedule adjustments, requiring modifications that must be analyzed for regulatory compliance.

Regulatory Impact Assessment & Policy Analysis would enable agencies to evaluate how new regulations or policy changes affect existing contracts and procurements. In the event of new regulatory requirements, agencies must determine the applicability of those changes to their contract portfolios and identify necessary compliance actions.

Small Business Compliance Support would enable more equitable federal procurement for small businesses facing regulatory complexity hurdles. A compliance assessment solution could provide small businesses with automated proposal reviews and notify users of compliance gaps and remediation suggestions before submission. This capability would “level the playing field” between large contractors with dedicated compliance teams and small businesses with limited resources.

5.3 Proof of Concept Pattern for Federal AI Adoption

This vendor proposal compliance solution offers a reusable architectural pattern that federal agencies can adapt for their unique use cases and workflows. Positioning this as a Well-Architected Framework (WAF) pattern for “AI-Powered Regulatory Compliance Assessment” can accelerate federal AI adoption by providing design principles and implementation guidance that address common concerns shared by federal customers about AI (i.e., reliability, transparency, etc.).

The core architectural components (document ingestion, vector-based knowledge retrieval, multi-agent orchestration, and generation of compliance status and recommendations) can serve as the building blocks for agencies’ respective compliance domains. Agencies could customize knowledge bases for their specific regulatory frameworks and adapt agent behaviors for domain-specific requirements. Any proof of concept provided would include guidance on knowledge base assembly, agent specialization, data privacy and residency, security compliance, FedRAMP/ GovCloud requirements, and audit trailing/monitoring among other considerations.

By establishing proven architectural patterns and providing reusable components, this initiative creates a foundation for accelerated AI adoption across federal agencies. This solution positions ATARC and AWS as leaders in federal AI innovation that serves the public interest while achieving the “art-of-the-possible” in AI technologies.



6. Conclusion

The ATARC Agentic AI Vendor Proposal Compliance Management use case demonstrates significant value in automating complex regulatory analysis. The system successfully identified critical compliance gaps that could prevent contract award, while also recognizing areas of strong performance. The convergent dual-analysis approach provides validation confidence, and the prioritized, actionable recommendations enable rapid remediation.

However, to maximize reliability and adoption, the implementation should address temporal challenges, incorporate confidence scoring, and add human validation checkpoints for high-stakes decisions. With these enhancements, this agentic AI solution could significantly reduce procurement risk while accelerating the proposal evaluation process across federal agencies.

The Vendor Proposal Compliance Management use case represents an advancement in applying agentic AI to federal procurement challenges. By combining sophisticated AI capabilities with robust security frameworks like MAESTRO, the initiative demonstrates how autonomous AI systems can improve government operations while maintaining necessary safeguards.

The solution's containerized, cloud-native architecture ensures scalability and adaptability, while its focus on transparency and explainability builds trust among stakeholders. As the inaugural use case for the ATARC Agentic AI Lab, this initiative establishes a foundation for broader adoption of agentic AI in government operations, with particular emphasis on supporting small businesses and improving procurement efficiency.

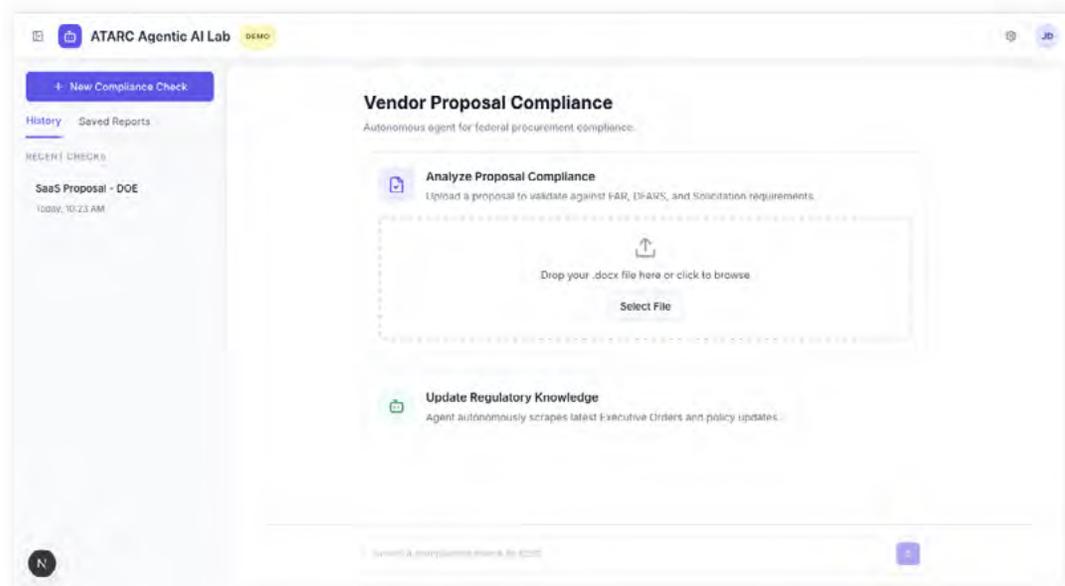
7. Appendices

Appendix A: User Interface

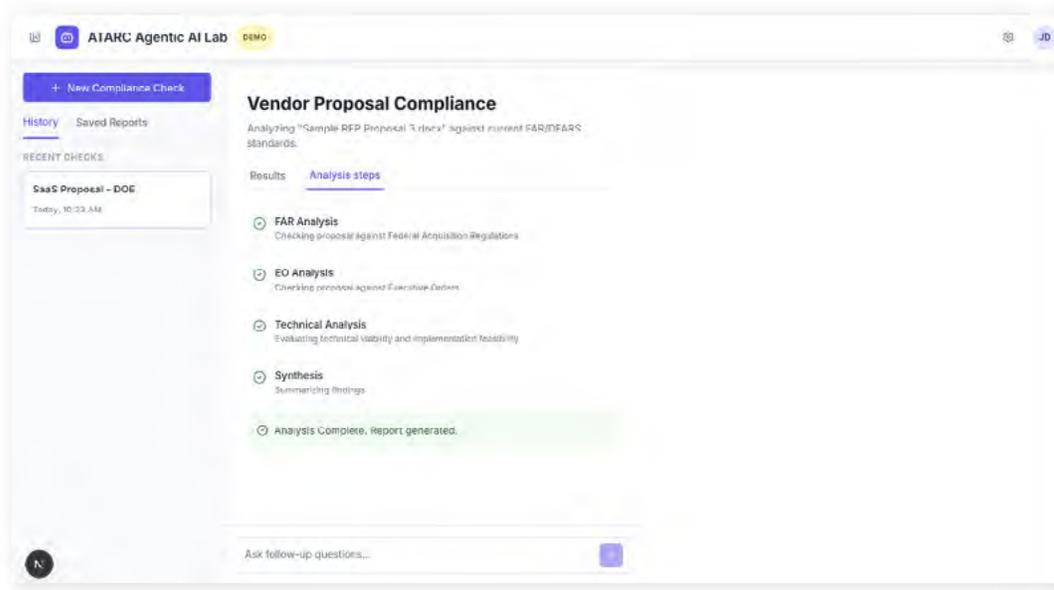
The user interface is built on Next.js 16 and React 19 for FAR compliance validation. The interface architecture prioritizes accessibility, performance optimization, and intuitive user workflows through a component-based design system. Key features include intelligent document upload with real-time validation, progressive analysis tracking with visual status indicators, and comprehensive compliance results display with severity-based categorization and AI confidence scoring. The development philosophy emphasizes type-safe development through TypeScript 5.9, automated testing across unit, integration, and end-to-end scenarios, and performance metrics targeting sub-2.5 second load times. This modular, accessible architecture demonstrates how agentic AI interfaces can balance sophisticated backend AI capabilities with user-centered design principles, providing government users with an efficient tool for ensuring regulatory compliance while maintaining the highest standards of code quality, security through Content Security Policy implementation, and extensibility for future enhancements including real-time updates, batch processing, and collaborative review capabilities.

Appendix B: Implemented Interface & Sample Output

Appendix B.1: Document Upload



Appendix B.2: Runtime Analyses



Appendix B.3: Sample Agent Outputs

EO Compliance

Federal Acquisition Regulation (FAR) Compliance Assessment Report

Solicitation Number: 70B1CR25R0000001 - Data and AI Modernization Initiatives

Submitted By: TechSecure Solutions, Inc.

Executive Summary

This compliance assessment evaluates TechSecure Solutions' proposal against applicable Federal Acquisition Regulation (FAR) requirements. While the proposal demonstrates strong technical capabilities and addresses many regulatory requirements, several critical compliance deficiencies have been identified that require immediate attention to ensure full regulatory adherence.

Priority Compliance Issues Identified

Priority 1: FAR 52.219-14 Small Business Subcontracting Limitations on Subcontracting

Issue: The proposal states TechSecure will perform "at least 50% of the work with our in-house team, per FAR 52.219-14," but fails to provide the required detailed analysis and documentation.

Regulatory Requirement: FAR 52.219-14 requires small business prime contractors to perform a minimum percentage of work themselves and provide specific documentation of how this requirement will be met.

Recommended Action: Provide a detailed work breakdown structure showing exactly which tasks will be performed by TechSecure personnel versus subcontractors, with dollar amounts and percentages clearly delineated.

Priority 2: FAR 52.219-9 Small Business Subcontracting Plan

Issue: While the proposal mentions "25% to small businesses" and references compliance with FAR 52.219-9, it lacks the comprehensive subcontracting plan required by regulation.

Regulatory Requirement: FAR 52.219-9 mandates a detailed subcontracting plan including specific subcontractor information, performance monitoring procedures, and compliance reporting mechanisms.

Recommended Action: Develop and include a complete subcontracting plan with named subcontractors, their certifications, assigned work scopes, and monitoring procedures.

Summary

Executive Summary - CBP Data Modernization Proposal Analysis

Executive Overview

TechSecure Solutions' proposal for CBP's Data and AI Modernization initiative demonstrates strong technical capabilities and competitive pricing but faces **critical compliance deficiencies** that pose significant award risk. While the technical approach is sound with moderate implementation risk, regulatory gaps—particularly in subcontracting documentation—require immediate remediation before contract award consideration.

Key Findings from Each Source

FAR Compliance Analysis - CRITICAL CONCERNS

Missing Subcontracting Plans: Both FAR 52.219-14 (limitations) and FAR 52.219-9 (small business plan) lack required detailed documentation

Security Documentation Gaps: CMMC Level 2 claims need comprehensive implementation details per FAR 52.204-21

Cost Justification Weaknesses: Insufficient detail for economic price adjustments and allowable cost determinations

Positive Elements: Proper business certifications (8(a), SDB, WOSB, HUBZone) and security framework references

Technical Viability Assessment - MODERATE RISK

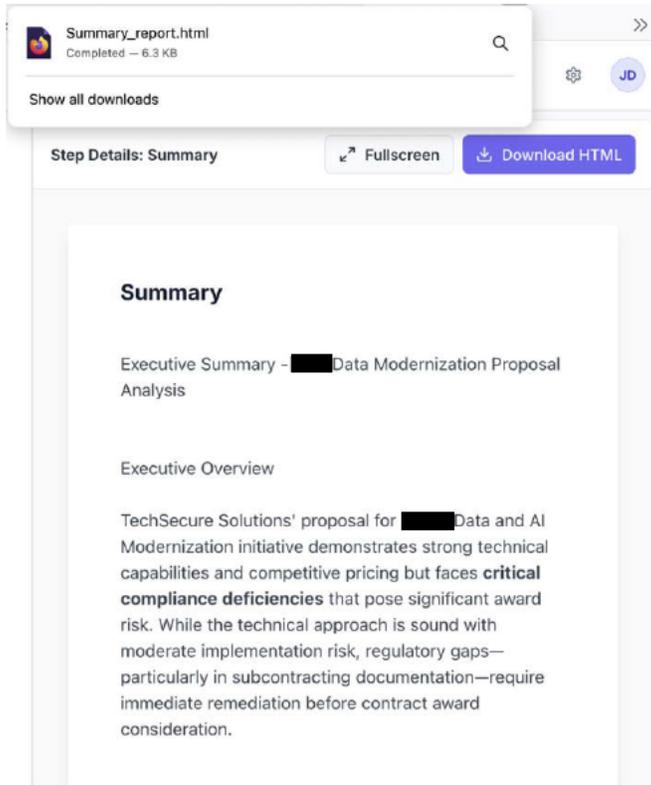
Sound Architecture: Well-designed data mesh using proven technologies (Apache Kafka, Delta Lake)

Appropriate Team: 15 FTEs with adequate experience levels and security clearances

Timeline Concerns: Aggressive 3-month MVP delivery may be unrealistic for system complexity

Integration Risks: Legacy system integration and 1M+ device scalability need validation

Appendix B.4: HTML Download & Sample Output



Executive Summary - [REDACTED] Data Modernization Proposal Analysis

Executive Overview

TechSecure Solutions' proposal for [REDACTED] Data and AI Modernization initiative demonstrates strong technical capabilities and competitive pricing but faces **critical compliance deficiencies** that pose significant award risk. While the technical approach is sound with moderate implementation risk, regulatory gaps—particularly in subcontracting documentation—require immediate remediation before contract award consideration.

Key Findings from Each Source

FAR Compliance Analysis - CRITICAL CONCERNS

- **Missing Subcontracting Plans:** Both FAR 52.219-14 (limitations) and FAR 52.219-9 (small business plan) lack required detailed documentation
- **Security Documentation Gaps:** CMMC Level 2 claims need comprehensive implementation details per FAR 52.204-21
- **Cost Justification Weaknesses:** Insufficient detail for economic price adjustments and allowable cost determinations
- **Positive Elements:** Proper business certifications (8(a), SDB, WOSB, HUBZone) and security framework references

Technical Viability Assessment - MODERATE RISK

- **Sound Architecture:** Well-designed data mesh using proven technologies (Apache Kafka, Delta Lake)
- **Appropriate Team:** 15 FTEs with adequate experience levels and security clearances
- **Timeline Concerns:** Aggressive 3-month MVP delivery may be unrealistic for system complexity
- **Integration Risks:** Legacy system integration and 1M+ device scalability need validation

Cross-Analysis Insights

Converging Concerns