

**Technical Response Document
Submitted To**



RECEIVED
2026 JUN -4 AM 9:51
WV DEPARTMENT OF ENVIRONMENTAL PROTECTION

**State of West Virginia
Department of Environmental Protection**

Agentic E-Permitting Platform powered by IBM Watsonx Orchestrate

Submitted by:

TrustNet Technologies, LLC

Prime Contractor

In strategic partnership with IBM Watsonx and Streebo Inc.

Sunday, 31 May 2026



Table of Contents

1. Company Profile and History..... 1

2. Executive Summary.....2

3. References.....4

4. Solution Architecture7

5. Project Goals, Objectives, and Proposed Approach..... 10

6. Mandatory Project Requirements..... 18

7. Implementation Timeline23

8. SaaS Compliance and Data Protection..... 24

1. Company Profile and History

1.1 TrustNet Technologies, LLC – Prime Contractor

TrustNet Technologies, LLC will serve as the Prime Contractor for this engagement. Headquartered in Atlanta, Georgia, TrustNet is a Service-Disabled Veteran-Owned Small Business (SDVOSB), SBA-verified and registered in SAM.gov (UEI TUJYSESM9BN3; CAGE 13RE5; NAICS 541519). As a strategic IBM Watsonx partner, the firm specializes in enterprise-grade Agentic AI platform engineering, AI governance, and managed technology services for government clients, and is led by a veteran team with decades of combined federal and state mission-support experience. As Prime, TrustNet will own the contract with WVDEP, lead program governance and delivery, and manage the integrated team — including subcontractor Streebo Inc. — as the single point of accountability for schedule, quality, security, and risk across the Agentic E-Permitting Platform.

TrustNet's core competencies map directly to the services in this proposal. The firm designs and deploys autonomous and semi-autonomous AI agents with mandatory human-in-the-loop controls, and orchestrates multi-agent workflows using LangGraph, the Model Context Protocol (MCP), and Agent-to-Agent (A2A) frameworks — the same patterns that underpin the UIC intake, completeness, technical compliance, GIS/AoR, and drafting agents described in Section 4. Its delivery platform is cloud-hosted and multi-tenant with container orchestration and cost monitoring, supports hybrid and on-premises deployment with no vendor lock-in, and is built on a FedRAMP-aligned security posture with role-based access controls, SSO, high-availability design, and fully attributable agent interactions traceable to individual agent identities — meeting the security, deployment, and observability requirements set out in Section 6.

TrustNet also brings the governance, observability, and rapid-deployment tooling that the WVDEP engagement requires. Integrated AI governance spans the full agent lifecycle — intake, risk assessment, deployment, monitoring, and retirement — aligned to the NIST AI Risk Management Framework, with real-time trajectory tracing, accuracy and bias evaluation, prompt-injection defense, PII-exposure prevention, and continuous vulnerability scanning of models and LLM endpoints. Pre-built audit documentation, standardized model factsheets, and automated regulatory reporting reduce audit burden for the agency, while a no-code agent builder and a pro-code Agent Development Kit (ADK) accelerate delivery from months to days. Together, these capabilities position TrustNet to deliver the Agentic E-Permitting Platform end-to-end while preserving the regulatory defensibility, security posture, and human oversight WVDEP requires for UIC Class I and Class VI permitting.

1.2 Streebo Inc.- Subcontractor

Streebo Inc., headquartered in Houston, Texas, is a global digital transformation and AI solutions provider, specializing in the design and implementation of enterprise-grade conversational AI platforms. Established in 2008, they have evolved into a full-service provider of AI-first digital engagement solutions, serving higher education institutions, government agencies, utilities, and enterprise organizations.

For this engagement, Streebo will serve as a key subcontractor, bringing proven expertise in deploying scalable, AI-driven support ecosystems with more than 10 years of experience. This includes the design of intelligent conversational interfaces/AI chatbots/voicebots, secure integration with institutional systems, and the implementation of structured human escalation workflows to ensure seamless user experience across all interaction scenarios. They operate with a global delivery model supported by a team of over 200 certified professionals across North America, Asia, and Europe. The team comprises solution architects, AI/ML engineers, conversational designers, and system integration specialists, enabling end-to-end delivery of complex AI implementations.

Additionally, the organization is a certified Minority Business Enterprise (MBE), aligning with the RFP's diversity and inclusion objectives and contributing toward MWBE participation goals.

2. Executive Summary

2.1 Understanding of WVDEP's Objectives

The West Virginia Department of Environmental Protection, Division of Water and Waste Management, administers primary enforcement authority (primacy) for UIC Class I and Class VI injection wells under the Safe Drinking Water Act. The Agency is seeking to modernize its permit review process through AI-enabled automation while preserving robust human oversight, regulatory defensibility, and data security.

We have carefully reviewed the RFP, including all specifications, and fully understand the following core objectives:

- Automate the complex, multi-stage review of Class I and Class VI permit applications using Agentic AI
- Enforce mandatory Human-in-the-Loop (HITL) decision gates at six critical regulatory milestones, AI may never autonomously issue a permit
- Integrate advanced GIS analysis for Area of Review (AoR) calculations and subsurface risk identification
- Process and extract data from large volumes of PDFs, engineering drawings, and CAD files using OCR and Computer Vision
- Generate AI-drafted permits, Notices of Deficiency, public notice documents, and response to comments
- Operate independently from legacy systems (ERIS, ESS, AppEnhancer) while providing human-mediated data synchronization
- Deliver on FedRAMP Moderate, NIST 800-53, SOC 2 Type II, Section 508, and all applicable state/federal compliance requirements
- Ensure complete data sovereignty within the continental United States with WVDEP retaining full data ownership

2.2 Proposed Solution Overview

We propose the WVDEP UIC AI Permit Intelligence Platform ("UIC-AIPI") a fully cloud-hosted, SaaS-delivered, IBM Watsonx powered agentic AI system that transforms the UIC permitting workflow from a manual, document-heavy process into an intelligent, auditable, and human-supervised regulatory review engine.

The platform consists of five integrated layers:

- I. **Agentic AI Core:** IBM Watsonx Orchestrate powers autonomous, multi-step permit review workflows for Class I and Class VI applications, routing between sub-agents for completeness checks, technical analysis, GIS evaluation, and document drafting.
- II. **Knowledge & RAG Engine:** IBM Watsonx orchestrate provides GraphRAG-style knowledge graph retrieval, grounding all AI responses in UIC regulations, WVDEP permit history, and geological databases. Eliminates hallucination through citation-enforced responses.
- III. **Document Processing:** AI-powered OCR, NLP, and Computer Vision agents extract data from PDFs, engineering schematics, CAD/DWG drawings, and well logs. Blueprint Vision Agents are trained on UIC-specific construction diagrams.
- IV. **GIS Intelligence:** Integrated GIS hub ingests Shapefile, GeoJSON, KML, CAD, DWG formats. Automates AoR calculation, risk identification, and cross-database spatial correlation.
- V. **HITL Workflow Interface:** Standalone, browser-based interface where all AI-to-human interactions occur. Enforces six mandatory decision gates. Fully independent from legacy ERIS, ESS, and AppEnhancer systems.

2.3 Key Value Proposition

Our solution is built around the T.R.U.S.T. Framework, ensuring WVDEP receives a system designed for regulatory confidence:

- **Transparent AI (T):** Every AI decision includes citations to source documents, confidence scores, and full Chain-of-Thought reasoning ensuring legal defensibility.
- **Regulatory Grounded (R):** All AI analysis is grounded in 40 CFR 146, 47 CSR 13/64, EPA Class I/VI standards, and WVDEP-specific templates through Watson's knowledge graph.
- **Uncompromised Security (U):** IBM Cloud for Government (FedRAMP Moderate), AES-256 encryption, TLS 1.3, RBAC, MFA, SAML 2.0 SSO, annual SOC 2 Type II and penetration testing.
- **Supervised Always (S):** Six mandatory HITL gates enforce human control at every critical milestone. AI cannot approve, issue, or deny a permit autonomously under any circumstance.
- **Total Data Sovereignty (T):** All WVDEP data remains within IBM Cloud US data centers. WVDEP retains full ownership.

2.4 Key Differentiators

The proposed solution is distinguished by the **S.A.F.E. Framework**, supported by enterprise-grade architecture:

- **Secure (S):** Built on IBM's enterprise cloud ecosystem, the solution leverages the full security capabilities of IBM, including identity management, encryption at rest and in transit, role-based access controls, audit logging, conditional access policies, and continuous threat monitoring. Data residency remains within U.S. boundaries, ensuring compliance with institutional requirements.
- **Accurate (A):** The AI agent is trained using structured institutional knowledge, historical ticket data, and validated documentation. Through supervised tuning, retrieval grounding, and iterative validation cycles, the solution is engineered to achieve up to 99% accuracy in defined service domains. Accuracy is continuously monitored using precision, recall, and containment analytics.
- **Fast Rollout (F):** Leveraging pre-configured AI frameworks and scalable cloud infrastructure, implementation timelines are measured in days and weeks rather than months. Modular onboarding allows phased deployment by channel (chat → mobile → voice) or by service category, minimizing operational disruption.
- **Ease of Use (E):** The conversational interface is intuitive for end users and requires minimal training. For administrators and human agents, guided workflows and AI-assisted interfaces reduce onboarding time and accelerate productivity. Structured enablement programs ensure rapid adoption across both technical and non-technical stakeholders.

3. References

We have extensive experience delivering AI agents, digital assistants, and end-to-end web solutions for higher education institutions, non-profit organizations, municipalities, utilities, and other public-sector entities. Our engagements span the design, development, integration, and deployment of intelligent conversational platforms, coupled with web-based content management and digital engagement solutions.

The following selected projects are particularly relevant, demonstrating our capability to handle complex technical requirements, large-scale knowledge management, and public-facing service objectives in institutional environments similar to this RFP.

In accordance with our confidentiality agreements and NDAs, we are unable to disclose certain client identities and contact details within this proposal. However, we will be pleased to furnish full contact information and arrange direct references **at the appropriate juncture** upon request or during contract negotiations.

1. AI powered virtual assistant for a County in Florida

- **Client:** County in Florida
- **Project Date:** August 2025

- **Project Location:** Florida, USA

Description of Services Provided: Delivered an AI-powered citizen service assistant that provides accurate, real-time responses using a large corpus of legal, policy, and archival documents. The solution enables seamless citizen interactions across web, mobile, voice, SMS, and email. Key capabilities include intelligent natural-language conversations, automated FAQs, ticket creation with document attachments, AI cognitive search, transactional workflows, omnichannel reach, backend system integration, and intelligent handoff to human agents. The assistant supports both English and Spanish, improving accessibility and inclusivity while reducing service turnaround time and operational costs.

Outcomes:

- Faster resolution of citizen service requests.
- Reduced workload on live agents, allowing focus on complex cases.
- Estimated 25–30% reduction in support workflow costs.
- Inclusive, bilingual support for English and Spanish speakers.
- 24x7 availability, reducing reliance on traditional office hours.
- Scalable deployment across additional county departments.
- Enhanced citizen satisfaction and trust through intelligent, streamlined engagement.

2. AI-Powered Cognitive Assistant for a Major Municipality in Ontario, Canada

- **Client:** A Leading Municipality in Ontario, Canada
- **Project Date:** November 2024
- **Project Location:** Ontario, Canada

Description of Services Provided:

The municipality sought to enhance its digital engagement capabilities by adding an AI-powered cognitive assistant. The objective was to enable citizens to access municipal services through intuitive, natural language interactions while reducing reliance on traditional support channels.

We designed and implemented an AI powered chatbot, integrating Cognitive Search and advanced Natural Language Processing (NLP) capabilities to deliver accurate, context-aware responses across multiple service domains.

The solution was designed to:

- Seamlessly crawl, index, and synchronize content from the Drupal Content Management System (CMS)
- Ensure real-time alignment between website updates and chatbot responses
- Enable consistent and accurate knowledge retrieval directly from structured and unstructured web content

The chatbot successfully ingested and indexed all Drupal-managed web content, demonstrating robust capability in handling large-scale, CMS-driven knowledge environments.

Outcome:

The implemented solution delivered measurable improvements across key performance indicators:

- **Enhanced Citizen Experience:** Delivered intuitive, human-like interactions, improving accessibility and engagement across diverse user groups
- **Increased Digital Adoption:** Enabled higher utilization of self-service channels, reducing dependency on live agents
- **Operational Efficiency Gains:** Automated high-volume queries, leading to faster response times and optimized resource utilization
- **Improved Content Utilization:** Leveraged Drupal-based web content effectively, ensuring accurate and up-to-date responses

3. AI Assistant for Public Utility Provider in the USA

- **Client:** Public Utility Provider in the USA
- **Project Date:** July 2024
- **Project Location:** United States

Description of Services Provided:

The utility provider, serving over 2 million customers, needed to reduce redundant support tickets and provide multilingual support. We deployed a smart customer facing AI assistant capable of handling high traffic volumes and performing real-time Spanish-to-English translation for escalated queries. The solution was integrated with a third-party live agent platform.

We also deployed an employee assistant AI chatbot which is a cloud-based built on Microsoft Azure infrastructure, using Azure Cognitive Services and Large Language Models for natural language chat and voice support. It integrates with SharePoint for knowledge access, uses Azure Active Directory for security, and is accessible through an employee portal and desktop app. It provides users secure access to LLMs, web search, file search and internal information search capability. All services and data remain secured within the organization’s Azure cloud environment, supporting enterprise compliance and scalability

Outcome:

- Reduced live agent load and improved response accuracy
- Multilingual coverage with seamless escalation
- Enhanced customer satisfaction, especially for Spanish-speaking communities
- Enhances Employee productivity
- Efficient handling of thousands of queries across multiple channels

Metrics:

- ~2 million customers served
- Designed for high concurrency and multilingual automation

4. Solution Architecture

4.1 Solution Architecture Overview

The UIC AI Permit Intelligence Platform will be delivered as a cloud-native SaaS solution hosted on IBM Cloud for Government (US data centers). The architecture is organized into four principal tiers that maintain clean separation between the AI processing environment, the legacy WVDEP systems, and external data sources.

Tier	Components
Tier 1: Presentation	E-Permitting Applicant Portal (web-based, Section 508 compliant), HITL Reviewer Workflow Interface, Admin Dashboard & Monitoring Console – all browser-based, SSO-enabled via SAML 2.0.
Tier 2: AI Processing	IBM Watsonx Orchestrate (agentic workflow engine), IBM watsonx.ai (LLM/NLP/Computer Vision), Graph RAG knowledge retrieval, GIS Analysis Engine, Document OCR & Blueprint Vision Agents, Regulatory Compliance Engine.
Tier 3: Data & Storage	IBM Cloud Object Storage (FedRAMP), PostgreSQL on IBM Cloud (structured permit data), Vector database for RAG embeddings, GIS data store, Audit log repository, Document management integration.
Tier 4: Integration & Legacy Bridge	Secure outbound data connectors to WVGES, EPA SDWIS, WV Oil and Gas records, WVDEP databases. Human-mediated export packages for manual legacy system (ERIS/ESS/AppEnhancer) updates. No direct AI-to-legacy system interaction.

4.2 Agentic Workflow Engine

IBM Watsonx Orchestrate serves as the central intelligence layer, orchestrating the entire permit review lifecycle through a network of specialized sub-agents. Each sub-agent is a purpose-configured AI module with defined scope, input/output contracts, and mandatory HITL suspension points.

- **Intake & Routing Agent:** Receives applications from the E-Permitting portal, assigns tracking numbers, creates electronic case files, and classifies applications as Class I or Class VI.
- **Completeness Review Agent:** Validates all required application sections against regulatory checklists using NLP and Computer Vision. Triggers Pre-NoD HITL gate.
- **Technical Compliance Agent:** Evaluates injection zone properties, confining zone integrity, USDW proximity, and well construction design against 40 CFR 146 and 47 CSR standards.
- **GIS & AoR Agent:** Performs spatial analysis, calculates Area of Review, identifies artificial penetrations, faults, and USDWs within the AoR. Triggers Technical Screening HITL gate.

- **Document Drafting Agent:** Generates formatted draft permits, Notices of Deficiency, public notice documents, and Response to Comments documents using WVDEP templates.
- **Regulatory Watchdog Agent:** Monitors Federal Register, EPA website, and WV Legislature feeds 24/7 for regulatory changes and alerts agency staff.

4.3 GraphRAG Knowledge Engine

The solution is powered with the Retrieval-Augmented Generation (RAG) layer that grounds all AI responses in verified regulatory and geological information. Watson Discovery's knowledge graph maps complex relationships between regulations, geological data, permit history, and applicant-submitted documentation.

- **Regulatory Knowledge Base:** 40 CFR Parts 144-148, 47 CSR 13 (Class I), 47 CSR 64 (Class VI), EPA UIC Program Guidance, WVDEP permit history and precedents
- **Geological Database Integration:** WVGES well logs, subsurface formation data, USDW aquifer maps, fault and fracture databases
- **Document Indexing:** All submitted application documents (PDF, CAD, GIS) are indexed and made retrievable with passage-level citations
- **Hallucination Mitigation:** Every AI response must cite its source document and passage. Responses without a grounded citation are blocked by the compliance guardrails.
- **Confidence Scoring:** All AI-generated determinations include a confidence score (0-100%) and the specific regulatory provisions supporting the determination.

4.4 Data Flow Architecture

The data flow is designed to maintain strict separation between the AI platform and WVDEP's legacy systems, in full compliance with the RFP:

Inbound Flow (Application Intake):

- Applicants submit permit applications through the WVDEP E-Permitting web portal (hosted at dep.wv.gov/WWE/PERMIT/UIC/ per RFP)
- The AI platform ingests documents via secure e-form API. A unique tracking number is assigned and an electronic case file is created immediately.
- Submitted documents (PDF, GIS, CAD, engineering drawings) are processed through the OCR, Computer Vision, and document extraction pipeline.
- The agentic workflow is initiated, progressing through defined processing stages with HITL gates at each regulatory milestone.

Outbound Flow (Permit Processing):

- At each HITL gate, the AI platform generates a reviewer task in the HITL Workflow Interface with full AI reasoning, confidence scores, and source citations.

- The human reviewer completes the review action in the HITL interface and approves progression.
- The system generates an export package (PDF documents, status updates, data summaries in open formats).
- The reviewer manually updates the corresponding records in ERIS, ESS, or AppEnhancer. The AI system does not write to, read from, or interface with legacy systems directly.

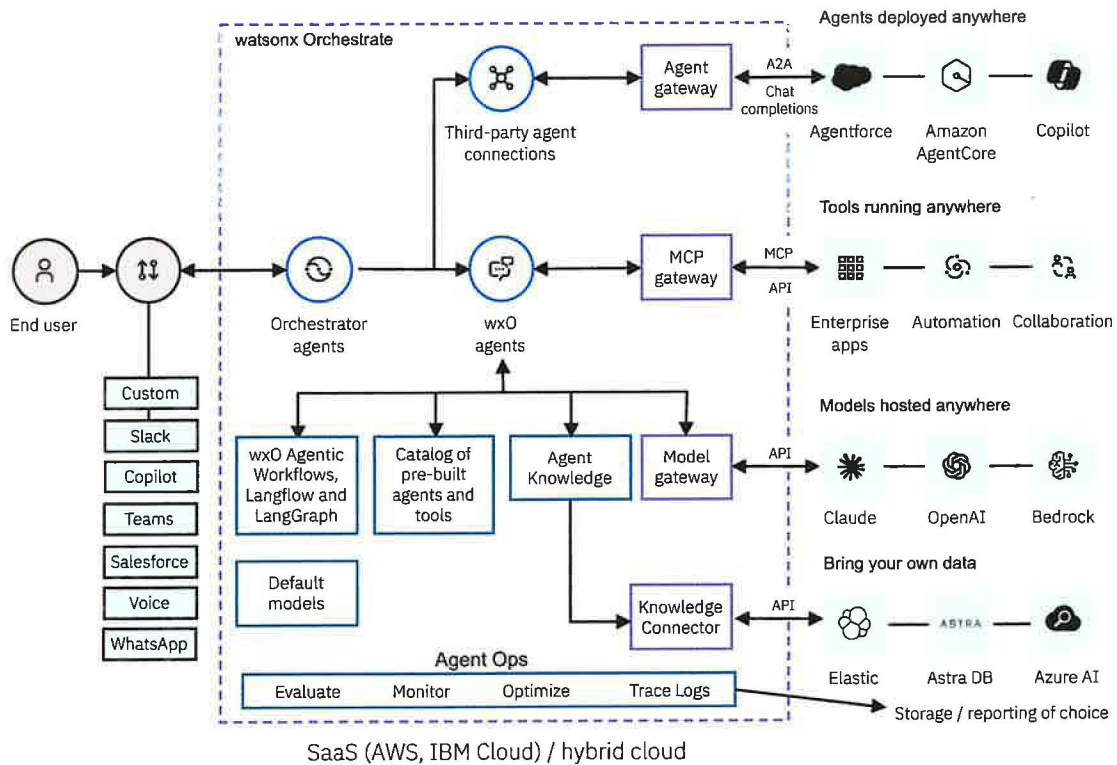


Figure 1 – Orchestration of the solution with Watsonx Orchestrate

5. Project Goals, Objectives, and Proposed Approach

This section describes our approach to each functional requirement. Our multiagent solution will meet the stated mandatory requirements.

5.1 General Automation and Dashboard Integration

5.1.1 System Automation and Integration

The UIC-AIPI platform integrates with WVDEP's operational environment through a carefully designed integration architecture that respects the integrity of legacy systems while enabling full AI-powered automation:

- **WVGES Integration:** Scheduled synchronization pulls geological survey data, well records, and formation data for cross-referencing against permit applications.
- **EPA SDWIS Integration:** Direct API calls to retrieve USDW identification data and public water system records for the Area of Review analysis.
- **WV Office of Oil and Gas:** Automated data retrieval for well records within the AoR, enabling conflict detection and artificial penetration identification.
- **WVDEP Internal Databases:** Secure outbound queries for existing permit records, inspection history, and compliance data.
- **AppEnhancer (Document Management):** Inbound document retrieval via secure API for historical documents. The AI does not modify AppEnhancer records; export packages are provided for human-mediated updates.

5.1.2 Dashboard Development

A centralized, web-based HTML dashboard is provided for all automation monitoring and control. The dashboard is hosted on the WVDEP domain or as a subdomain and offers:

- Secure role-based login via WVDEP SSO (SAML 2.0) with MFA enforcement
- Real-time job status board showing all active, pending, and completed permit review workflows
- Full automation logs with timestamped action entries, AI decisions, confidence scores, and HITL gate statuses
- HITL notification center: reviewers receive in-dashboard alerts and email notifications when human action is required
- Regulatory Watchdog alerts panel for regulatory change notifications
- Analytics: application volumes, processing times, AI accuracy metrics, and reviewer workload reporting
- Admin panel for authorized staff to adjust automation rules, update regulatory knowledge base, and configure HITL thresholds

5.2 UIC Class I and Class VI Agentic AI Processing

5.2.1 Digital Intake Specialist Functionality

The Intake & Routing Agent functions as a Digital Intake Specialist, executing the multi-stage regulatory review process autonomously from submission receipt through permit draft generation, with mandatory HITL gates at each regulatory milestone.

- Class I / Class VI Routing: Upon application receipt, the agent analyses the permit application type field, injection fluid characterization, and CO2 stream indicators to determine Class I vs. Class VI classification and routes to the appropriate agentic sub-workflow.
- Sub-Workflow Orchestration: Beyond classification, the system dynamically triggers granular sub-routines for AoR analysis, financial responsibility verification, corrective action planning (Class VI), and injection well construction standards based on specific data detected.
- EPA Form 7520-6 Verification: Computer Vision agents verify the presence, completeness, and correct execution of EPA Form 7520-6, including signature verification by Responsible Corporate Officers using NLP signature block detection.

5.2.2 Administrative Completeness Review

The Completeness Review Agent validates all required sections of the permit application against a configurable regulatory checklist:

- Class I Checklist: Project plans, site maps, AoR calculations, EPA Form 7520-6, facility location/legal description, operator identification, well construction details (casing/cementing), injection zone identification, confining zone characterization, proposed injection rates/volumes/pressures, injection fluid characterization.
- Class VI Additional Requirements: Geologic characterization, Testing and Monitoring Plan (TMP), Injection Well Plugging Plan, PISC Plan, Emergency and Remedial Response Plan, Financial Responsibility documentation. Computer Vision verifies signatures by Responsible Corporate Officers and mandatory certification statements.
- Missing Item Identification: Any missing or insufficient items are flagged with specific regulatory citations (e.g., 40 CFR 146.22(a)(3)) and queued for the Pre-NoD HITL review gate.

5.2.3 Technical Compliance Review

The Technical Compliance Agent performs deep technical evaluation of all submitted technical data:

- Injection Zone Evaluation: Depth, thickness, lithology, porosity, permeability, and formation pressure extracted from well logs and geological reports and compared against regulatory standards.
- Confining Zone Integrity: Analysis of confining zone characterization data to verify its ability to prevent fluid migration into USDWs.
- USDW Identification: Cross-referenced against EPA SDWIS and WVGES data to identify all USDWs within the AoR.

- Well Construction Design Review: Casing materials, cementing programs, and mechanical integrity test data evaluated against 40 CFR 146 Subpart B/G/H requirements.
- Operational Parameters: Maximum injection pressure, rate, and volume evaluated against formation fracture pressure and USDW protection standards.

5.2.4 RAG and Source Grounding

The proposed solution employs IBM Watson's knowledge graph as the RAG engine. This approach is specifically superior to standard vector-similarity RAG because it maps structured relationships between regulatory provisions, geological concepts, permit conditions, and submission documents not merely surface-level text similarity.

- Knowledge Graph Structure: Nodes represent regulatory provisions, geological formations, permit types, and compliance requirements. Edges represent relationships such as 'applies_to', 'requires', 'supersedes', and 'cross_references'.
- Multi-Hop Retrieval: When evaluating a Class VI AoR, the system can traverse: CO2 injection → geologic sequestration → AoR requirements (40 CFR 146.84) → computational modeling → USDW protection → related WVGES formation data – in a single retrieval pass.
- Staleness Prevention: The knowledge base is updated through the Regulatory Watchdog Agent upon detection of Federal Register amendments. Outdated regulatory provisions are flagged and retired from active retrieval. The Agent alerts WVDEP staff to any regulatory changes requiring manual knowledge base updates.
- Irrelevance Filtering: Watson's relevance scoring ensures that only passages with a confidence score above a configurable threshold are returned for AI reasoning.

5.2.5 Hallucination Mitigation

We implement a multi-layered hallucination mitigation framework specific to regulatory AI systems:

- Citation Enforcement: Every AI-generated determination must include a reference to a specific document, page, and paragraph from the submitted application or the regulatory knowledge base. Responses without grounded citations are blocked.
- Confidence Thresholds: AI determinations below a configurable confidence threshold (default: 80%) are automatically escalated to human review rather than progressed through the workflow.
- Constitutional AI Guardrails: IBM watsonx.ai's guardrails prevent the AI from generating regulatory determinations outside its trained scope or inventing regulatory provisions not present in the knowledge base.
- Cross-Agent Validation: Critical determinations (e.g., AoR delineation, USDW presence) are independently evaluated by two AI agents and reconciled before proceeding.
- Human Override: At every HITL gate, the human reviewer has full visibility into the AI's reasoning chain and can override, correct, or reject any AI determination.

5.2.6 Citations and Explainability

All AI-generated interactions provide full source attribution and explainability:

- Document Citations: Every fact extracted from the application is linked to the specific page, section, and paragraph in the source PDF with a direct reference link.
- Regulatory Citations: Every compliance determination cites the specific regulatory provision (e.g., '40 CFR 146.84(a)(2)') that the determination is based on.
- Confidence Scores: All automated decisions include a confidence percentage (0-100%) displayed to human reviewers.
- Chain-of-Thought Logging: Full step-by-step reasoning is logged for every decision and accessible in the observability/analytics dashboard.

5.3 Geospatial Analysis and GIS

5.3.1 Automated Risk Assessment and AoR

The GIS & AoR Agent provides automated spatial analysis capabilities:

- Class I AoR (Fixed Radius): Automated 1/4-mile radius delineation from the well bore, with overlay of all artificial penetrations, faults, fractures, USDWs, sensitive receptors, and surface water bodies.
- Class VI AoR (Computational Modeling): Pressure front modeling using reservoir simulation parameters from the application. Validates applicant-submitted computational model inputs, boundary conditions, and assumptions against WVGES formation data.
- Artificial Penetration Identification: Automated query of WV Office of Oil and Gas records, mine records, and borehole databases for all penetrations within the calculated AoR.
- Fault and Fracture Analysis: Cross-reference against WVGES fault and fracture databases to identify potential fluid migration pathways.
- USDW Mapping: Automated identification and mapping of all USDWs within the AoR using EPA SDWIS and WVGES aquifer data.
- Topology Validation: Automated topology checks on all submitted GIS data to verify geometric integrity and compliance with state coordinate standards (NAD83/UTM Zone 17N).

5.3.2 Data Integration and Correlation

- Property & Plats: Overlay of surface ownership tracts, digital courthouse plat records, and pore space ownership documentation to verify legal rights.
- State & Federal Database Synchronization: Scheduled and on-demand synchronization with WVGES, EPA SDWIS, WV Office of Oil and Gas, and WVDEP databases.
- Conflict Detection: Automated flagging of overlapping subsurface mineral rights, conflicting injection permits, and proximity violations.
- GIS Format Support: Shapefile, GeoJSON, KML, CAD, DWG ingested and auto-projected to NAD83/UTM Zone 17N.

- 2D/3D Visualization: Subsurface stratigraphy visualization in both 2D cross-section and 3D spatial views, accessible to human reviewers through the HITL dashboard.

5.4 Document Processing and AI Drafting

5.4.1 Blueprint Vision Agents

Beyond standard OCR, the solution is specifically trained on UIC Class I and Class VI engineering schematics:

- Well Construction Diagrams: Automatic extraction of casing depths, cement thickness, tubing packer placement, and casing shoe depth from engineering drawings.
- CAD/DWG Processing: Technical drawings in CAD and DWG format are parsed to extract dimensioned data points critical for construction standards compliance.
- Cross-Section Diagrams: Subsurface cross-section diagrams are analyzed to extract formation identifications, depth markers, and confining zone boundaries.
- Pressure-Depth Charts: Automated extraction of injection pressure profiles and formation pressure data from graphical chart submissions.

5.4.2 AI Draft Generation

The solution will generate permit documents using WVDEP's standard templates:

- Draft Permits: Facility Information, Well Construction Requirements, Operating Requirements (maximum injection pressure/rate/volume), Monitoring and Reporting Requirements, Plugging and Abandonment Requirements, General Conditions – all populated from AI-extracted application data.
- Class VI-Specific Permit Conditions: Approved AoR with reevaluation schedule, corrective action requirements, CO2 stream specifications, TMP conditions, Emergency Response requirements, and PISC requirements.
- Template Fidelity: All draft documents would strictly follow WVDEP formatting and numbering conventions. Human reviewers can modify AI-drafted text before final issuance.

5.4.3 Completeness, NoD, Public Notice, and Response to Comments

- Completeness Determination: AI cross-checks submitted e-form data and attachments against Class I and Class VI regulatory checklists. Missing items trigger the Pre-NoD HITL gate.
- Notice of Deficiency Generation: AI drafts NoD email/letter specifying exact missing or insufficient items with regulatory citations. Document is reviewed and approved by human staff before transmission.
- Public Notice Documents: Platform generates public notice documents compliant with 40 CFR 124, including fact sheets, notification recipient lists, and comment period deadlines (30 days Class I, 45 days Class VI).

- Response to Comments: Public comments are ingested, categorized, and AI drafts responses using a WVDEP-specific regulatory response library. Substantive technical comments are flagged for detailed human review. Final Response to Comments document is compiled for public release.

5.5 Workflow Integration

5.5.1 Secure Submission Handling

- Applications are received through the WVDEP E-Permitting web portal (dep.wv.gov/WWE/PERMIT/UIC/). A unique application tracking number will be automatically generated and an electronic case file will be created upon submission.
- All submitted documents are encrypted at rest (AES-256), and stored in IBM Cloud Object Storage within US boundaries.

5.5.2 Agency Logs and Deep Observability

- Comprehensive audit logs are maintained for all AI actions, decisions, and workflow transitions. Logs are stored for the period meeting WVDEP's retention policy.
- All log entries include timestamps, agent identifiers, input data references, output decisions, confidence scores, and regulatory citations.
- An adjacent agency folder is maintained in the HITL interface where WVDEP staff can review AI comments and actions for quality control at any time.

5.5.3 AgentOps Observability

The platform integrates IBM watsonx's native AgentOps capabilities for full Chain-of-Thought recording:

- Every decision is recorded with its complete reasoning chain
- Reasoning chains are accessible to human reviewers in the HITL interface and downloadable as audit documents for legal proceedings.
- All Chain-of-Thought logs are immutably archived and are retained per WVDEP's retention policy.

5.5.4 Mandatory HITL Decision Gates

The six mandatory HITL gates are non-negotiable control points enforced by the workflow engine. The AI system physically cannot progress past any HITL gate without recorded human action:

#	HITL Gate	AI Provides to Reviewer	Required Human Action
1	Pre-NoD Review	Completeness checklist with missing items, citations, and confidence scores	Verify missing items; approve or modify draft NoD before transmission
2	Admin Compliance Approval	Fee verification status, ownership validation results, public notice requirements identified	Supervisor approves compliance or issues deficiency; initiates public notice process
3	AoR & Risk Validation	AI-generated AoR map with identified penetrations, faults, USDWs, risk ratings, and confidence scores	Human review of AoR and fault analysis; determines expedited vs. standard review track
4	Technical Analysis Approval	Full technical review findings, corrective action needs, permit condition recommendations	Senior technical reviewer validates findings; approves permit conditions; may request additional analysis
5	Draft Permit Approval	Complete AI-drafted permit document with all conditions, citations, and data sources	Human review and modification of AI-drafted permit before document is locked for public comment
6	Final Decision	Complete application record, public comment summary, Response to Comments draft	Human administrator executes final Issue or Deny command. AI cannot autonomously issue a permit.

5.6 AI Token Usage and Cost Management

- **Token Cost Model:** Costs are priced per 1 Million Tokens (MTok)
- **Cost Predictability:** IBM watsonx.ai provides real-time token consumption dashboards. Budget alerts are configurable at 70%, 85%, and 95% of the contracted allocation. All token usage is billed in arrears based on actual consumption.
- **Token Optimization:** We employ prompt compression, document chunking, caching of repeated regulatory passages, and model routing (smaller models for completeness checks; larger models for technical analysis) to minimize token consumption without sacrificing output quality.
- **Estimated Token Usage:** See Attachment B for detailed estimated token consumption per processing activity.
- **Usage Transparency:** Monthly token usage reports are provided in WVDEP's preferred format, broken down by processing activity, application type, and agent.
- **Model Flexibility:** IBM's model-agnostic watsonx.ai architecture allows model substitution as newer, more efficient foundation models become available – without architectural changes to the platform.
- **Cost Guarantees:** Any model efficiency improvements that reduce per-application token costs are passed to the client.

5.7 HITL Workflow Interface and Legacy System

The standalone HITL Workflow Interface is a self-contained, browser-based application that serves as the complete workspace for human-AI interaction, entirely independent of ERIS, ESS, and AppEnhancer:

- **Standalone Interface:** The HITL interface is a separate web application accessible via WVDEP SSO. It does not share database connections, APIs, or infrastructure with legacy systems.
- **AI Request Routing:** When an AI processing milestone generates a HITL request, the watsonx Orchestrate workflow engine creates a task record in the HITL interface database, attaches all AI-generated documents and reasoning, and notifies the assigned reviewer via dashboard alert and email. No legacy system is involved.
- **Human-Mediated Synchronization:** After completing a HITL review task, the reviewer manually inputs relevant information and status updates into ERIS, ESS, or AppEnhancer. The AI platform generates a structured export package (pre-populated data summaries, formatted documents) to minimize manual data entry burden.
- **Data Flow Separation:** The AI platform never reads from or writes to legacy systems. All data flows through the human reviewer as the intermediary, maintaining the integrity and audit compliance of both systems.

6. Mandatory Project Requirements

6.1 Data Integration and Regulatory Compliance

Mandatory Requirement	Compliance	How We Meet It
Format Support: XML, CSV, PDF, GIS (Shapefile, GeoJSON, KML, CAD, DWG), well logs, permit forms	FULLY MET	IBM Watsonx orchestrate and document processing pipeline supports all listed formats natively. CAD/DWG processed by Blueprint Vision Agents.
Compliance Engine & Regulatory Watchdog Agent	FULLY MET	Configurable compliance engine encodes Class I and Class VI requirements. Dedicated Watchdog Agent monitors Federal Register, EPA, and WV Legislature feeds 24/7.
External System Integration for data cross-checking	FULLY MET	API integrations with WVGES, EPA SDWIS, WV Office of Oil and Gas, and WVDEP databases. Synchronization schedules configurable per source.

6.2 Security and Deployment

Mandatory Requirement	Compliance	How We Meet It / IBM Evidence
TLS 1.3 in transit, AES-256 at rest	FULLY MET	IBM Cloud for Government enforces TLS 1.3 for all data in transit and AES-256 encryption for all data at rest. IBM FIPS 140-2 validated cryptographic modules.
RBAC + MFA	FULLY MET	RBAC aligned with WVDEP organizational structure (applicant, reviewer, senior reviewer, admin, manager roles). IBM Cloud Identity and Access Management (IAM) enforces MFA for all users.
PII and CBI Protection	FULLY MET	PII/CBI fields are identified at ingestion using IBM Watson NLP entity recognition. Stored in designated encrypted fields with access restricted to authorized roles. CBI data not included in public-facing outputs.
FedRAMP Moderate – AWS GovCloud or Azure Government equivalent	FULLY MET	IBM Cloud for Government holds FedRAMP Moderate Authorization (Authorization to Operate). All data residency within the continental United States. IBM Cloud ATO package available to WVDEP.
SSO via SAML 2.0 or OpenID Connect	FULLY MET	IBM Cloud IAM supports SAML 2.0 and OpenID Connect federation. Integration with WVDEP's existing identity provider is a standard configuration task.
Annual penetration testing; patching within 30 days for critical vulnerabilities	FULLY MET	Annual third-party penetration testing results shared with WVDEP. IBM Cloud continuous vulnerability scanning. Critical vulnerability patching within 30 days contractually committed. IBM PSIRT process for zero-day disclosures.

6.3 Support and Maintenance

Warranty Support

- We are committed to providing a comprehensive 30-days warranty period after "Project Go-Live". During this period, we will promptly address any application issues that may arise, ensuring a stable and reliable solution. Please note that the warranty period starts from the date when the solution is signed off for "Go-Live".

Annual Support & Maintenance

- After the 30-day warranty period, Vendor will transition to comprehensive application support and maintenance services, offering flexible engagement for security and functional enhancements, ongoing bot training, and regular maintenance.
- Vendor will provide L2 and L3 support services post-go-live using an offshore support model for maintaining technical and functional capabilities.
- For the benefit of the client, Level 2, and Level 3 application support along with their severity, what classifies them to be in a particular severity level and the estimated RPO & RTO for the same have been described below:

Service Type	Processes	Maintenance & Support Activities
L2 Support	Service Help Desk	Ticket Management activities by responding to customer bug reporting
		Monitored email support 8*5 for customers to email incidents/issues to support engineering team
		Web based support incident tracking by allowing customer contact to access our secure support portal and log incidents/problems
		Tracking and escalation of all activities/processes
		Chatbot Solution monitoring - Proactive monitoring of the applications
	Incident Management	User incidents and queries
Incident management for chatbot solution support		
Detect, classify, record and provide initial support of incidents		
L3 Support	Problem Management	Technical Support and advice
		Providing incident support reports to customer team on a set frequency to keep the team updated on the progress of the support
		Analysis of chatbot solution data for chalking out preventive actionable plans
		Support chatbot solution installations
		Proactive maintenance, root cause analysis and permanent fixes
		Preventive maintenance – chatbot solution tuning, code restructuring, and performance tuning
		Trouble shooting
	Notification of planned and unplanned outage	
Change Management	Request	Analyse the impact of all change and enhancement requests across all systems, and accordingly authorize or deny change requests based on feasibility and criticality

		Estimate and ensure that CR's are properly scoped, estimated, and high-level project plans for each item are created for customer
		Research peers and competition and suggest suitable enhancements and modifications to the existing functionality of the digital properties.

Support Scope - Severity Definition:

Severity Levels	Support Category	Criteria
1	Urgent	<ul style="list-style-type: none"> Critical Business Impact the Incident has caused a complete and immediate work stoppage affecting a critical function such that a primary business process or a broad group of users such as an entire department, floor, branch, line of business, or external customer. No Workaround available.
2	High	<ul style="list-style-type: none"> A business process is affected in such a way that business functions are severely degraded, multiple users are impacted, a key customer is affected, or a Critical Function is operating a significantly reduced capacity or functionality. A Workaround may be available; however, the Workaround is not easily sustainable.
3	Medium	<ul style="list-style-type: none"> A business process is affected in such a way that certain functions are unavailable to end users, or a system and/or service is degraded. A Workaround may be available.
4	Low	<ul style="list-style-type: none"> Minimal Business Impact: an Incident that has little impacts on normal business processes and can be handled on a scheduled basis. A Workaround is available or there is minimal negative impact on a user's ability to perform their normal daily work.

Support Scope - Response & Resolution Times based on Severity of issue:

Priority Level	Estimated Response time
Priority 1 (Urgent)	2 business hours for issues classified as Urgent priority.
Priority 2 (High)	4 business hours for issues classified as High priority.

Priority 3 (Medium)	8 business hours for issues classified as medium priority.
Priority 4 (Low)	24 business hours for issues classified as Low priority.

6.4 Training post Go-live

The proposed solution includes a comprehensive and structured training program designed to ensure effective adoption, operational independence, and long-term sustainability of the chatbot platform. Training activities will be planned, developed, and delivered in alignment with the client's operational structure, covering both functional usage and administrative management of the solution.

Training will encompass platform navigation, content management, conversation design, workflow configuration, reporting and analytics, and basic troubleshooting. Special emphasis will be placed on enabling non-technical users to confidently manage knowledge bases, review conversations, and optimize chatbot performance without dependency on technical teams.

To ensure scalability and knowledge continuity, a "Train-the-Trainer" model will be adopted. Key personnel will be identified as master trainers and provided with in-depth, hands-on training sessions. These trainers will then be equipped to extend training to broader user groups across. Supporting materials such as user manuals, quick reference guides, and recorded sessions will be provided to reinforce learning and enable ongoing reference.

The effectiveness of training will be continuously evaluated through feedback mechanisms and practical assessments, ensuring that client's teams are fully prepared to manage, maintain, and enhance the solution independently while maximizing its value across institutional functions.

6.5 Regulatory Compliance

Compliance Requirement	Status	Evidence / Approach
FedRAMP Moderate Authorization	FULLY MET	IBM Cloud for Government holds current FedRAMP Moderate ATO. IBM's FedRAMP package (System Security Plan, POA&M, continuous monitoring reports) available for WVDEP's ATO review.
NIST 800-53 Compliance	FULLY MET	IBM Cloud for Government's FedRAMP authorization is based on NIST 800-53 Rev. 5 controls. All applicable control families (AC, AU, IA, SC, SI, etc.) are satisfied.
Auditability – 5 Year Retention	FULLY MET	All AI actions, decisions, and Chain-of-Thought reasoning are immutably logged in IBM Cloud Object Storage with required retention. Logs are exportable in open formats.
Section 508 Accessibility	FULLY MET	All user-facing components (E-Permitting portal, HITL interface, dashboard) are developed to WCAG 2.1 AA / Section 508 standards. Accessibility testing conducted pre-launch.
AI Governance	FULLY MET	We document model training data, methodologies, and bias testing. Annual bias testing reports will be provided to WVDEP. Model drift detection alerts are active. Human override available at all decision points.

Annual SOC 2 Type II Audit

FULLY MET

IBM Cloud maintains SOC 2 Type II certification. Annual audit reports provided to WVDEP. Right-to-audit clause for WVDEP or designated third-party security assessments included in contract.

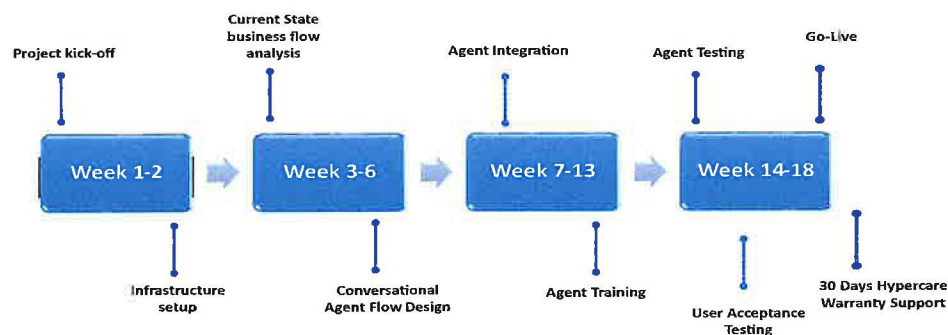
6.6 Data Ownership and Exit Strategy

- All WVDEP data, application materials, permit documents, GIS data, AI-generated content, and audit logs are and remain the sole property of the State of West Virginia.
- WVDEP data is never used to train AI models that benefit other customers. Data is not transmitted outside the United States. All AI model training uses WVDEP-authorized data only with explicit written consent.
- Compliance with WV Technology Office AI policies (<https://technology.wv.gov/policy-governance/artificial-intelligence>) is maintained throughout the contract, including the prohibition of FCC-banned AI models and software.
- Complete data export within the decided days of contract termination in PDF, CSV, JSON, XML, and standard document formats at no additional cost.
- Post-termination transition assistance to support migration to a replacement system.
- All WVDEP data securely deleted from all vendors' systems within the expected days of confirmed data transfer, with a written NIST-approved certificate of destruction if applicable.

7. Implementation Timeline

In alignment with the RFP, we propose an initial implementation timeline of approximately 14–18 weeks to deliver a Minimum Viable Solution (MVS) based on a prioritized set of high-impact use cases. This approach ensures early value realization while establishing a scalable foundation for continuous expansion. It can be followed by **subsequent enhancement phases** where additional use cases, integrations, and capabilities are incrementally introduced.

Estimated Timeline



Given timeline is tentative and subject to finalization at the project kick-off

Figure 2 – Estimated timeline

We will work closely with the client's team throughout the project to ensure smooth collaboration, quick decision-making, and mitigation of potential delays due to dependencies.

While we have designed this plan to deliver within a focused 14-18 week's window, it is based on our prior experience.

We remain flexible and adaptive, and will recalibrate our timeline as needed based on:

- Final scope confirmation during Discovery
- Availability of access to third-party systems and knowledgebase
- Any changes to go-live dependencies

8. SaaS Compliance and Data Protection

FedRAMP and NIST Compliance via IBM Cloud for Government

A key differentiator of the proposed solution is the use of IBM Cloud for Government as the hosting environment, which carries a current FedRAMP Moderate Authorization to Operate (ATO). This means WVDEP can leverage IBM's existing FedRAMP package (including System Security Plan, Control Implementation Summary, and continuous monitoring artifacts) to support its own Agency ATO, significantly reducing the time and cost of security authorization.

- FedRAMP Moderate Impact Level authorization covers the 325+ security controls defined in NIST SP 800-53 Revision 5
- IBM's FedRAMP package is available through the FedRAMP Marketplace and can be provided directly to WVDEP upon contract award
- IBM maintains a dedicated FedRAMP Program Management Office (PMO) for continuous monitoring, annual assessments, and control maintenance
- All WVDEP data is stored exclusively in IBM Cloud US-East and US-South data centers – never outside the continental United States

AI Governance

The proposed solution meets the AI governance requirements:

- Model Training Documentation: AI Factsheets document the training data sources, data lineage, preprocessing steps, and model architecture for every AI model deployed in the UIC-AIPI platform
- Bias Testing: Annual bias analysis reports are generated using automated fairness monitoring, covering disparate impact testing across applicant demographics and geographic regions
- Model Drift Detection: Continuous monitoring alerts WVDEP staff when AI model accuracy or prediction confidence drifts beyond configurable thresholds, triggering retraining review
- Human Override: Every AI determination in the system has a documented human override path – no AI decision is irreversible without human action

Prohibited AI and Software Compliance

We confirm full compliance with the West Virginia Office of Technology AI Policy. The proposed solution exclusively uses IBM watsonx-family models and services, none of which appear on the FCC-prohibited software list. No data is transmitted to prohibited jurisdictions or third-party AI services outside the approved platform.