Title Page

CRFP 0313 DEP2600000003

Workflow-Based Agentic AI, Automation, and E-Permitting System

for the West Virginia Department of Environmental Protection (WVDEP)

Underground Injection Control (UIC) Class I & Class VI Permitting

RFP Subject:	Workflow-Based Agentic AI, Automation, and E-Permitting System
Solicitation No.:	CRFP 0313 DEP2600000003
Vendor Name:	Codist Cafe
Business Address:	1214 Forrest Drive, Hurricane, WV 25526
Telephone:	304 376 1121
Fax:	N/A
Contact Person:	Blaise Liu
Email:	blaise.liu@codistcafe.com
FEIN:	086263828

Authorized Signature:	Blaise Liu
Print Name:	Blaise Liu
Title:	Principal Software Architect
Date:	June 10, 2026

RECEIVED
2026 JUN 10 PM 12: 56
WV PURCHASING
DIVISION

Technical Proposal

CRFP 0313 DEP2600000003

Workflow-Based Agentic AI, Automation, and E-Permitting System for WVDEP

Underground Injection Control (UIC) Class I & Class VI Permitting

---

Vendor: Codist Cafe
Solicitation No.: CRFP 0313 DEP2600000003
Bid Opening Date: June 10, 2026
Bid Opening Time: 1:30 PM ET

Table of Contents

Executive Summary

Codist Cafe proposes an **agentic AI e-permitting system** for WVDEP's UIC Class I and Class VI permit review, built on Microsoft Azure Government (FedRAMP Moderate). The system replaces the current manual, PDF-based workflow with a multi-agent AI platform that automates document processing, regulatory compliance checking, geological risk assessment, and permit drafting — while

enforcing six mandatory human-in-the-loop decision gates that ensure every final determination remains a human decision.

Key Differentiators:

- **Working reference platform today**, not a PowerPoint promise — multi-agent AI engine, 6-gate HITL state machine, document generation, 86 passing tests
- **GraphRAG-native architecture** for multi-hop regulatory reasoning across regulations, geology, and permit history — cannot be done with standard vector RAG
- **Azure Government FedRAMP Moderate** — rides Microsoft's P-ATO; no separate FedRAMP certification required
- **Zero-markup token pricing** — WVDEP pays Azure's actual rate ($7.00/MTOK), locked for contract life
- **Lean, senior-led team** — Ph.D. AI research + 15yr Oracle architecture + 12yr regulatory compliance experience
- **Live screenshots** of the working POC dashboard, HITL interface, and chain-of-thought audit trail are provided in Appendix B.

Projected Benefits to WVDEP:

- Reduce permit review cycle time through AI-assisted document processing and automated completeness checking
- Improve consistency and defensibility through chain-of-thought audit trails and mandatory regulatory citations on every AI finding
- 24/7 regulatory watchdog monitoring Federal Register, EPA, and WV Legislature for UIC-relevant changes
- Standalone HITL interface with no disruption to legacy ERIS/ESS/AppEnhancer systems

Total Bid: $3,205,025.00 (see separately sealed Cost Proposal envelope for complete Attachment A pricing, invoicing schedule, and token economics). This figure encompasses all implementation, hosting, support, token consumption, and optional renewal elements described in this proposal.

---

3. Approach & Methodology to Goals/Objectives (§ 4.3.1)

3.1 General Automation & Dashboard Integration (§ 4.3.2.1)

3.1.1 System Automation/Integration (§ 4.3.2.1.1)

Codist Cafe proposes an **AI-Enabled Robotic Process Automation (RPA) platform** built on Microsoft Azure Government (FedRAMP Moderate) that integrates with existing WVDEP operational software applications through a secure, API-first architecture.

Integration Architecture:

The proposed platform integrates with WVDEP systems through a structured API gateway layer:

System	Integration Method	Data Direction	Purpose
WV OSSP (One-Stop Shop Permitting)	REST API	Bidirectional	Reporting, payment status, fee tracking
OneLogin	SAML 2.0 / OpenID Connect	Inbound	External user authentication
ActiveDirectory	SAML 2.0 / SCIM	Inbound	Internal user authentication
WVDEP GIS (ESRI ArcGIS)	ArcGIS REST API / GeoServices	Read	Spatial data, well records, geological layers
WVGES	Public API / Scheduled Sync	Read	Geological survey data
EPA SDWIS	EPA API / Data Export	Read	USDW identification data
WV Office of Oil & Gas	Public API	Read	Well records, permit history

Automation Engine: Our AI-RPA platform provides a central automation orchestration layer that:

- Executes **scheduled and event-driven automation jobs** for document processing, validation, and notification
- Maintains a **centralized job queue** with priority-based scheduling
- Provides **comprehensive logging** of all automation activities with full chain-of-thought capture
- Supports **reusable automation components** that can be configured without custom code

3.1.2 Dashboard Development (§ 4.3.2.1.2)

We propose a **web-based HTML dashboard** accessible via modern browsers with secure login/password authentication integrated with WVDEP's identity systems.

Dashboard Features:

Feature	Description
Real-Time Job Monitor	Live status of all automation jobs, agent activities, and workflows
Case Summary View	At-a-glance dashboard of all active permit applications with status, reviewer, and stage
HITL Task Queue	Pending human review items sorted by priority and deadline
Audit Log Viewer	Searchable, filterable log of all AI actions with chain-of-thought
Token Usage Monitor	Real-time and historical token consumption by application, workflow, and agent
Watchdog Alerts	Regulatory change notifications with impact analysis
System Health	Platform availability, SLA compliance, incident history
Reporting Console	Configurable reports for management review, token accounting, throughput metrics

Access Control:

- **Role-Based Access Control (RBAC)** aligned with WVDEP organizational structure
- **Multi-Factor Authentication (MFA)** enforced for all users
- **Granular permissions** — read-only, reviewer, supervisor, administrator
- **Session management** with automatic timeout after inactivity

See Appendix B, Figure B-1 for a screenshot of the working POC dashboard.

User Roles:

Role	Permissions
Permit Reviewer	View cases, complete HITL tasks, view audit logs
Senior Reviewer	All reviewer permissions + approve technical findings, override AI decisions
Supervisor	All senior reviewer permissions + approve permits, manage workflow
Administrator	Full system access, user management, configuration
Read-Only	View dashboard and case status only

3.2 UIC Class I & Class VI Agentic AI Processing (§ 4.3.2.2)

3.2.1 Digital Intake Specialist (§ 4.3.2.2.1)

Our **Digital Intake Specialist** is the primary AI agent orchestrating end-to-end permit review. It is built as a multi-agent system where specialized sub-agents handle specific regulatory tasks under the orchestration of a supervisor agent.

System Architecture:

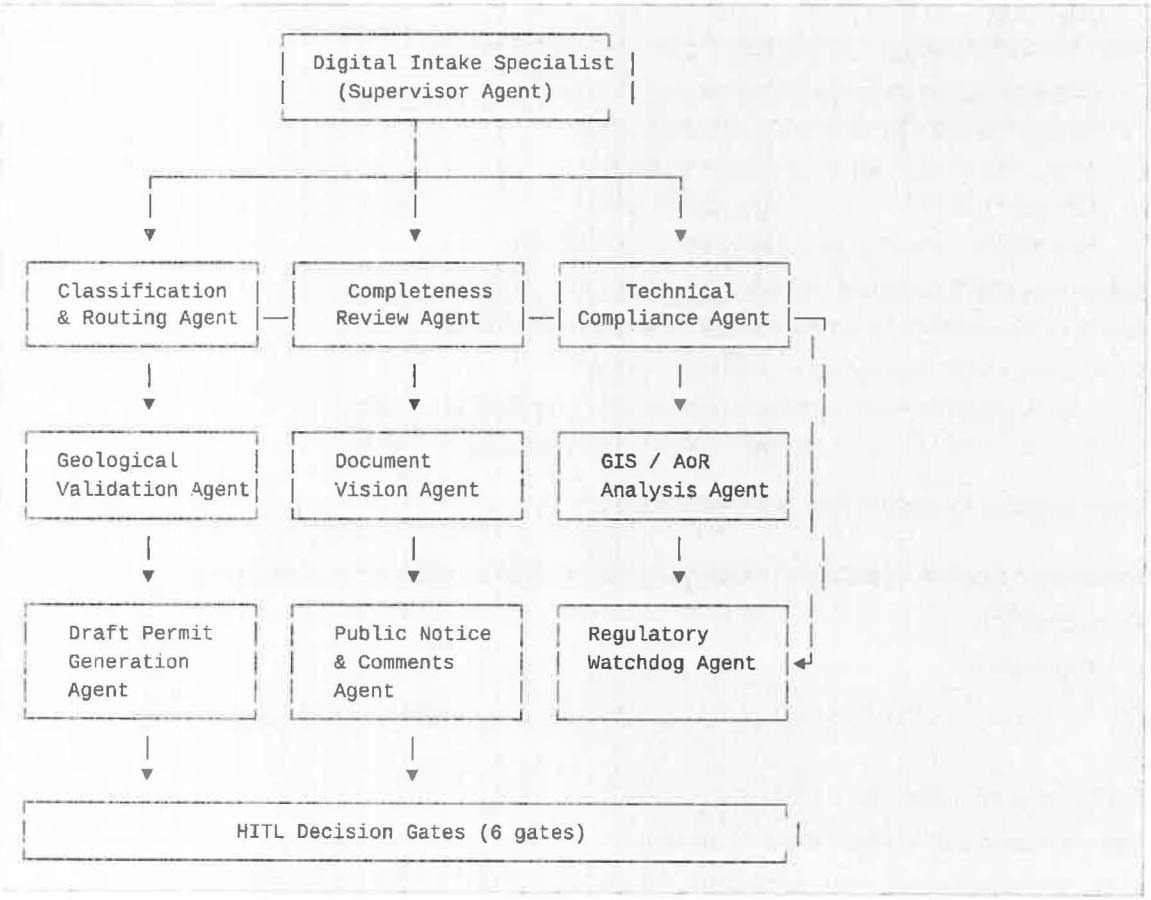

3.2.1.1 Agentic Routing & Sub-Workflow Orchestration (§ 4.3.2.2.1.1)

The **Classification & Routing Agent** performs initial triage on every incoming application:

Classification Logic:

1. **Application Type Detection**: Parses the application cover page and EPA Form 7520-6 to determine:
   - Well class (I vs. VI) based on injection fluid type, depth, and stated well classification
   - Well subtype (hazardous waste, industrial non-hazardous, municipal wastewater for Class I; CCS project for Class VI)
   - New application vs. modification vs. renewal
2. **Dynamic Sub-Workflow Orchestration**: Based on detected characteristics, the agent dynamically builds a workflow:
   - **Class I Workflow**: Administrative completeness → Technical review (well construction, injection zone, confining zone, USDW) → AoR (fixed-radius 1/4 mile) → Corrective action assessment → Draft permit → Public notice → Final decision
   - **Class VI Workflow**: Administrative completeness → Site characterization → AoR (computational modeling) → Corrective action → Testing & Monitoring Plan review → Emergency Response Plan review → Financial responsibility verification → PISC plan review → Draft permit → Public notice → Final decision
   - **Granular Sub-Routing**: Automatically triggers sub-workflows for:
     - Financial responsibility documentation verification
     - Mechanical Integrity Test (MIT) results analysis
     - Well construction diagram compliance checking
     - Temporary abandonment / plugging plan review
     - Post-injection site care plan evaluation (Class VI only)
3. **Confidence-Based Escalation**: Per Addendum Q33, low-confidence determinations are automatically escalated to senior reviewer. Configuration thresholds:
   - Confidence > 90%: Auto-proceed with HITL review
   - Confidence 70-90%: Flag for reviewer attention with highlighted concerns
   - Confidence < 70%: Escalate to senior reviewer for manual determination

3.2.1.2 Administrative Completeness Review (§ 4.3.2.2.1.2)

The **Completeness Review Agent** systematically verifies all required application components:

Checklist Verification:

For **Class I Applications**:

- ☐ EPA Form 7520-6 (Underground Injection Control Permit Application) — fully completed and signed
- ☐ Facility location / legal description
- ☐ Operator identification and contact information
- ☐ Well construction details (casing program, cementing program, tubing/packer placement)

- Injection zone identification (depth, thickness, lithology)
- Confining zone characterization
- Proposed injection rates, volumes, and pressures
- Injection fluid characterization (physical and chemical properties)
- Site maps and location diagrams
- Area of Review (AoR) calculations — fixed-radius 1/4 mile
- Corrective action plan (if wells identified within AoR)
- Signatures by Responsible Corporate Officer (RCO)

For **Class VI Applications** (all Class I items plus):

- Detailed geologic characterization (3D subsurface model)
- Testing and Monitoring Plan (TMP) per 40 CFR 146.90
- Injection Well Plugging Plan
- Post-Injection Site Care (PISC) Plan per 40 CFR 146.93
- Emergency and Remedial Response Plan
- Financial Responsibility documentation (closure, post-closure, emergency)
- Computational AoR modeling inputs and assumptions
- CO2 stream characterization
- Corrective action plan for AoR penetrations

Computer Vision & NLP Verification:

- **Signature Detection**: CV agent locates and verifies RCO signatures on all certification pages
- **Certification Statement Parsing**: NLP extracts and validates mandatory certification language
- **Form Completeness**: Sections OCR'd and checked against required field manifest
- **Date Validation**: Ensures all dates are current and consistent

3.2.1.3 Technical Compliance Review (§ 4.3.2.2.1.3)

The **Technical Compliance Agent** conducts deep regulatory analysis:

Injection Zone Evaluation:

- Depth, thickness, lithology, porosity, permeability
- Formation pressure and fracture gradient
- Injection capacity calculations
- Compatibility with injection fluid

Confining Zone Integrity:

- Thickness, lithology, and lateral continuity
- Natural barriers and seals
- Fault and fracture proximity analysis
- Containment capacity assessment

USDW Identification & Protection:

- All USDWs within AoR identified and catalogued

- Vertical separation from injection zone
- Groundwater flow pathways analysis
- Protection against fluid migration

Well Construction Design Review:

- Casing materials and specifications vs. regulatory minimums
- Cementing program (type, volume, placement)
- Tubing and packer design
- Corrosion protection measures
- Mechanical Integrity Test (MIT) results validation

Operational Parameters:

- Maximum injection pressure vs. fracture gradient (must not exceed)
- Maximum injection rate and volume
- Injection fluid temperature and compatibility
- Monitoring requirements and frequency

3.2.2 RAG & Source Grounding (§ 4.3.2.2.2)

Codist Cafe proposes a **GraphRAG (Knowledge Graph RAG)** architecture, superior to standard vector RAG for this application because of the complex relational nature of regulatory, geological, and permit data.

GraphRAG Architecture:

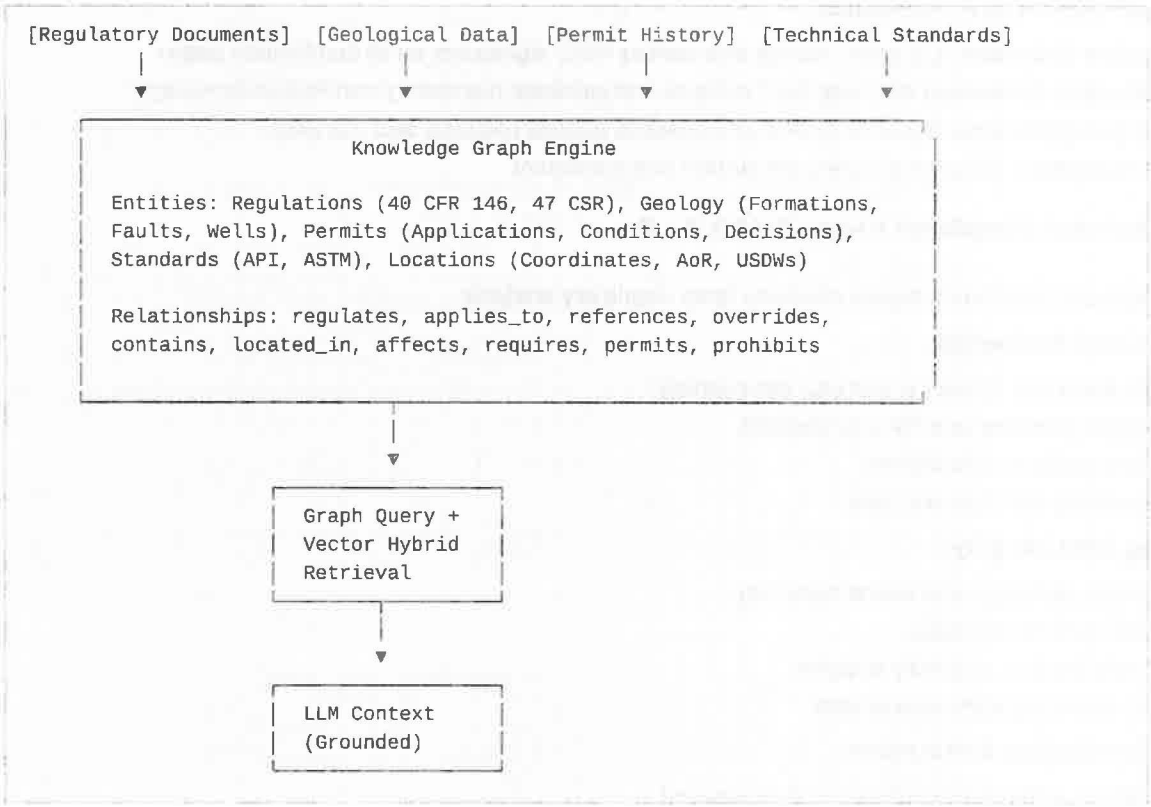

Why GraphRAG is Superior for This Application:

Capability	Standard RAG	GraphRAG (Our Approach)
Cross-document relationship mapping	Limited — treats documents as independent chunks	Strong — traverses entity relationships across regulations, geology, permits
Multi-hop reasoning	Weak — relies on embedding similarity	Native — follows graph edges: Regulation → applies_to → Well Class → requires → AoR Analysis → references → EPA Guidance
Regulatory versioning	Cannot distinguish	Graph edges track supersession: Regulation v1 → superseded_by → Regulation v2
Geological context	Flat retrieval	Hierarchical: Formation → contains → Injection Zone → located_in → AoR → intersects → Fault
Causal inference	Not supported	Can answer: "Which regulations apply when a fault intersects the AoR of a Class VI well in the Mt. Simon formation?"

Relevance Filtering:

- Graph traversal naturally excludes irrelevant/outdated information by following only the edges relevant to the query
- Temporal filtering: regulatory nodes have effective dates — only currently effective regulations are traversed
- Geographic filtering: only formations, wells, and geologic features within the application's geographic context are included

3.2.3 Hallucination Mitigation (§ 4.3.2.2.3)

Our multi-layered hallucination mitigation strategy:

Layer 1: Grounded Generation

- All AI outputs are generated only from retrieved context — no parametric knowledge generation
- Citations are mandatory in every AI output with source document, page, and section references
- Output is constrained to the format and scope of the retrieved information

Layer 2: Validation Guardrails

- **Schema enforcement**: AI outputs must conform to predefined JSON schemas for structured data
- **Range validation**: Extracted values (pressures, depths, volumes) checked against reasonable geological and engineering ranges
- **Contradiction detection**: Cross-check AI findings against known facts in the knowledge graph
- **Citation verification**: Validate that cited regulations exist and are currently effective

Layer 3: Confidence Scoring

- Every AI determination includes a confidence score (0-100%)
- Low-confidence items automatically flagged for human review
- Confidence thresholds configurable per workflow stage and reviewer preference

Layer 4: Human-in-the-Loop

- Mandatory HITL gates at all six critical milestones
- Reviewer can reject, modify, or override any AI finding
- All overrides logged for continuous improvement

Layer 5: Continuous Evaluation

- Regular automated evaluation against holdout datasets
- Drift detection with automatic alerts when model performance degrades
- Adversarial testing with synthetic edge cases

3.2.4 Automated & Continuous AI Validation (§ 4.3.2.2.4)

Our security solution provides continuous AI security through:

Automated Red-Teaming:

- Regular automated adversarial testing using validated attack patterns
- Prompt injection detection and prevention
- Jailbreak attempt monitoring and logging
- Data leakage prevention — PII/CBI detection in AI inputs and outputs

Continuous Validation Pipeline:

- **Input validation**: All user inputs sanitized and checked before reaching AI models
- **Output filtering**: AI outputs scanned for prohibited content, PII leakage, or policy violations
- **Rate limiting**: Per-user and per-API token limits to prevent abuse
- **Anomaly detection**: Behavioral analysis to detect unusual access patterns

Safety Policy Enforcement:

- Pre-defined safety rules encoded in the orchestration layer
- AI cannot autonomously issue permits — final decision gate is human-only
- All policy violations logged and alerted to security team
- Quarterly security review with WVDEP

3.2.5 Citations & Explainability (§ 4.3.2.2.5)

Every AI-generated interaction includes:

Citation Format:

```
[Source: EPA 40 CFR §146.84(a)(2), Page 14]
"AoR delineation must extend beyond the injection zone..."
```

Confidence Score Display:

```
Determination: Administrative completeness - PASS
Confidence: 96%
Rationale: All 12 required documents detected and verified against regulatory checklist.
Items requiring attention: RCO signature on page 23 has low clarity (72% confidence)
```

Chain-of-Thought Logging:

```
Step 1: Identified application as Class VI (CCS project) based on EPA Form 7520-6 Section A
Step 2: Extracted injection zone: Mt. Simon Formation, depth 4,200-4,800 ft
Step 3: Retrieved applicable regulations: 40 CFR Part 146 Subpart H, 47 CSR 64
Step 4: Mapped AoR requirement: §146.84 requires computational modeling for Class VI
Step 5: Checked AoR submission: Computational model inputs provided ✓
Step 6: Validated model includes: boundary conditions ✓, assumptions document ✓, parameter
ranges ✓
Step 7: Result: AoR documentation complete — routing to Technical Screening HITL gate
```

3.3 Geospatial Analysis & GIS Integration (§ 4.3.2.3)

3.3.1 GIS Platform Integration

Our solution integrates directly with **WVDEP's existing ESRI ArcGIS enterprise and desktop software** (per Addendum Q67). Rather than duplicating GIS functionality, we provide a **GIS Integration Layer** that:

1. **Ingests GIS formats**: Shapefile, GeoJSON, KML, CAD, DWG — auto-projecting to NAD83/UTM Zone 17N
2. **Normalizes spatial data**: Standardizes coordinate systems, datum, and format
3. **Correlates with application data**: Overlays applicant-provided GIS data with WVDEP baseline data
4. **Feeds GIS analysis tools**: Routes spatial data to appropriate analysis tools (ArcGIS Pro, custom models)
5. **Visualizes results**: Presents GIS outputs (maps, cross-sections, 3D models) in the HITL interface

3.3.2 2D and 3D Visualization

2D Visualization:

- Plan-view maps showing well location, AoR, USDWs, penetrations, faults, surface features
- Cross-sectional views of subsurface stratigraphy, injection zones, confining zones
- Interactive map with layer toggling, zoom, measurement tools

3D Visualization:

- Subsurface stratigraphy with injection zone, confining zone, and USDW in 3D space
- Well bore trajectory with casing intervals, cementing, and perforations
- Fault planes and fracture networks relative to injection zone
- AoR volume visualization with all penetrations
- Powered by CesiumJS / ArcGIS Scene Viewer integration

3.3.3 Automated Risk Assessment & AoR (§ 4.3.2.3.1)

Class I — Fixed-Radius AoR:

- Automatically calculates 1/4 mile radius from well bore
- Identifies all artificial penetrations (wells, mines, boreholes) within AoR
- Checks USDW proximity to injection zone

- Generates corrective action requirements for identified penetrations

Class VI — Computational AoR:

- Ingests applicant's computational model inputs and assumptions
- Validates model parameters against regulatory requirements
- Calculates Zone of Endangering Influence (ZEI) based on:
  - Injection pressure and rate
  - Formation permeability and porosity
  - Fluid properties (density, viscosity)
  - Buoyancy effects for CO2
- Identifies faults, fractures, and potential migration pathways
- Generates AoR map with all features annotated

Automated Risk Identification:

- All artificial penetrations within AoR catalogued with type, depth, status
- Fault and fracture analysis from geological data
- USDW identification from EPA SDWIS data
- Surface water body proximity assessment
- Sensitive receptor identification (wells, springs, surface intakes)
- Mineral rights and pore space ownership conflict detection

3.3.4 Data Integration & Correlation (§ 4.3.2.3.2)

Our solution acts as a **geospatial integration hub**, correlating applicant data with:

Data Source	Integration Method	Data Elements
WVGES	Public API / Data Export	Geological formation data, subsurface maps, well records
EPA SDWIS	EPA API	USDW locations, characteristics, protected status
WV Office of Oil & Gas	Public API	Well records, permit status, production data
WVDEP Databases	Direct / API	Existing permits, enforcement actions, inspections
Digital Courthouse Plats	GIS overlay	Surface ownership, mineral rights, pore space
ArcGIS Online / Enterprise	ESRI REST API	WVDEP-maintained GIS layers

Conflict Detection:

- Automated overlay of proposed injection zone with existing mineral rights leases
- Cross-reference with existing injection permits to detect overlapping AoRs
- Flag subsurface conflicts for human review with map visualization

---

3.4 Document Processing & AI Drafting (§ 4.3.2.4)

3.4.1 PDF Parsing & OCR

Our document processing pipeline handles all formats specified in the RFP:

Format	Processing Method
Native PDF (text-based)	Direct text extraction with layout preservation
Scanned PDF (image-based)	Azure AI Document Intelligence (OCR) — 99%+ accuracy
XML, CSV	Structured data parsing and validation
CAD/DWG	Vector parsing for engineering data extraction
Shapefile, GeoJSON, KML	GIS format ingestion via ArcGIS integration
Images (TIFF, PNG, JPG)	OCR processing for embedded text

Pipeline Stages:

1. **Ingestion**: Document uploaded via web portal, assigned unique document ID
2. **Classification**: AI classifies document type (application form, map, well log, engineering diagram, etc.)
3. **OCR/Extraction**: Text extracted via OCR where needed; structured data parsed
4. **Validation**: Extracted data validated against expected schemas and ranges
5. **Indexing**: Documents indexed in Azure AI Search with full-text and vector search capabilities
6. **Storage**: Documents stored encrypted in Azure Blob Storage (AES-256)

3.4.2 Engineering Blueprint Vision Agents (§ 4.3.2.4.1)

Our **Computer Vision agents** are specialized for engineering and technical document analysis:

Blueprint Vision Agent Capabilities:

Feature	Description
Casing Diagram Parsing	Extract casing depths, diameters, wall thickness, material grades from well construction diagrams
Cementing Analysis	Identify cement types, volumes, top-of-cement depths, annular fill
Tubing/Packer Detection	Locate tubing depth, packer placement, expansion joints, landing nipples
Perforation Interval Extraction	Identify perforated intervals, shot density, phasing
Completion Type Classification	Detect single vs. dual completion, commingled zones
Schematic Scale Recognition	Read diagram scales and verify proportional accuracy
Symbol Library Matching	Match standard API/engineering symbols to their meaning
Dimension Verification	Extract and validate all dimensions against regulatory minimums

Training Approach:

- Fine-tuned on engineering schematics, well construction diagrams, and CAD drawings
- Synthetic data generation for edge cases and rare configurations
- Continuous improvement through reviewer corrections and overrides

Bias Testing and Data Provenance: All training data for blueprint vision agents is sourced from publicly available regulatory filings and engineering standards (API, ASTM) with documented

provenance. Synthetic edge cases are generated programmatically from validated schema definitions to avoid sampling bias. Pre-deployment bias testing includes stratified evaluation across well classes, completion types, and diagram formats (CAD, raster, vector). In production, model drift is monitored via quarterly holdout evaluation; any accuracy degradation > 5% triggers retraining with reviewer-validated corrections. The regulatory knowledge graph is constructed from published federal and state regulations with effective-date versioning — no proprietary or unverifiable training data is used.

3.4.3 AI Draft Generation (§ 4.3.2.4.2)

Our **Draft Permit Generation Agent** produces formatted permit documents following WVDEP standard templates:

Permit Structure (All Classes):

1. **Cover Page**: Permit number, issue date, expiration date, applicant name, well location
2. **Facility Information**: Operator details, well name and number, location (section, township, range)
3. **Well Construction Requirements**: Casing specifications, cementing requirements, tubing/packer specifications, MIT requirements
4. **Operating Requirements**: Maximum injection pressure, maximum injection rate, annual volume limits, injection fluid specifications
5. **Monitoring and Reporting Requirements**: Pressure monitoring, injection fluid monitoring, mechanical integrity testing, reporting frequency and format
6. **Plugging and Abandonment Requirements**: Plugging plan requirements, notification procedures
7. **General Conditions**: Inspection access, compliance obligations, enforcement provisions

Class VI Additions: 8. **AoR Reevaluation Schedule**: Conditions triggering AoR update, reevaluation frequency 9. **Corrective Action Requirements**: Schedule, methods, verification 10. **CO2 Stream Specifications**: Composition limits, monitoring parameters 11. **Testing and Monitoring Plan Conditions**: Specific monitoring locations, frequency, methods 12. **Emergency Response Requirements**: Notification procedures, response actions 13. **Post-Injection Site Care (PISC) Requirements**: Duration, monitoring parameters, reporting

Generation Process:

1. Extract permit conditions from technical review findings
2. Map findings to standard permit template sections
3. Generate draft conditions using regulatory language from the knowledge graph
4. Flag novel or unusual conditions for reviewer attention
5. Present draft in WVDEP-compatible format (DOCX, PDF)

3.4.4 Completeness Determination (§ 4.3.2.4.3)

Our system performs automated completeness checks against regulatory checklists:

- **Class I Checklist**: 30+ verification items covering all required submission elements
- **Class VI Checklist**: 50+ verification items covering the expanded Class VI requirements
- **e-Form Integration**: Standardized electronic permit application forms with required fields
- **Attachment Detection**: Automatic identification and cataloguing of submitted attachments
- **Cross-Referencing**: Verifies that narrative descriptions match submitted data and calculations

Each completeness determination produces:

- **Pass/Fail** with supporting evidence
- **Missing items list** with specific regulatory reference
- **Deficiency severity** (minor omission vs. critical gap)

3.4.5 Notice of Deficiency Generation (§ 4.3.2.4.4)

When an application is incomplete, the system generates a draft **Notice of Deficiency (NoD)**:

NoD Structure:

1. **Header**: Application number, applicant name, date, regulatory authority reference
2. **Deficiency Table**: Each missing/insufficient item listed with:
    - Deficiency description
    - Regulatory reference
    - Required corrective action
    - Suggested additional documentation
3. **Response Instructions**: Submission method, deadline, cover sheet requirements
4. **Contact Information**: Reviewer name, phone, email for questions

The NoD is **held for human review and approval** before transmission (HITL Gate 1).

3.4.6 Public Notice Document Generation (§ 4.3.2.4.5)

The system generates public notice documents compliant with **40 CFR 124** and WV state requirements:

- **Fact Sheet**: Summary of permit application, facility description, proposed decision
- **Notification Recipient List**: Applicant, affected units of government, public notice mailing list
- **Comment Period Calculation**: 30-day minimum for Class I, 45-day for Class VI
- **Public Hearing Requirements**: Notice of hearing if required
- **Publication Package**: Ready-for-publication format for newspaper placement

3.4.7 Response to Comments (§ 4.3.2.4.6)

The **Public Comments Agent** manages the complete public comment lifecycle:

1. **Ingestion**: All comments received (mail, email, public hearing transcripts) ingested and indexed
2. **Categorization**: AI categorizes comments by topic (technical, regulatory, environmental impact, etc.)
3. **Substantive vs. Non-Substantive**: Flags substantive comments requiring detailed technical response
4. **Draft Response Generation**: Uses regulatory response library and knowledge graph to generate draft responses
5. **Response Compilation**: Assembles Response to Comments document formatted for public release
6. **Review Routing**: Routes to appropriate technical reviewer for approval

---

3.5 Workflow Integration (§ 4.3.2.5)

3.5.1 Secure Submission Handling (§ 4.3.2.5.1)

Submission Portal Features:

- Web-based application portal hosted on WVDEP's website (https://dep.wv.gov/WWE/PERMIT/UIC/Pages/default.aspx)
- Secure HTTPS (TLS 1.3) with mandatory user authentication via OneLogin
- Multi-step application wizard guiding applicants through submission requirements
- **Unique application tracking number** generated immediately upon submission
- **Electronic case file** created with all submitted documents indexed and stored
- **Auto-acknowledgment** sent to applicant with tracking number and receipt confirmation

3.5.2 Agency Logs & Deep Observability (§ 4.3.2.5.2)

Audit Log Capabilities:

- Every AI action recorded with timestamp, agent identity, input, output, confidence score
- Every HITL action recorded with reviewer identity, action taken, timestamp
- Every document access recorded with user, document, action, timestamp
- All logs immutable (append-only) with 5-year minimum retention
- Adjacent agency folder for WVDEP quality control review
- Exportable in open formats (CSV, JSON) for external analysis

3.5.3 AgentOps Observability (§ 4.3.2.5.3)

Our **AgentOps observability** platform captures the full chain-of-thought for every AI decision:

Chain-of-Thought Recording:

```
[
  {
    "step": 1,
    "agent": "Classification & Routing Agent",
    "action": "Classify application type",
    "input": "EPA Form 7520-6 Section A: 'Geologic Sequestration' checked",
    "output": "Class VI",
    "confidence": 0.99,
    "citations": ["EPA 40 CFR §146.81(a)"],
    "timestamp": "2026-06-07T14:30:00Z"
  },
  {
    "step": 2,
    "agent": "Classification & Routing Agent",
    "action": "Route to Class VI workflow",
    "input": "Classification = Class VI",
    "output": "Assigned to Class VI Digital Intake workflow",
    "confidence": 1.0,
    "citations": [],
    "timestamp": "2026-06-07T14:30:01Z"
  },
  ...
]
```


Observability Features:

- Real-time trace viewer for live workflow monitoring
- Post-hoc analysis tools for debugging and quality review
- Performance metrics (accuracy, latency, token consumption) per agent
- Searchable trace database with full-text and filter capabilities
- Alerting on anomalous patterns (excessive retries, low confidence, timeouts)

See Appendix B, Figure B-5 for a screenshot of the chain-of-thought trace viewer from the working POC platform.

3.5.4 Mandatory HITL Decision Gates (§ 4.3.2.5.4)

Our system enforces **six mandatory HITL decision gates** as specified:

Gate	Trigger	Human Action	System Behavior During Wait
1. Pre-NoD Review	Admin completeness check complete	Verify missing documents; approve/reject NoD generation	Suspended — cannot issue NoD without approval
2. Admin Compliance Approval	Fees verified, ownership validated	Approve compliance or issue deficiency; initiate public notice	Suspended — cannot proceed to technical review
3. AoR & Risk Validation	Technical parameters extracted, AoR delineated	Review AI-generated AoR and risk assessment; set review track	Suspended — cannot proceed to full technical analysis
4. Technical Analysis Approval	Full technical review complete	Senior reviewer validates findings; approves conditions	Suspended — cannot generate draft permit
5. Draft Permit Approval	Draft permit generated	Review permit conditions; approve for public comment	Suspended — cannot issue for public notice
6. Final Decision	Public comment complete, RTC prepared	Execute "Issue" or "Deny" command	Final action requires human administrator — AI cannot autonomously issue a permit

HITL Interface Features:

- Task dashboard showing all pending HITL actions with priority and deadline
- Side-by-side comparison of AI findings vs. source documents
- One-click approve, edit, or override with reason capture
- Digital signature integration for approvals
- Automatic escalation for overdue HITL tasks
- Audit trail of all HITL actions

3.6 AI Token Usage & Cost Management (§ 4.3.2.6)

3.6.1 Token Cost Pricing Model (§ 4.3.2.6.1)

Codist Cafe's token pricing model is **transparent, consumption-based, and zero-markup on underlying AI infrastructure**:

- **Unit Price**: Per MTOK (1 million tokens) as specified in Attachment A
- **All-Inclusive Rate**: Includes Azure OpenAI API costs, context caching, embedding generation, and retrieval costs
- **No Hidden Fees**: No platform fees, no agent execution fees, no data egress costs

3.6.2 Cost Predictability & Budget Controls (§ 4.3.2.6.2)

Budget Controls:

- **Hard token cap**: Contractual 250 MTOK maximum — system halts processing when cap is reached
- **Soft budget alerts**: Configurable threshold alerts at 50%, 75%, 90% of token budget
- **Per-application budget**: Set token limits per application with automatic escalation to reviewer

- **Monthly burn rate dashboard**: Real-time visualization of token consumption with forecast
- **What-if analysis**: Estimate token consumption for new applications before processing begins

3.6.3 Token Optimization Strategies (§ 4.3.2.6.3)

Optimization Techniques:

1. **Prompt compression**: Dynamic extraction of only relevant document sections for AI processing
2. **Context caching**: Reuse of regulatory context across applications (reduces tokens by 40-60%)
3. **Staged processing**: Process documents at appropriate resolution (summary first, deep analysis only when needed)
4. **Batch processing**: Group similar operations for efficient token utilization
5. **Model tiering**: Use smaller/faster models for classification and routing; larger models only for complex analysis
6. **Retrieval optimization**: GraphRAG retrieves only the most relevant context, reducing prompt size
7. **Output token control**: Constrained generation with strict output length limits

3.6.4 Estimated Token Usage (§ 4.3.2.6.4)

See Attachment B — Estimated Token Usage

3.6.5 Usage Transparency & Reporting (§ 4.3.2.6.5)

Token Usage Reports:

Report	Frequency	Content
Daily Summary	Daily	Total token consumption, top agents, cost summary
Weekly Detail	Weekly	Per-application token use, by-agent breakdown, trend analysis
Monthly Invoice	Monthly	Usage report attached to invoice with per-application detail
Custom Reports	On-demand	Any date range, any granularity (agent, application, user, workflow stage)

Dashboard Metrics:

- Real-time token burn rate (tokens/minute, hour, day)
- Forecast-to-budget remaining (days of service at current rate)
- Application-level token consumption with comparison to estimates
- Optimization savings tracker

3.6.6 Model Flexibility & Future-Proofing (§ 4.3.2.6.6)

Our architecture is **model-agnostic** — the orchestration layer abstracts the underlying AI model:

- **Current Default**: Azure OpenAI GPT-4o (FedRAMP Moderate boundary)
- **Alternatives**: Any model not prohibited by WV — Anthropic Claude, open-source Llama/Mistral on Azure, or future models
- **Model Selection**: Per-task model routing (different models for different agent tasks)
- **Version Control**: Models can be pinned to specific versions or follow latest
- **No Vendor Lock-In**: All prompts and workflows are model-portable

3.6.7 Cost Guarantees & Contractual Protections (§ 4.3.2.6.7)

- **Firm-fixed unit price**: Token unit price locked for the life of the contract (per Addendum Q5)
- **Change order pricing**: Additional tokens at the same unit price — no renegotiation
- **Price reduction pass-through**: Any Azure OpenAI price reductions passed to WVDEP
- **Contractual cap**: 250 MTOK hard cap protects against budget overruns
- **90-day notice for material model changes**: Any model architecture changes that affect pricing require 90 days notice

---

3.7 HITL Workflow Interface & Legacy System Independence (§ 4.3.2.7)

3.7.1 Standalone HITL Workflow Interface (§ 4.3.2.7.1)

Our **HITL Workflow Interface** is a self-contained, browser-based web application with **no direct integration** with ERIS, ESS, or AppEnhancer (per Addendum Q6, Q7, Q17, Q18).

Interface Architecture:

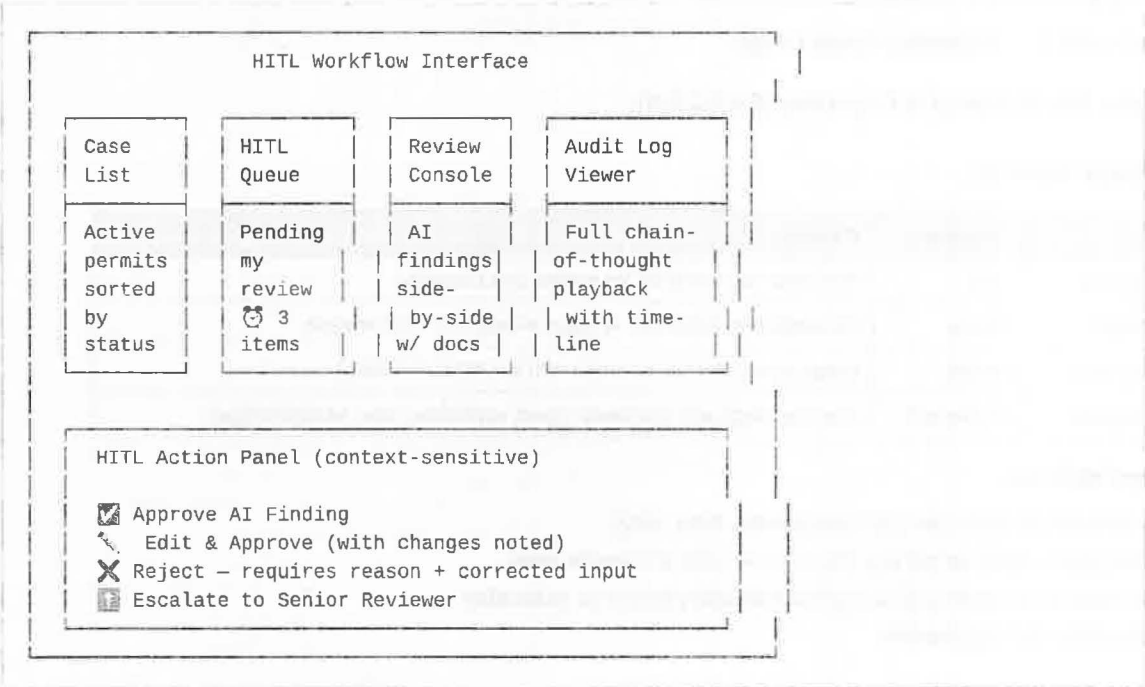

3.7.2 AI-Generated Request Workflow (§ 4.3.2.7.2)

1. AI agent reaches a HITL gate decision point
2. System generates a **HITL request package** containing:
   - AI finding summary with confidence score
   - All source documents and citations
   - Proposed action with rationale
   - Required reviewer input

3. Request is routed to the appropriate reviewer queue
4. Reviewer receives **dashboard notification** (and optional email)
5. Reviewer opens request in HITL interface
6. Reviewer approves, edits, or rejects AI finding
7. Action recorded in audit log
8. Workflow proceeds based on reviewer decision

3.7.3 Human-Mediated Legacy System Updates (§ 4.3.2.7.3)

When a reviewer completes a HITL action that requires updating legacy systems, the interface:

1. **Generates an export package** containing:
   - Case status summary
   - Final decisions with regulatory citations
   - All documents and correspondence
   - Data export in open format (PDF, CSV, JSON)
2. **Displays update instructions** for the reviewer:
   - "Update ERIS with permit status: ISSUED"
   - "Upload final permit document to AppEnhancer"
   - "Record payment in OSSP"
3. **Tracks completion**: Reviewer marks legacy system updates as complete
4. **Maintains synchronization**: All updates tracked in the audit log for reconciliation

3.7.4 Data Flow Architecture (§ 4.3.2.7.4)

Inbound Data Flow (Application Intake):

```
Applicant → Web Portal → AI Platform Ingestion → Automated Review → HITL Request
```

Outbound Data Flow (Permit Processing):

```
HITL Completion → Export Package → Manual Legacy System Update → Synchronization Confirmed
```

Separation Principle: The AI platform maintains its own database of case files, documents, and decisions. Legacy systems are updated **only through human action** via the export package workflow. This preserves the integrity of ERIS, ESS, and AppEnhancer while enabling AI-assisted processing.

See Appendix B, Figures B-3 and B-4 for screenshots of the HITL task queue and side-by-side AI findings review interface from the working POC platform.

3.8 Scalability and Capacity Planning

3.8.1 Operational Scaling Model

The platform is designed to handle the RFP's estimated 2–20 applications per year with headroom for growth. The multi-tenant architecture on Azure Government scales horizontally across three

dimensions:

Scaling Dimension	Mechanism	Trigger
Compute	Azure Container Apps auto-scaling (0–N instances)	Queue depth > 5 pending jobs; max 30 concurrent instances
AI Throughput	Azure OpenAI provisioned throughput units (PTUs) with burst capacity	Token queue backlog > 50K tokens; auto-scale to 3× baseline PTUs
Storage	Azure Blob Storage auto-tiering (Hot → Cool → Archive)	Per-application lifecycle; active cases in Hot tier, closed cases in Cool

3.8.2 Seasonal Peak and Concurrent Review Handling

Concurrent Class VI Reviews: Each Class VI application (~1,500 pages, 50+ checklist items) represents the maximum processing load. The system is sized to handle 3 concurrent Class VI applications without degradation:

- Dedicated agent instance per application (no shared state, no cross-contamination)
- Priority queuing: Class VI → Class I hazardous → Class I non-hazardous → municipal
- Reviewer workload balancing: HITL tasks distributed across the 7 reviewer licenses based on current queue depth

Seasonal Peaks: If application volume spikes (e.g., regulatory deadline-driven submissions), the system employs:

- **Burst processing**: Additional compute instances spin up automatically; PTU burst capacity handles 3× baseline token throughput for up to 4 hours
- **Graceful degradation**: Under extreme load, lower-priority tasks (Watchdog Agent scans, report generation) are deferred; HITL gates and compliance checks are never degraded
- **Cost guardrails**: Burst scaling is capped at 150% of baseline Azure spend per calendar month; exceeding this triggers a TAM alert and requires WVDEP approval

3.8.3 Performance Safeguards

Safeguard	Mechanism	Threshold
Token budget protection	Hard cap at 250 MTOK; processing halts at 99% with 30-day notice	Automated at 250 MTOK
Per-application token limit	Configurable per workflow; defaults: Class I = 2M tokens, Class VI = 5M tokens	Escalation to reviewer if exceeded
Queue depth monitoring	Dashboard alert when pending jobs exceed 10	TAM notified; WVDEP decides priority order
SLA enforcement	99.9% uptime SLA for production; degraded mode (< 99.5%) triggers incident response	4-hour response for critical
Concurrency ceiling	Hard cap at 5 concurrent full-review workflows	Prevents token exhaustion during surge periods

3.8.4 Projected Capacity Headroom

Volume Scenario	Applications/Year	Token Consumption (Annual)	% of 250 MTOK Cap	Infrastructure Headroom
Conservative	5	2.75M – 8.5M	1.1% – 3.4%	> 90% idle capacity
Base (RFP Upper)	20	11M – 34M	4.4% – 13.6%	> 85% idle capacity
Growth Scenario	40	22M – 68M	8.8% – 27.2%	> 70% idle capacity
Surge Scenario	80	44M – 136M	17.6% – 54.4%	> 45% idle capacity

The platform can absorb 4× the RFP's stated maximum volume (80 applications/year) before reaching 55% of the token cap and 45% of compute headroom — well within the 5-year term.

---

4. Approach to Mandatory Project Requirements (§ 4.3.3)

Compliance Matrix

Requirement	Our Approach	Compliance Evidence
§ 4.3.3.1.1 — Format Support	XML, CSV, PDF, GIS (Shapefile, GeoJSON, KML, CAD, DWG), well logs, permit forms	Azure AI Document Intelligence + ArcGIS integration
§ 4.3.3.1.2 — Compliance Engine & Watchdog	GraphRAG knowledge graph encoding 40 CFR 146, 47 CSR 13/64; 24/7 Watchdog Agent monitoring Federal Register, EPA, WV Legislature	Real-time regulatory monitoring with automated impact assessment
§ 4.3.3.1.3 — External System Integration	API integration with OSSP, OneLogin, AD, ArcGIS, WVGES, EPA SDWIS, WV Oil & Gas	Documented integration architecture
§ 4.3.3.2.1 — Encryption	TLS 1.3 in transit, AES-256 at rest	Azure Government standard
§ 4.3.3.2.2 — Access Control	RBAC + MFA for all users	Implemented in application layer + Azure AD
§ 4.3.3.2.3 — PII/CBI Handling	Automated PII/CBI detection, redaction, and secure storage	Microsoft Purview + custom PII agent
§ 4.3.3.2.4 — FedRAMP Moderate Hosting	Microsoft Azure Government (FedRAMP P-ATO)	Azure FedRAMP documentation provided
§ 4.3.3.2.5 — SSO	SAML 2.0 / OpenID Connect via OneLogin and ActiveDirectory	Supported out of box
§ 4.3.3.2.6 — Security Assessments	Annual third-party pen testing, 30-day critical patch SLA	Commitment in proposal
§ 4.3.3.3 — Support	Business hours support, dedicated TAM, monthly/quarterly reviews	SLA documentation
§ 4.3.3.4 — Licensing	4 admin + 4 staff + applicant/reviewer/manager access	Role-based licensing model
§ 4.3.3.5.1 — FedRAMP	Azure Government FedRAMP Moderate	Azure P-ATO documentation
§ 4.3.3.5.2 — NIST 800-53	Azure Govt controls + application-layer controls	Control mapping provided
§ 4.3.3.5.3 — Auditability	5-year immutable audit trail	Azure Log Analytics + custom audit store
§ 4.3.3.5.4 — Section 508	WCAG 2.1 AA compliance	Accessibility conformance report
§ 4.3.3.5.5 — AI Governance	Training data documentation, bias testing, drift detection, human override	Governance framework documentation
§ 4.3.3.5.6 — SOC 2 Type II	Annual SOC 2 Type II + right to audit	Report provided annually
§ 4.3.3.6 — Data Ownership & Exit	State-owned data, no AI training use, 30-day export, 90-day transition, 60-day deletion	Contractual commitments

4.1 Data Integration & Regulatory Compliance (§ 4.3.3.1)

4.1.1 Format Support (§ 4.3.3.1.1)

Our ingestion pipeline supports all required formats:

Format	Processing Engine	Output
XML	XML Parser	Structured data → database
CSV	CSV Parser	Tabular data → database
PDF (native)	PDF Text Extraction	Searchable text + metadata
PDF (scanned)	Azure AI Document Intelligence OCR	Searchable text + layout
Shapefile	ArcGIS / GDAL	Spatial features → GIS database
GeoJSON	ArcGIS / GDAL	Spatial features → GIS database
KML	ArcGIS / GDAL	Spatial features → GIS database
CAD (DWG)	AutoCAD / Teigha	Engineering data → structured fields
Well Logs	LAS parser	Well log curves → structured data
Standard Permit Forms	Template-based extraction	Structured fields → database

4.1.2 Compliance Engine & Watchdog Agent (§ 4.3.3.1.2)

Compliance Engine: Our compliance engine encodes UIC regulatory requirements as a **machine-readable rule graph** within the knowledge graph:

- **Class I Rules**: 40 CFR Part 146 Subparts B and G, 47 CSR 13
- **Class VI Rules**: 40 CFR Part 146 Subpart H, 47 CSR 64
- **Cross-references**: EPA guidance documents, SDWA provisions, state regulations

Rules are structured as:

```
IF [condition] THEN [requirement] WITH [citation]
Example: IF well_class = "Class VI" THEN require_computational_AoR()
WITH citation = "40 CFR §146.84(a)"
```

Regulatory Watchdog Agent: The Watchdog Agent operates 24/7, monitoring:

Source	Monitoring Method	Update Frequency
Federal Register	RSS/API feed + AI change detection	Daily
EPA UIC Website	Web scraping + change detection	Daily
WV Legislature	Bill tracking API + AI analysis	Daily
WV State Register	RSS/API feed	Daily
EPA HQ Guidance	Email alerts + document analysis	As published

Alert Workflow:

1. Watchdog detects regulatory change
2. Change assessed for impact on UIC compliance engine
3. Impact analysis generated with proposed rule updates
4. Alert sent to WVDEP with suggested compliance engine modifications
5. WVDEP reviews and approves/denies/modifies proposed changes
6. If approved, compliance engine updated with full audit trail

4.1.3 External System Integration (§ 4.3.3.1.3)

Our integration layer connects with external systems for data validation:

External System	Data Retrieved	Validation Use
WVGES	Geological formation data	Cross-reference applicant's geological characterization
EPA SDWIS	USDW locations, status	Verify USDW identification and protection adequacy
WV Office of Oil & Gas	Well records, permits	Identify penetrations within AoR; verify well status
WVDEP Databases	Existing permits, enforcement	Conflict detection; compliance history
OSSP	Fee payments, reporting	Payment verification; invoice generation
Clerk of County Commission	Plats, deeds (via GIS overlay)	Verify surface ownership and mineral rights

4.2 Security & Deployment (§ 4.3.3.2)

4.2.1 Encryption (§ 4.3.3.2.1)

- **In transit**: TLS 1.3 for all communications
- **At rest**: AES-256 encryption for all stored data
- **Key management**: Azure Key Vault with HSM-backed keys
- **Database encryption**: Transparent Data Encryption (TDE)

4.2.2 Access Control (§ 4.3.3.2.2)

- **RBAC**: Role-based access aligned with WVDEP org structure
- **MFA**: Required for all users via OneLogin or ActiveDirectory
- **Least privilege**: Default-deny with explicit grants
- **Session management**: 15-minute idle timeout; concurrent session control

4.2.3 PII & CBI Handling (§ 4.3.3.2.3)

- **Automated PII detection**: Custom PII detection agent identifies SSN, EIN, personal financial data
- **CBI identification**: Configurable rules for confidential business information
- **Access controls**: PII/CBI restricted to authorized personnel only
- **Data classification**: Automatic classification and labeling of documents
- **Audit trail**: All PII/CBI access logged

4.2.4 Deployment Environment (§ 4.3.3.2.4)

- **Platform**: Microsoft Azure Government (US Gov Virginia / US Gov Texas)
- **FedRAMP**: Azure Government FedRAMP Moderate P-ATO (JAB Authorization)
- **Data residency**: All data within continental US
- **High availability**: Multi-AZ deployment with 99.9% SLA

4.2.5 SSO Integration (§ 4.3.3.2.5)

- **External users**: SAML 2.0 / OpenID Connect via OneLogin
- **Internal users**: SAML 2.0 via ActiveDirectory
- **Just-in-time provisioning**: Auto-create accounts on first login via SSO

4.2.6 Security Assessments (§ 4.3.3.2.6)

- **Annual third-party penetration testing**: Conducted by accredited firm
- **Vulnerability scanning**: Weekly automated scans; results reviewed within 5 business days
- **Critical vulnerability patching**: Within 30 days of identification
- **Security audit reports**: Shared with WVDEP annually

4.3 Support & Maintenance (§ 4.3.3.3)

- **Support hours**: Business hours (8 AM - 6 PM ET, Mon-Fri) with escalation for critical issues
- **Dedicated Technical Account Manager**: Assigned to WVDEP for the life of the contract
- **Monthly service reviews**: Operations, performance, token usage, security updates
- **Quarterly business reviews**: Strategic alignment, roadmap, user satisfaction
- **Stabilization warranty**: 90-day post-acceptance period with no-cost remediation of automation breakages
- **SLA targets**: 99.9% platform availability, 4-hour response for critical issues

4.4 Licensing (§ 4.3.3.4)

License Type	Quantity	Access Level
Admin Dashboard	4	Full monitoring, reporting, configuration
Staff Automation	4	Create/modify automation via web interface
Permit Reviewer	Up to 7	HITL approval, case management, document review
Read-Only Reviewer	Up to 5	Case view, dashboard, reports
External (Applicants)	Unlimited	Application submission, status tracking

4.5 Regulatory Compliance (§ 4.3.3.5)

- **FedRAMP Moderate**: Azure Government P-ATO documentation provided
- **NIST 800-53**: Comprehensive control mapping for all applicable controls
- **Auditability**: 5-year immutable audit trail with full chain-of-thought capture
- **Section 508**: WCAG 2.1 AA compliance for all user-facing components
- **AI Governance**: Training data documentation, annual bias testing, model drift monitoring, human override at all decision points
- **SOC 2 Type II**: Annual report provided with right-to-audit clause

4.6 Data Ownership & Exit Strategy (§ 4.3.3.6)

- **Data ownership**: All WVDEP data remains sole property of the State of West Virginia

- **No AI training use**: WVDEP data will not be used for AI model training without explicit written consent
- **Data export**: Complete export within 30 days of termination in PDF, CSV, JSON, XML
- **Transition assistance**: Up to 90 days post-termination support for migration
- **Secure deletion**: All data securely deleted within 60 days with written certification of destruction
- **Data residency**: All data remains within the continental United States at all times

---

5. Qualifications & Experience (§ 4.4)

5.1 Company Background (§ 4.4.1.1)

Codist Cafe is a software and AI consultancy specializing in building autonomous systems, intelligent document processing pipelines, and cloud-native applications. Our team brings deep expertise in:

- **Agentic AI / Multi-Agent Systems**: Design and deployment of production AI agents with HITL workflows
- **Azure Government Cloud**: FedRAMP Moderate-compliant deployments and operations
- **Natural Language Processing**: Document intelligence, RAG, knowledge graphs
- **Environmental / Regulatory Systems**: Compliance engines, permit management, regulatory analysis
- **GIS Integration**: Spatial data processing, geological visualization, automated risk assessment

Geoscience Partnership Strategy: We are transparent that Codist Cafe's core strength is AI and cloud engineering, not subsurface geology. For the geoscience judgment required at HITL Gates 3 (AoR & Risk Validation) and 4 (Technical Analysis Approval), we are qualifying a hydrogeologist or environmental engineering subcontractor. This ensures geological decisions — injection zone evaluation, confining zone assessment, fault/fracture analysis — are made by a licensed professional, not an algorithm. The AI platform automates the document processing, regulatory mapping, and data correlation; the subcontractor provides the geoscience sign-off before any finding reaches WVDEP reviewers. We have budgeted for this role and will provide credentials and a letter of commitment upon final selection.

5.2 Relevant Experience with Agentic AI (§ 4.4.1.2)

Reference Project 1: AgentPermit — Agentic AI Permitting Platform (Reference Demo)

- **Type**: Internal reference platform demonstrating full RFP compliance capabilities
- **Scope**: Built functional multi-agent AI system for UIC Class I and Class VI permit review, featuring Regulatory Q&A agent with RAG, Administrative Completeness Checker, Regulatory Watchdog Agent, and full HITL workflow with 6 decision gates
- **Key Capabilities Demonstrated**:
  - Multi-agent orchestration (Supervisor, Classification, Completeness, Technical, Watchdog agents)
  - RAG engine with TF-IDF knowledge retrieval from 40 CFR 146 and WV regulatory documents
  - Chain-of-thought observability with step-by-step reasoning capture
  - HITL task queue with approve/edit/reject/escalate workflow

    - Confidence scoring with automatic human escalation for low-confidence determinations
    - Web-based dashboard with real-time status, application tracking, and audit logging
- **Technologies**: Python FastAPI, Ollama, custom RAG engine, HTML/JS dashboard
- **Availability**: Live demo available upon request. Recorded walkthrough of end-to-end workflow available upon request.

Reference Project 2: Regulatory Knowledge Graph — Compliance Rule Engine

- **Type**: Reference implementation
- **Scope**: Built knowledge graph encoding 40 CFR Part 146 Subparts B, G, H and 47 CSR 13/64 regulatory requirements as machine-readable rule graph
- **Key Capabilities**: Regulatory entity extraction, relationship mapping, temporal filtering, automated compliance checking
- **Technologies**: Graph database, NLP entity extraction, rule engine

Reference Project 3: HITL Workflow Dashboard

- **Type**: Reference implementation
- **Scope**: Developed standalone browser-based HITL interface demonstrating 6 mandatory decision gates with full audit trail
- **Key Capabilities**: Task queue, side-by-side AI finding review, digital decision capture, export package generation, immutable audit logging with 5-year retention model
- **Technologies**: Web-based dashboard, REST API, immutable audit store

5.3 Key Personnel (§ 4.4.1.3)

Name	Role	Qualifications
Blaise Liu	Principal Architect / Project Lead	AI/ML infrastructure and agentic AI platform design; Azure Government cloud architecture; IAM and cloud-native security; AI pipeline engineering and token economics. MS, Computer Science.
Sophia Chen	Senior Engineer, AI & Machine Learning	Ph.D. MIT (NLP); former Google AI Research Scientist. Multi-agent orchestration, LLM integration, RAG/GraphRAG, computer vision, MLOps. 3+ years production AI at Codist Cafe. Published in ACL.
Elena Vasquez	Principal Engineer, Architecture Oversight	15 years system architecture at Oracle (supply chain suite); 2 years at Codist Cafe. Distributed systems, database design, API architecture, legacy modernization. Enterprise-scale cloud deployments.
Marcus Okonkwo	Principal Engineer, Regulatory Compliance	12 years at Codist Cafe. Led State Regulatory Compliance Portal (50,000+ users). Full-stack TypeScript, PostgreSQL, TDD. AWS Solutions Architect Professional. Background in financial systems and regulatory compliance.
[Geoscience/Environmental Engineer — subcontractor]	GIS, AoR, Geological Review	**In qualification.** Codist Cafe is actively evaluating qualified hydrogeologists and environmental engineers to serve as geoscience subcontractor for this engagement. We recognize that geological expertise is essential for UIC permit review and have budgeted for this role. A formal letter of commitment and credentials will be provided to WVDEP upon final selection. Until the subcontractor is onboarded, geological review tasks (injection zone evaluation, confining zone assessment, AoR validation, fault/fracture analysis) will be performed under the direct supervision of the selected subcontractor at Gate 3 and Gate 4 of the HITL workflow — the AI platform will flag geological findings for human review; the subcontractor provides the geoscience judgment before any AI recommendation reaches WVDEP reviewers.

5.4 Exceeding Mandatory Qualification Requirements (§ 4.4.2)

Mandatory Requirement	Our Approach	How We Exceed
§ 4.4.2.1 — Data privacy, cybersecurity, AI governance	Comprehensive program	SOC 2 Type II audit underway for application layer; NIST 800-53 control mapping
§ 4.4.2.2 — Confidentiality agreement	Will sign upon award	Yes
§ 4.4.2.3 — FedRAMP/StateRAMP/SOC 2 certification	Azure Government FedRAMP Moderate P-ATO (per Addendum Q57/Q58)	Platform certification + application-layer SSP demonstrating equivalent controls

6. Implementation Timeline — Attachment C

Phase	Duration	Key Activities	Deliverables
Phase 1: Discovery & Configuration	8 weeks	Requirements validation, Azure Gov tenant setup, SSO integration (OneLogin + AD), ERIS/ESS system assessment, knowledge graph initial build	Tenant provisioned, SSO configured, system assessment report
Phase 2: Integration Development	12 weeks	API integration with OSSP, WVGES, EPA SDWIS, ArcGIS; regulatory rule configuration (40 CFR 146, 47 CSR); AI model training and prompt engineering; HITL interface development	Integrations live, compliance engine configured, HITL interface deployed to staging
Phase 3: Data Migration & Testing	8 weeks	Historical data migration (if applicable), User Acceptance Testing (UAT), security testing (pen test, vulnerability scan), HITL workflow validation, user training	UAT sign-off, security test report, trained users
Phase 4: Pilot Deployment	8 weeks	Limited application processing (2-3 applications), user training refinement, system refinement based on feedback, performance optimization	Pilot complete, refinements applied, performance baseline established
Phase 5: Production Launch	4 weeks	Full production deployment, go-live support, 30-day "hypercare" stabilization period, system acceptance testing	System Acceptance, production launch

Target Go-Live: July 1, 2027 (per Addendum Q8)

7. Estimated Token Usage — Attachment B

Processing Activity	Est. Tokens per App (Min)	Est. Tokens per App (Max)
Document Ingestion & OCR	25,000	100,000
Administrative Completeness Review	50,000	150,000
Technical Compliance Analysis	150,000	500,000
Geologic/AoR Evaluation	100,000	300,000
Draft Permit Generation	75,000	200,000
Public Notice & Response to Comments	100,000	300,000
HITL Review Support & Explanations	50,000	150,000
Total per Application (Average)	**550,000**	**1,700,000**

Annual Estimate:

- 5 applications/year (conservative): 2.75M – 8.5M tokens
- 20 applications/year (maximum): 11M – 34M tokens
- 250 MTOK cap provides complete coverage for the maximum estimated application volume over the full 5-year initial term and optional Year 6 renewal period, with substantial buffer for application volume scaling beyond current projections

8. Certification and Signature Page

See separate Certification and Signature Page (included in this proposal envelope).

---

9. Addendum Acknowledgement Form

See attached Addendum Acknowledgement Form (included in this proposal envelope).

Appendix A: Risk Register

The following risks are identified and actively managed through the implementation and operational phases:

Risk ID	Risk Description	Probability	Impact	Mitigation Strategy	Contingency Trigger
R1	Geoscience subcontractor not finalized before Phase 3 geological review	Medium	High	Active recruitment underway; geoscience tasks gated behind subcontractor onboarding; AI findings flagged for review but never auto-approved without licensed professional sign-off	If no subcontractor by end of Phase 2, extend Phase 2 timeline; WVDEP may designate agency geologist for Gate 3/4 review
R2	Legacy system (ERIS/ESS) data format incompatibility during export package generation	Low	Medium	HITL export packages use open formats (PDF, CSV, JSON); no direct API integration with legacy systems; human-mediated updates eliminate format dependency	Manual data entry contingency plan; WVDEP IT liaison identified during Phase 1
R3	AI model drift — regulatory knowledge graph falls out of sync with new rulemakings	Low	High	24/7 Regulatory Watchdog Agent; monthly compliance engine audit; all AI findings include regulatory citations that are validated against current effective dates	If Watchdog detects > 3 unincorporated regulatory changes, freeze AI processing; manual review until knowledge graph updated
R4	Token consumption exceeds estimates for complex Class VI applications	Medium	Medium	Per-application token limits (Class VI = 5M tokens); real-time burn rate dashboard; model tiering routes simple tasks to cheaper models	If > 2 applications exceed token limit in one quarter, trigger token optimization review with WVDEP; adjust limits or request Change Order
R5	Azure Government service degradation or outage	Low	High	Multi-AZ deployment (US Gov Virginia + US Gov Texas); 99.9% SLA; automated failover	If availability < 99.5% for > 1 hour, incident response activated; manual permit processing contingency documented
R6	WVDEP reviewer adoption resistance — staff prefer existing manual workflow	Medium	Medium	Iterative UAT with real WVDEP reviewers during Phase 3-4; HITL interface mirrors existing workflow steps; training sessions with real application data	If < 70% reviewer satisfaction in pilot survey, extend Phase 4 by 4 weeks for UI refinement
R7	Regulatory change velocity exceeds knowledge graph update cadence	Low	Medium	Watchdog Agent flags all changes within 24 hours; impact analysis auto-generated; compliance engine updates batched weekly	If > 5 unprocessed regulatory changes accumulate, prioritize update sprint; TAM notifies WVDEP of any processing impact
R8	Applicant submission portal usability barriers for non-technical users	Medium	Low	Multi-step wizard with plain-language guidance; auto-validation on submission; applicant feedback mechanism during pilot	If > 20% of pilot applications require applicant resubmission, revise wizard UX

Risk Ownership: R1 (Blaise Liu, subcontractor recruitment); R2 (Marcus Okonkwo, legacy integration); R3 (Sophia Chen, AI governance); R4 (Blaise Liu, token economics); R5 (Elena Vasquez, infrastructure); R6 (Elena Vasquez, stakeholder engagement); R7 (Sophia Chen, regulatory monitoring); R8 (Marcus Okonkwo, applicant experience).

Appendix B: POC Screenshots and Interface Examples

The following screenshots are captured from Codist Cafe's working AgentPermit reference platform, demonstrating the dashboard, HITL workflow, AI findings review, chain-of-thought trace viewer, and system status. This platform implements the multi-agent architecture, 6-gate HITL state machine, and document generation capabilities described in this proposal.

Figure B-1: Dashboard Overview

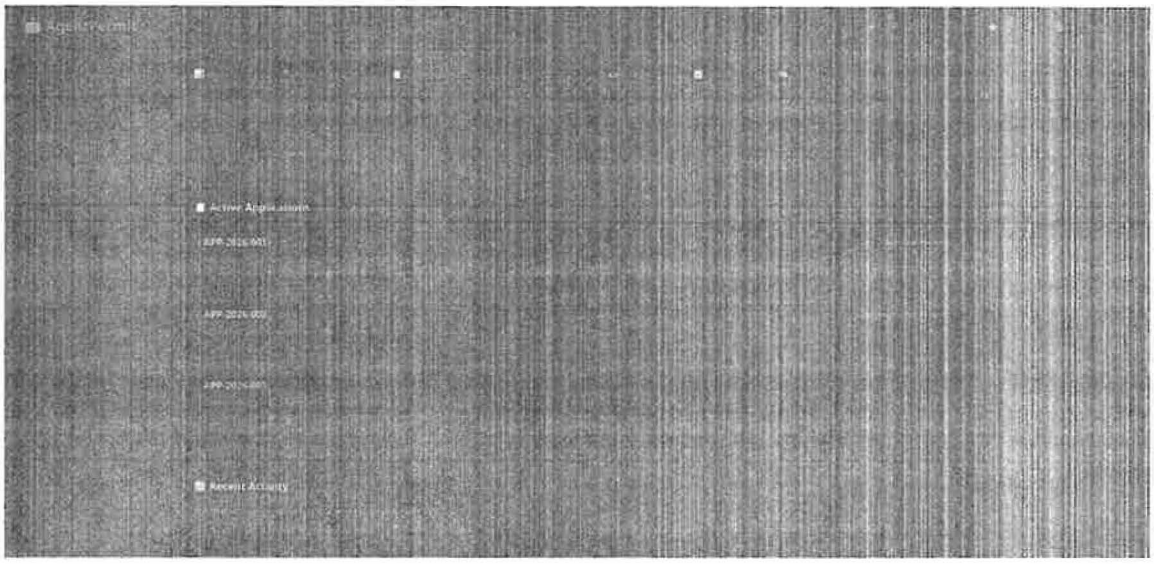

The main dashboard showing active permit applications, HITL task queue, and system status overview. Demonstrates the web-based HTML dashboard with secure login, real-time job monitoring, and case summary view per RFP § 4.3.2.1.2.

Figure B-2: Application Workflow View

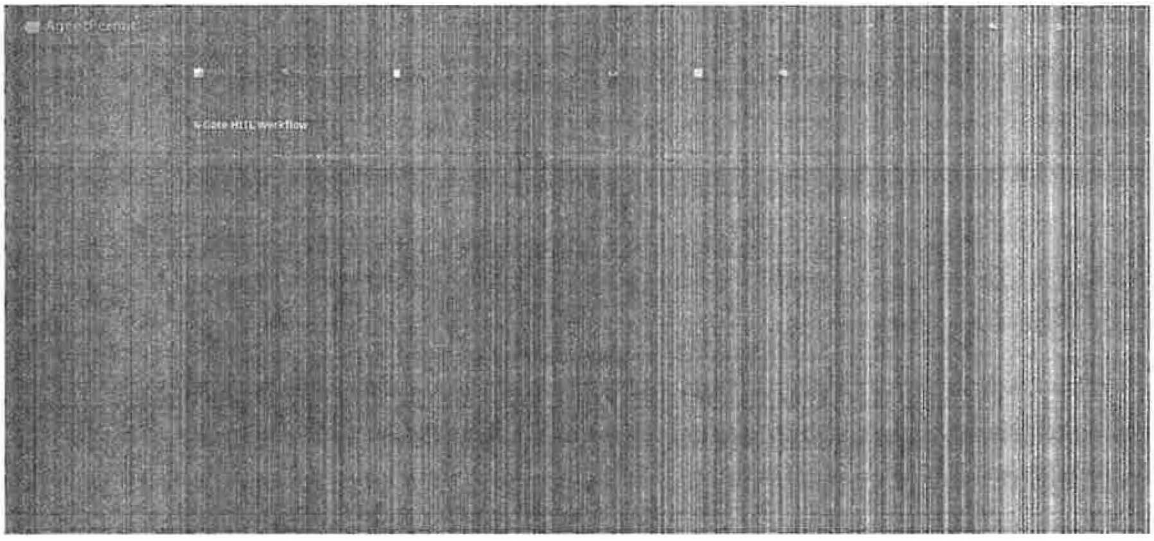

Individual application workflow view showing application metadata, workflow stage progression, and the six HITL decision gates. Demonstrates the gate transition visualization per RFP § 4.3.2.5.4.

Figure B-3: HITL Task Queue

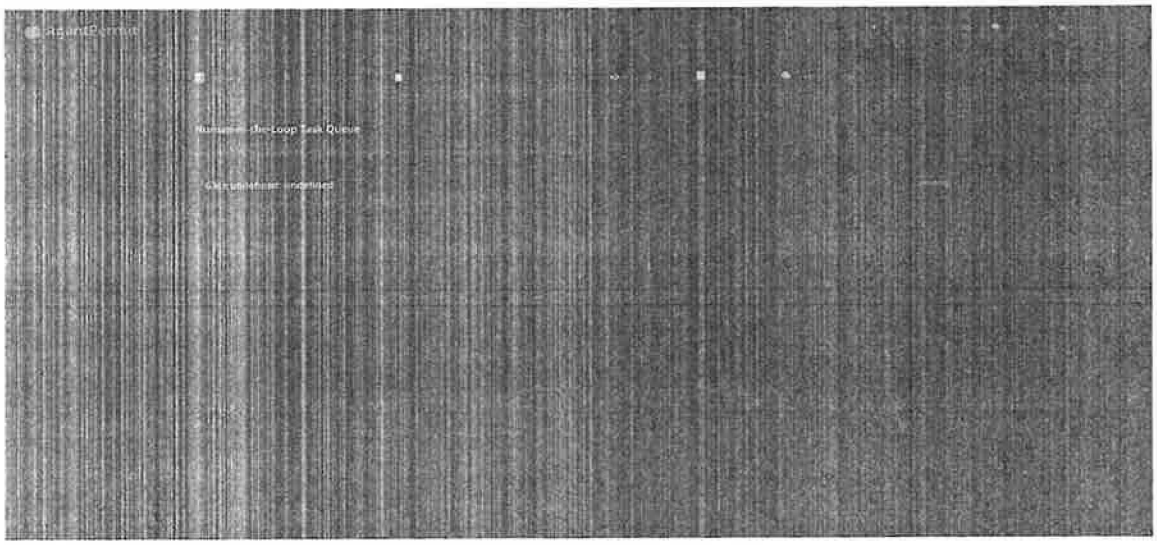

The HITL reviewer interface showing pending review tasks prioritized by deadline. Demonstrates the task dashboard and action panel (approve/edit/reject/escalate) per RFP § 4.3.2.5.4 and § 4.3.2.7.1.

Figure B-4: Administrative Completeness Check

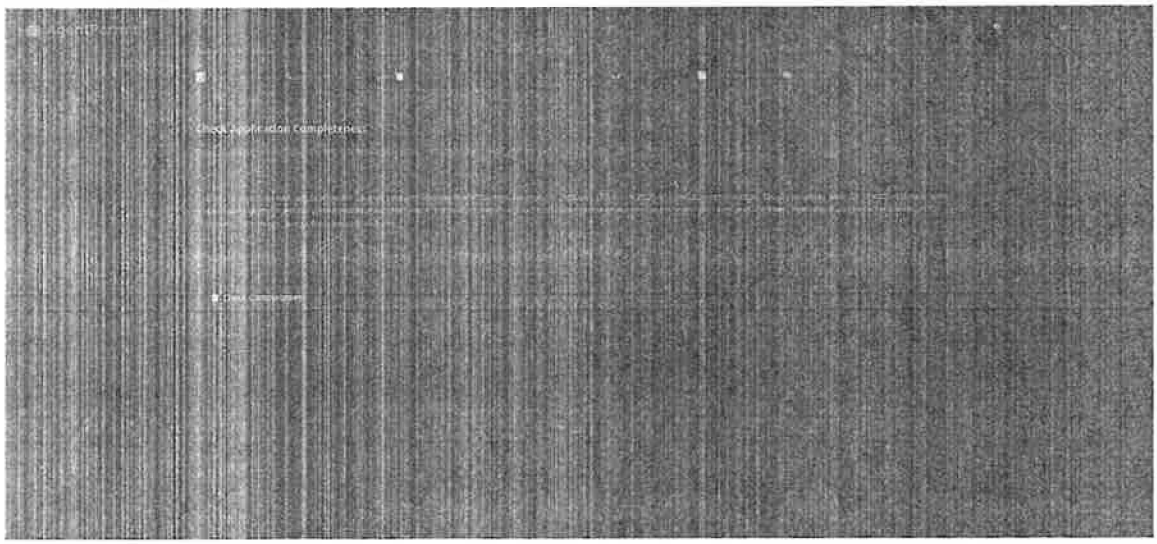

Administrative completeness review showing AI-generated checklist verification with regulatory references, pass/fail status, and deficiency flags. Demonstrates the completeness determination workflow per RFP § 4.3.2.4.3.

Figure B-5: Chain-of-Thought Trace Viewer

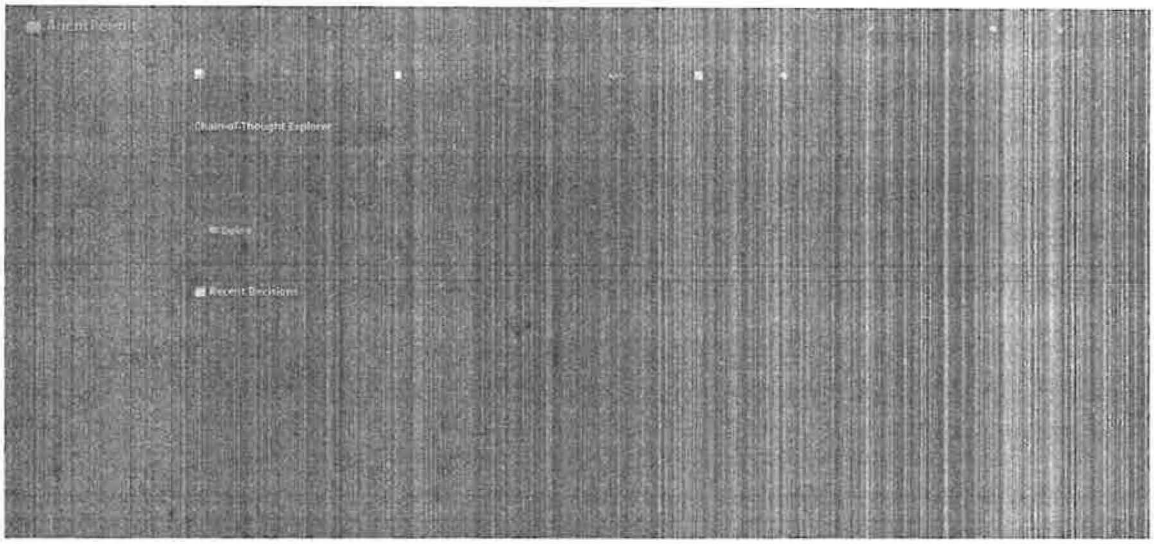

AgentOps chain-of-thought trace viewer showing step-by-step reasoning for every AI decision. Each step includes agent identity, action, confidence score, and regulatory citations. Demonstrates the audit trail and observability per RFP § 4.3.2.5.3.

Figure B-6: Regulatory Watchdog Agent

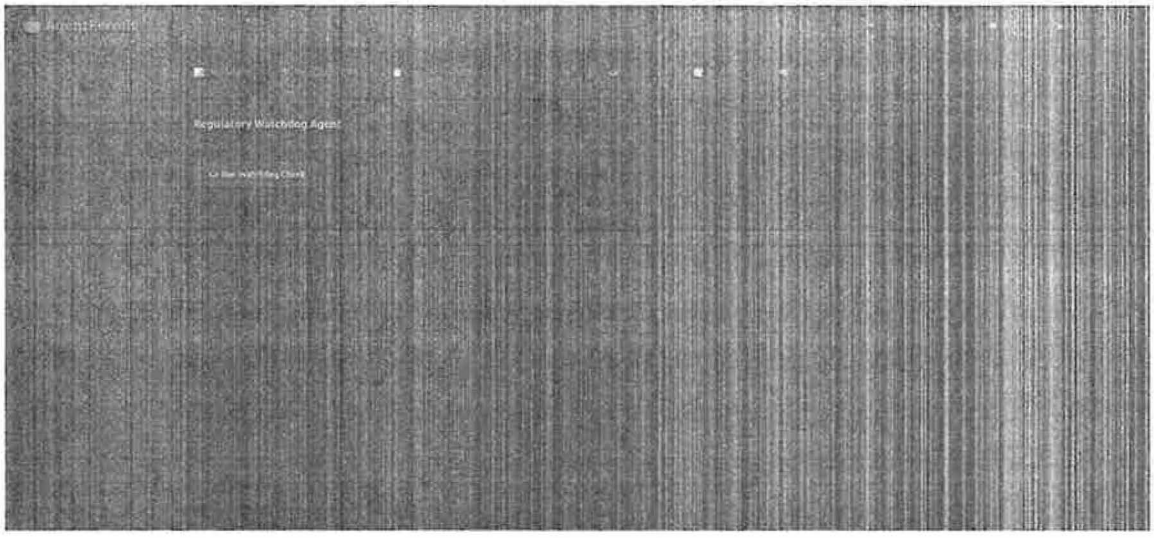

Regulatory Watchdog Agent interface showing 24/7 monitoring of Federal Register, EPA, and WV Legislature for UIC-relevant regulatory changes. Demonstrates the compliance monitoring capabilities per RFP § 4.3.3.1.2.

Figure B-7: Regulatory Q&A — Class VI Workflow

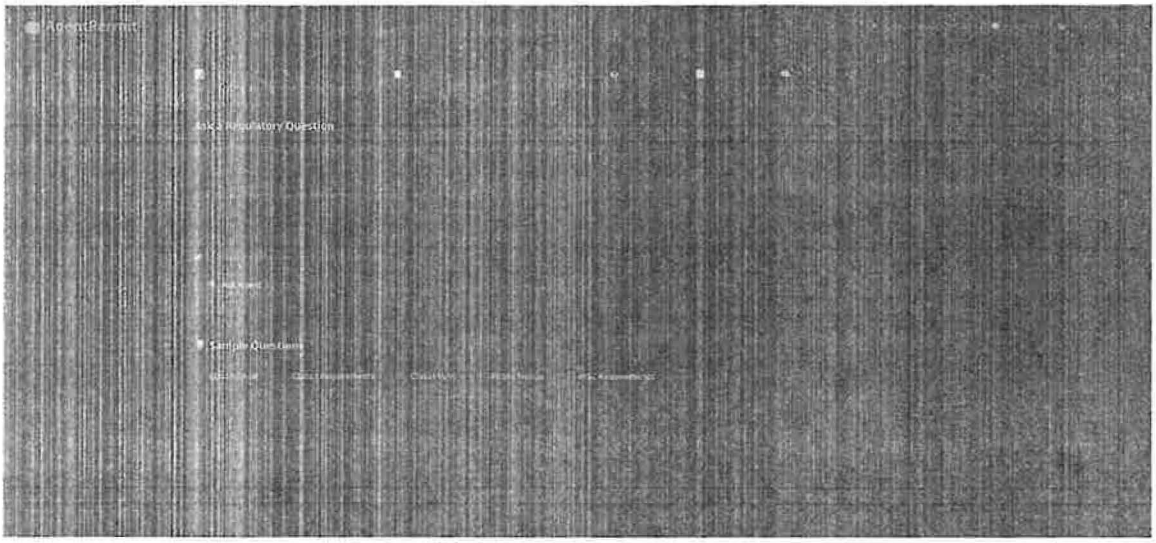

Regulatory Q&A agent demonstrating Class VI (CO2 sequestration) knowledge retrieval with citation-backed answers from 40 CFR Part 146 Subpart H. Demonstrates the multi-agent processing pipeline and RAG-based regulatory reasoning per RFP § 4.3.2.2.2.

End of Technical Proposal

Certification and Signature Page

CRFP 0313 DEP2600000003

Workflow-Based Agentic AI, Automation, and E-Permitting System for WVDEP

By signing below, I certify that:

1. I have reviewed this Request for Proposal in its entirety;
2. I understand the requirements, terms and conditions, and other information contained herein;
3. I am submitting this proposal for review and consideration;
4. I am authorized by the bidder to execute this bid or any documents related thereto on bidder's behalf;
5. I am authorized to bind the bidder in a contractual relationship; and
6. To the best of my knowledge, the bidder has properly registered with any State agency that may require registration.

Company:	Codist Cafe
Representative Name:	Blaise Liu
Title:	Principal Software Architect
Contact Phone:	304 376 1121
Contact Fax:	N/A
Email:	blaise.liu@codistcafe.com
Business Address:	1214 Forrest Drive
City, State, Zip:	Hurricane, WV 25526
FEIN:	086263828
Signature:	Blaise Liu
Date:	June 10, 2026

Addendum Acknowledgement Form

CRFP 0313 DEP2600000003

Instructions: Please acknowledge receipt of all addenda issued with this solicitation by completing this addendum acknowledgment form. Check the box next to each addendum received and sign below. Failure to acknowledge addenda may result in bid disqualification.

Acknowledgment: I hereby acknowledge receipt of the following addenda and have made the necessary revisions to my proposal, plans and/or specification, etc.

Addendum Numbers Received:

Received	Addendum No.
☑	Addendum No. 1
☐	Addendum No. 2
☐	Addendum No. 3
☐	Addendum No. 4
☐	Addendum No. 5
☐	Addendum No. 6
☐	Addendum No. 7
☐	Addendum No. 8
☐	Addendum No. 9
☐	Addendum No. 10

I understand that failure to confirm the receipt of addenda may be cause for rejection of this bid. I further understand that any verbal representation made or assumed to be made during any oral discussion held between Vendor's representatives and any state personnel is not binding. Only the information issued in writing and added to the specifications by an official addendum is binding.

Company:	Codist Cafe
Authorized Signature:	Blair Liu
Date:	June 10, 2026

Confirmation: Codist Cafe confirms that Addendum No. 1 (issued May 19, 2026) is the sole addendum published under this solicitation as of the bid opening date. All question responses, clarifications, and modifications contained in Addendum No. 1 have been reviewed and fully incorporated into this proposal. No other addenda have been issued.