



TECHNICAL PROPOSAL

Solicitation No. CRFP 0313 DEP2600000003

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

UIC Class I and Class VI Permitting - West Virginia Department of Environmental Protection

Submitted By	Digitamize, Inc. dba Supervity
Address	11921 Freedom Drive, Ste. 570, Reston, VA, 20190
Principal Contact	Brett Fraser, Vice President of Public Sector
Phone	571-524-9125
Email	brett@supervity.ai
FEIN	82-0687505
Submitted To	WV Dept. of Administration, Purchasing Division
Buyer	Joseph (Josh) E Hager III
Solicitation Closes	June 10, 2026 at 13:30
Date of Submission	June 10, 2026
Document Type	TECHNICAL PROPOSAL ONLY — Cost Proposal submitted separately

RECEIVED
 2026 JUN -9 AM 8:55
 WV PURCHASING
 DIVISION

Supervity acknowledges that this proposal is a public document subject to disclosure under West Virginia Code §§5A-3-1 et seq. and the West Virginia Freedom of Information Act. No claim of confidentiality or proprietary designation is asserted over its contents. Personnel names and contact information are disclosed solely to satisfy mandatory solicitation requirements. Upon contract award, data privacy obligations will be governed by the SaaS Addendum incorporated into this solicitation, which Supervity agrees to in full.

Authorized Signature:

Date: June 10, 2026

Table of Contents

Executive Summary.....	4
Section 1: Company Background and Introduction	4
1.1 Supervity Differentiators (§4.3).....	5
Section 2: Project Goals and Proposed Approach (§4.3).....	7
2.1 General Automation and Dashboard Integration (§4.3.2.1)	7
2.1.1 System Automation / Integration (§4.3.2.1.1).....	7
2.1.2 UIC Permitting Command Center Dashboard (§4.3.2.1.2).....	7
2.2 UIC Class I and VI Agentic AI Processing (§4.3.2.2).....	8
2.2.1 Digital Intake Specialist Functionality (§4.3.2.2.1).....	8
2.2.2 Intelligent Context Engine, GraphRAG and Source Grounding (§4.3.2.2.2).....	9
Dynamic AI Policies for WVDEP Regulatory Governance	9
2.2.3 Grounded Regulatory Execution and Hallucination Mitigation (§4.3.2.2.3).....	9
2.2.4 Automated and Continuous AI Validation (§4.3.2.2.4).....	10
2.2.5 Citations and Explainability (§4.3.2.2.5).....	10
2.3 Geospatial Analysis and GIS Integration (§4.3.2.3).....	10
2.3.1 Automated Risk Assessment and AoR (§4.3.2.3.1)	10
2.3.2 Data Integration and Correlation (§4.3.2.3.2).....	10
2.4 Document Processing and AI Drafting (§4.3.2.4)	11
2.4.1 Engineering Blueprint Vision Agents (§4.3.2.4.1).....	11
2.4.2 AI Draft Generation (§4.3.2.4.2)	11
2.4.3 Completeness Determination (§4.3.2.4.3)	11
2.4.4 Notice of Deficiency Generation (§4.3.2.4.4).....	11
2.4.5 Public Notice Document Generation (§4.3.2.4.5)	11
2.4.6 Response to Comments (§4.3.2.4.6).....	11
2.5 Workflow Integration and HIC Architecture (§4.3.2.5).....	11
2.5.1 Secure Submission Handling (§4.3.2.5.1).....	11
2.5.2 Agency Logs and Deep Observability (§4.3.2.5.2)	12
2.5.3 Decision Traceability and AgentOps Observability (§4.3.2.5.3).....	12
2.5.4 Mandatory HIC Decision Gates (§4.3.2.5.4)	12
2.6 AI Token Usage and Cost Management Approach (§4.3.2.6)	13
2.7 HIC Workflow Interface and Legacy System Independence (§4.3.2.7).....	13
2.7.1 Standalone HIC Workflow Interface (§4.3.2.7.1).....	13
2.7.2 AI-Generated Request Workflow (§4.3.2.7.2)	13
2.7.3 Human-Mediated Legacy System Updates (§4.3.2.7.3).....	13
2.7.4 Data Flow Architecture (§4.3.2.7.4)	13
Section 3: Mandatory Project Requirements Compliance (§4.3.3).....	14

3.1 Data Integration and Regulatory Compliance (§4.3.3.1)..... 15

3.2 Security and Deployment (§4.3.3.2)..... 15

3.3 Support and Maintenance (§4.3.3.3)..... 16

3.4 Licensing (§4.3.3.4)..... 16

3.5 Regulatory Compliance and AI Governance (§4.3.3.5) 16

3.6 Data Ownership and Exit Strategy (§4.3.3.6) 16

Section 4: Qualifications and Experience (§4.4) 16

 4.1 Company Background (§4.4.1.1)..... 16

 4.2 Relevant Agentic AI Experience and References (§4.4.1.2)..... 17

 4.3 Key Personnel (§4.4.1.3) 17

Section 5: Implementation Timeline (§4.6 / Attachment C) 18

Section 6: Mandatory Compliance Summary 18

Conclusion 19

Certification and Signature Page 20

 Designated Contract Administrator 20

Addendum Acknowledgement Form 21

Executive Summary

The West Virginia Department of Environmental Protection (WVDEP) requires a secure, auditable, and regulator-controlled system to modernize UIC Class I and Class VI permitting without compromising human judgment, legal defensibility, or state data sovereignty. These permitting workflows are document-heavy, technically complex, geospatially sensitive, and governed by overlapping federal and state regulatory requirements. Reviewers must evaluate application completeness, geological and engineering evidence, AoR risks, USDW protection, public notice requirements, deficiency responses, and final permit conditions while maintaining a complete record of every decision.

Supervity proposes a UIC Permitting AI Employee Command Center that operates as a governed digital permitting unit for WVDEP. The AI Employee ingests application packages, extracts structured and unstructured content, applies Class I and Class VI regulatory logic, correlates geospatial and historical records, generates source-cited review packages, drafts deficiency notices and permit materials, and routes every consequential action through Human-in-Command approval gates.

While the RFP specifies Human-in-the-Loop (HITL) as the governance framework, Supervity's solution delivers Human-in-Command (HIC), a more rigorous and regulator-appropriate standard that goes beyond passive checkpoint approvals to give WVDEP reviewers active, policy-governed authority over every AI action, with full decision traceability, reviewer-defined escalation logic, and the ability to modify, override, or redirect AI behavior at any point in the permitting workflow, ensuring that WVDEP staff are not merely consulted at predefined stops, but remain in command of the entire permitting process from intake to final determination.

The proposed solution is powered by Supervity's Intelligent Context Engine, which enables the AI Employee to operate with WVDEP-specific regulatory context, application context, geospatial context, precedent context, and human reviewer feedback. This allows the system to support faster, more consistent permit review while preserving WVDEP's authority over every regulatory determination.

Key outcomes include reduced manual review burden, faster completeness screening, stronger source-grounded decision support, improved visibility into application status and bottlenecks, more consistent regulatory application, and a complete audit trail of AI-assisted and human-approved actions.

Section 1: Company Background and Introduction

Supervity is an enterprise AI company headquartered in Reston, Virginia, purpose-built to deploy governed AI Employees and AI Command Centers for regulated, audit-sensitive operations across government, banking, insurance, healthcare, and shared services. Trusted by more than 300 global enterprise teams and recognized with the prestigious Technology of the Year award at the SSON Impact Awards, Supervity's platform enables organizations to deploy AI Employees that execute complex, document-intensive workflows while keeping humans in command of every consequential decision. Supervity's multi-modal platform combines Agentic RAG, Agentic OCR, Agentic Workflow orchestration, Computer Vision, and the Intelligent Context Engine into a unified operating layer governed by Dynamic AI Policies, Human-in-Command approval gates, AgentOps observability, and full decision traceability, making it uniquely suited to the regulatory precision, auditability, and human oversight requirements of the WVDEP UIC permitting program. Supervity has direct experience deploying AI Employees for government permitting and compliance workflows, including a prior engagement with the West Virginia Department of Environmental Protection, demonstrating not only proven public-sector capability but specific familiarity with WVDEP's operational environment, data sources, and regulatory requirements. Supervity's flagship platform is SOC 2 Type II certified, SaaS-delivered system that orchestrates multi-agent AI workflows for complex regulatory permitting, compliance determination, and document intelligence tasks.

Our proposal responds directly to WVDEP's need for a Workflow Based Agentic AI, Automation, and E-Permitting System for UIC Class I and Class VI permitting. We understand that these permits govern hazardous waste injection and geologic sequestration of CO₂ - activities with profound implications for underground sources of drinking water (USDWs). The Supervity solution ensures that no permit is issued without explicit human authorization, every AI determination is traceable to its source, and the State's data sovereignty is fully protected.

Supervity proposes a purpose-built, UIC-specialized system that encodes West Virginia's Class I (40 CFR Part 146 Subparts B & G; 47 CSR 13) and Class VI (40 CFR Part 146 Subpart H; 47 CSR 64) regulatory requirements directly into the intelligent context layer and compliance engine, integrates with WVGES, EPA SDWIS, and WVDEP's existing systems (ERIS, ESS, AppEnhancer), and delivers a 5-year production system backed by contractual SLAs, annual adversarial validation, and FedRAMP-level security controls.

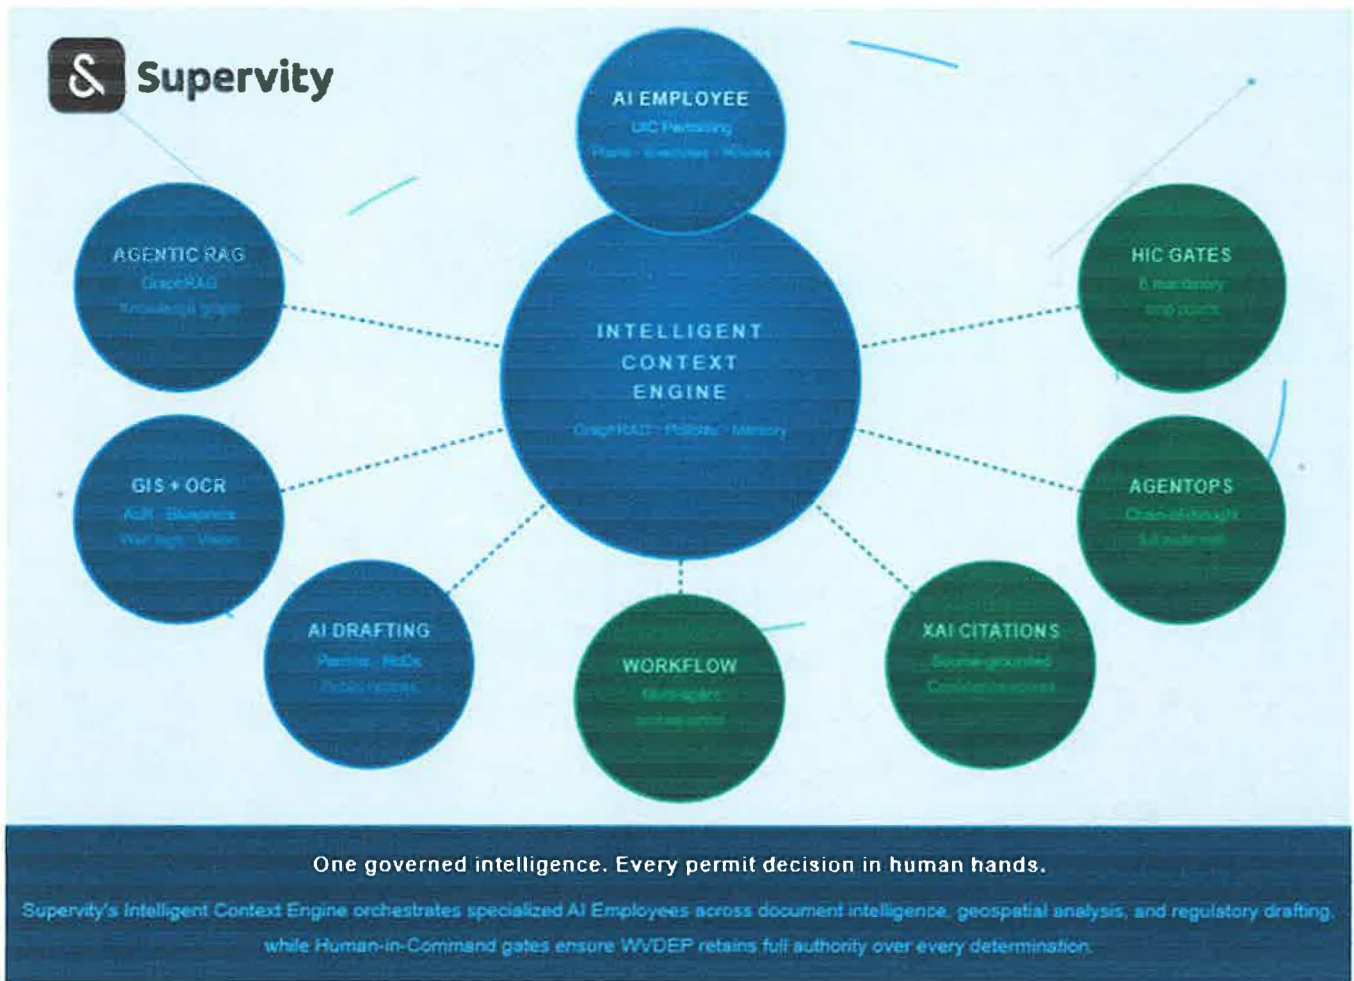


Image 1. The Supervity AI Employee for WVDEP Permitting

1.1 Supervity Differentiators (§4.3)

The Workflow Based Agentic AI, Automation, and E-Permitting System RFP requires vendors to describe how their proposed approach is superior or inferior to other possible approaches. The following differentiators distinguish Supervity's solution from competing agentic AI and e-permitting platforms.

1.1.1 AI Employee Intelligent Context Engine

Most agentic AI platforms treat each permit application as an isolated inference task, losing institutional knowledge between sessions and relying on generic retrieval to ground AI responses. Supervity's AI Employee Intelligent Context Engine fundamentally changes this model by maintaining a persistent, structured understanding of WVDEP's regulatory environment, permit history, geological data, and agency preferences across every interaction. Rather than re-learning context from scratch with each application, the Intelligent Context Engine continuously enriches its understanding, meaning the system becomes demonstrably smarter and more accurate over time as it processes more WVDEP permit decisions. For evaluators, this translates directly into higher-quality AI determinations, fewer false deficiency notices, and a compliance engine that grows more precisely calibrated to West Virginia's specific regulatory interpretations rather than generic federal standards.

Through Dynamic AI Policies, WVDEP can govern how the AI Employee evaluates completeness, routes deficiencies, flags technical risk, escalates ambiguous determinations, and prepares review packages for human approval. Context is carried across the permitting workflow so that administrative completeness findings, technical review outcomes, AoR risk flags, public notice requirements, and final decision support remain connected inside one governed operating layer.

This is what differentiates Supervity from generic AI or automation tools. The AI Employee does not simply retrieve documents or generate drafts. It operates from a persistent, governed permitting context that allows

WVDEP staff to supervise, audit, and continuously improve the permitting process while retaining full authority over every regulatory decision. Competing approaches that rely solely on static RAG or session-scoped context cannot replicate this compounding accuracy advantage.

1.1.2 AI Employee Orchestration and Workflow Management

Generic automation platforms require significant custom engineering to coordinate multiple AI agents across a complex, multi-stage regulatory workflow. Supervity's AI Employee Orchestration and Workflow Management framework is purpose-built for exactly this challenge - natively managing the handoffs between specialized agents (classification, completeness review, technical compliance, GIS analysis, document drafting, public notice generation) without custom middleware or brittle point-to-point integrations.

Each AI Employee operates within a defined scope, escalates to human reviewers at the appropriate Human-in-Command gate, and passes structured context packages to the next agent in the chain ensuring no information is lost between processing stages and no agent operates outside its authorized function. Competing platforms typically offer either monolithic single-model approaches (which sacrifice specialization) or loosely coupled agent frameworks (which require extensive custom orchestration logic and introduce failure points at every handoff). Supervity's native orchestration layer eliminates both risks while providing WVDEP with the full AgentOps Chain of Thought observability required for legal defensibility.

1.1.3 Cost Optimization for AI Token Usage

Token costs are the hidden variable in most AI permitting proposals. Vendors who price on raw token consumption expose agencies to unpredictable and potentially significant cost escalation as application volumes or document complexity increase. Supervity's token optimization architecture addresses this at every layer of the system.

At the retrieval layer, the Intelligent Context Engine dramatically reduces redundant token consumption by serving pre-enriched context rather than re-retrieving and re-processing the same regulatory reference material for every application. At the inference layer, Supervity's tiered model routing applies smaller, specialized models to discrete tasks, such as document classification, completeness checking, and form field extraction. The Supervity platform reserves larger, more expensive models only for tasks that genuinely require their capability, such as complex technical compliance analysis and draft permit generation. At the output layer, prompt compression templates encode regulatory context efficiently, reducing input token volume without sacrificing analytical quality.

The combined effect is estimated peak annual consumption of approximately 94 MTok against a 250 MTok contractual ceiling providing WVDEP with over 60% cost headroom and contractually capped exposure, rather than the open-ended token billing models offered by competing platforms.

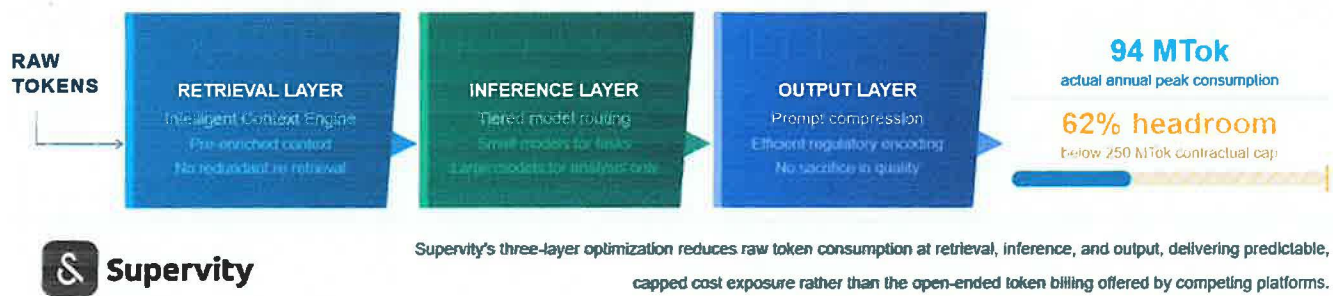


Image 2. Supervity Token Optimization

Supervity brings to this engagement a platform that was not retrofitted for government permitting, it was designed from the ground up to operate in environments where regulatory precision, human accountability, and institutional defensibility are non-negotiable. The Intelligent Context Engine, the Human-in-Command governance architecture, the six mandatory decision gates, and the AgentOps observability framework are not features added to a general-purpose AI tool. They are the foundational design decisions of a platform purpose-built for exactly the kind of high-stakes, document-intensive, legally consequential workflow that UIC Class I and Class VI permitting represents.

Every finding the AI Employee produces traces to a source document, a regulatory citation, or an approved policy rule. Every consequential determination waits for a human reviewer to authorize it. Every action the system takes

is logged, time-stamped, and exportable in a format that survives legal challenge and regulatory audit. Supervity's prior engagement with WVDEP demonstrates that this is not a theoretical capability, it is a proven operational reality, deployed in a West Virginia regulatory environment, by the same team proposing to serve WVDEP here. That combination of purpose-built architecture, proven government deployment experience, and contractually capped, predictable cost structure is what Supervity offers WVDEP as its partner for the UIC permitting mission.

Section 2: Project Goals and Proposed Approach (§4.3)

2.1 General Automation and Dashboard Integration (§4.3.2.1)

2.1.1 System Automation / Integration (§4.3.2.1.1)

The UIC Permitting AI Employee Command Center serves as the central orchestration and governance layer, integrating with WVDEP's existing operational ecosystem through a non-invasive, read-then-human-update pattern. The AI Employee retrieves, classifies, analyzes, and prepares work packages while WVDEP personnel retain authority over all regulatory determinations and legacy system updates.

- **WV One-Stop Shop Permitting (OSSP):** The AI Employee integrates with the WV OSSP system for reporting and payment processing. Upon receipt of a new application, it retrieves the submission package, assigns a unique tracking number (format: UIC-[CLASS]-[YEAR]-[SEQUENCE]), classifies the application, and creates the initial review package for WVDEP staff. External users access the system via OneLogin; internal WVDEP staff authenticate via Active Directory.
- **Identity and Access Management:** The Command Center integrates with OneLogin for external applicant access and Active Directory for internal WVDEP staff authentication, providing role-based access control aligned to WVDEP organizational structure. All legacy system updates (ERIS, ESS) are performed manually by authorized WVDEP staff based on AI-generated export packages, preserving human authority over all records entries.
- **External Regulatory Data Sources:** The Command Center supports scheduled synchronization with WVGES geological data, EPA SDWIS USDW records, and WV Office of Oil and Gas well records via public API access. WVDEP has confirmed minimal legacy UIC data exists given that Class I and Class VI are new programs in West Virginia.
- **External Regulatory Sources:** The Command Center supports scheduled synchronization with WVGES geological data, EPA SDWIS USDW records, WV Office of Oil and Gas well records, and relevant regulatory update sources, subject to WVDEP-approved integration design.

2.1.2 UIC Permitting Command Center Dashboard (§4.3.2.1.2)

Supervity will deliver a secure UIC Permitting Command Center that serves as the Human-in-Command operating environment for WVDEP reviewers, supervisors, administrators, and technical specialists. The Command Center provides a unified workspace, with consolidated dashboard data, to manage permit queues, review AI-generated packages, inspect evidence and citations, approve or reject recommendations, monitor bottlenecks, and maintain a complete audit trail of AI-assisted and human-approved actions.

- Centralized permit queue showing active, pending, completed, and exception-state permit review workflows with real-time status and ownership.
- AI Workbench for reviewer action, where staff can inspect evidence, source citations, confidence scores, extracted application data, regulatory policies applied, recommended next steps, and prior human decisions before approving, rejecting, or requesting revision.
- Comprehensive decision trace and AgentOps observability logs capturing source evidence, policies evaluated, extracted data, recommended actions, human decisions, timestamps, and final dispositions. (step-by-step reasoning capture per §4.3.2.5.3).
- Role-based secure login with MFA; RBAC aligned to WVDEP organizational structure and reviewer authority levels.
- AI Insights and regulatory update feed surfacing relevant regulatory changes, workflow bottlenecks, recurring deficiency categories, AoR risk flags, and exception patterns for WVDEP leadership review.

2.2 UIC Class I and VI Agentic AI Processing (§4.3.2.2)

Supervity's UIC Permitting AI Employee executes end-to-end Class I and Class VI processing through specialized agents operating under Human-in-Command (HIC) governance. This section outlines the architecture that transforms raw submissions into structured, source-grounded, regulator-ready review packages. From intake and classification through compliance evaluation and draft generation, each stage preserves regulatory fidelity, enforces policy-driven decisions, and maintains auditability. All AI-supported actions remain transparent, evidence-based, and validated, while WVDEP reviewers retain full control over determinations and outcomes.

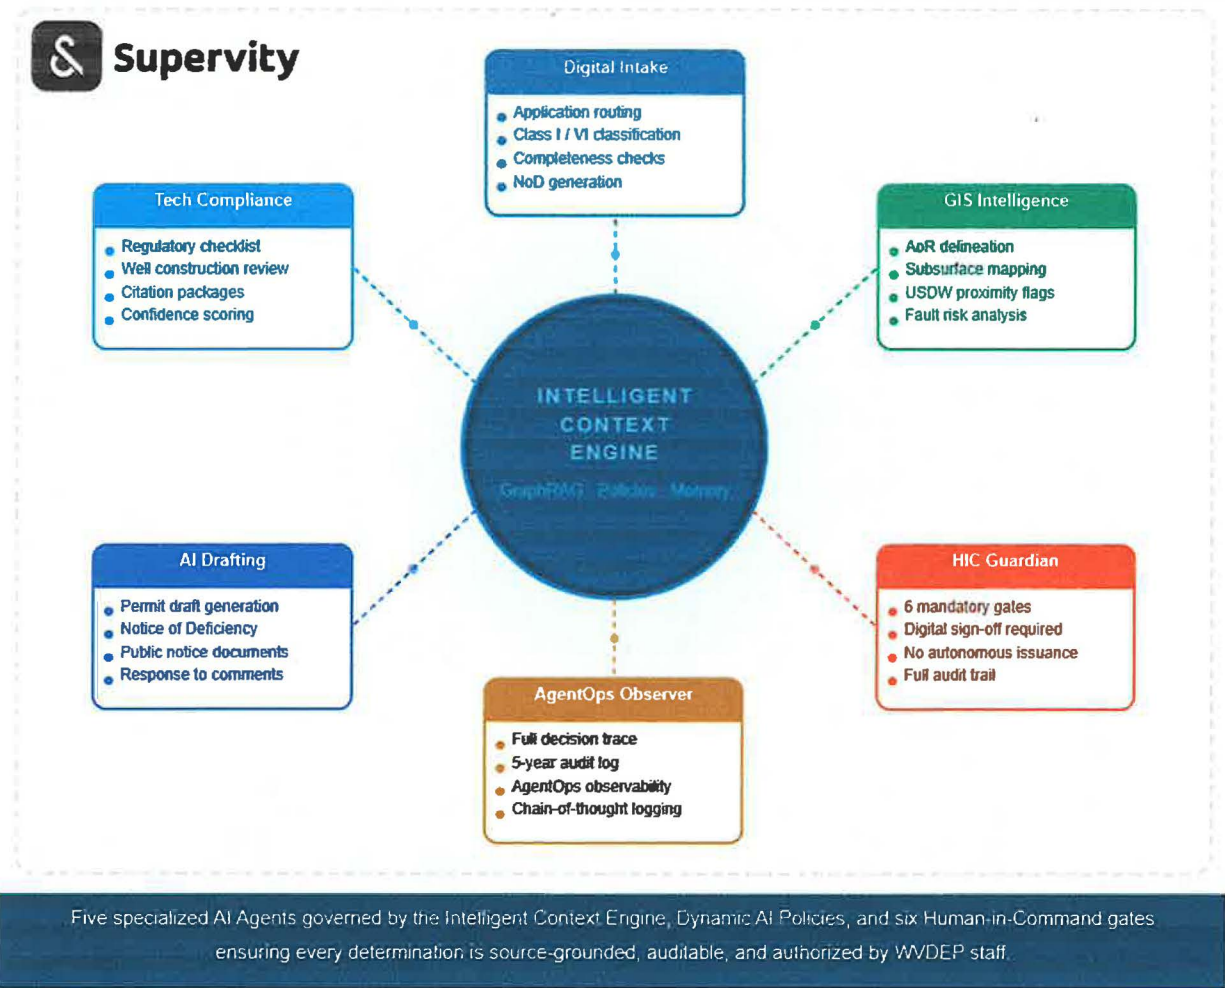


Image 3. Agentic AI Processing for WVDEP

2.2.1 Digital Intake Specialist Functionality (§4.3.2.2.1)

The Digital Intake Specialist is the first Operator within the UIC Permitting AI Employee. It converts application submissions into structured permitting work items, classifies the application type, identifies missing or incomplete materials, and initiates the appropriate review workflow under WVDEP-defined Dynamic AI Policies, including the following components:

Agentic Routing and Sub-Workflow Orchestration (§4.3.2.2.1.1): Upon receipt, the Classification Agent analyzes submitted forms, facility descriptions, and injection zone characterization to route the application to the appropriate Class I or Class VI workflow. Within each primary workflow, Agentic AI processing dynamically spawns sub-agents for:

- Area of Review (AoR) analysis with fixed radius (1/4 mile) for Class I; for Class VI, accounting and validation of modeling parameters directly stated within the permit application, consistent with the EPA reference tool and applicable WV regulations. Active reservoir simulation or plume modeling is not performed by the AI system; advanced GIS analysis is conducted by WVDEP reviewers in ESRI tools.

- Financial responsibility verification against 40 CFR 144.52 and 146.85 requirements.
- Corrective action planning identifying artificial penetrations within the AoR.
- Injection well construction standards validation against WV and federal requirements.
- PISC, TMP, and CO₂ stream specification review (Class VI only).

Administrative Completeness Review (§4.3.2.2.1.2): The Administrative Completeness Agent confirms presence and adequacy of all required elements using a digitized regulatory checklist.

- **For Class I:** project plans, site maps, AoR calculations, EPA Form 7520-6, facility location and legal description, operator identification, well construction details, injection zone identification, confining zone characterization, proposed injection rates/volumes/pressures, and injection fluid characterization.
- **For Class VI:** detailed geologic characterization, TMP (40 CFR 146.90), Injection Well Plugging Plan, PISC Plan (40 CFR 146.93), Emergency and Remedial Response Plan, and Financial Responsibility documentation. Computer Vision agents validate signatures by Responsible Corporate Officers and mandatory certification statements.

Technical Compliance Review (§4.3.2.2.1.3): The Technical Compliance Operator evaluates injection zone properties, formation pressure, confining unit integrity, well construction details, corrective action requirements, monitoring plans, and other technical evidence, then generates a review package for WVDEP technical staff approval.

2.2.2 Intelligent Context Engine, GraphRAG and Source Grounding (§4.3.2.2.2)

Supervity's Intelligent Context Engine combines GraphRAG, regulatory knowledge graphs, source-grounded retrieval, Dynamic AI Policies, and human reviewer feedback into one governed context layer, allowing the UIC Permitting AI Employee to reason across application materials, regulatory requirements, geospatial records, and WVDEP-specific review practices without relying on disconnected retrieval or generic prompting.

The knowledge graph is populated from WV Class I and Class VI regulations (47 CSR 13, 47 CSR 64, 40 CFR Part 146), WVDEP-provided SOPs, federal EPA guidance, WVGES geological survey data, EPA SDWIS data, WV Office of Oil and Gas records, GIS layers, and submitted application materials. As Class I and Class VI are new program types in West Virginia, training data will consist of regulatory documents and example materials provided by WVDEP.

Supervity prevents reliance on outdated information through temporal versioning, source confidence scoring, and policy-version tracking. When source material is missing, outdated, or ambiguous, the AI Employee routes the issue to Human-in-Command review rather than inferring unsupported conclusions.

Dynamic AI Policies for WVDEP Regulatory Governance

WVDEP's regulatory and procedural requirements will be encoded as Dynamic AI Policies inside the UIC Permitting AI Employee Command Center. These policies define how the AI Employee evaluates completeness, applies Class I and Class VI review logic, routes deficiency items, escalates technical risks, applies citation requirements, triggers HIC gates, and determines when human review is mandatory.

Dynamic AI Policies allow authorized WVDEP users to govern the AI Employee in plain language while maintaining version control, auditability, and approval controls. When regulations, internal procedures, review thresholds, or agency preferences change, those updates can be reviewed and incorporated into policy logic without requiring a full rebuild of the workflow. Every policy evaluation is logged as part of the audit trail, including the policy version applied, the source evidence considered, the recommended action, and the human reviewer's final decision.

2.2.3 Grounded Regulatory Execution and Hallucination Mitigation (§4.3.2.2.3)

Supervity mitigates hallucination risk by designing the UIC Permitting AI Employee to operate as a source-grounded regulatory assistant, not an open-ended generative system. The AI Employee may only produce findings, deficiency recommendations, draft language, or review packages when those outputs are grounded in retrieved source material, extracted application data, approved WVDEP policy logic, or applicable regulatory text. Where evidence is missing or ambiguous, the system flags the item and routes it to Human-in-Command review rather than inferring unsupported conclusions through the following controls:

1. **Source-grounding enforcement** - all AI-supported determinations must be traced to a retrieved document, database record, application field, GIS layer, or approved regulatory source; gaps are flagged as deficiencies or review items rather than inferred.
2. **Independent verification Operator** - re-queries the knowledge base and source package to confirm key factual claims before presenting findings to WVDEP reviewers.
3. **Citation and confidence package** - all review packages include source citations, extracted fields, confidence scores, applicable policy logic, and recommended next action.
4. **Human-in-Command gates** - authorized WVDEP reviewers make all consequential determinations, including deficiency issuance, AoR/risk findings, draft permit approval, and final decision support.
5. **Adversarial validation and monitoring** - validation scenarios, WV-specific geological/regulatory test cases, and reviewer feedback are used to identify, monitor, and improve systematic failure modes.

2.2.4 Automated and Continuous AI Validation (§4.3.2.2.4)

Supervity's AI Safety and Validation Module operates continuously, enforcing prompt-injection guardrails that prevent applicant-submitted content from altering system-level instructions, output filters that block unsupported regulatory determinations, and reviewer-feedback loops that improve policy handling and exception routing over time.

Annual third-party adversarial validation exercises are contractually committed, with findings shared with WVDEP and remediation timelines agreed in writing.

2.2.5 Citations and Explainability (§4.3.2.2.5)

Every AI-supported determination includes a citation panel with source documents, page numbers or record references, applicable regulatory provisions, confidence scores, policies evaluated, and reviewer disposition. These determinations are presented as review packages for authorized WVDEP staff rather than autonomous regulatory decisions.

2.3 Geospatial Analysis and GIS Integration (§4.3.2.3)

2.3.1 Automated Risk Assessment and AoR (§4.3.2.3.1)

The AI Employee ingests GIS data in standard formats (Shapefile, GeoJSON, KML, and ESRI File Geodatabase), reprojects all layers to NAD83/UTM Zone 17N to align with West Virginia state standards, and performs geometry validation before analysis proceeds. CAD/DWG drawings are routed to Computer Vision agents described in §2.4.1 rather than treated as native GIS data.

For Class I wells, the AI Employee applies the fixed 1/4-mile radius AoR buffer and identifies all artificial penetrations, faults and fractures, USDWs, and sensitive receptors within the zone. For Class VI wells, the AI Employee extracts and validates modeling parameters from the application package, validates input boundary conditions against stated assumptions, and supports 2D map visualization of injection zones and confining units. Reservoir simulation, plume migration modeling, and complex subsurface analysis are performed by WVDEP reviewers in their existing ESRI desktop and Enterprise environment. The AI Employee prepares the geospatial review package, ingests reviewer ESRI outputs back into the permitting workflow, and records reviewer disposition in the audit trail. WVDEP will provide baseline GIS and stratigraphic mapping data for system access. Automated risk flags are generated for artificial penetrations requiring corrective action, faults or fractures serving as potential fluid migration pathways, USDW proximity within the Zone of Endangering Influence, and surface water body proximity.

2.3.2 Data Integration and Correlation (§4.3.2.3.2)

The geospatial context layer correlates applicant data with authoritative external sources including:

- **Property and Plats:** Surface ownership tract overlays and digital courthouse plat references to verify legal rights and pore space ownership.
- **WVGES:** Scheduled ingestion via published data services for geological formation data, well logs, and subsurface characterization, subject to confirmed access patterns established during discovery.
- **EPA SDWIS:** Scheduled synchronization for USDW identification and public water system proximity analysis.

- **WV Office of Oil and Gas:** Well record retrieval for AoR artificial penetration identification.
- **WVDEP Systems:** Integration with OSSP for reporting and payment; OneLogin for external user authentication; Active Directory for internal staff authentication. WVDEP will provide technical documentation, data dictionaries, and API specifications during the discovery phase.
- **Conflict Detection:** Automated checking for overlapping existing injection permits within the AoR.

2.4 Document Processing and AI Drafting (§4.3.2.4)

2.4.1 Engineering Blueprint Vision Agents (§4.3.2.4.1)

Beyond standard OCR, the AI Employee uses Computer Vision agents trained in engineering schematics, well construction diagrams, and CAD/DWG files. These agents extract technical specifications such as casing depths, cement thickness, tubing packer placement, perforation intervals, and well construction geometries, mapping extracted values against Class I and Class VI construction standards for automated compliance flagging.

2.4.2 AI Draft Generation (§4.3.2.4.2)

The AI Drafting Operator generates formatted permit draft documents following WVDEP standard templates, including Facility Information, Well Construction Requirements, Operating Requirements, Monitoring and Reporting Requirements, Plugging and Abandonment Requirements, and General Conditions. Class VI drafts additionally include approved AoR with reevaluation schedule, corrective action requirements, CO₂ stream specifications, TMP conditions, Emergency Response requirements, and PISC requirements. All draft content requires human review and approval at HIC Gate 5 before finalization.

2.4.3 Completeness Determination (§4.3.2.4.3)

The Completeness Operator processes applications submitted through the WV OSSP system, checking each submission against the full regulatory checklist for the applicable permit class. Completeness determinations are presented to WVDEP staff at HIC Gate 1 for human verification before any action is taken.

2.4.4 Notice of Deficiency Generation (§4.3.2.4.4)

When an application is incomplete, the AI Employee generates a draft Notice of Deficiency specifying each missing or insufficient item with precise regulatory citations, routed to WVDEP staff for review and approval at HIC Gate 1 before any communication is transmitted to the applicant. The AI Employee maintains a NoD tracking log linking each deficiency item to the applicant's subsequent response.

2.4.5 Public Notice Document Generation (§4.3.2.4.5)

The AI Employee generates draft public notice documents compliant with 40 CFR Part 124 and West Virginia requirements, including fact sheets summarizing permit applications in plain language, notification recipient lists, and comment period deadline calculations (30 days minimum; 45 days for Class VI with automatic calendar computation and holiday adjustment).

2.4.6 Response to Comments (§4.3.2.4.6)

The Response to Comments Operator ingests and categorizes public comments, which may range from 50 to 1,000 per application. Draft responses are generated using a WVDEP regulatory response library seeded at go-live. Substantive technical comments are flagged for detailed human-authored responses. The AI Employee compiles a formatted Response to Comments document for public release, with all AI-drafted content subject to human review and approval before finalization.

2.5 Workflow Integration and HIC Architecture (§4.3.2.5)

2.5.1 Secure Submission Handling (§4.3.2.5.1)

The AI Employee retrieves applications via secure, authenticated API connections. Upon receipt, it automatically assigns a unique application tracking number (format: UIC-[CLASS]-[YEAR]-[SEQUENCE]) and creates a comprehensive electronic case file containing all submitted documents, extracted data, and a complete audit trail.

2.5.2 Agency Logs and Deep Observability (§4.3.2.5.2)

The Command Center maintains an adjacent agency log folder for each application, accessible to WVDEP reviewers through the dashboard. Logs capture every AI action, data retrieval, determination, and human interaction. Logs are stored in tamper-evident, append-only format and are exportable in CSV and JSON for the full 5-year retention period (or longer per WVDEP policy).

2.5.3 Decision Traceability and AgentOps Observability (§4.3.2.5.3)

Supervity's AgentOps Observability Framework captures a complete decision trace for every AI-supported workflow step. Rather than exposing internal model decision trace, the system records the externally auditable elements required for regulatory defensibility: source documents retrieved, data fields extracted, policies evaluated, regulatory citations applied, confidence scores, recommended actions, human reviewer decisions, and final workflow outcomes.

For example, where the AI Employee flags an AoR issue, the audit trail may show:

1. the application document and GIS layer reviewed;
2. the fault, well, USDW, or sensitive receptor identified;
3. the measured distance or extracted technical parameter;
4. the applicable WV or federal regulatory requirement;
5. the reason the item was routed to human review; and
6. the reviewer's final approval, rejection, or request for additional analysis.

This approach provides WVDEP with legal defensibility, transparency, and auditability without exposing non-auditable internal model reasoning. Every decision package remains tied to evidence, policy, citation, confidence, and human authorization.

2.5.4 Mandatory HIC Decision Gates (§4.3.2.5.4)

The UIC Permitting AI Employee Command Center enforces six mandatory Human-in-Command stop points, each requiring digital sign-off by authorized WVDEP personnel before workflow progression. The AI Employee cannot autonomously issue, deny, transmit, or finalize any regulatory action.

These HIC gates are enforced as Dynamic AI Policies within the Command Center. The AI Employee cannot proceed, transmit determinations, issue notices, or finalize permit materials without WVDEP reviewer approval. Each gate records the AI recommendation, source evidence, policies evaluated, reviewer decision, timestamp, and final disposition.

#	Human-in-Command Milestone	Trigger Conditions	Supervity Required Actions
1	Pre-NoD Review (Administrative Completeness)	Application intake complete; all required documents checked	Supervity Workflow Engine suspends pipeline; reviewer verifies missing documents via HIC console before NoD is finalized or transmitted
2	Administrative Compliance Approval	Fees verified; ownership validated; public notice requirements identified	Supervisor reviews AI findings; approves compliance or issues deficiency notice; initiates public notice process
3	AoR and Risk Validation (Technical Screening)	Initial technical parameters extracted; AoR delineated; risk assessment complete	Human review of AI-generated AoR and fault analysis; determines expedited vs. standard review track
4	Technical Analysis Approval	Full technical review complete; corrective action needs identified	Senior technical reviewer validates findings; approves permit conditions; may request additional analysis
5	Draft Permit Approval	Draft permit document generated with all conditions	Human review of AI-drafted permit conditions before document is locked for public comment
6	Final Decision	Public comment period complete; Response to Comments prepared	Human administrators execute the final Issue or Deny command; the AI cannot autonomously issue a permit. All HIC gates include mandatory override capability, with reviewer decisions captured in the audit trail and fed back into the Dynamic AI Policy refinement cycle.

2.6 AI Token Usage and Cost Management Approach (§4.3.2.6)

Supervity's token management approach is designed to give WVDEP full transparency, predictability, and control over AI inference costs. The UIC Permitting AI Employee uses deterministic workflow logic where appropriate and reserves model calls for tasks requiring document understanding, regulatory reasoning, summarization, extraction, or source-grounded drafting. Detailed cost figures are provided in the separately sealed Cost Proposal.

Token Optimization: Supervity uses semantic chunking, document preprocessing, prompt compression via domain-specific templates, caching of frequently accessed regulatory context, and targeted model routing to reduce unnecessary LLM calls across workflow steps.

Model Flexibility: Supervity's multi-model architecture supports evolving AI capabilities without application redesign. Model updates are deployed under governance controls and regression-tested against WVDEP-approved scenarios before production use.

Supervity commits to fixed per-MTok pricing for the initial 5-year term, with Year 6 renewal escalation capped at the prior year's CPI index.

2.7 HIC Workflow Interface and Legacy System Independence (§4.3.2.7)

2.7.1 Standalone HIC Workflow Interface (§4.3.2.7.1)

The UIC Permitting Command Center is a self-contained, browser-based AI Workbench where all AI-to-human and human-to-AI interactions occur. Reviewers access AI-generated packages, review decision traces, source citations, extracted data, policy evaluations, confidence scores, and recommended actions, then approve, reject, or request revision.

2.7.2 AI-Generated Request Workflow (§4.3.2.7.2)

As each HIC gate is reached, the AI Employee generates a structured HIC Request Package containing a plain-language summary of findings, source citations and confidence scores, recommended action options, required reviewer action, and downstream legacy-system update instructions.

2.7.3 Human-Mediated Legacy System Updates (§4.3.2.7.3)

Upon completing a HIC review, the Command Center generates a structured Export Package containing reviewer-approved status update summaries for manual staff entry into agency systems, AI-generated draft documents in standard formats for WVDEP records retention, reviewer-approved deficiency or permit materials, and audit-log references for the agency record. Export packages are formatted for integration with the WV OSSP system for reporting and payment as required.

2.7.4 Data Flow Architecture (§4.3.2.7.4)

Supervity's UIC Permitting AI Employee Command Center is purpose-built for the regulatory complexity, document intensity, and auditability requirements of Underground Injection Control Class I and Class VI permitting under the Safe Drinking Water Act. The solution is built on three core design principles.

1. Every AI action is source-grounded; outputs are only generated when tied to retrieved application data, regulatory text, or WVDEP-approved policy logic.
2. Every consequential determination requires human authorization; six Human-in-Command gates ensure WVDEP staff, not the AI, make all critical decisions.
3. The system operates alongside WVDEP's existing infrastructure without replacing legacy systems, integrating directly with OSSP, OneLogin, and Active Directory.

The platform's Intelligent Context Engine unifies regulatory knowledge, geospatial records, and live application data into a governed context layer, giving the AI Employee institutional-grade awareness while preserving transparency, traceability, and auditability for regulatory defensibility.



AI Employee Command Center Data Flow Architecture

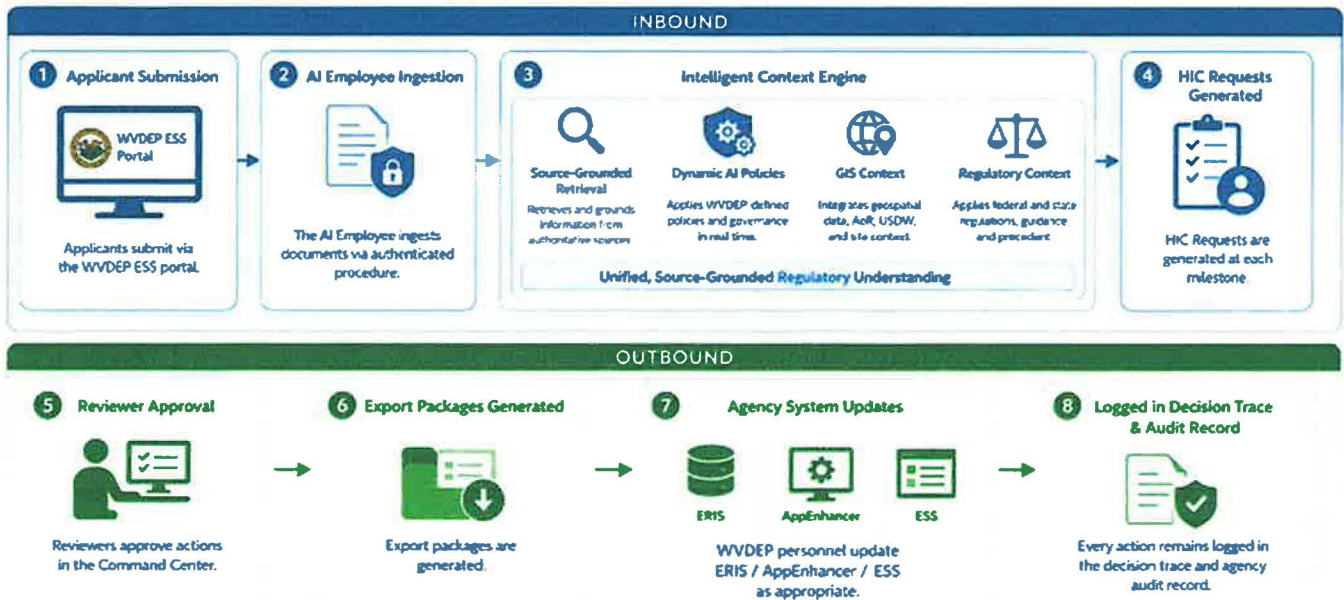


Image 4. Command Center Inbound and Outbound Data Flow

Section 3: Mandatory Project Requirements Compliance (§4.3.3)

Supervity will validate final compliance language during contract negotiations and security review. Where the RFP requires evidence of certification, authorization, insurance, audit reports, or security controls, Supervity will provide supporting documentation through the State’s designated procurement and security process. Any deployment-specific control inheritance from cloud infrastructure, third-party services, or agency systems will be documented in the implementation security package.

The table below summarizes Supervity’s intended compliance approach for each mandatory project requirement identified in §4.3.3 of the RFP.

Mandatory Requirement	Status	Supervity Approach
FedRAMP Moderate Authorization	✓ Met	Deployment on AWS GovCloud (US) or Azure Government (both FedRAMP High authorized). FedRAMP authorization documentation from the cloud platform provider is accepted per Addendum #1, Q57-58. Entire application stack, including all AI model inference, operates within the FedRAMP boundary per Q22 and Q26.
NIST 800-53 Compliance	✓ Met	Supervity will provide applicable control mapping and deployment-specific documentation, including inherited cloud controls where applicable.
TLS 1.3 Encryption in Transit	✓ Met	All API calls and browser sessions enforce TLS 1.3 minimum.
AES-256 Encryption at Rest	✓ Met	All stored data encrypted with AES-256; key management via AWS KMS.
RBAC and MFA	✓ Met	Role-based access aligned with WVDEP org structure; MFA enforced for all users.
SSO via SAML 2.0 / OpenID Connect	✓ Met	SSO integration with WVDEP identity management systems.
Section 508 Accessibility	✓ Met	All user-facing components conform to Section 508 requirements.
SOC 2 Type II Certification	✓ Met	Annual SOC 2 Type II audit; report provided to WVDEP; right-to-audit clause included.

Mandatory Requirement	Status	Supervity Approach
FedRAMP, StateRAMP, or SOC 2 Type II (§4.4.2.3)	✓ Met	Holds SOC 2 Type II certification.
Annual 3rd-Party Penetration Testing	✓ Met	Annual pen test results shared with WVDEP; critical vulnerabilities patched within 30 days.
Data Residency – Continental United States	✓ Met	All data stored and processed exclusively within CONUS AWS GovCloud or WVDEP-approved deployment environment regions.
AI Governance Documentation	✓ Met	Training data provenance; bias testing; model drift detection; human override at all decision points.
5-Year Audit Log Retention	✓ Met	All AI actions, HIC decisions, and workflow events retained for minimum 5 years.
Confidentiality Agreement	✓ Met	Signed upon contract award as required.
Data Ownership – State of West Virginia	✓ Met	All WVDEP data remains sole property of the State; not used to train models benefiting other customers.
30-Day Data Export on Termination	✓ Met	Complete data export in PDF/CSV/JSON/XML within 30 days; 90-day transition assistance; certified data destruction.
99.9% Platform Uptime SLA	✓ Met	SLA contractually guaranteed excluding scheduled maintenance windows.

3.1 Data Integration and Regulatory Compliance (§4.3.3.1)

From a format support perspective, the AI Employee can ingest XML, CSV, PDF (including scanned/OCR), GIS layers (Shapefile, GeoJSON, KML, ESRI File Geodatabase), CAD/DWG drawings (processed via Computer Vision agents per §2.4.1), well logs (LAS, DLIS), and standard UIC permit forms, subject to final integration and sample-data validation during implementation.

With the Intelligent Context Engine and Dynamic AI Policies, Supervity encodes the relevant Class I and Class VI regulatory framework as governed review logic that can be versioned, audited, and refined with WVDEP-approved changes.

Confirmed system integrations include WV OSSP (reporting and payment), OneLogin (external user authentication), Active Directory (internal staff authentication), WVGES (geological data API), EPA SDWIS (USDW records), and WV Office of Oil and Gas (well records via public API). AppEnhancer integration is not in scope. WVDEP will provide technical documentation, data dictionaries, and API/interface specifications during discovery. All integrations will use WVDEP-approved authentication and data-access patterns. At oral presentations, Supervity will demonstrate representative workflows using sample or synthetic data as permitted by Addendum #1, Q73.

3.2 Security and Deployment (§4.3.3.2)

Supervity's security posture is designed to align with the §4.3.3.2 requirements and will be documented through the State's security review process:

- TLS 1.3 encryption in transit
- AES-256 encryption at rest via AWS KMS
- RBAC aligned to WVDEP organizational structure
- MFA enforced for all users
- PII and CBI identified at ingestion using NLP classifiers and stored in isolated, access-controlled vaults
- Deployment in AWS GovCloud (US) or Azure Government, both holding FedRAMP High authorization, satisfying the FedRAMP Moderate hosting requirement. FedRAMP authorization documentation from the SaaS platform provider (AWS or Microsoft) is accepted by WVDEP per Addendum #1, Q57. The entire application, including all agentic workflows, data processing tasks, and all model inference, resides and operates entirely within the FedRAMP Moderate boundary. No processing or data transmission occurs through external commercial API endpoints outside the FedRAMP boundary.

- SSO integration via SAML 2.0 and OpenID Connect
- Annual third-party penetration testing with results shared with WVDEP and critical vulnerabilities patched within 30 days

3.3 Support and Maintenance (§4.3.3.3)

Support commences only upon written System Acceptance by WVDEP following successful demonstration and training. Supervity provides technical support during standard business hours with escalation paths for critical issues. A Dedicated Technical Account Manager will be assigned exclusively to WVDEP with monthly service review meetings and quarterly business reviews. Supervity provides a stabilization warranty covering all AI agent breakages caused by minor updates or environment changes at no additional cost along with continuous platform monitoring with automated incident alerting.

3.4 Licensing (§4.3.3.4)

Supervity's UIC Permitting AI Employee Command Center is delivered exclusively as a Software-as-a-Service subscription, with no perpetual license fees, no per-seat charges for individual user accounts, and no additional licensing costs tied to application volume or workflow complexity. The subscription covers unlimited named user access for all WVDEP personnel required to interact with the system in any capacity, including technical reviewers, administrative staff, supervisors, system administrators, and any designated agency personnel responsible for HIC gate approvals, audit log review, or policy governance. Supervity's subscription model is inclusive by design. All platform capabilities, updates, security patches, model governance controls, AgentOps observability features, Dynamic AI Policy management tools, and the Human-in-Command workflow interface are available to all licensed users under the base subscription without additional module fees or tiered access restrictions. Licensing will be provisioned and maintained for the full duration of the initial five-year contract term and any exercised renewal period, with no lapse in access during transitions between contract years. All subscription pricing is provided in the separately sealed Cost Proposal in accordance with Section 5.3.1 of the solicitation.

3.5 Regulatory Compliance and AI Governance (§4.3.3.5)

Supervity will provide available SOC 2 Type II audit documentation, security control documentation, and AI governance materials to WVDEP through the designated security review process. Section 508 compliance is verified through automated accessibility testing on all user-facing components prior to each release.

3.6 Data Ownership and Exit Strategy (§4.3.3.6)

All WVDEP data remains the sole property of the State of West Virginia. Supervity will not use WVDEP data to train AI models benefiting other customers without explicit written consent. Supervity's model selection approach uses commercial and/or open-source models that are not prohibited by the State of West Virginia, deployed and operated entirely within the required FedRAMP Moderate boundary. The selected model architecture will be approved by WVDEP/WVOT prior to production deployment, consistent with Addendum #1, Q1, Q21, and Q59. Data retention, export, deletion, and exit requirements will be finalized in the contract and implementation security package. Upon contract termination, Supervity will provide complete data export within 30 days in an agreed format, 90 days of transition assistance, and certified data destruction from vendor systems within 60 days of confirmed data transfer.

Section 4: Qualifications and Experience (§4.4)

4.1 Company Background (§4.4.1.1)

Supervity was founded with the mission of making enterprise-grade AI Employees accessible and governable for regulated industries and complex enterprise operations. Supervity's approach combines AI Employees, AI Command Centers, the Intelligent Context Engine, Dynamic AI Policies, Human-in-Command workflows, and source-grounded execution to support real operational work with auditability and control.

Supervity's core competencies directly relevant to this RFP include AI Employee and Command Center design, regulated workflow automation, document intelligence, source-grounded retrieval, geospatial data integration, Human-in-Command governance, decision traceability, and legacy-system-adjacent execution.

4.2 Relevant Agentic AI Experience and References (§4.4.1.2)

Supervity has deployed enterprise AI Employees and Command Centers across regulated, audit-sensitive environments, including public sector, defense, healthcare, finance, shared services, and enterprise operations. For WVDEP, Supervity will apply the same governed AI Employee model with Human-in-Command review, source-grounded execution, auditability, and policy-driven control to the UIC permitting workflow. Reference contact information will be provided in the required reference attachment or directly to the State upon request, subject to customer approval and confidentiality requirements.

Project	Description and Outcomes	Reference Contact
West Virginia DEP Application Processing	Supervity deployed an Application Processing AI Employee Command Center, hosted on the Azure Public Cloud, to automate application intake, data extraction, completeness checks, signature validation, rule-based assignment, exception handling, and audit visibility. Eliminated a 3-month backlog, increased straight-through processing from 30% to 75%, reduced manual data entry by 100%, and delivered an estimated 5 FTE equivalent productivity gain.	Todd Cooper 304-926-0499 x43855 Todd.Cooper@wv.gov
Daikin Accounts Payable Command Center	Supervity deployed an Accounts Payable AI Employee Command Center to automate invoice ingestion, extraction, validation, matching, exception handling, and audit tracking across global shared service centers. Decreased invoice cycle times by 88%, improved straight-through processing by 75%, reduced manual data entry by 92%, and achieved zero data-entry errors.	Rajiv Bali 91-98991-28368 rajiv.bali@daikinindia.com
NEC Sales Command Center	Supervity deployed a Sales AI Employee Command Center to automate sales research, quote generation, technician note interpretation, asset matching, and quoting policy enforcement. Reduced sales research from weeks to minutes, quote generation from days to minutes, improved quote accuracy to zero-error output, and accelerated deal cycles with faster, compliant quotes.	Ram Menghani 972-672-2052 Ram.menghani@necam.com

4.3 Key Personnel (§4.4.1.3)

The following personnel will be assigned to the WVDEP UIC project. All roles include personnel who contributed to the referenced experience projects.

Role	Responsibilities
Project Manager (Vikash Chandra)	Overall delivery accountability; primary Agency point-of-contact; PMP-certified
Lead Solutions Architect (Ramu Yarru)	System design; FedRAMP compliance; legacy integration strategy
AI/ML Engineering Lead (Vennela Lanka)	Agentic workflow design; LLM fine-tuning; RAG/GraphRAG implementation; hallucination mitigation
Security and FedRAMP Lead	FedRAMP Moderate ATO; NIST 800-53; penetration testing coordination; SOC 2 oversight
Data Engineer	Historical data ingestion; ETL pipelines; GIS layer normalization; OSSP, OneLogin, and Active Directory integration
Dedicated Technical Account Manager (Anusha P)	Ongoing Agency liaison; monthly/quarterly business reviews; SLA monitoring

Supervity commits that the key personnel listed above will remain assigned to this project for its duration. Any proposed substitution requires WVDEP's prior written approval. All personnel with access to WVDEP data are subject to criminal background checks compliant with W. Va. Code §15-2D-3.

Supervity recommends that WVDEP on staff GIS and Geoscience Specialists and Environmental Regulatory SMEs contribute to the construction of the Intelligent Context Engine, Orchestrated Workflows, user acceptance testing, and support.

Resumes for named key personnel listed above are provided in Appendix A: Key Personnel Resumes.

Section 5: Implementation Timeline (§4.6 / Attachment C)

The implementation plan below assumes timely access to WVDEP SMEs, system documentation, test environments, representative permit application samples, integration credentials, security review materials, and identified reviewers for UAT. Final sequencing may be refined during kickoff and BRD confirmation to ensure the scope, integration design, HIC gates, and acceptance criteria are aligned before production deployment.

Supervity proposes the following phased implementation timeline from contract award to production deployment, targeting a go-live date of July 1, 2027 as specified by WVDEP (Addendum #1, Q8). The total estimated duration is 36 weeks (approximately 9 months), followed by 30 days of hypercare. Full system functionality across all agents will be confirmed prior to go-live; a phased production rollout of individual agents is not proposed (Q29). Final timeline will be validated during kickoff and BRD confirmation based on integration access, sample data availability, security review, and WVDEP SME participation.

Phase	Name	Duration	Key Activities
Phase 1	Discovery and Configuration	Weeks 1 - 6	Requirements validation, tenant setup, OSSP integration design, OneLogin and Active Directory SSO configuration. WVDEP user role mapping and Dynamic AI Policy baseline configuration. Go-live target: July 1, 2027.
Phase 2	Integration Development	Weeks 7-16	API integration with OSSP, OneLogin, and Active Directory; WVGES and WV Office of Oil and Gas data source connections; WV-specific regulatory rule configuration; GIS layer ingestion; and AI model training on WVDEP-provided regulatory documents, SOPs, and applicable federal guidance. Priority focus: Digital Intake Specialist agent configuration for administrative review.
Phase 3	Data Migration and Testing	Weeks 17 - 24	Knowledge graph population using WVDEP-provided regulatory documents, SOPs, federal EPA guidance, WVGES data, and available example materials (no historical Class I/VI UIC records exist in WV as these are new program types). User Acceptance Testing (UAT), security and penetration testing, and HITL workflow validation with WVDEP reviewers.
Phase 4	Pilot Deployment	Weeks 25 - 30	Pilot launch prioritizing the Digital Intake Specialist agent for administrative review as the primary production agent, consistent with WVDEP's stated priority (Addendum #1, Q24). User training for WVDEP staff, HIC gate validation, and system refinement based on reviewer feedback. Full system functionality confirmed prior to go-live per Q29.
Phase 5	Production Launch	Weeks 31 - 36	Full production deployment, go-live support, 30-day hypercare stabilization period, SLA monitoring activation, and AgentOps observability dashboard handover to WVDEP.

Note: Timeline assumes timely access to WVDEP system credentials, API documentation for OSSP, OneLogin, and Active Directory, and availability of WVDEP subject matter experts for UAT and HITL validation sessions. Target production go-live: July 1, 2027. Phase durations may be adjusted by mutual written agreement. All milestone completions are subject to WVDEP sign-off.

Section 6: Mandatory Compliance Summary

Supervity confirms compliance with all mandatory requirements and qualification requirements of this RFP:

- §4.4.2.1 - Compliance with applicable data privacy, cybersecurity, and AI governance procedures: **CONFIRMED**, subject to final security documentation and control mapping during procurement and implementation review.

- §4.4.2.2 - Willingness to sign confidentiality agreement upon contract award: CONFIRMED.
- §4.4.2.3 - Current FedRAMP, StateRAMP, or SOC 2 Type II certification: Supervity holds SOC 2 Type II certification and will provide available certification and audit documentation through the State's required process.
- §4.3.3.2.4 - FedRAMP Moderate authorized cloud environment: SUPPORTED THROUGH AWS GovCloud or another WVDEP-approved environment, with control inheritance and supporting documentation to be provided during security review.
- §4.3.3.5.2 - NIST 800-53 compliance: ALIGNMENT TO BE DOCUMENTED through the implementation security package and applicable cloud-control inheritance.
- §4.3.3.5.6 - Annual SOC 2 Type II audit with right-to-audit clause: CONFIRMED WHERE APPLICABLE, with supporting report access and contractual audit language to be handled through the designated procurement/security process.
- §4.3.3.6 - Data ownership, US data residency, and exit strategy: CONFIRMED.
- Israel Boycott Prohibition (W. Va. Code §5A-3-63): CONFIRMED. Supervity does not and will not engage in any boycott of Israel during the contract term.
- No Debt Certification: Supervity confirms that neither the Vendor nor any related party owes a debt to the State of West Virginia or any political subdivision as defined under W. Va. Code §5A-3-10a.

Conclusion

Supervity's UIC Permitting AI Employee Command Center represents the only purpose-built, Human-in-Command governed solution designed from the ground up for the precise regulatory demands of UIC Class I and Class VI permitting under the Safe Drinking Water Act. Where competing platforms offer general-purpose AI tools adapted for government use, Supervity delivers a system whose architecture, from the Intelligent Context Engine and GraphRAG knowledge layer to the six mandatory HIC gates and AgentOps audit trail, was engineered specifically for the document intensity, source-grounding requirements, and human accountability standards that WVDEP's permitting mission demands. Critically, Supervity has already demonstrated this capability in practice: the reference engagements detailed in Section 4 include prior AI-assisted application processing work performed directly for WVDEP, establishing a proven operational track record in West Virginia's regulatory environment that no competing vendor can claim. The platform's three-layer token optimization architecture holds projected annual consumption to 94 MTok against a 250 MTok contractual ceiling, delivering cost predictability that open-ended AI billing models cannot match.

WVDEP deserves an AI partner that treats human authority as a design requirement, not an afterthought. Every finding the Supervity AI Employee produces traces to a source document, a regulatory citation, or an approved WVDEP policy rule. Every consequential determination waits for an authorized WVDEP reviewer to act on it. Every action the system takes is logged, time-stamped, and exportable in a format that survives legal challenge, regulatory audit, and public records request. With a committed go-live target of July 1, 2027, a phased implementation that prioritizes the Digital Intake Specialist for early production value, and pricing that is firm-fixed for the full five-year term, Supervity offers WVDEP not just a technology solution but a long-term operational partnership - one governed by the same principles of transparency, accountability, and regulatory defensibility that define WVDEP's mission.

Certification and Signature Page

By signing below, I certify that I have reviewed this Request for Proposal in its entirety. I understand the requirements, terms and conditions, and other information contained herein. I am submitting this proposal for review and consideration, and I am authorized by the bidder to execute this bid on the bidder's behalf. I am authorized to bind the bidder in a contractual relationship, and that, to the best of my knowledge, the bidder has properly registered with any State agency that may require registration.

I further certify that I understand this Contract is subject to the provisions of West Virginia Code §5A-3-62, which automatically voids certain contract clauses that violate State law, and that pursuant to W. Va. Code §5A-3-63, the entity entering into this contract is prohibited from engaging in a boycott against Israel.

Company: Supervity



Signature of Authorized Representative:

Printed Name and Title: Brett Fraser, Vice President of Public Sector

Date: 6/10/2026

Phone Number: 571-524-9125

Fax Number: N/A

Email Address: brett@supervity.ai

FEIN: 82-0687505

Designated Contract Administrator

Printed Name and Title: Siva Moduga

Address: 11921 Freedom Drive, Ste. 570, Reston, VA, 20190

Phone / Fax: 337-764-3131

Email: siva@supervity.ai

Addendum Acknowledgement Form

Solicitation No.: CRFP 0313 DEP2600000003

- Addendum No. 1 - Extension of close date to June 10, 2026; publication of agency responses to vendor questions
- 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. Only information issued in writing and added to the specifications by an official addendum is binding.

Company: Supervity



Authorized Signature:

Date: June 10, 2026

Appendix A: Key Personnel Resumes

Project Manager – Vikash Chandra



VIKASH CHANDRA

Enterprise AI Engagement Leader · Strategic Account Growth · Platform GTM

20+ years · Mumbai · +91 9158937256 · vchandra@supervity.ai

Vikash brings together the strategic clarity of a business leader and the execution depth of a hands-on program owner. He is trusted by CXOs to lead the most complex enterprise engagements — navigating organisational dynamics, aligning multi-stakeholder teams, and converting ambitious AI roadmaps into real, measurable business outcomes.

Across two decades, he has operated at the intersection of technology transformation, commercial growth, and operational scale — in roles spanning enterprise consulting, platform GTM, AI deployment, and P&L leadership. His instinct is to build for lasting impact: scalable systems, trusted relationships, and replicable playbooks.

20+

Years of Enterprise Leadership

5x

Business scaled in 18 months

50+

Enterprise improvement mandates delivered

10+

End-to-end transformations

CURRENT ENGAGEMENT

Adani Group Enterprise AI Deployment *via Supervity Agentic AI Platform* — 2024 – Present

Leading end-to-end program ownership for one of India's most ambitious enterprise AI deployments — multi-module Agentic AI agents rolled out across Adani's business functions. Vikash owns CXO-level relationships, delivery governance, adoption strategy, and commercial expansion, with a 4-module roadmap targeting ~\$1M ARR.

WHAT VIKASH DOES BEST

Enterprise Program Ownership

AI Platform GTM & Adoption

CXO & Stakeholder Management

Account Growth & Expansion

Cross-Functional Orchestration

Scaling Teams & Operating Models

WHERE HE HAS OPERATED

Cutshort.io

COO — Built and scaled an AI-powered tech hiring marketplace (3M candidates, 18K recruiters) to \$1M ARR. Built multi-channel GTM engines and a leadership team capable of reaching \$10M.

Merkle Sokrati

SVP — P&L owner for a 300+ person digital marketing organisation. Scaled operations 5x in 18 months, delivered across global brand accounts spanning search, social, and performance platforms.

Tata Sons – TBExG

Practice Consultant & Key Account Manager — Delivered 50+ process improvement projects across 6+ Tata Group companies, working directly with CXOs on strategy deployment and operational excellence. Programmes spanned Voice of Customer research that drove the relaunch of automobile brands, Key Account Management redesign for sales effectiveness, OPMS operational dashboards for a steel manufacturer, balanced scorecard deployment, and GTM re-segmentation that uncovered a segment growing at ~100% CAGR. TBEM Gold Assessor; trained 450+ senior leaders.

Infosys Technologies

Sr. Consultant, Enterprise Solutions — Led 10+ end-to-end CRM transformation programmes for major global telecom clients, advising on business strategy, regulatory compliance, and technology adoption. Led a large-scale UK telecom migration within strict regulatory timelines, earning the CIO 'Right First Time' award. Drove significant capability uplift — including a campaign management transformation that expanded campaign throughput from 2 per year to 7 per month.

XLRJ Jamshedpur (PGDGM) · AIM Manila (International Exchange) · JSSATE Noida (Mechanical Engineering) · ASQ Certified
 [CQIA & CMO/OE] · TBEM Gold Assessor · IFC Sustainability Leadership

Lead Solutions Architect – Rama Yarru



Ramu Yarru

SUMMARY

Forward Deployed AI Architect with 20 years of experience in delivering enterprise-scale Generative AI transformation for Fortune-tier clients across infrastructure, international services, and global enterprise sectors. Specialized in Voice AI, Agentic AI, and RAG on AWS & Azure, with end-to-end ownership from POC to production. Delivered 4,000+ daily automated interactions, 40%+ reduction in manual AP operations, and 90% faster legal research over 500K+ documents.

CORE SKILLS

Generative AI & LLMs: Agentic AI, RAG, Voice AI, Prompt Engineering, LangChain, LangGraph, LangSmith, OpenAI, AWS Bedrock, Speech AI (Whisper, Transcribe, Polly), OCR Workflows

Cloud & DevOps: AWS (Bedrock, EC2, S3, Transcribe, Polly), Azure (App Services, DevOps), GCP (Cloud Run), Docker, Kubernetes, CI/CD

Backend & Data: Python, SQL, FastAPI, Django, REST APIs, Microservices, PostgreSQL, MySQL, FAISS, Milvus

Delivery & Engineering: System Design, Evaluation Frameworks, Performance & Cost Optimization, Monitoring, Stakeholder Management, Agile Delivery

PROFESSIONAL EXPERIENCE

Senior Director & Forward Deployed AI Architect @ Supervity, Hyderabad (Hybrid) Feb 2018 - Present

- **Enterprise Client Delivery:** Led end-to-end delivery of GenAI programs for Fortune-tier enterprise clients — primary technical point of contact across Finance, Logistics, and IT from discovery through production.
- **Agentic AI — Financial Operations Automation:** Delivered an AI-powered AP Command Center — OCR invoice processing, contract clause extraction, 3-way Invoice/PO/GRN matching, and SAP posting — cutting manual AP effort 40%+ and improving operational visibility 30%+ at enterprise GCC scale.
- **Voice AI Platform:** Architected a production Voice AI platform for a global visa & consular services enterprise, handling 4,000+ daily interactions via Asterisk (SIP/VoIP), AWS Bedrock, Transcribe, and Polly — sub-second latency, high availability.
- **Enterprise RAG Platform:** Built production RAG on AWS Bedrock, Titan Embeddings, LangChain/LangGraph, FAISS, and Milvus — retrieval accuracy +25%, latency -30% via chunking, prompt, and inference tuning.
- **Team Leadership:** Led a delivery pod of 16–20 engineers and interns, owning architecture reviews, onboarding standards, and GenAI mentorship.

PRINCIPAL ARCHITECT @ Supreme Netsoft Pvt Ltd | May 2017 to Feb 2018

- Business Process Consultant for enterprise portfolio optimization
- Architecture Recommendation and implementation

SENIOR SOLUTION ARCHITECT @ Tech Mahindra Ltd | Jul 2014 – May 2017

- Solution Architecture design using Oracle SOA, OSB and Oracle B2B technologies
- Provide response to RFPs, estimations and implementation plan

SENIOR CONSULTANT @ Oracle India Pvt Ltd | Dec 2010 – Jun 2014

- Project Lead & Integration Layer Architecture design and development using Oracle SOA, B2B and AIA technologies

SENIOR SOFTWARE ENGINEER @ Satyam Computer Services Ltd | Feb 2006 -Dec 2010

- Integrations development using Oracle SOA
- Web Application development using Oracle Portal with PL/SQL Portlets

ANALYST PROGRAMMER @ Syntel Ltd | Aug 2005 -Feb 2006

- Manual Testing

Lead Solutions Architect – Rama Yarru (Continued)



KEY ENGAGEMENTS & PROJECTS

Real-Time Voice AI Agent | Asterisk, Python, FastAPI, Generative AI, Docker, Kubernetes, AWS

- Production Voice AI with Asterisk SIP/VoIP and real-time STT - LLM - TTS pipeline & Speech to Speech models
- Sub-second latency, 4,000+ daily production calls, deployed on AWS EKS with horizontal scalability.

AI Accounts Payable (AP) Command Center | Generative AI, Python, FastAPI, Docker, Kubernetes, Azure Devops

- Enterprise finance transformation at GCC scale — contract clause extraction, OCR invoice processing, 3-way matching (Invoice/PO/GRN-SES), SAP posting, logistics gate-in, and deductions.
- Architected for enterprise scale via Azure DevOps CI/CD, Dockerized microservices, and Kubernetes orchestration.

Speech AI Analytics Tool | Python, Django, Generative AI, LangChain, LangSmith, Kubernetes, Azure Devops

- Multi-input (audio/video/text/cloud URL) platform with transcription, AI analysis, and multilingual translation — +20% accuracy via LangSmith observability.

EDUCATION

B. Tech from Jawahar Lal Nehru Technological University, Hyderabad

ACHIEVEMENT & CERTIFICATIONS

- Received Key Associate award in Tech Mahindra (Only top 5% of 35000 employees)
- Received Best Project Lead award in Satyam Computer Services Ltd in small projects category

AI/ML Engineering Lead – Vennela Lanka



Vennela Lanka

Lead / Project Executive

SUMMARY AND ACCOMPLISHMENTS

With over 11 years of professional experience, I am an AI Technical Lead bridging the gap between academic depth and industrial innovation. My 8-year IT tenure is defined by a diverse technical portfolio: 3 years developing Prompt-based AI Agents, 3 years specializing in RPA (Supervity.ai and Techforce.ai), and 2 years in application development. This is complemented by 3 years as an Assistant Professor, where I refined my ability to translate complex technical concepts into actionable strategies.

Currently leading AI and automation at Supervity.ai, I excel at managing end-to-end software lifecycles and mentoring high-performing teams. From my roots as an Automation Engineer to my current leadership role, I have a proven track record of delivering AI-driven workflows and chatbot solutions (Camunda, Dialogue Maker) that consistently achieve key project milestones across multiple industries.

Team Lead - AI FDE/ RPA Developer

- **Training s Knowledge Sharing:** Conducted training for internal teams, clients, and partners on Techforce.AI, Supervity.ai, and chatbot applications.
- **Built and Led RPA Team:** Established the RPA framework, managed production support, and developed automation solutions to streamline business processes.
- **End-to-End Automation:** Designed, developed, tested, and deployed RPA solutions for clients in Automobile, Finance, Accounting, BPO, HR Operations, and Insurance domains.
- **Technical Expertise:** Hands-on experience with Desktop, Web, Excel, JSON, PDF, and Email Automation, as well as OCR and Image Recognition.
- **Migration s Database Management:** Led environment migrations and worked with SQL databases and collections.
- **Testing s Deployment:** Conducted exception handling, UAT support, and production deployment using Techforce EM Console Orchestrator.
- **Chatbot Development:** Designed, tested, and deployed conversational AI chatbots with robust error handling using Camunda and Dialog maker.
- **Cross-Organizational AI Solutions:** Collaborated with multiple organizations to ideate, scope, and implement tailored Generative AI solutions.
- **Prompt-based AI Agents:** Developed and deployed Prompt-based AI Agents, specializing in crafting effective prompts and integrating them into business workflows to drive efficiency and automation.
- **Client Communication:** Acted as a bridge between developers and clients, efficiently handling functional and technical inquiries.
- **Project s Workflow Management:** Oversaw the entire RPA development lifecycle, including workflow design and process optimization.

AI/ML Engineering Lead – Vennela Lanka (Continued)



- **Leadership s Problem-Solving:** Managed and guided a team, resolving technical challenges and ensuring smooth project execution.
- **Process Optimization:** Identified automation opportunities and improved process efficiency, reducing manual efforts.
- **Cross-Platform Integration:** Integrated RPA solutions with third-party tools and enterprise applications.
- **Compliance s Security:** Ensured all automation solutions adhered to industry standards and security policies.
- **Continuous Learning:** Adapted to new technologies, upskilled the team, and implemented best practices.
- **Documentation s Reporting:** Maintained detailed project documentation, reports, and knowledge-sharing resources.

WORK EXPERIENCE

AI/RPA Tech Lead | Supervity.ai September 2015 – Current | Hyderabad

- Leading AI and RPA development using Supervity.ai platform with Camunda, Dialog maker, JavaScript, Python, and MySQL.
- Established RPA framework and managed production support across multiple business domains.
- Deployed automation solutions for Automobile, Finance, Accounting, BPO, HR Operations, and Insurance sectors.

Asst. Professor | Techforce.ai April 2015 – April 2018 | Korangi

- I served at KIET Engineering College, Korangi.
- Developed technical curriculum and conducted training programs.
- Contributed to technical strategy development and educational innovation.

Software Engineer | Achanta Ace Engr Pvt Ltd August 2013 – March 2015 | Kakinada

- Developed software applications and participated in full development lifecycle.
- Gained foundational experience in project management and technical problem-solving.

CORE COMPETENCIES

- **Built and Led RPA Team:** Established RPA frameworks and managed production support.
- **End-to-End Automation:** Designed and deployed solutions for Finance, HR, and BPO.
- **Prompt-based AI Agents:** Specialized in crafting effective prompts for business workflows.
- **Chatbot Development:** Deployed conversational AI using Camunda and Dialog maker.
- **Technical Expertise:** Hands-on experience with OCR, Image Recognition, and SQL.

AI/ML Engineering Lead – Vennela Lanka (Continued)



TECHNICAL SKILLS

- **Automation Tools:** Supervity.ai, Techforce.ai
 - **Chatbot Tools:** Camunda, Dialog Maker
 - **Programming Languages:** JavaScript, Basic Python, PHP
 - **Database:** Oracle DB, MySQL
 - **Web Applications:** AJAX, CSS, HTML, JSON, jQuery, Web Services
-

EDUCATION

- **M Tech - Computer Science | KIET College of Engineering (2017)**
- **B Tech - Computer Science s Engineering | NOVA Engineering College (2012)**

Dedicated Technical Account Manager – Anusha P



ANUSHA PAPANI

SENIOR TECHNICAL PRODUCT MANAGER

GENAI SAAS | API INTEGRATIONS | ENTERPRISE IMPLEMENTATIONS | CUSTOMER TO PRODUCT STRATEGY

SUMMARY

Senior Product and Customer Success professional with 13+ years of experience across GenAI SaaS, product management, enterprise implementations, web platforms, and software engineering. Experienced in translating customer needs into BRDs, user stories, workflows, dashboards, release feedback, and implementation plans. Strong technical exposure across APIs, integrations, QA/UAT, troubleshooting, data analysis, and deployment monitoring. Proven ability to work closely with engineering, customer success, and enterprise stakeholders to improve adoption, reduce operational friction, and deliver measurable product value.

WORK EXPERIENCE

Senior Technical Product Manager, Supervity.ai (Digitamize Inc)

Jan 2024 - Present

Supervity is a GenAI-powered AI Agent platform that enables enterprises to automate workflows, improve digital adoption, and drive operational efficiency.

Technical Product Management & Requirements

- Translate enterprise customer needs into BRDs, user stories, workflow diagrams, implementation documentation, and product requirements.
- Partner with engineering teams to define workflows, prioritise bugs, validate product behaviour, and improve release readiness.
- Gather structured feedback from customers and internal teams to identify product gaps, usability issues, and automation opportunities.
- Support product backlog discussions by escalating high-impact customer requirements, technical blockers, and adoption pain points.

Implementation, QA/UAT & Release Validation

- Conduct QA/UAT for new features and product enhancements to ensure releases meet customer use cases and business requirements.
- Monitor enterprise implementations, integrations, dashboards, and deployment progress to identify risks and reduce rollout delays.
- Define and improve implementation workflows to reduce deployment time and improve customer onboarding efficiency.

Customer Adoption & Product Value Realisation

- Manage 8 enterprise accounts, supporting onboarding, adoption, implementation monitoring, and long-term product value realisation.
- Introduce customers to Supervity's automation capabilities and help identify manual business processes suitable for AI-agent-led automation.
- Track product usage patterns, adoption trends, customer health, and feedback signals to identify opportunities for product and process improvement.

Data, Dashboards & Operational Improvements

- Build dashboards and reports to monitor customer health, usage trends, adoption metrics, implementation status, and enablement impact.

Dedicated Technical Account Manager – Anusha P (Continued)



- Analyse customer feedback and usage data to recommend product improvements, workflow changes, and adoption initiatives.
- Helped reduce ticket response time(-40%) and deployment timelines(-50%) by improving escalation processes, implementation tracking, and cross-functional communication.

Technical Product Manager, Supervity.ai (Digitamize Inc)

Strategy & Vision:

Jul 2020 - Dec 2023

- Defined and documented the product vision, outlining the target market, value proposition, and long-term roadmap.
- Conducted market research and competitive analysis to identify opportunities and white space for product differentiation.

Customer Research & Feedback:

- Conducted customer interviews, surveys, and usability testing to understand needs, pain points, and desired functionalities.
- Analysed the data and feedback to inform product roadmap prioritisation and feature development, increasing user satisfaction by 48%.

Product Development & Execution:

- Collaborated with cross-functional teams (engineering, design, marketing) to translate product vision into a cohesive and user-centric experience.
- Owned product backlog prioritisation, defining user stories and acceptance criteria for development.
- Managed product launches and go-to-market strategies, ensuring a smooth user experience and successful adoption.
- A/B tested features and functionalities to measure user response and optimise product design.

Website Manager, Mersey Eventures Pvt Ltd (teestory.in)

Jul 2016 - Jun 2020

Website Operations & Performance:

- Managed a team of 5 with web designers, developers and testers to troubleshoot technical issues and implement website improvements.
- Interpreted google analytics data to understand user behaviour and implemented fixes from coding enhancements to loading speed, increasing website conversion rate by 1.6% and user retention by 3x.
- Implemented A/B testing to determine customers top choices and suggested solutions to improve under performing products, surpassing revenue targets by 20%.

Marketing & Promotions:

- Managed website content (e.g., product listings, blog posts, banners) to promote products and sales events.
- Collaborated with marketing and merchandising teams on product launch strategies and promotional activities.

Additional Responsibilities:

- Ensured compliance with e-commerce regulations and data security standards.
- Analysed website traffic and user behavioural data to identify customer trends and inform

Dedicated Technical Account Manager – Anusha P (Continued)



strategic decision-making.

- Built and maintained merchandise partnerships with leading brands like Disney and Warner media to improve product catalogue and user retention.

Senior Systems Engineer, Infosys Ltd

Mar 2013 - Jun 2016

.NET Development & Architecture:

- Designed, developed, and implemented complex .NET applications for retail clients at Infosys.
- Possess strong expertise in core .NET technologies (C#, ASP.NET MVC/Web API, .NET Framework/Core).
- Implemented scalable and maintainable code adhering to best practices and design patterns.
- Experience with integrating .NET applications with various APIs and third-party systems relevant to the retail domain (e.g., payment gateways, inventory management systems).

EDUCATION

Bachelor of Technology

Aug 2008 - Jun 2012

Sreenidhi Institute of Science and Technology

- Branch: Electronics and Computer Engineering
- Final Project: Built an automatic resume generator for HR function in Tech Mahindra.

CERTIFICATIONS AND COURSES

Microsoft AI Product Manager Professional Certification

May 2026

Generative AI s Prompt Engineering , Udemy

Nov 2025

Scrum Certification , Udemy

Sep 2024

Product and Growth Program, GrowthX

Jan 2023

Runners up in GX Demo day 11 , presented growth strategies for [cult.fit](#) to a crowd of 2000+ product and growth enthusiasts. Built a strong product and growth mindset during the course duration and got to interact with an amazing community.

CORE SKILLS & TOOLS

- **Product Management:** Product Road mapping, BRDs, User Stories, Workflow Design, Backlog Prioritization, Release Testing, QA/UAT, Customer Feedback Loops
- **Technical Product:** APIs, Integrations, Implementation Monitoring, Troubleshooting, Dashboards, Data Analysis, SQL, Postman
- **Tools:** Jira, MS Planner, Notion, GitHub, HubSpot, Google Analytics, Miro, Figma, Whimsical, Canva
- **Domains:** GenAI SaaS, AI Agent Platforms, Automation, Enterprise Implementations, E-commerce, Retail Tech
- **Programming Languages:** C#, PHP, Python, HTML, JS, SQL, MySQL