

Agentic AI

Enabling Multi-Agent Solutions



Super Major Oil and Gas Company

Audit.AI DocIntelligence

Enhancing Finance Audit Efficiency with
AI-Powered Data Management



Challenges

- Difficulty managing vast volumes of heterogeneous data from multiple sources (financial records, contracts, material data).
- Manual audit checks were time-consuming and increased the risk of errors.
- Existing data extraction tools included irrelevant or redundant information, complicating analysis.
- Inconsistent data relevance and quality impacting decision-making.

Our Solution

- Developed and implemented DocIntel, an AI-powered audit automation framework using Azure OpenAI.
- Automated data cleaning and management for finance audits, handling both structured and unstructured data.
- Integrated the solution seamlessly with existing client systems for easy configuration and monitoring.
- Enabled scalable, high-availability operations using Azure's cloud infrastructure.

Tech Stack

Python, Azure SDK | Azure Cloud Platform | Blob Storage, Azure functions, Azure OpenAI services, Azure SQL DB, Azure AI search, Azure Key Vault, API Management, GitHub | Visual Studio code (for Python and Azure functions), Azure Machine Learning Studio (for ML model development) | Scikit Learn, Azure Cognitive Services API, Azure Logic apps

Outcome Delivered

Achieved **100%** automated document processing, eliminating the need for sampling and manual intervention

Increased audit coverage and accuracy, scaling from **25 to 100–1000** documents per audit

Improved data quality and completeness, supporting faster and more agile audit cycles

Freed up auditors to focus on insights and exceptions, strengthening risk detection and decision-making

US Energy Basin Operator

Smarter Hazard Analysis with Agentic AI

Adaptive AI framework delivering real-time compliance, safety, and cost efficiency

Challenges

- Ensuring comprehensive and accurate hazard identification remains a challenge.
- Integrating JHA processes into daily operations without disrupting workflows is difficult.
- Keeping JHA processes aligned with new hazards and regulatory changes requires continuous effort.

Our Solution

- The solution involves utilizing advanced AI models to analyze historical JHA data and identify potential gaps in new submissions.
- Implemented an intelligent, adaptive agentic AI framework to power the JHA Gap Analysis Engine.
- Generated detailed reports on identified gaps, historical trends, and compliance status.



Tech Stack

Azure OpenAI Service | Azure Data Factory | Agentic AI
| Azure Blob Storage | Azure DevOps

Outcome Delivered

Enhanced
safety and
compliance

Improved efficiency by
reduction in manual
effort and allowing
foremen to focus
on critical tasks

Significant cost savings with
reduced risks of accidents

Super Major Oil & Gas Company

Audit.AI – Material Recon Solution

Enhancing Material Mapping Accuracy

Challenges

- Difficulty managing & extracting data from inconsistent, un-structured, & diverse sources (vendor invoices, internal documents).
- Inconsistent naming, descriptions, & properties across documents make accurate data extraction and standardization challenging.
- Manual processes were time-consuming, error-prone, and resource-intensive.

Our Solution

- Developed an AI-powered solution for automated material mapping and reconciliation across diverse document sources.
- Automated detection of document types and templates, ensuring accurate extraction and mapping of materials.
- Generated comprehensive mapping reports, tagging materials and maintaining consistency across vendor and internal records.
- Provided an intuitive interface for auditors, simplifying review & approval.



Tech Stack

Programming: Python (ML models, Azure SDK, scripting), SQL | **Cloud and Data:** Azure Cloud Platform, Blob Storage, Functions, OpenAI Services, SQL Database, AI Search, Key Vault, API Management, GitHub | **Development Tools:** Visual Studio Code, Azure Machine Learning Studio | **Libraries & Integration:** Scikit Learn, Azure Cognitive Services API, Azure Logic Apps

Outcome Delivered

Automation enabled each auditor to save approximately six weeks per audit

Achieved annual savings of around **18** person-weeks previously spent on manual reconciliation and validation

Accelerated audit timelines and improved accuracy and consistency in data validation and reporting

Freed up auditor capacity to focus on deeper analysis and strategic insights

Audit.AI – Risk Analyzer

AI-Driven Risk Analysis for Automated Compliance



Challenges

- Manual risk categorization and rating were time-consuming, inconsistent, and dependent on individual expertise.
- Large and growing volumes of structured and unstructured data from diverse sources were difficult to process consistently, leading to inefficiencies, errors, and delays.
- Existing processes relied heavily on manual effort from compliance analysts to collate, aggregate, & verify data, making the approach inefficient, hard to scale, & unsustainable as data volumes increased.

Our Solution

- Built a reusable Gen AI-powered component for automated data profiling, categorization, and risk rating.
- Leveraged natural language processing and historical data mapping to assign risk categories and ratings dynamically.
- Integrated the solution with existing systems for scalable, efficient, and proactive risk management.
- Automated the process of defining categories by learning from past data patterns.

Tech Stack

Programming: Python (ML models, Azure SDK, scripting), SQL (database queries, data management) | **Cloud and Data:** Azure Cloud Platform, Blob Storage, Functions, OpenAI Services, SQL Database, AI Search, Key Vault, API Management, GitHub, Azure Machine Learning Studio | **Libraries and Integration:** Scikit Learn, Azure Cognitive Services API, Azure Logic Apps

Outcome Delivered

Automated risk scoring & calibration, reducing manual effort & turnaround time from days to minutes

Delivered consistent, data-driven risk assessments across all audit findings

Saved over **1,400** person-days annually, enabling teams to focus on strategic analysis and remediation

Enhanced audit quality, coverage, & agility without increasing headcount

Agentic AI-driven solutions for Energy

Driving precision, productivity, and operational excellence



Challenges

- Manual, rule-based production processes were limiting operational efficiency, reducing throughput, and negatively impacting Overall Equipment Effectiveness (OEE).
- High scrap rates and quality variability driven by limited, labor-dependent defect detection and inspection accuracy were increasing costs and reducing sustainability.
- Lack of integrated human-machine collaboration resulted in fragmented monitoring, slower diagnostics, & delayed issue resolution across operations.

Our Solution

- Virtual line and digital twin for end-to-end process modeling.
- Vision AI for accurate material picking and fewer errors.
- Real-time monitoring to optimize quality and reduce waste.
- Precision & predictive AI for better punching accuracy & early failure detection.
- Agentic AI for instant quality checks and corrective actions.
- Unified data analysis for smarter, faster decisions to improve productivity, quality, and efficiency.

Tech Stack

Agentic AI systems for process automation and quality control | Digital twin technology for virtual line simulation | Vision AI for material handling and inspection | Predictive analytics for equipment health and maintenance | Real-time data integration & monitoring platforms

Outcome Delivered

5,000
hours/year saved
through automation and
process optimization

~99%
quality and availability
with advanced
AI-driven monitoring
and control

70%
improvement in labor
productivity through
reduced manual effort
and errors

Over USD
2Mn
annual ROI from cost
savings, efficiency,
and improved output

Super Major Oil & Gas Company

Agentic AI for Smarter ServiceNow Ticket Management

Automate prioritization and routing, improving response times

Challenges

- Manual, inconsistent prioritization and categorization caused delays and misrouted tickets.
- Limited real-time visibility into recurring issues and potential incident spikes.
- Reactive operations hindered proactive planning & resource allocation.

Our Solution

- Embedded an AI-enabled workflow that triggers on ticket creation to analyze content, assign priority, and categorize it into predefined groups.
- Implemented auto-routing to the right teams, reducing misclassification rework.
- Enabled trend dashboards and predictive signals to identify recurring issues and forecast ticket spikes.
- Delivered measurable improvements in response time, ticket handling accuracy, and user satisfaction.



Tech Stack

Agentic AI integrated with ServiceNow workflows | AI Foundry's OpenAI models for automated priority & category assignment

Outcome Delivered

Faster resolution:
Reduced incident resolution time by
30–50%
through AI-driven auto-prioritization

Lower manual effort:
Cut manual ticket triage effort by
40–60%,
allowing teams to focus on high-value tasks

Proactive ops:
Improved incident spike prediction accuracy by
25–40%,
enabling proactive resource planning and stability

Quality uplift:
Improved routing accuracy to
90–95%
increased end-user satisfaction scores by
15–25%

Super Major Oil & Gas Company

Claimify.AI: Optimizes Claims Substantiation

Processing time reduced by 50%, Saving \$144Mn

Challenges

- Time-consuming manual data search for claims substantiation.
- Inefficient, repetitive legal research processes.
- Difficulty in generating actionable insights from past substantiated claims.
- Limited decision-making capability and high operational costs due to lack of timely, data-driven insights.

Our Solution

- Implemented AI-based semantic search solution using an LLM to find previous claims data.
- Enabled accurate, context-aware search and evidence retrieval using AI/NLP.
- Provided legal analysts and corporate teams with an easy-to-navigate evidence repository.
- In-built design to improve search accuracy-based user interactions.
- Delivered actionable insights and improved decision support for corporate affairs attorneys with excel-like UI.



Tech Stack

AI/NLP | Power automate | Data lake | Azure AI

Outcome Delivered

\$1.44Mn
in yearly savings
approved

50%
reduction in claim
processing time

9,600
business
hours saved annually

600
claims processed
annually with improved
efficiency in the
claim process

Improved productivity,
insights generation, and decision-making

Super Major Oil and Gas Company

TAG AI Assistant – Intelligent Terminal Operations & Support Copilot

Agentic AI-powered knowledge assistant enabling faster incident resolution, self-service support, and scalable terminal operations across a global network

Challenges

- Supporting 50+ terminals globally required rapid access to operational knowledge and support guidance.
- Critical information was scattered across SharePoint, ServiceNow, SOPs, and multiple systems.
- Manual searches and tribal knowledge increased incident resolution times and operational inconsistencies.
- Frequent L1-L2 escalations and lengthy onboarding created additional support overhead.

Our Solution

- Developed the TAG AI Assistant, an agentic AI-powered copilot for terminal operations and support teams.
- Leveraged enterprise knowledge sources to deliver conversational, role-based, and context-aware assistance.
- Enabled unified access to procedures, troubleshooting guides, and operational documentation.
- Automated incident guidance, routing, and escalation support to improve issue resolution.
- Established a scalable knowledge platform supporting terminal expansion, workforce enablement, and operational standardization.

Tech Stack

Agentic AI | Generative AI | SharePoint | ServiceNow |
Knowledge Management Platform | Role-Based Access Controls |
Enterprise Support Automation



Outcome Delivered

Accelerated incident triage and resolution, reducing downtime and improving support responsiveness across terminal operations.

Reduced L1-to-L2 escalations through AI-driven guidance, self-service capabilities, and contextual troubleshooting support.

Enabled conversational knowledge access, delivering accurate and rapid answers from enterprise knowledge repositories.

Provided 24x7 self-service support, improving operational continuity and reducing dependence on specialist resources.

Streamlined onboarding and knowledge transfer, helping new support engineers become productive faster.

Established a scalable digital support model capable of supporting new terminals, facilities, and operational growth without proportional increases in support overhead.

AI-Powered Intelligent Document Processing for Financial Reconciliation

Transforming financial reconciliation through intelligent document processing, accelerating data extraction, improving accuracy, and reducing manual effort at scale.

Challenges

- Thousands of credit and debit note documents arrived in diverse PDF and Word formats.
- Non-standard templates made data extraction complex and inconsistent.
- Manual processing increased effort, turnaround time, and data quality risks.
- Growing document volumes limited scalability and increased operational costs.

Our Solution

- Implemented an AI-powered Intelligent Document Processing (IDP) solution using Microsoft Copilot Studio.
- Created a unified extraction framework capable of handling multiple document formats and layouts.
- Leveraged OCR, AI-based extraction, and business-rule validation to automate critical data capture.
- Enabled scalable processing while maintaining consistency and accuracy across documents.
- Delivered a reusable, production-ready solution to streamline reconciliation workflows.

Tech Stack

Microsoft Copilot Studio | Intelligent Document Processing (IDP) |
OCR | AI-Powered Field Extraction | Business Rules Engine |
Workflow Orchestration



Outcome Delivered

Automated processing of 1,000+ documents per month-end cycle, significantly increasing operational throughput.

Reduced manual effort by 60-70%, minimizing dependency on human validation and repetitive processing activities.

Improved processing turnaround time by 50-60%, accelerating reconciliation cycles and financial close activities.

Enhanced extraction accuracy and consistency, reducing reconciliation exceptions and improving overall data quality.

Delivered a production-ready solution within 2-3 days, enabling rapid business value realization and accelerated adoption.

Reduced operational costs through automation, improved productivity, and streamlined financial workflows.

Strengthened compliance and audit readiness through improved data reliability, standardization, and process governance.

Fortune 1000 Public Utility

Rates & Regulatory Intelligence Bot

AI-powered regulatory intelligence assistant accelerating research, response generation, and filing preparation through instant access to historical regulatory knowledge



Challenges

- Manual regulatory research required significant effort to locate documents and analyze historical filings.
- Over 2,500 annual data requests created lengthy turnaround cycles and operational overhead.
- Response template creation and document retrieval relied heavily on manual effort.
- Limited access to historical context and precedents weakened regulatory response readiness.

Our Solution

- Developed a custom conversational AI assistant for Rates & Regulatory Affairs teams.
- Automated regulatory research, response generation, and report preparation workflows.
- Enabled intelligent search and summarization of historical filings, precedents, and supporting documentation.
- Automated template creation for recurring regulatory data requests.
- Integrated cloud-based publishing and archiving for streamlined document management.

Tech Stack

Front-end technologies: React, JavaScript, TypeScript, Tailwind CSS, CSS | **Backend Technologies:** AWS Services – Cloud Base (Lambda, Stepfunction, ECR, Glue, Sagemaker, AWS Amplify) | **Backend Technologies:** AWS AI Services (Bedrock (Model – Claude Sonnet 3.5 & Amazon Titan), Kendra (SharePoint Data source, S3 Data Source, WebCrawler Data Source) | **Backend Technologies:** AWS Security Services (Cognito, AI Gateway, WebSocket, Macie, Guardrail, System Manager (parameter store), Secret Manager) | **Database:** Snowflake, Dynamodb, SQL Server | **Cloud resources:** Azure (application management), Web App | **DevOps:** Jenkins, AWS CloudFormation.

Outcome Delivered

5%
improvement
in regulatory
outcomes

\$1Mn-\$2Mn
annual cost recovery enabled
through improved regulatory
positioning and filing quality.

Reduced template
creation time from
10-15 minutes to
seconds

Significantly reduced
document search and
research effort

Legal Intelligence Bot

AI-powered legal intelligence solution accelerating case preparation through automated policy discovery, keyword identification, and evidence analysis

Challenges

- Manual identification of relevant keywords and reference policies was time-consuming and inefficient.
- Legal teams faced challenges locating applicable regulations and supporting documentation.
- Building strong legal cases required extensive research, leading to delays in legal proceedings.

Our Solution

- Developed an AI-powered legal intelligence engine to analyze case documents and supporting evidence.
- Automated identification of relevant keywords, policies, and regulatory references.
- Ranked keywords and references based on relevance and case context.
- Enabled faster legal research and evidence discovery through intelligent document analysis.



Tech Stack

Front-end technologies: React, JavaScript, TypeScript, Tailwind CSS, CSS | **Backend Technologies:** AWS Services – Cloud Base (Lambda, Stepfunction, ECR, Glue, Sagemaker, AWS Amplify) | **Backend Technologies:** AWS AI Services (Bedrock (Model – Claude Sonnet 3.5 & Amazon Titan), Kendra (SharePoint Data source, S3 Data Source, WebCrawler Data Source)) | **Backend Technologies:** AWS Security Services (Cognito, AI Gateway, WebSocket, Macie, Guardrail, System Manager (parameter store), Secret Manager) | **Database:** Snowflake, Dynamodb, SQL Server | **Cloud resources:** Azure (application management), Web App | **DevOps:** Jenkins, AWS CloudFormation.

Outcome Delivered

Reduced keyword identification time from days to under a minute.

Reduced policy and reference research effort from 4+ hours to less than 1 hour.

Accelerated case preparation

Improved access to relevant policies and precedents

Fortune 1000 Public Utility

AI-Powered Marketing Campaign Assistant

Conversational AI-powered marketing assistant enabling personalized campaign creation, faster audience targeting, and improved marketing productivity

Challenges

- Manual campaign creation and customer data retrieval were time-consuming and resource intensive.
- Marketing teams struggled to quickly identify the right audience and develop personalized content.
- Limited automation impacted campaign efficiency and creativity.
- Reliance on external resources increased execution costs and turnaround times.

Our Solution

- Developed a custom conversational AI assistant for marketing campaign planning and execution.
- Enabled real-time access to customer data through natural-language interactions.
- Automated creation of personalized campaign content and email drafts.
- Centralized customer insights and campaign intelligence through a unified interface.



Tech Stack

Front-end technologies: React, JavaScript, TypeScript, Tailwind CSS, CSS | **Backend Technologies:** AWS Services – Cloud Base (Lambda, Stepfunction, ECR, Glue, Sagemaker, AWS Amplify) | **Backend Technologies:** AWS AI Services (Bedrock (Model – Claude Sonnet 3.5 & Amazon Titan), Kendra (SharePoint Data source, S3 Data Source, WebCrawler Data Source) | **Backend Technologies:** AWS Security Services (Cognito, AI Gateway, WebSocket, Macie, Guardrail, System Manager (parameter store), Secret Manager) | **Database:** Snowflake, Dynamodb, SQL Server | **Cloud resources:** Azure (application management), Web App | **DevOps:** Jenkins, AWS CloudFormation.

Outcome Delivered

Improved marketing productivity

Accelerated campaign development with AI-generated personalized email drafts and marketing content.

Enhanced campaign effectiveness

Reduced dependency on external contractors

Fortune 1000 Public Utility

Talent Acquisition AI Assistant

GenAI-powered recruitment assistant accelerating hiring workflows through automated job description creation, interview preparation, and candidate evaluation support

Challenges

- Creating job descriptions and interview questionnaires required significant manual effort.
- Recruiting teams spent considerable time tailoring hiring content based on candidate profiles and role requirements. Manual hiring processes increased turnaround times and reduced recruiter productivity.

Our Solution

- Developed a conversational AI assistant tailored for Talent Acquisition teams.
- Integrated with the Career Gallery portal to access live candidate and job data.
- Automated generation of role-specific job descriptions using multiple hiring inputs.
- Created interview question guides based on job requirements and candidate profiles.



Tech Stack

Front-end technologies: React, JavaScript, TypeScript, Tailwind CSS, CSS | **Backend Technologies:** AWS Services – Cloud Base (Lambda, Stepfunction, ECR, Glue, Sagemaker, AWS Amplify) | **Backend Technologies:** AWS AI Services (Bedrock (Model – Claude Sonnet 3.5 & Amazon Titan), Kendra (SharePoint Data source, S3 Data Source, WebCrawler Data Source) | **Backend Technologies:** AWS Security Services (Cognito, AI Gateway, WebSocket, Macie, Guardrail, System Manager (parameter store), Secret Manager) | **Database:** Snowflake, Dynamodb, SQL Server | **Cloud resources:** Azure (application management), Web App | **DevOps:** Jenkins, AWS CloudFormation.

Outcome Delivered

Reduced candidate hiring cycle by up to 3 days

Improved Talent Acquisition team productivity

Enabled faster creation of job descriptions and interview guides.

~\$200K annual savings through greater hiring efficiency and process optimization

Fortune 1000 Public Utility

AI-Powered Work Order Design Review

Automating design validation through AI-driven review intelligence

Challenges

- Manual work order design reviews were time-consuming and resource intensive.
- Review standards varied across reviewers, resulting in inconsistent outcomes.
- Prolonged approval cycles delayed project execution and delivery timelines.

Our Solution

- Developed an AI-powered review assistant to automatically validate work order sketches and designs.
- Applied rule-based intelligence to assess designs against predefined engineering standards and checklists.
- Standardized the review process through automated validation and compliance checks.
- Enabled faster review cycles while improving consistency and accuracy.



Tech Stack

Front-end technologies: React, JavaScript, TypeScript, Tailwind CSS, CSS | **Backend Technologies:** AWS Services – Cloud Base (Lambda, Stepfunction, ECR, Glue, Sagemaker, AWS Amplify) | **Backend Technologies:** AWS AI Services (Bedrock (Model – Claude Sonnet 3.5 & Amazon Titan), Kendra (SharePoint Data source, S3 Data Source, WebCrawler Data Source)) | **Backend Technologies:** AWS Security Services (Cognito, AI Gateway, WebSocket, Macie, Guardrail, System Manager (parameter store), Secret Manager) | **Database:** Snowflake, Dynamodb, SQL Server | **Cloud resources:** Azure (application management), Web App | **DevOps:** Jenkins, AWS CloudFormation.

Outcome Delivered

Reduced design review time from
30 minutes to **4** minutes per design

Accelerated project delivery through faster approval cycles and increased review throughput

Improved design quality and consistency

Reduced rework and review effort

Fortune 1000 Public Utility

AI-Powered Expense Audit Bot

Automating expense compliance and audit workflows through AI-driven policy validation

Challenges

- Manual expense audits were difficult to scale across growing transaction volumes.
- High risk of non-compliant expense claims being approved and processed.
- Policy validation relied heavily on manual review, increasing effort and inconsistency.

Our Solution

- Developed an AI-powered expense audit assistant to automatically validate employee expense claims.
- Embedded reimbursement rules and policy guidelines into an intelligent review framework.
- Automated compliance checks across submitted expense claims before approval.
- Enabled policy-driven auditing and exception identification at scale.



Tech Stack

Front-end technologies: React, JavaScript, TypeScript, Tailwind CSS, CSS | **Backend Technologies:** AWS Services – Cloud Base (Lambda, Stepfunction, ECR, Glue, Sagemaker, AWS Amplify) | **Backend Technologies:** AWS AI Services (Bedrock (Model – Claude Sonnet 3.5 & Amazon Titan), Kendra (SharePoint Data source, S3 Data Source, WebCrawler Data Source) | **Backend Technologies:** AWS Security Services (Cognito, AI Gateway, WebSocket, Macie, Guardrail, System Manager (parameter store), Secret Manager) | **Database:** Snowflake, Dynamodb, SQL Server | **Cloud resources:** Azure (application management), Web App | **DevOps:** Jenkins, AWS CloudFormation.

Outcome Delivered

100% of expense claims up to \$7K audited automatically

90%+ accuracy in identifying non-compliant expense submissions

Improved policy adherence across employee reimbursement programs

Reduced manual audit effort

Fortune 1000 Public Utility

AI-Powered Procurement Compliance Bot

Automating purchase order audits through AI-driven policy validation

Challenges

- Manual PO audits covered only 1–2% of total purchase orders.
- High risk of non-compliant POs being approved and processed.
- Policy violations often created downstream supply chain and operational disruptions.

Our Solution

- Developed an AI-powered procurement review assistant to automatically audit purchase orders against procurement policies and business rules.
- Embedded policy controls into an intelligent validation framework for real-time compliance checks.
- Automated PO review and exception detection prior to approval and processing. Enabled scalable procurement governance through continuous, policy-driven auditing.



Tech Stack

Front-end technologies: React, JavaScript, TypeScript, Tailwind CSS, CSS | **Backend Technologies:** AWS Services – Cloud Base (Lambda, Stepfunction, ECR, Glue, Sagemaker, AWS Amplify) | **Backend Technologies:** AWS AI Services (Bedrock (Model – Claude Sonnet 3.5 & Amazon Titan), Kendra (SharePoint Data source, S3 Data Source, WebCrawler Data Source) | **Backend Technologies:** AWS Security Services (Cognito, AI Gateway, WebSocket, Macie, Guardrail, System Manager (parameter store), Secret Manager) | **Database:** Snowflake, Dynamodb, SQL Server | **Cloud resources:** Azure (application management), Web App | **DevOps:** Jenkins, AWS CloudFormation.

Outcome Delivered

100%
of purchase orders
audited automatically

90%+
accuracy in identifying
non-compliant
purchase orders

Reduced
procurement risk

Minimized downstream
supply chain disruptions

Fortune 1000 Public Utility

AI-Powered Requisition Compliance Bot

Automating requisition reviews through AI-driven policy validation



Challenges

- Low-risk and low-value requisitions lacked capacity for manual review.
- Policy violations could pass undetected, increasing compliance risk.
- Unchecked requisitions created potential downstream procurement and operational issues.

Our Solution

- Developed an AI-powered requisition review assistant to automatically validate purchase requisitions against business rules and policies.
- Embedded compliance checks into an intelligent review framework for consistent evaluation.
- Automated requisition auditing regardless of risk level or value.
- Eliminated manual reviews for low-risk requisitions while maintaining policy compliance.

Tech Stack

Front-end technologies: React, JavaScript, TypeScript, Tailwind CSS, CSS | **Backend Technologies:** AWS Services – Cloud Base (Lambda, Stepfunction, ECR, Glue, Sagemaker, AWS Amplify) | **Backend Technologies:** AWS AI Services (Bedrock (Model – Claude Sonnet 3.5 & Amazon Titan), Kendra (SharePoint Data source, S3 Data Source, WebCrawler Data Source) | **Backend Technologies:** AWS Security Services (Cognito, AI Gateway, WebSocket, Macie, Guardrail, System Manager (parameter store), Secret Manager) | **Database:** Snowflake, Dynamodb, SQL Server | **Cloud resources:** Azure (application management), Web App | **DevOps:** Jenkins, AWS CloudFormation.

Outcome Delivered

640+ hours saved annually through reduced manual review effort.

Improved requisition audit accuracy through automated policy validation.

Enhanced compliance coverage across all purchase requisitions.

Reduced labor costs by eliminating manual checks for low-risk transactions.