

Domain Meets AI

Energizing Efficiencies and Powering Possibilities

American Multinational Company

Driving Digital Field Transformation

AI-powered insights and real-time decision support for safer, smarter operations

Challenges

- Fragmented and hard-to-integrate data across OT and IT networks.
- Manual, time-consuming processes slowed decision-making and productivity.
- Lack of user-friendly, real-time support for field and office staff.

Our Solution

- Intelligent agents extract and process data from multiple sources for timely, precise insights.
- Amazon Bedrock powers robust embeddings and daily reporting with advanced LLMs.
- RAG system contextualizes information, optimizing responses for relevance and accuracy.
- GenAI delivers daily briefings, real-time Q&A, and optimized route planning for field teams.



Tech Stack

AWS Bedrock (Titan embedding, Claude LLMs), Anthropic Claude Sonnet, Hiaku LLMs | Data services: DBT, Snowflake, Stonebrach | Business apps: SCADA, SAP, GIS | AWS Transcribe, AWS Polly for audio/transcription, DynamoDB, S3 for storage | Streamlit for UI framework

Outcome Delivered

35–40%

Increase in productivity due to Gen AI-powered daily briefings & streamlined workflows

Significant reduction in non-value-added manual effort—engineers spend less time searching for information

Real-time, accurate insights enable safer and more efficient completion of field tasks

Automated task prioritization and optimal route selection improve team's efficiency

Supermajor Oil & Gas Company

AI-Driven Anomaly Detection for Unconventional Wells

Challenges

- Manual surveillance and alert configuration consumed significant expert time.
- Scattered data systems delayed anomaly detection and response.
- Difficulty in classifying and acting on short- and long-term well anomalies.
- Limited scalability for monitoring a growing number of wells.

Our Solution

- Developed an automated anomaly detection solution using a heuristic Z-score model.
- Leveraged Azure AI/ML and Seeq for real-time monitoring of 1,000+ gas-lift wells.
- Implemented rule-based and pattern-recognition engines for anomaly classification.
- Integrated real-time data for proactive well management and decision-making.



Tech Stack

Azure AI/ML, Seeq, Python, PETEX, Relational Databases | Z-Score Statistical Model for anomaly detection | Real-time data integration and scalable analytics infrastructure | SQL for data storage and dashboarding

Outcome Delivered

Saved
800+
man hours per month
by automating
anomaly detection and
eliminating manual
surveillance

Enabled real-time
operational insights
for over
1,000
gas-lift wells, improving
responsiveness

Accelerated root cause
analysis and anomaly
classification, reducing
downtime and
operational delays

Scalable analytics
infrastructure
supports seamless
monitoring as new
wells are added

Smarter Wells, Faster Decisions: AI-Driven IPR Estimation at Scale

Challenges

- Traditional inflow performance models failed for rapidly declining reservoir pressures.
- Manual, error-prone IPR estimation processes.
- Difficulty in forecasting key subsurface metrics for unconventional wells.
- Need for more proactive, data-driven well management.

Our Solution

- Built AI/ML models on Azure for intelligent IPR and reservoir pressure estimation.
- Integrated physics-based and machine learning inputs for robust predictions.
- Developed modules for real-time decision support and production optimization.
- Enabled scalable, automated workflows for subsurface operations.



Tech Stack

Azure AI/ML Workspace, Python, PETEX, Harmony Models | Machine learning for historical and predictive analytics | Physics-informed input integration | Relational databases for data management

Outcome Delivered

Delivered advanced, ML-driven IPR estimation for unconventional wells, improving production planning

Enabled proactive well management by forecasting pressure and drawdown trends

Enhanced accuracy by integrating physics-informed inputs with AI models

Supported quicker, data-driven decisions, reducing operational overhead

Zero Guesswork: AI-Driven Calibration for Optimized Well Production

Challenges

- Manual calibration of production networks was time-consuming and error-prone.
- Difficulty in minimizing production gaps and optimizing network performance.
- Complex, multi-constraint optimization requirements.
- Need for real-time data integration and scalable scenario processing.

Our Solution

- Automated production network calibration using the PSwarm algorithm and PETEX suite.
- Developed a production gap minimization engine with real-time feedback loops.
- Enabled multi-constraint optimization and cluster-based parallel processing.
- Integrated real-time data for continuous calibration and updates.



Tech Stack

PETEX, PSwarm Model, Cluster Management, Relational Databases Real-time data integration platform | Distributed computing for parallel scenario processing | Machine learning for optimization & calibration

Outcome Delivered

Saved
1,800+
hours per field annually by
automating calibration
of complex production
networks

Real-time rate
prediction
minimized
unplanned
downtime and
production losses

Improved model
reliability and reduced
model-to-field error
percentage for
consistent forecasts

Cluster-based
parallel processing
accelerated
calibration cycles
and production
adjustments

No More Blind Spots: AI-Driven, Real-Time Well Rate Estimation with ML NET

Challenges

- Lack of multiphase flow meters in many wells limited visibility.
- Manual, delayed production allocation cycles.
- Need for scalable, cost-effective alternatives to traditional evaluation methods.
- Difficulty in tracking performance across diverse asset types.

Our Solution

- Developed ML-based rate estimation engine using ML NET.
- Automated model lifecycle management for continuous accuracy.
- Enabled high-frequency output and integrated accuracy monitoring.
- Provided alternative production evaluation methodology for non-instrumented wells.



Tech Stack

ML NET, Python, Relational Databases | Automated pipelines for model training and validation | High-frequency data processing and analytics

Outcome Delivered

Provided real-time well rate estimates for wells lacking multiphase flow meters, enabling faster decisions

Automated model lifecycle management ensured sustained accuracy and adaptability

Offered a scalable, cost-effective alternative to traditional production evaluation methods

High-frequency output and transparent analytics improved stakeholder trust and auditability

American Oil & Gas Company

Accelerating Production Insights with Generative AI

Automated HSE data analysis, enabled faster decisions, and minimized manual effort with XPERT AI-powered platform

Challenges

- Substantial time and effort required to locate and analyze relevant production and HSE data.
- Manual processes slow down business decisions and reduce overall efficiency.
- Reports often fail to provide actionable insights promptly, leading to delays in decision-making.
- Deriving insights required reference multiple reports which took significant time resulting in inefficiencies.

Our Solution

- Developed and deployed XPERT, a generative AI-powered platform.
- Automated the process of retrieving & analyzing production & HSE data.
- Enabled business and technology teams to make faster, more informed decisions with minimal manual effort.
- Customized and consolidated reports delivered based on user-specific requests.
- Performance tracking insights made readily accessible to all stakeholders.



Tech Stack

Custom-built AI solution leveraging generative AI technology | Automated retrieval, synthesis, and analysis of HSE documentation and production/deferment data

Outcome Delivered

Quick access to data and insights enables faster business decisions

Minimized time spent searching for information across multiple reports

\$1.44Mn
business value
Delivered

Provided actionable insights and recommendations through consolidated reports, driving production improvement and operational efficiency

Microsoft ADME

From Silos to Insights

Redefining Data Discovery with Gen AI and OSDU



Challenges

- Fragmented data silos across proprietary E&P systems.
- Complex data integration and inconsistent data quality.
- Limited collaboration and slow decision-making due to legacy infrastructure.
- Manual, time-consuming model updates and data preparation.
- Security, governance, & skill gaps in adopting open standards & AI.

Our Solution

- Implemented Gen AI-powered Virtual Data Analyst (VDA) on the OSDU™ Data Platform to enable natural language search & interpretation of subsurface data.
- Integrated Microsoft Azure Data Manager for Energy (ADME) and OSDU APIs for seamless, secure data access.
- Enabled advanced analytics and natural language queries using Azure OpenAI and Copilot. Delivered a user-friendly interface for geoscientists and business users.
- Unified diverse datasets—well logs, seismic, rock, and fluid—into a single subsurface model for better decision-making.

Tech Stack

Microsoft Azure Data Manager for Energy (ADME) | OSDU™ Data Platform and APIs | Azure Data Lake and Vector Store | Microsoft Search and AI Document Intelligence

Outcome Delivered

Centralized
subsurface
data to
eliminate silos

Accelerated
data discovery
cutting
interpretation
cycles by
20–30%

Automated
data quality
control and
governance

Delivers savings
USD
~2–3Mn
annually
By time saving
from manual prep

Enhanced
insight
generation
and supports
new
discoveries

Scalable to
non-
conventional
energy
domains

Fortune 1000 Public Utility

Hydro Forecasting Reinvented with AI

ML-powered inflow predictions achieved 10–15% MAPE accuracy, saving \$750K annually and boosting operational agility



Challenges

- Inaccurate inflow forecasts led unnecessary water spill & lost value.
- Frequent pond level approaches to maximum limits, reducing operational flexibility.
- Suboptimal generation scheduling and energy market participation due to unreliable forecasts.
- Reliability concerns, including risk of uneconomic dispatch and potential load shedding.

Our Solution

- Data Science: Developed a machine learning model for 7-day hourly inflow forecasts, achieving 10–15% MAPE accuracy. Forecasts updated hourly, available 75 minutes before each hour for real-time operations.
- Data Engineering: Automated data ingestion from multiple sources into Snowflake for model training and reporting. Enabled direct comparison of forecast accuracy (MAPE/MAE) between DSM and DOPD models.
- Reporting: Tableau dashboard visualizes forecast performance, error metrics, and financial benefits. Supports trading and strategy decisions with clear comparisons of actuals vs forecasts.

Tech Stack

Machine learning, Python (forecasting model) | S3, nowflake, Lambda (data engineering) | Tableau (front-end/dashboard)

Outcome Delivered

Improved forecast accuracy and operational planning

Enhanced real-time decision-making and resource management

Estimated hard savings of **\$750,000** per year since 2024 by reducing avoidable spill

Optimized hydroelectric operations and better compliance with regulatory and environmental requirements

US Midstream Major

Digital Twin for Compressors: Edge-to-Cloud Intelligence

Unified analytics and automation deliver \$3Mn savings, cut emissions by \$05Mn per year

Challenges

- Manual and siloed compressor onboarding and data definitions.
- Lack of predictive analytics and real-time operational insights.
- High unscheduled downtime and suboptimal asset performance.
- Complex data aggregation & reporting across distributed assets.

Our Solution

- Deployed the Insight NXT platform for unified edge-to-cloud analytics and orchestration.
- Enabled compressor onboarding, predictive analytics, and remote monitoring.
- Automated data aggregation, alerting, and emissions reporting.
- Integrated with enterprise systems for seamless data flow and compliance.
- Delivered rapid, scalable deployment across 130 sites.



Tech Stack

Insight NXT Platform (IIoT, AI/ML, Edge and Cloud) |
Real-time analytics at the edge and in the cloud |
Low-code platform with self-service analytics |
Integration with SCADA, PI, sensors, and enterprise systems (EMS, EHS, HRIS)

Outcome Delivered

Over
\$3Mn
in optimal
operations
savings

Reduced
emissions by
~\$0.5Mn
per year

6–8%
reduction in
reporting time

~4
hours
saved per alert

Improved asset visibility,
performance, and useful
life; reduced downtime and
operational costs

Super Major Oil and Gas Company

AI-Powered Predictive Maintenance for ESPs

Boosted MTBF by 10% and saved \$500K annually



Challenges

- Frequent Electric Submersible Pump (ESP) failures led to high downtime & costs.
- High failure rates due to operational complexity and limited diagnostic capabilities and low ROI from traditional physics-based models requiring continuous calibration.
- Difficulty in monitoring and optimizing performance due to limited diagnostics and frequent cavitation events.

Our Solution

- Implemented an ML-based predictive maintenance solution tailored for ESPs in dry gas wells.
- Streamed and preprocessed real-time sensor data (pressure, temperature, current, vibration, flow, and GLR).
- Developed and trained SOM models to detect anomalies, predict failures, and optimize pump operation.
- Enabled anomaly detection, diagnostics, and prescriptive actions for proactive maintenance.
- Integrated alerts and recommendations into operator workflows.

Tech Stack

Data Collection: Rockwell PLCs (SCADA systems, IoT gateways) | Preprocessing: Python (Pandas, NumPy) | Modeling: Self-Organizing Map (SOM) Model | Alerting & Integration: Alerts integrated into the application layer above Rockwell PLCs (SCADA) | Database/Storage: SQL Server

Outcome Delivered

Increased mean time between failures (MTBF) of ESPs by an average of **10%**

Achieved estimated annual savings of **~\$500,000**

Enabled earlier detection of anomalous pump operations for proactive maintenance

Improved equipment reliability, reduced downtime, & enhanced operational efficiency

US Midstream Major

AI-Powered Supply Chain Risk Sensing

Enhanced visibility, predicted supplier risks early, and reduced production downtime with automated mitigation

Challenges

- Lack of visibility on supply parts delivery led to delayed shipments and incorrect inventory.
- Short shipments and breakdowns caused production delays.
- Manual tracking and coordination for planners and business users.
- High costs associated with even brief production downtime.

Our Solution

- Developed an AI/ML tool to sense and mitigate supplier risks using current data points.
- Built machine learning models to predict and flag supplier risks before they impact production.
- Automated notifications and escalation to responsible teams.
- Enabled proactive mitigation actions to minimize delays.
- Provided continuous model evaluation and improvement for sustained results.



Tech Stack

AI/ML tool for supply chain risk sensing and mitigation | Machine learning models for risk prediction | Automated person-in-charge assignments and notifications | Model metrics evaluation for continuous improvement

Outcome Delivered

Enhanced visibility
and resiliency
in the supply chain

Early identification
and mitigation
of supplier risks

Reduced production
downtime and
associated costs

Streamlined
communication and
faster response to
disruptions

Maximizing Reliability with AI-Driven Uncrewed Operations

Enabled real-time simulation, predictive analytics, and automated workflows to reduce failures and optimize asset performance

Challenges

- Frequent asset failures and unpredictable equipment downtime hindered operational efficiency.
- High maintenance and repair costs, with suboptimal asset usage, limit equipment lifespan.
- Ensuring safety and environmental compliance was challenging due to operational risks and the need for continuous monitoring.
- Accelerating root cause analysis was difficult without effective knowledge capture and advanced pattern recognition.

Our Solution

- Established a unified namespace to organize and contextualize data for advanced analytics. Enhanced data enrichment and processing for pattern recognition and alarm management.
- Deployed AI/ML models for actionable insights and forecasting.
- Implemented a Gen AI-powered compressor bot for real-time, natural language data access.
- Enabled comprehensive workflow management for efficient root cause analysis and recommendations.



Tech Stack

Unified Namespace (UNS) | Data enrichment and processing pipelines | AI/ML models | Gen AI-powered compressor bot | Digital twin models for real-time simulation and scenario planning

Outcome Delivered

Real-time simulation of compressor stations using digital twin models for scenario planning and optimization

Advanced analytics deliver actionable insights from continuous real-time data analysis

Generative AI-powered compressor bots enable natural language interaction with real-time and historical data

Automated alert and alarm management ensures timely, proactive responses

Comprehensive workflow management streamlines root cause analysis and recommendations across sites

American Multinational Company

Optimizing Gas Lift Compression for Production Wells

AI-powered real-time monitoring and automation for efficiency and performance

Challenges

- Manual processes for managing compressor energy use and gas injection were inefficient.
- Lack of real-time alerting and workflow automation led to delayed responses and suboptimal performance.
- Difficulty in identifying the optimal injection point and amount of gas for maximum efficiency.

Our Solution

- Designed and implemented an AI solution to optimize compressor operations in production wells.
- Enabled real-time alerting, alarm management, and workflow automation.
- Delivered a system that efficiently manages energy use & maximizes liquid flow, supporting continuous production improvements.



Tech Stack

Custom-built AI solution leveraging Palantir platform |
AI-driven optimization for compressor energy, gas injection volume, and injection point selection |
Real-time monitoring and predictive analytics for operational decision support

Outcome Delivered

Real-time alerting, alarm management, and workflow automation for proactive operations

Efficient energy use, reducing operational costs and environmental impact

Maximized liquid flow to the surface, improving overall production rates

Super Major Oil and Gas Company

Smarter Turnarounds with AI: Driving Safety and Cost Efficiency

Save up to \$10Mn, cut planning time by weeks, & reduce risk with predictive analytics and automated workflows

Challenges

- High cost of unplanned downtime at large facilities.
- Scope creep, contractor availability issues, and fragmented visibility disrupt schedules and budgets.
- Manual, time-consuming shift handovers and log reviews led to errors and delays.
- Data scattered across spreadsheets, emails, and legacy systems hampered fast decision-making.

Our Solution

- TAR Event Manager Copilot: A conversational partner that understands and draws from real-time data, historical trends, and even shift logs to provide clear, actionable answers instantly.
- Predictive Analytics Engine with S-Curve Tracking.
- NLP Work Analyzer: This tool uses natural language processing to clean, structure, and standardize them work packages.
- Shift Handover Intelligence: Compiles updates from logbooks, emails, and voice notes into a single, prioritized summary.



Tech Stack

AI and Machine Learning: Azure Machine Learning, NLP models, predictive analytics |

Data Management: ADF, Azure Data Lake, Synapse
Analytics | User Interface: Power BI dashboards, Power Apps, web apps for end-user consumption

Outcome Delivered

~\$10Mn
saved
per day by
minimizing
unplanned
downtime

12%
reduction in
turnaround
scope through
AI-driven task
validation

48 hour
delay avoided by
early detection
and
resequencing of
critical tasks

Weeks of manual
planning reduced
to hours with
NLP-based job
plan

50%
faster shift
handover process
through
automated,
AI-generated
summaries

Multi-million
dollar cost
overruns
prevented by
real-time risk
detection

Supermajor Oil & Gas

Optimizing Maintenance with Generative AI for Australia Asset Class

Saved \$400K–\$800K annually, cut manual effort, and enabled scalable outcome classification across equipment classes

Challenges

- Reviewing lengthy text in work orders to accurately classify test outcomes (pass/fail) was manual, time-intensive, and error-prone.
- The existing text analytics solution relied on manual keyword identification, which varies across equipment classes, making the process non-scalable and inconsistent.

Our Solution

- Implement Generative AI (GenAI) models to analyze work order text automatically.
- Classify outcomes (failures and passes) at scale without manual intervention.
- Eliminate manual keyword mapping, reducing human effort & errors.
- Enable a scalable, intelligent approach for outcome classification across all equipment classes.



Tech Stack

Tailwind CSS | AWS API Gateway | AWS Lambda | AWS Step Functions | AWS Bedrock | Claude 3.5 Sonnet | Amazon Titan | Bedrock Guardrails

Outcome Delivered

\$400K–\$800K
per year savings by
extending planned
work order frequency
from 1 year to 2 years
(50–100% optimization)

\$8Mn–\$15Mn
savings per year from
10–20% reduction in
annualized maintenance
spend

\$10Mn–\$30Mn
per year additional value from improved uptime
(1–3 extra production days at **\$10Mn/day**)
enabled by optimized maintenance programs

Super major Oil & Gas

Preventing Equipment Failures with AI-Driven Valve Positioner Monitoring

Enabled early anomaly detection, avoided costly outages, and prevented \$32Mn+ production losses through predictive insights

Challenges

- The Wheatstone LNG plant experienced frequent equipment and trip events primarily caused by valve positioner failures.
- Historical data analysis revealed early indicators of valve positioner faults that could have prevented these failures if detected in time.
- However, due to the absence of regular monitoring, these early warning signs were missed, resulting in unanticipated equipment trips and significant production losses.

Our Solution

- Implemented advanced anomaly detection techniques and integrated the results into PI Vision.
- Automate continuous monitoring of valve positioner performance.
- Provide real-time alerts to onsite personnel.
- Enable anomaly verification annotation for proactive maintenance actions.



Tech Stack

Data Ingestion and Integration | Anomaly Detection and Analytics | Visualization and Monitoring | Alerting and Notifications | Storage and Data Management | Security and Access Control | CI/CD and Deployment | Observability and Logging

Outcome Delivered

Early detection of valve positioner failures prevented unexpected plant or equipment trips

Minimized downtime and avoided costly production losses through predictive maintenance

Improved operational reliability and safety by addressing anomalies before escalation

Avoided major incidents like the 2023 failure that caused **\$32Mn** Lost Production Opportunity (LPO)

Demonstrated success: August 2024 detection of Turbo Expander IGV issue, preventing **~\$0.9Mn** day LPO

Fortune 500 Utility Leader

Cut Risk Reporting Time by 40% with AI-Driven Insights

Optimized trading decisions, reduced compliance risk, and boosted efficiency through predictive models and automation

Challenges

- Disconnected platforms and manual processes hinder end-to-end automation and operational efficiency.
- Static systems struggle to learn and adapt to rapidly changing market and regulatory environments.
- Absence of conversational, natural-language interfaces makes risk insights difficult to access and act upon in real time.

Our Solution

- Predictive Models: Developed AI/ML models for predictive risk, VaR/ES, and anomaly detection.
- Simulations: Implemented AI-driven simulations for advanced scenario analysis with Azure Knowledge Fabric.
- Strategy Optimization: Created AI-powered strategy optimization with dynamic parameter adjustment.
- Risk Monitoring, Regulatory Reporting, and testing.
- Cognitive Automation: Built cognitive automation for routine tasks, including chatbot and Azure Copilot integration.
- Unified Platform: Integrated components into a unified, cloud-based platform.



Tech Stack

GitHub copilot | Azure OpenAI | Azure Foundry | Machine Learning models

Outcome Delivered

Optimized Returns:

Improved decision turnaround time by **30–50%** through AI-driven predictive analytics and simulations. (Indicative)

Reduced Risk:

Reduced potential risk exposure by **20–35%** using predictive risk models and scenario-based simulations. (Indicative)

Increased Efficiency:

Freed up **40–60%** of expert analyst effort by automating reporting, analysis, query resolution via conversational AI. (Indicative)

Enhanced Compliance:

AI-powered tools streamlined regulatory reporting

World's Leading Integrated Energy Company (LNG)

AI-Powered Turbine Intelligence for Faster Resolution

Enabling real-time access to critical turbine OEM data, reducing downtime risk, and improving maintenance accuracy through conversational AI

Challenges

- Unplanned turbine trips caused significant lost production opportunity (LPO) and plant downtime.
- Field teams required rapid, trusted access to turbine manuals and part information, but existing OEM documentation was extremely large (1,300–2,000+ pages), difficult to search onsite, and slow to cross-reference with JDE material data, increasing repair time and risk of incorrect part ordering.

Our Solution

- Chat-based LLM solution developed to help the Turbine Team identify turbine part numbers, lookup operating and troubleshooting instructions for Turbine operations and maintenance, as well as the ability to cross reference part numbers with JDE item numbers, retrieving suppliers' details for JDE parts using agentic AI.



Tech Stack

Databricks | SharePoint | Generative AI

Outcome Delivered

Reduced time to resolution during turbine events

Improved first-time-right part identification, minimized ordering errors and material losses

Provided a mobile-first experience for onsite technicians

Helped mitigating high-impact production losses

Fortune 1000 Energy Company

Improving Outage Restoration Accuracy with Agentic AI

Dynamic, data-driven Estimated Restoration Time (ERT) predictions leveraging real-time grid conditions, historical data, and crew progress

Challenges

- Poor customer communication during outages.
- Lack of dynamic updates as field conditions changed.
- Unreliable Estimated Restoration Times (ERTs).
- Manual estimation based on limited historical data.

Our Solution

- Designed a Smart Restoration Predictor Agent, an Agentic AI solution that continuously predicts outage restoration timelines.
- The agent leverages historical outage data, real-time grid conditions, and crew progress updates to dynamically adjust ERTs and provide accurate, real-time insights across operations.



Tech Stack

GridVerse™ | Agentic AI | MLOps

Outcome Delivered

50%+
improvement in
restoration time
accuracy

Real-time ERT
updates improving
customer
communication

Reduced manual
estimation effort

Increased trust in
outage resolution
timelines

Fortune 1000 Energy Company

Proactive Outage Forecasting with Agentic AI

Advanced AI agents forecast outages using real-time insights from weather and asset conditions, enabling faster, preventive decision-making.

Challenges

- Reactive outage response model.
- Inability to anticipate outages due to weather/grid conditions.
- Poor preparedness for large-scale outage events.

Our Solution

- Implemented a Proactive Outage Forecasting Agent, an Agentic AI solution that predicts potential outages before occurrence.
- The agent analyzes weather patterns, grid vulnerabilities, and asset conditions to identify high-risk areas and trigger preventive actions.



Tech Stack

GridVerse™ | Agentic AI | MLOps

Outcome Delivered

Enabled proactive
outage management
strategy

Improved
preparedness before
major weather
events

Reduced outage
occurrences in
vulnerable grid zones

Optimized resource
pre-positioning

British Water Supply and Treatment Utility Company

Agentic AI for Intelligent Field Operations

Agentic, multi-agent solution enabling contextual knowledge access, automated workflows, and seamless field operations



Challenges

- Fragmented SOPs, manuals, and alarm-fix records made on-site troubleshooting slow and inconsistent.
- Heavy reliance on manual search and back-office support increased resolution delays.
- Lack of real-time, contextual knowledge access reduced first-time fix rates and operational efficiency.

Our Solution

- Built SSAM Buddy, an agentic AI-powered, conversational field assistant for real-time knowledge access.
- Deployed a multi-agent architecture.
- Conversational agent for issue understanding.
- RAG-based knowledge retrieval from SharePoint SOPs/FAQs.
- Intent routing for resolution/escalation workflows.
- Ticketing agent for ITSM (Samanage) integration.
- Learning agent for continuous KB updates.
- Delivered a mobile-enabled PWA, ensuring seamless field accessibility.

Tech Stack

Azure OpenAI (GPT-4.1) | Azure AI Search | Azure Functions | SharePoint | Samanage ITSM

Outcome Delivered

Enabled technicians to scan and access solutions across 500+ technical documents instantly
30–40% reduction in time spent searching for relevant document

20–25% improvement in issue resolution time with faster access to SOPs and guidance

Responses generated in **<20** seconds, enabling real-time decision support in the field

15–20% improvement in first-time fix rate through AI-driven troubleshooting
Reduced dependency on back-office support

Super Major Oil and Gas Company

AI-Powered Equipment Failure Intelligence & Predictive Maintenance

Generative AI transforms decades of maintenance records into actionable failure intelligence

Challenges

- 115,000+ historical work orders contained valuable reliability insights trapped in unstructured records.
- Traditional methods made failure pattern and trend analysis difficult and time-consuming.
- Limited visibility into recurring failure modes impacted predictive maintenance efforts.
- Reliability teams lacked scalable mechanisms to proactively improve asset performance.

Our Solution

- Developed a GenAI-powered failure intelligence solution to analyze unstructured maintenance and work order data.
- Automated identification of recurring failure patterns, root causes, and reliability trends.
- Applied AI-driven classification with contextual explanations to improve insight accuracy and usability.
- Processed historical maintenance records at enterprise scale, accelerating knowledge extraction.
- Established a data-driven foundation for reliability improvement and predictive maintenance.

Tech Stack

Generative AI | Predictive Maintenance Analytics | JD Edwards |
Failure Intelligence Engine | Asset Reliability Analytics |
Maintenance Knowledge Mining



Outcome Delivered

95% classification accuracy achieved during proof-of-concept validation.

Processed **115,000+** historical work orders in approximately 2.5 hours, compared with an estimated 16,000 manual person-hours.

Avoided approximately AUD **\$2.8Mn** in labor costs through large-scale automation of maintenance intelligence activities.

Unlocked previously inaccessible operational knowledge embedded within unstructured work order records, accelerating maintenance insight generation.

Strengthened predictive maintenance programs through improved visibility into recurring failure patterns and asset performance trends.

Supported Failure Modes and Effects Analysis (FMEA) initiatives, contributing to an estimated **\$100Mn** potential business value through enhanced reliability and proactive maintenance

Super Major Oil and Gas Company

SmartSignal AI-Powered Predictive Maintenance & Anomaly Detection

AI-driven early warning intelligence enables proactive maintenance, improves asset reliability, and reduces production losses

Challenges

- Reactive monitoring and threshold-based alarms limited early visibility into equipment degradation.
- Emerging asset issues often went undetected, increasing outage and production risks.
- Reactive maintenance resulted in higher repair costs and unplanned interventions.
- Maintenance teams lacked predictive insights to proactively manage asset health.

Our Solution

- Implemented SmartSignal, an AI-powered predictive maintenance platform for continuous monitoring of critical assets.
- Leveraged historical and real-time sensor data to identify anomalies before threshold violations occur.
- Enabled intelligent anomaly detection against expected performance baselines.
- Integrated root-cause diagnostics and failure prediction to support proactive maintenance decisions.
- Delivered early-warning alerts and established a predictive reliability framework for long-term asset performance.

Tech Stack

SmartSignal Predictive Analytics Platform | AI/ML Models |
Real-Time Sensor Integration | Predictive Maintenance Engine |
Root Cause Diagnostics | Anomaly Detection Framework



Outcome Delivered

Enabled proactive detection of equipment anomalies before operational failures occurred, reducing reliance on reactive maintenance practices.

Avoided maintenance and production losses through early identification of equipment degradation and emerging failure conditions.

Improved uptime, reliability, and production continuity across monitored field assets.

Reduced risk of unplanned outages through advanced predictive analytics, anomaly detection, and early intervention strategies

Enhanced maintenance decision-making by providing data-driven insights, root cause diagnostics, and failure forecasting capabilities.

Improved workforce productivity by reducing firefighting activities and enabling more effective maintenance planning and resource utilization.

Demonstrated significant cost avoidance, including scenarios where early warnings helped prevent compressor failures with an estimated avoidance value of approximately **\$250K** per event.

Established a scalable predictive maintenance capability with potential applicability across power generation, manufacturing, chemical, and process industries.

Global Energy Technology Provider

DrillMiner: AI-Powered Operations Intelligence for Drilling

Transforming drilling operations and enabling engineers to access complex well data, uncover insights and make faster decisions.

Challenges

- Critical drilling and well-construction data was stored across complex EDM databases, requiring specialized knowledge to access and analyze.
- Engineers spent significant time navigating hundreds of interconnected data tables and building manual queries.
- Valuable operational insights remained locked within technical systems, limiting accessibility across teams.
- Fragmented workflows and data-quality challenges slowed decision-making and reduced operational efficiency.

Our Solution

- Developed DrillMiner, a GenAI-powered operations intelligence platform for drilling and well construction environments.
- Enabled engineers to query complex EDM databases using natural language without SQL expertise.
- Automated data retrieval, analysis, visualization, and insight generation through conversational AI.
- Implemented built-in analytical capabilities including drilling performance analysis, NPT insights, hazard identification, and data-health monitoring.
- Delivered real-time dashboards, AI-generated summaries, and collaborative workspaces for knowledge sharing and decision support.



Tech Stack

Azure OpenAI | GPT Models | LangGraph | LangChain | Python | Node.js | Angular | SQL Server (EDM) | PostgreSQL | MongoDB Atlas | Redis | ECharts

Outcome Delivered

Eliminated dependency on SQL specialists

Improved visibility into drilling performance, operational risks, and well data quality

Accelerated engineering analysis and decision-making through conversational AI and automated data exploration

Reduced time spent on data retrieval and reporting

Enabled self-service analytics through real-time dashboards, AI summaries, and automated visualizations

Automated Well Log Interpretation Using Machine Learning

AI-powered electrofacies classification that enables faster reservoir characterization

Challenges

- Manual well log interpretation is time-consuming and highly dependent on specialist expertise.
- Large volumes of petrophysical data make electrofacies identification difficult to scale.
- Inconsistent interpretation methodologies can impact reservoir evaluation outcomes.
- Delayed subsurface insights slow exploration and development decisions.

Our Solution

- Developed an ML-powered solution to automate electrofacies classification from well log datasets.
- Trained machine learning models using historical well data to identify reservoir facies patterns.
- Automated facies prediction, visualization, and interpretation workflows.
- Enabled rapid generation of well log strip logs and subsurface insights.
- Delivered an interactive platform to support faster reservoir evaluation and geological analysis.



Tech Stack

Machine Learning | Electrofacies Modeling | Well Log Analytics | Reservoir Characterization | Python | Interactive Visualization Platform

Outcome Delivered

Accelerated well log interpretation through automated electrofacies classification

Reduced dependency on manual analysis

Improved reservoir characterization accuracy through ML-driven pattern recognition

Enhanced scalability by processing large well log datasets efficiently

Enabled faster subsurface decision-making

Well Reports Insights: GenAI-Powered Drilling Intelligence

Transforming drilling and completions reporting through AI-powered insights.

Challenges

- Critical drilling and completions information was distributed across large volumes of daily well reports and operational records.
- Engineers spent significant effort reviewing reports to identify operational issues, trends, and performance insights.
- Manual analysis delayed decision-making and reduced responsiveness to field conditions.
- Valuable knowledge remained buried within unstructured operational documentation.

Our Solution

- Developed a GenAI-powered Well Reports Insights assistant to analyze drilling and completions reports using natural language interactions.
- Automated extraction of operational summaries, events, risks, and performance trends from daily reports.
- Enabled engineers to ask questions and receive contextual insights without manually reviewing extensive documentation.
- Delivered conversational search, report summarization, and drilling intelligence through a unified interface.
- Improved access to operational knowledge and accelerated insight generation across drilling programs.



Tech Stack

Generative AI | Natural Language Processing (NLP) |
Operational Intelligence Platform | Conversational AI
| Drilling & Completions Data Repository

Outcome Delivered

Accelerated report analysis through AI-generated summaries and contextual insights.

Reduced manual effort spent reviewing daily drilling and completions reports.

Improved operational visibility into drilling activities, performance trends, and field events.

Enabled faster decision-making through real-time access to critical operational intelligence.

Unlocked knowledge from unstructured reports, improving information accessibility across teams.

Enhanced collaboration through a shared, conversational interface for drilling intelligence

Fortune 1000 Public Utility Company

Smart Restoration Predictor Agent

Agentic AI-powered outage intelligence delivering dynamic restoration predictions, improving customer communication, and enabling data-driven outage management



Challenges

- Estimated Restoration Times (ERTs) were often inaccurate and based on limited historical data.
- Manual estimation processes reduced operational efficiency during outage events.
- Customers received limited visibility into changing restoration timelines.
- Field condition changes were not reflected dynamically in outage communications.

Our Solution

- Designed a Smart Restoration Predictor Agent to continuously forecast outage restoration timelines.
- Leveraged historical outage data, real-time grid conditions, and crew progress updates to improve prediction accuracy.
- Enabled dynamic adjustment of ERTs as field conditions changed.
- Delivered real-time outage insights to support operational planning and customer communication.

Tech Stack

AWS Lambda | Amazon S3 | Amazon SageMaker |
Python ML Models | Snowflake | Real-Time Outage
Data Integration | Weather Data Enrichment |
Crew & Outage Feature Engineering | Tableau
Analytics & Reporting

Outcome Delivered

50%+
improvement in
restoration time
accuracy

Real-time ERT
updates improving
customer
communication

Reduced manual
estimation effort

Increased trust in
outage resolution
timelines

Fortune 1000 Public Utility Company

Proactive Outage Forecasting Agent

Agentic AI-powered outage forecasting enabling utilities to anticipate disruptions, improve preparedness, and proactively protect grid reliability

Challenges

- Outage management relied primarily on reactive response models.
- Limited ability to predict outages caused by weather and grid conditions.
- Large-scale outage events often resulted in inadequate operational preparedness.
- Resource deployment decisions were made with limited predictive visibility.

Our Solution

- Implemented a Proactive Outage Forecasting Agent to predict potential outages before occurrence.
- Leveraged weather patterns, grid vulnerabilities, and asset conditions to identify high-risk zones.
- Enabled early detection of outage risks and recommended preventive actions.
- Supported proactive planning and resource allocation through AI-driven forecasting



Tech Stack

AWS Lambda | Amazon S3 | Amazon SageMaker | Python ML Models | Snowflake | Real-Time Outage Data Integration | Weather Data Enrichment | Crew & Outage Feature Engineering | Tableau Analytics & Reporting

Outcome Delivered

Enabled proactive outage management strategy

Improved preparedness before major weather events

Reduced outage occurrences in vulnerable grid zones

Optimized resource pre-positioning

Fortune 1000 Public Utility Company

Comprehensive Duration Analyzer Agent

Agentic AI-powered outage duration intelligence enabling accurate restoration planning, faster decision-making, and optimized resource allocation

Challenges

- Inaccurate outage duration estimates impacted restoration planning and customer communications.
- Limited data-driven prioritization led to inefficient response management.
- Multiple outage variables made restoration forecasting complex and difficult to manage.
- Misaligned planning often resulted in delays and operational inefficiencies.

Our Solution

- Developed a Comprehensive Duration Analyzer Agent to predict outage duration using multiple contextual parameters.
- Analyzed outage type, severity, location, weather conditions, and resource availability to improve forecast accuracy.
- Enabled data-driven prioritization of restoration activities across multiple outage events.
- Delivered intelligent outage planning insights to support faster and more effective restoration operations.



Tech Stack

AWS Lambda | Amazon S3 | Amazon SageMaker |
Python ML Models | Snowflake | Real-Time Outage
Data Integration | Weather Data Enrichment |
Crew & Outage Feature Engineering | Tableau
Analytics & Reporting

Outcome Delivered

25%
reduction in average
outage duration

Enabled data-driven
prioritization of
restoration efforts

Reduced delays
caused by misaligned
planning

Improved operational
planning across
multiple outages

Fortune 1000 Public Utility Company

Crew Optimization & Simulation Agent

Agentic AI-powered workforce optimization enabling smarter crew deployment, faster outage response, and improved field productivity

Challenges

- Inefficient crew allocation impacted outage response effectiveness.
- Limited scenario planning capabilities hindered preparedness for large-scale outages.
- Poor utilization of field resources increased operational costs and response delays.
- Manual planning made it difficult to optimize crew size, skills, and routing strategies.

Our Solution

- Implemented a Crew Optimization & Simulation Agent to dynamically model outage scenarios and optimize crew deployment.
- Leveraged AI-driven simulation to determine optimal crew size, skill requirements, and dispatch strategies.
- Enabled intelligent routing and resource allocation based on outage conditions and operational priorities.
- Delivered scenario-based planning capabilities to improve response coordination and workforce utilization.



Tech Stack

AWS Lambda | Amazon S3 | Amazon SageMaker | Python ML Models | Snowflake | Real-Time Outage Data Integration | Weather Data Enrichment | Crew & Outage Feature Engineering | Tableau Analytics & Reporting

Outcome Delivered

Optimized crew allocation across outage incidents

Improved field productivity and response coordination

Reduced resource wastage & manual planning effort

Enabled scenario-based strategic planning for outages

25%
Increased in crew utilization

5X
Faster dispatch time