FEATURE ARTICLE | AI IN O&G SUPPLY CHAIN

From Fragmented Data to Decision Intelligence: A Practical AI Roadmap for O&G Supply Chain Teams

By Nitin | *Enterprise AI & Digital Transformation Lead, Aays*

Data fragmentation is a near-universal challenge in oil and gas supply chain operations. Across upstream, midstream, and downstream environments, procurement and planning teams routinely contend with SAP exports that conflict with field system records, material master data with duplicate or free-text entries, and maintenance histories distributed across multiple generations of software.

This is not an exceptional situation — it is the norm. And it need not be a barrier to AI adoption. Organisations that have made meaningful progress with AI in procurement and supply chain did not wait for a clean data environment. They treated data quality as an ongoing workstream and built capability incrementally, with measurable value at each stage.

This article outlines that progression — from data foundation to an agentic AI decision intelligence platform — and the practical considerations at each stage.

Stage 1: Establishing a Governed Data Foundation

Data governance in O&G tends to stall when treated as a discrete project rather than operational infrastructure. Progress requires embedding cleansing, standardisation, and lineage tracking into existing pipelines — not periodic remediation exercises.

A practical entry point is to scope to a single high-value data domain: MRO material master, procurement spend, or maintenance work orders. Common data quality issues at this stage include:

- Free-text SAP MM material descriptions that are inconsistent and unsearchable at scale
- Supplier names recorded in multiple formats across procurement and finance systems
- Maintenance data siloed across IBM Maximo, SAP PM, and legacy CMMS platforms with no shared reference
- Field consumption data submitted manually via spreadsheet, outside core ERP workflows
- Absence of a unified cost centre or asset hierarchy linking spend to operational context

Addressing one domain end-to-end — cleansed, standardised, and connected to a governed pipeline — delivers more actionable insight than a broader, less disciplined initiative across multiple domains simultaneously.

Stage 2: Building Decision-Oriented BI and Semantic Models

With a governed data foundation in place, the next priority is structured visibility. Semantic models translate raw operational data into business-relevant constructs — cost centres, asset hierarchies, vendor performance tiers — that support consistent interpretation across teams.

Dashboard design should be anchored to decisions rather than data availability. A procurement dashboard showing spend variance against budget by asset — with drill-through to the underlying purchase orders — supports a different and more actionable set of decisions than one showing aggregate spend by vendor. The semantic layer produced at this stage also serves as the foundation for the AI capabilities described below.

Stages 3 & 4: Agentic AI and Decision Intelligence

Once a governed semantic layer is established, organisations are positioned to extend their data investment into agentic AI — tools that operate continuously, reason across existing models, and surface relevant decisions to appropriate stakeholders without requiring manual querying or dashboard navigation.

One approach to this capability tier is a two-component architecture: a query-driven intelligence layer for broad stakeholder access, and a proactive recommendation engine for decision support. Implementations of this type — such as the AADI platform developed by Aays — illustrate what is achievable when a well-governed semantic layer is in place:

AADI Conversa — Query-Driven Intelligence	AADI Recommender — Proactive Decision Support
Handles 3 categories of questions:	**Illustrative alert scenario:**
1. Descriptive — What & Which 'What was our MRO spend at Facility 3 last quarter vs. budget?' 'Which vendors had the most emergency POs in Q2?'	Alert: Pump seal consumption at Facility 3 will exceed reorder point in 18 days. Impact: ~$340K in lost production uptime if unaddressed.
2. Root Cause Analysis — What & Why 'Why did maintenance spend spike in March?' 'What drove the emergency PO increase at the Houston site?'	Recommendation: Raise PO for 12 pump seal units (critical tier). Financial upside: Avoiding this stockout preserves an estimated $340K–$520K in deferred production value.
3. Simulation & Recommendation 'If we reduce safety stock on rotating parts by 15%, what is the stockout risk over the next quarter?' 'Which contracts are up for renewal in 90 days and what is the spend exposure?'	Action options: Notify Procurement Lead | Escalate to T&I Manager | Log feedback | Defer with reason All from a single interface — no dashboard navigation required.

The query layer addresses a structural limitation of conventional BI: most of an organisation's workforce lacks the tools or training to extract insight from dashboards without analyst support. A natural language interface supporting descriptive queries, root cause analysis, and scenario simulation extends data access to any stakeholder — in the terminology they already use.

The proactive recommendation layer addresses a different gap. Standard dashboards are retrospective — they report conditions that have already occurred. A well-implemented recommendation engine identifies developing KPI changes before they become operational issues, quantifies each in financial terms, and routes structured actions to the relevant decision-maker. In practice, this means recommendations are grouped into defined categories and sub-layers by report and time period, each accompanied by a dollar-denominated impact estimate.

The Recommender in Practice: An Illustrative Scenario

A supply chain manager opens their morning briefing. Rather than reviewing multiple dashboards,
they receive the following structured alert for the Refinery Operations report:

1. MRO stockout probability — rotating equipment parts: HIGH (14 items below reorder point)
   Financial exposure: ~$340K–$520K in deferred production value
   → Action: Expedite PO for 6 critical items | Notify: Procurement Lead

2. Contractor maintenance spend: tracking 18% above monthly budget
   Financial exposure: ~$210K projected overage by month-end
   → Action: Review Work Order #4471 for scope variance | Notify: Finance Controller

3. Turnaround readiness: 61% against 80% target — 21 days to T&I commencement
   Financial exposure: ~$1.2M per day in lost throughput if start date is affected
   → Action: Escalate 4 long-lead procurement items | Notify: T&I Manager

Each recommendation links to the underlying data, states the dollar exposure, and provides one-click options to notify stakeholders, log feedback, or defer with a reason.

The Four-Stage Progression

The following summarises the capability and operational value available at each stage of the journey:

Stage	Capability	Operational Value
Stage 1	Data Foundation	Cleanse and unify O&G data from SAP, historians, and field systems into a governed data lake

Stage 2	BI & Dashboards	Build semantic models and decision-oriented dashboards for procurement, operations, and finance teams
Stage 3	AADI Conversa	Natural language Q&A across semantic models — descriptive, root cause, and simulation queries for any stakeholder
Stage 4 ★	**AADI Recommender**	Proactive KPI monitoring with dollar-quantified recommendations, structured by category, routed to the relevant stakeholder

Each stage delivers standalone value and creates the preconditions for the next. Organisations do not need to reach Stage 4 to realise meaningful returns. The critical discipline is treating each stage as a production deployment rather than a pilot, and maintaining continuity of data investment across the progression.

A Note on Implementation Experience

The progression outlined in this article reflects an approach I have applied across enterprise engagements in energy, CPG, and financial services — working with operators at varying stages of data and AI maturity. The use cases and scenarios referenced are drawn from those implementations, and the AADI platform described in Stages 3 and 4 is in active production deployment across several of those organisations.

In my experience, the most consistent predictor of success is not the sophistication of the technology selected, but the discipline applied to the earlier stages — particularly data governance and semantic model design. Organisations that invest adequately in Stages 1 and 2 reach the agentic AI tier faster, with fewer rework cycles and more reliable outcomes.

About the Author

Nitin is an Enterprise AI & Digital Transformation Lead at Aays, an AI-first analytics consultancy that delivers production-grade data and AI platforms for Fortune 500 companies. Aays provides end-to-end capability — from data engineering and semantic modelling to agentic AI deployment — across energy, CPG, financial services, and healthcare.