VIEW POINT

FROM MONTH-END FIRE DRILLS TO REAL-TIME CONFIDENCE

APPLIED AI FOR HYDROCARBON ACCOUNTING IN UPSTREAM OIL AND GAS

Infosys | CONSULTING

Introduction

Upstream Oil and Gas industry has spent billions digitizing the oilfield – connecting wells, automating SCADA, and deploying digital twins. Nevertheless, the critical bridge between the field operations and the financial truth, i.e., Production Accounting, remains fragile.

Production accounting, often referred to as Hydrocarbon Accounting or Allocation (HCA), is one of the most important activities in upstream operations. It's how companies turn raw field data, such as, meter readings, tank levels, well tests, lab results and tickets into the official numbers used for regulatory reporting, entitlements, and financial close. It is the cash register, the value realization engine – but we run it like a back-office report filing cabinet. For many years, companies have relied on established third party software to manage these workflows. For a long time, these tools were the only practical way to handle the complexity of allocations and reconciliations. Yet, when the allocation logic that only one person understands and buried in a spreadsheet fails in a new scenario, when a minor rounding error manifests months later as a Prior Period Adjustment (PPA) triggering joint venture disputes and regulatory fines, the business is exposed. Many operators often carry a backlog of PPAs stretching back 12-24 months and a single misallocated offshore well can result in $5M-$10M in disputed revenue. In many companies, people who wrote these calculations have moved out and companies spend anywhere between $200K to $500K just to maintain these systems, not including the license costs. Civil penalties can accrue daily and can be material depending on violation type and jurisdiction; for example, ONRR penalty schedules include daily maximums that can exceed tens of thousands of dollars per day for certain willful violations, and even inadvertent errors result in massive interest payments on unpaid royalties.

We are entering an era where AI has evolved beyond chatbots or fancy automation. We now have capabilities that can, write code, interpret data from binary document formats, spot patterns in time-series data, generate explanations and learn business rules, all while working with humans in each step.

As someone who has spent decades consulting in Production Accounting, and witnessed the recurring stress that production engineers and allocation analysts face every month end - I keep coming back to a simple question:

"Can some parts of production accounting be handled differently—especially in areas where we still rely heavily on manual effort or complex customizations?"

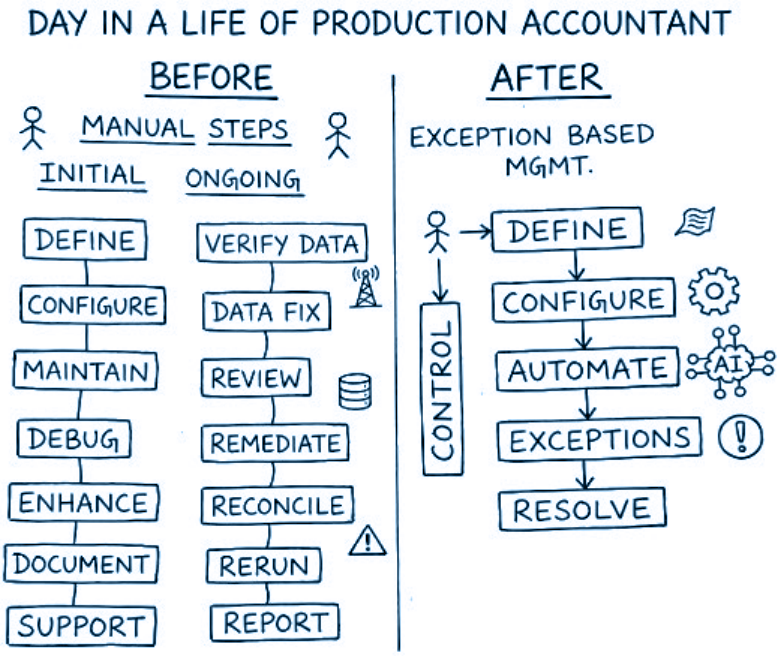

Image 1: Shifting Production Accounting from manual set up and verification to exception based management.

This is not about replacing existing product vendors. Instead, it's about considering whether AI driven, in house capabilities, could complement or simplify what companies already have and give them more flexibility and control; thereby turning production accounting from a forensic exercise to a real time strategic advantage.

Where Modern AI Might Help

Below are a few areas where early-stage patterns are observed in pilots, where the outcomes vary by the asset complexity and data quality.

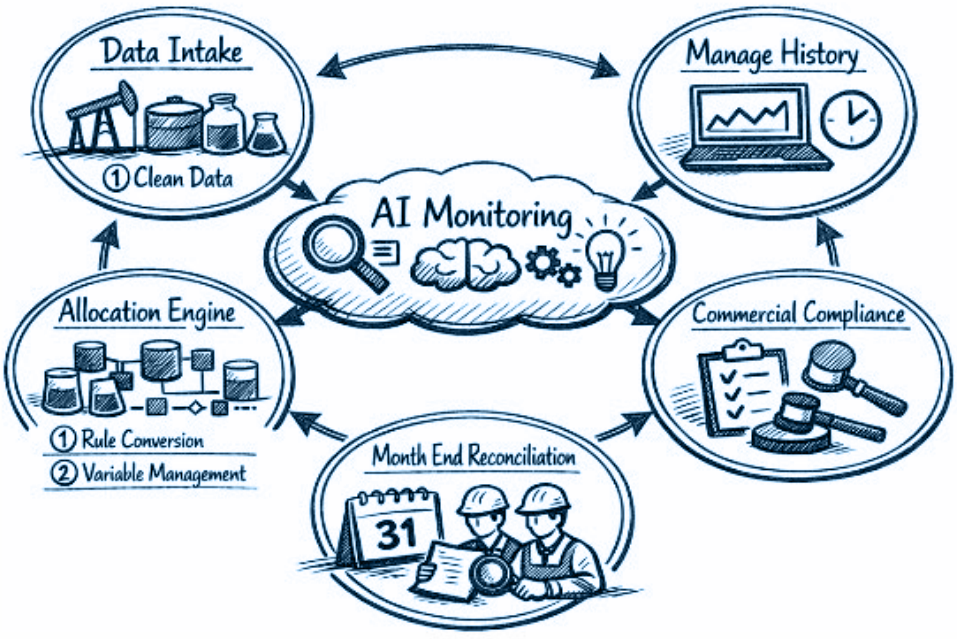

Image: Applied AI for Production Accounting Life Cycle

1. Clean up raw data before it enters the accounting process

Production accounting teams spend a surprising amount of time fixing stuck or incorrect meter readings from historian tags, identifying missing data, synthetic theoretical estimation, cross-checking field tickets and checking for inconsistencies between SCADA and LACT meters.

Modern AI can help here by detecting unusual patterns in the data, suggesting corrections, extracting structured data from PDFs, images, and handwritten notes, and highlighting only the issues that truly need human attention. This acts like a "smart filter" before the data flows into existing accounting software.

Here the input layer will need to have multi-modal ingestion with built-in connectors for standard protocols like OPC UA for SCADA, FTP for Lab reports, and APIs for other measurements. A normalization or a pre-processing agent will standardize the parameters like units and timestamps.

Agentic core of this 'Smart Filter' is a controlled agent that will require a Composite AI approach blending deterministic rule engine, probabilistic models, and generative reasoner. The deterministic engine can be generated through AI code generation, which will enforce hard constraints (physics and contracts) like – maximum water cut and zero flow rate for a 'closed' value status. The probabilistic model will detect anomalies based on history by applying techniques like LTSM for time-series forecasting. This will determine outliers or deviations and flag them. The Generative reasoner is the 'agent' layer that looks at the previous outputs and decides what to do. It will use LLM with RAG to observe non-correlations, retrieve context, reason the impacts, and take actions to either fix the data or mark the reasons. The agent layer also assigns confidence scores which will be used for making decisions about auto-correct or suggestion or just alert.

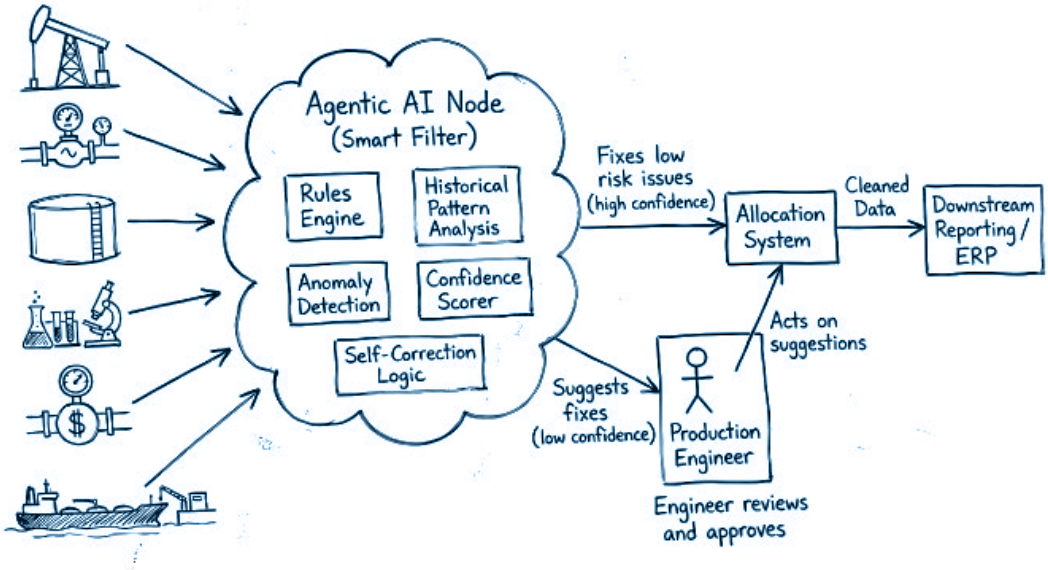

A self-correction loop will follow reinforcement learning. This will accept feedback from the production engineer and learns to either make or stop making suggestions accordingly. Not just learning, such an Agentic node will explain its work with a "why" and "what," and this transparency is what delivers business trust.

2. Accelerate Deployment of allocation engine

Configuring allocation logic is one of the hardest parts of production accounting. Major effort and cost are spent while configuring allocation networks, variables, and calculations. In many configurations, these variables and formulas live in Excel templates, MathML constructs, or SQL procedures. As assets grow and commingle, the complexity multiplies—and small setup issues may only surface years later, often during audits or commercial reviews.

2.1 Tame Configuration Complexity (Excel/MathML/SQL):

Allocation logic often lives in a mix of Excel templates, MathML expressions, and SQL procedures, making the setup hard to maintain, as assets and rules evolve. Modern AI can simplify this by turning constrained natural language policies into a canonical rule model to draft SQL/Python with built in tests, validating network structures as graphical schematics, and creating clean, centralized config files from canonical allocation models. Such claims were riskier in the past, but Modern AI code generators use Agentic Thought Process to self-critic and self-correct. With human oversight, this can be a powerful tool for coding and configuring complex calculations. It is critical to note the deterministic guardrails – AI suggests the logic using policies but the code it generates is deterministic SQL/Python that can be audited, and the review gates decide acceptance. AI can also autogenerate documentation and catch missing elements or inconsistent assumptions early, reducing configuration errors and improving long term

2.2 Manage Large Variable Sets:

Production accounting involves hundreds of input variables such as theoretical and measured values from wells and fields, storage, export/fiscal data, laboratory samples, well tests, fallback drivers, shrinkage factors, and water cuts, each with different schedules and rules. AI can help by automatically generating data dictionaries (measured, derived, lab, fallback, constant), identifying relationships between variables (lab sample applies only to certain streams and periods), and recommending standard rule patterns. This keeps variable handling consistent across assets and prevents subtle errors caused by outdated or mismatched inputs.

3. Ensure compliance

On top of configuration risk, production accounting often requires accuracy to the third decimal because even small volume differences have financial implications—especially once currency conversions and price volatility are included. So, results must be traceable, explainable, and reproducible.

3.1 Precision and Rounding (to the Third Decimal):

Production accounting requires precise math, often to the third decimal, as small rounding differences can translate into real financial impacts. AI can validate correct decimal handling, apply consistent rounding rules, and generate unit tests directly from written policy. It can also analyze where precision may be lost in the calculation chain and suggest safer ordering of operations. It can ensure standardized decimal handling across all calculations and rounding at the correct stages.

3.2 Currency, Pricing, and Valuation Volatility:

Entitlements and valuations depend on pricing formulas, FX rates, and contractual terms that can vary by date and source. An AI based valuation governance layer can help encoding rate hierarchies from the policy, detecting anomalies in market data, and pulling applicable contract clauses using retrieval augmented search. It can simulate valuation scenarios, thus improving the accuracy and defensibility of financial calculations tied to allocated volumes.

3.3 Full Traceability and Explainability:

Allocation results must always be explainable - internally, to partners, and to auditors. Audit-defensible AI can automatically build calculation lineage graphs showing every step from raw inputs to final volumes, generate evidence packs with all relevant data and assumptions, and explain results or variances in plain language. It can also provide transparency by showing exactly which rules applied and why each well received its allocated volumes, based on the recorded lineage or metadata.

4. Manage History and Reduce Errors

4.1 Reduce Setup Errors (and Find Them Early):

Because allocation rules are coded manually, errors can stay hidden for years making the fixes hard and time consuming. AI can improve reliability by generating stress tests, comparing results to known baselines, scanning code for risky constructs, and running "shadow mode" comparisons against existing systems. It can continuously verify logic over time and alerts users when the results drift unexpectedly, helping teams catch configuration defects much earlier.

4.2 Prior Period Adjustments (PPAs):

When new information arrives for the past month, production accounting often requires recalculating that month and all following months—a complex and time consuming process. AI can model the allocation process as a dependency graph, instantly determining which months need re runs, and maintain bitemporal (valid time vs transaction time) data structures to preserve history. It can automate PPA workflows, summarize impacts across wells and periods, and explain the financial effects clearly, making adjustments faster, cleaner, and easier to govern.

5. Assist with reconciliation and month-end close

Production engineers and reporting analysts spend a lot of time during month end reconciliation as they need to ensure all systems line up, remediate the discrepancies, and support all explanations with evidence.

AI could help by comparing different data sources automatically, finding the likely root causes for discrepancies, providing the right references (work orders, lab results, proving reports) and drafting clear explanations for review by partners and regulatory agencies.

Most production accounting tools are built around a month-end cycle. AI can layer on top of this by providing provisional allocation estimates during the month, summarizing trends, highlighting assets that need more attention and explaining what changed and why. This gives operations and accounting teams better visibility without altering the formal reporting process.

6. A balanced approach to incorporate AI into Production Allocation

A growing number of operators are considering building certain AI-driven capabilities internally, with support from technology partners. Their intent is not to replace existing products, but to incorporate unique business needs, reduce cycle time for changes and reduce costs to compliance. Most companies won't shift away from third party tools entirely, and they should not. These systems are stable, trusted, and required for compliance.

A more realistic path is to keep the core systems as-is and use AI as an additional layer around data quality, analysis, explanations, and workflow. Build pieces that offer unique value or reduce recurring effort, by working with tech partners to develop safe, well-governed AI components. To ensure safe and repeatable outcomes, AI agents used in production accounting must operate under explicit control mechanisms rather than fully autonomous. Agent behavior should be constrained through bounded scope of action and clearly defined inputs and outputs derived from written policy. Techniques such as rule-constrained execution, versioned prompts, and mandatory human approval for key actions help ensure agent recommendations are reproducible across runs and retain human leadership over the AI behavior.

This way, companies can explore the benefits of modern AI without taking unnecessary risks or disrupting vendor relationships.

Conclusion

Oil and Gas Production Accounting as a subject area is a suitable candidate for Applied and Explainable AI. This should be an area of interest to not just the production operators but also the software product firms that currently provide production accounting systems. While operating companies can incorporate these modern AI techniques into their production allocation workflow, the real acceleration at industry level can be brought in by the product companies through incorporating AI into their systems, and by technology consulting firms through designing an AI driven production accounting platform with built in components consumable by operating companies. While the core calculations remain deterministic in the form of scripts, AI's role lies in augmentation and pattern detection around the core allocation engine.

Technology is no longer the bottleneck; the opportunity lies in its application. The real competitive advantage will shift to the operators who start validating these on their own data. Small scale, domain specific pilots offer the clearest mechanism to unlock this value, turning production accounting into a proactive strategic asset.

About the Author:

Nandakumar Muthukrishnan
Business Consulting

Nandakumar Muthukrishnan is a Business Consulting leader specializing in the Upstream Oil and Gas sector. He began his upstream career developing scripts for Top Down Reservoir Modelling, and over the next two decades, he led large-scale technology programs across the entire upstream value chain. Today, he advises global Oil and Gas companies on technology-driven transformation and heads Infosys Consulting in India for the Energy, Utilities, and Services industries.

For more information, contact consulting@infosys.com

Infosys® | CONSULTING