

# The people and the money are in different buildings.

We scored every position in energy by AI exposure. The biggest workforce is where the smallest dollars are. The industry has been staring at the wrong building.

By [Raj Mistry](https://rajmistry.com) (https://rajmistry.com) · 404 positions · 6 layers · BLS/industry-sourced employment · Inspired  
by [Karpathy](https://www.youtube.com/watch?v=c3b-JASoP10) (https://www.youtube.com/watch?v=c3b-JASoP10)

---

## 2.5

### ROUGHNECK

Pipe, mud, steel. AI is not great at climbing things.

## 8.5

### PRODUCTION ACCOUNTANT

Same spreadsheet. Same format. Rebuilt from scratch every month. Nobody knows why.

---

### ABOVE THE FIELD

## ~90%

of every AI-exposed dollar in energy sits in people who never touch a wellhead. Analysts building reserve reports. Associates packaging lender decks. Coordinators assembling regulatory filings. The people everyone ignores when they talk about "AI in energy."

---

### FIVE THINGS WE FOUND (30 SECONDS)

- 
- 01 **The mismatch is the whole story.** About 29% of energy workers are in the field. Less than 10% of AI-exposed wage dollars are there. That's a 3:1 mismatch between where the people are and where the money is. In finance, they'd call that a mispricing.
- 
- 02 **Titles survive. Hours don't.** The board still signs. The engineer still attests. Nobody gets a new title. But the 70% of their week they used to spend assembling slide decks and reconciling spreadsheets? That compresses. The assemblers are exposed, not the decision-makers.
- 
- 03 **The upside isn't efficiency.** Cheaper analysis doesn't make the same memo faster. It lets you run the forty scenarios you've been skipping — because each one used to cost three engineer-weeks, and nobody has three spare engineer-weeks. Four scenarios become forty. That's not a faster decision. It's a different decision.
- 
- 04 **Three workflows already pencil.** Treasury/lender readiness, ownership and title, regulatory planning. You own all the inputs. You own all the document assembly. You don't need permission from the regulator, the lender, or the board to start. Just someone willing to try it. In energy, that's the hard part.
- 
- 05 **Energy is AI's own bottleneck.** 2,300 GW sitting in U.S. interconnection queues. An estimated 60–70% of that wait is document and regulatory work. Compress the paperwork and more data centers come online. More data centers mean more AI compute. AI needs energy, and energy is stuck behind a pile of paper that AI could compress. There's a joke in there somewhere.
- 


**You're thinking "just give me the bullets." We know.**

---

Eighteen months ago, you couldn't feed a 200-page reserve report into a model and get anything useful back. You can now. That's the whole timing argument. The gap between "models can do this" and "your competitor already is" tends to close faster than anyone expects.

---

ⓘ All scores are modeled estimates. Every claim is tagged by confidence level. [Full methodology below.](#)



Left: where the people are. Right: where the exposed dollars concentrate. The lines cross. Ninety cents of every AI-exposed wage dollar sits above the field – seventy-four cents even when restricted to BLS-anchored rows only.

**WHAT COULD BREAK THIS**

**Faster packets ≠ faster decisions.** A compressed borrowing-base package still waits for the VP's calendar, the lender committee, and the bank's own reserve engineers. Maybe 30% of cycle time is work product. The rest is organizational friction AI doesn't touch.

**Bad source data breaks everything.** Well logs from the 1970s were hand-transcribed. Conflicting spreadsheets, stale type curves, vintage assumptions. AI on bad data produces bad results faster.

**Weakly trusted outputs increase review burden.** If the model's work isn't trusted, reviewers spend more time checking than they saved. The net effect can be negative until confidence builds.

**How robust is the ~90% figure?**

Below: what happens when we stress field-layer headcount and compensation — the two inputs most likely to be underestimated. Non-field employment estimates and compensation proxies are held constant; stressing those would require a broader sensitivity test. The pattern holds across the range tested because the compensation gap between layers dominates the math.

Field emp ↓ / Field comp →	0.75× comp	Base comp	1.25× comp	1.5× comp
0.5× headcount	96%	95%	94%	93%
0.75× headcount	94%	93%	91%	89%
Base headcount	93%	90%	88%	86%
1.5× headcount	89%	86%	83%	80%
2× headcount	86%	82%	79%	76%

Rows: field employment multiplier (how much larger or smaller the physical layer headcount might actually be). Columns: field compensation adjustment. The base case uses the current model assumptions. Even at the extremes — 2× field headcount at 1.5× field pay — the above-field share still shows the majority of exposed dollars sitting outside physical operations.

**EXTENDED SENSITIVITY: ABOVE-FIELD PARAMETER STRESS**

What happens when we vary the parameters the *above-field* layers use? Rows: multiplier on all non-Physical layer bases. Columns: multiplier on all non-Physical compensation proxies. Both directions weaken the result by making above-field layers less dominant.

Above-field base ↓ / comp →	0.5× comp	0.75× comp	Base comp	1.25× comp
0.5× headcount	70%	78%	82%	85%
0.75× headcount	78%	84%	87%	90%
Base headcount	82%	87%	90%	92%
1.5× headcount	87%	91%	93%	95%

**What drives the ~90% — and what survives stress**

Three factors build the above-field concentration. The decomposition below shows how much each contributes (under the score-neutral employment model), and the stress tests show what happens when we deliberately weaken the most attackable assumptions.

**FACTOR DECOMPOSITION**



Even with uniform compensation and uniform compressibility, ~71% of estimated headcount sits above the field (under score-neutral employment). Compensation proxies add ~9 percentage points. Compressibility scores add another ~11. The direction is established by headcount distribution alone — the other factors amplify it.

**STRESS TESTS**

Scenario	Above field
Baseline — all 404 positions, score-neutral employment	90.3%
Flatten compensation to \$80K for all layers (removes the pay-gap amplifier between Physical at \$58K and Governance at \$180K)	85.2%
Cap compressibility at 7.0 for all rows (assumes above-field roles are less compressible than scored)	89.0%

Both: flat compensation + capped scores (harshes combination — removes both amplifiers)	83.4%
Anchor rows only (60 of 404, BLS/industry-sourced) (excludes all formulaic employment estimates)	74.2%
Roles only — exclude workflows and artifacts (373 of 404) (addresses the mixed-ontology objection)	89.4%

The above-field share ranges from 74–96% across all stress tests (parametric Monte Carlo: 75.5–96.1%; anchor-row-only test: 74.2%). Even under the harshes tested specification (BLS-anchored rows only, excluding all formulaic estimates), roughly three of every four exposed wage dollars still sit above the field. The directional finding is not an artifact of the employment model, the compensation proxies, or the scoring method.

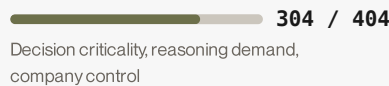
**WHAT COULD BREAK THIS**

344 of 404 employment estimates use a deterministic formula (layer base × tier multiplier) rather than observed data. Three tier bands (0.75×, 1.0×, 1.25×) assigned by stable hash provide structural spread, not an economic model. The compensation proxies are layer-level averages, not role-specific wages. If physical-layer headcount is substantially larger than modeled, or if above-field compensation proxies are too high, the ~90% figure compresses — but the stress grid shows it stays above 74% even when physical headcount doubles and physical comp increases 50%. The anchor-row test (74.2%, using 60 BLS/industry-sourced rows only) provides the hardest floor. The direction holds across all specifications. The exact percentage is scaffold-dependent — treat it as an interval (74–96%), not a point estimate.

**COMPRESSIBILITY SCORES**



**AUGMENTED AXES**



**EMPLOYMENT ANCHORING**



**WHAT THIS DOES NOT CLAIM**

*(a.k.a. the section most research puts in an appendix nobody reads)*

**Modeled, not observed.** Every score is informed estimation. No before/after workflow data exists yet. Validation against deployed systems is planned for a future edition.

**Exposure is not replacement.** A high compressibility score means the prep stack compresses. The title, the signoff, and the judgment remain human.

**Scores rank directionally.** A 0.4-point difference between two roles is noise, not signal. Read tiers, not decimals.

**Employment and comp are proxies.** 60 of 404 positions use BLS-sourced employment. The remaining 344 use a formulaic estimate (layer base × tier multiplier) — structural scaffolding, not survey data. Compensation uses six layer-level proxies (\$58K–\$180K), not role-specific wages. The layer-level pattern is the claim; individual role numbers are not precise.

**External approvals stay external.** Compressing a borrowing-base packet does not make the lender approve faster. Compressing a rate-case filing does not make the PUC rule sooner. Company-controlled loops compress. Externally governed outcomes do not.

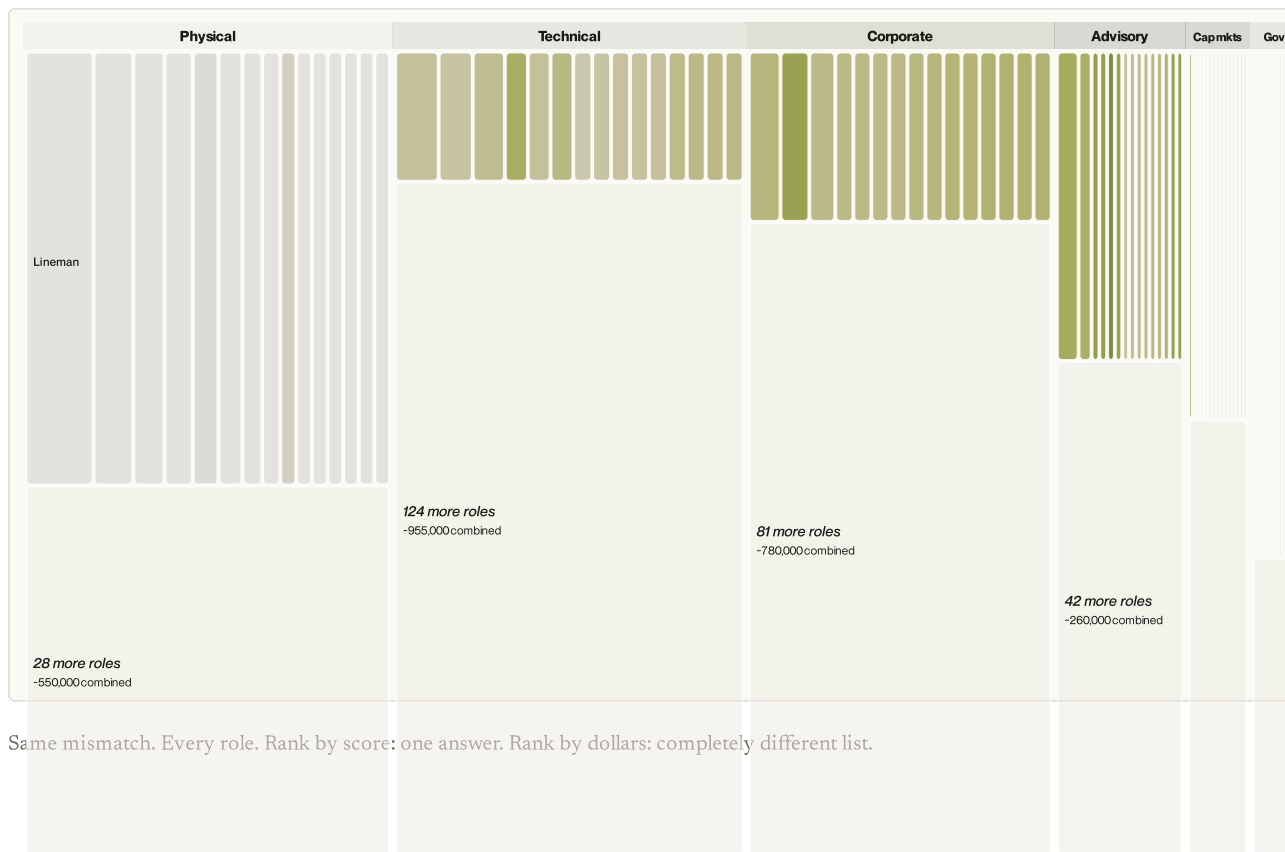
*Energy is finance in 1975.*

\$3.3 trillion in annual investment,  
almost all of it driven by individual  
judgment, individual spreadsheets,  
individual analysts running  
four scenarios when they should  
run forty.

(The firms doing this in finance in 1975 — Kidder Peabody, Drexel, Salomon Brothers — don't exist anymore.)

Beeswarm

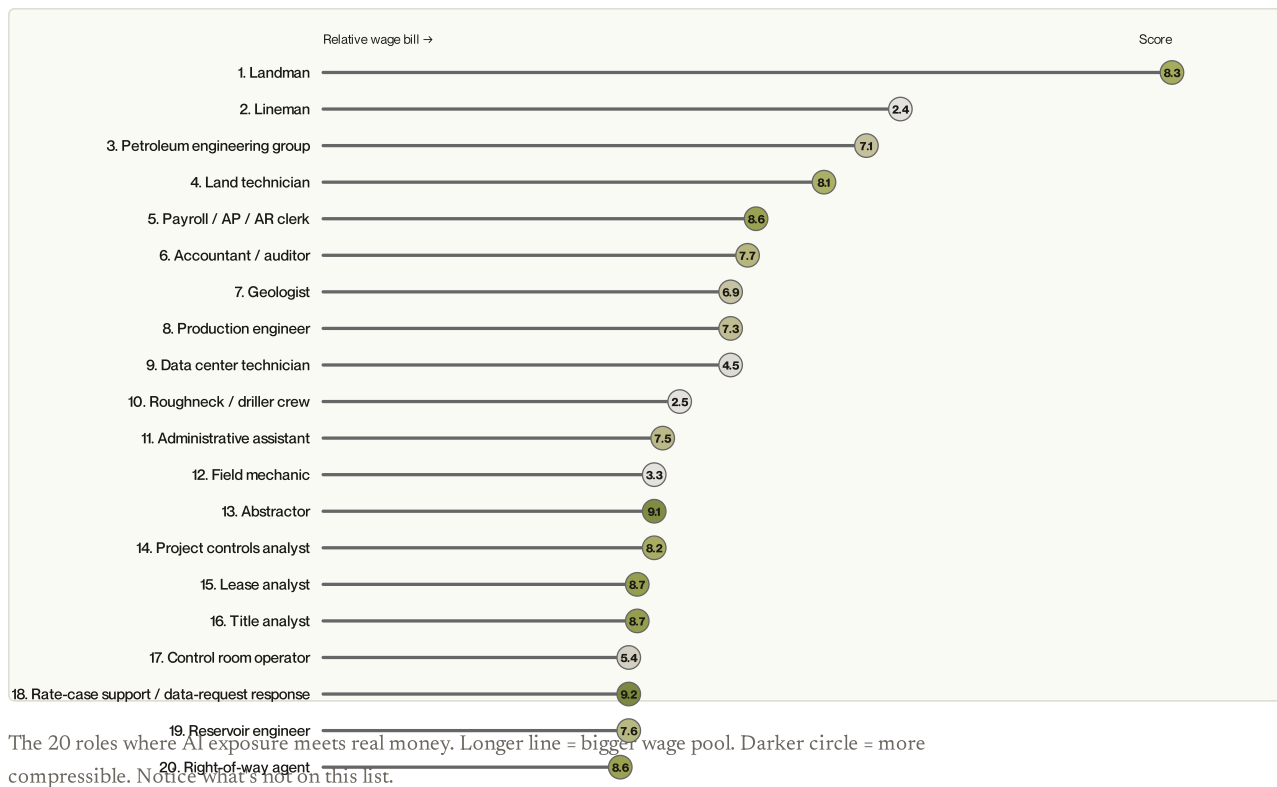
Treemap



Same mismatch. Every role. Rank by score: one answer. Rank by dollars: completely different list.



- By wage bill
- By score
- By employment



Karpathy's framing applies here directly: *"To get the most out of the tools that have become available now, you have to remove yourself as the bottleneck."* In energy, the bottleneck today is evidence assembly. When that compresses, the constraint shifts to judgment – which is where you want engineers spending their time.

**TWO PATHS (HYPOTHESIS — EACH REQUIRES VALIDATION)**

**Path A · Copilot**

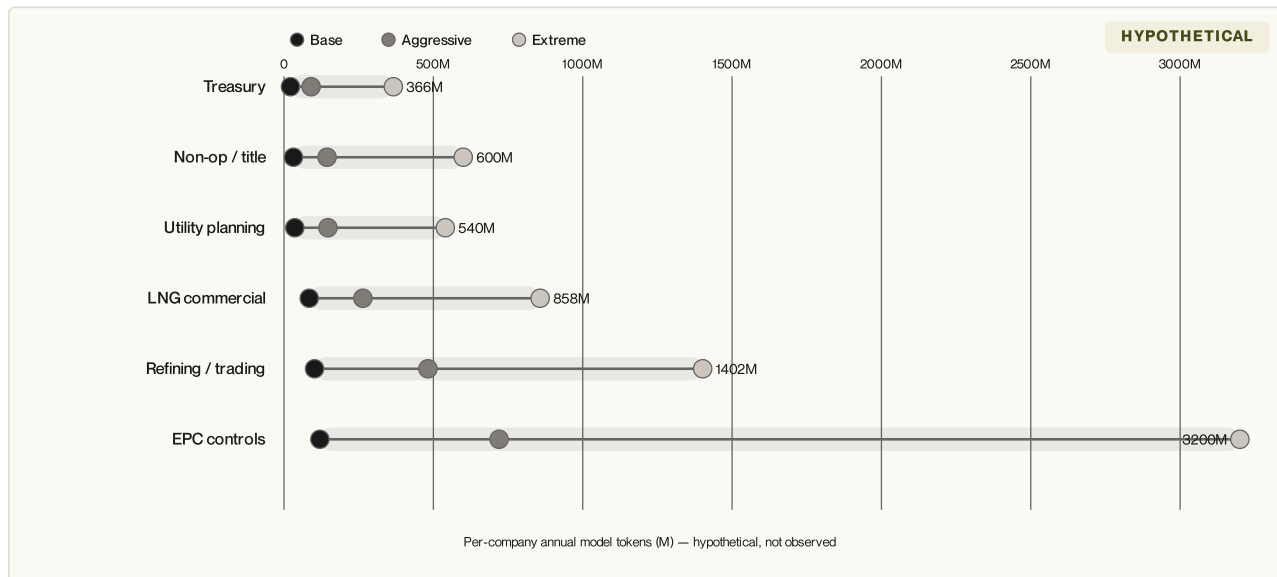
Same workflow, faster. Analysts run 4 scenarios in 3 days instead of 3 weeks. The decision stays the same. Cost savings: 10–20% on analytical labor.

**Path B · Intelligence**

Different structure entirely. The system composes 40 scenarios from capabilities, flags the 3 that change the decision, and delivers them before the

cycle starts. The decision changes. Value: 10–100× the labor savings.

**The gap today:** No published evidence for either path in energy. Financial trading precedent supports Path A. A future edition targets the question: do additional scenarios actually change the final decision, or does the team anchor on the first? (If you've sat in a management committee, you already know the answer.)



Annual model-token scenarios per operating unit. Hypothetical — not observed. But the shape of the curve is the argument: demand doesn't flatten when analysis gets cheap. It compounds.

**THE JEVONS TRAP — AND WHY G&A IS THE WRONG DENOMINATOR**

Here's what most energy companies do first with AI: email drafts. Meeting summaries. Formatting. Which makes sense — nobody gets fired when a model hallucinates a meeting recap. But it misses the point by about 40:1.

A mid-size E&P spends \$50–100M on G&A. It spends \$2–6B on development capital. That's a 30-to-1 ratio. The entire AI-for-energy pitch is aimed at the smaller number. It's like hiring a personal shopper to optimize your grocery bill while your investment portfolio picks itself.

Jevons noticed this in 1865 with coal: make it cheaper and people don't use less of it, they use more. Same thing happens with analysis. Make it cheap and the company doesn't just do the old analysis faster — it asks questions it couldn't previously afford to ask. Four scenarios become forty. The team runs the sensitivity work it's been skipping because each additional case costs three engineer-weeks and nobody has three spare engineer-weeks. Total demand for analytical work goes *up*, not down.

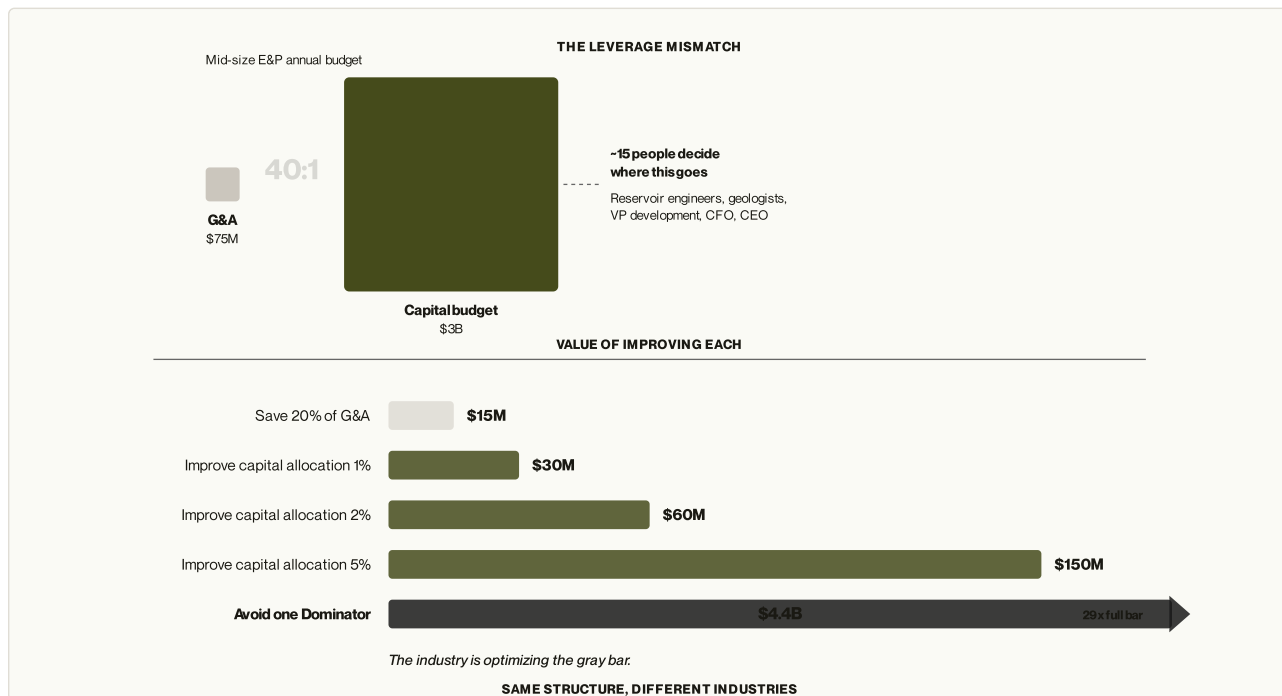
Which is why this is not an efficiency story. Saving 20% of G&A on \$75M is \$15M. A 2% improvement in capital allocation on a \$3B development program is \$60M — four times the G&A savings the industry's writing press releases about. The industry is optimizing the small number and ignoring the big one.

In 2019, Concho Resources drilled 23 wells at 230-foot spacing on the Dominator pad — roughly a third of what the rest of the Delaware Basin was running. The thesis was more wells per section, more resource recovery. What happened was well interference: production was tracking 38% below the rest of Concho's Wolfcamp projects — and headed toward 45% below pre-drill estimates. The stock dropped 22% in a single day — \$4.4 billion in market cap. The team that selected the spacing — reservoir engineers, geologists, a VP, a CEO — earned a combined \$3–5M a year. They destroyed a thousand times their own compensation in one decision with inadequate scenario coverage.

— CONCHO RESOURCES, DELAWARE BASIN, 2019

That's the leverage structure of this industry. A handful of knowledge workers control a capital budget that dwarfs their salaries by orders of magnitude. The Dominator wasn't a failure of intelligence. It was a failure of scenario coverage — and scenario coverage was rationed because analysis was expensive. Make analysis cheap, and the question changes from "can we afford to run more cases?" to "can we afford not to?"

This leverage structure isn't unique to oil. A frontier AI lab decides which \$100M+ training run to fund — five people in the room. TSMC allocates \$52–56B in fab capex for 2026. Hyperscalers are deploying \$600B+ in data center infrastructure this year, 75% of it for AI — and the analysis that determines whether that capacity comes online in 2028 or 2031 runs through the same energy permitting workflows mapped in this benchmark. Same structure every time: small team, massive capital budget, decision quality as the binding constraint. Energy is the industry where we counted the roles.



Mid-size E&P (\$75M G&A, \$3B development capital). G&A savings assume 20% compression. Capital allocation improvement assumes same acreage, better scenario coverage. Dominator reference: Concho Resources, Delaware Basin, 2019. Cross-industry capex figures from company guidance and analyst

estimates (2025–2026).

### THE OWNERSHIP PROBLEM

If the value is in the capital allocation decision – not the meeting summary, not the email draft, not the reformatted slide deck – then the person who owns that decision has to own the AI deployment. Not "approve it." Not "sponsor it." Own it. Be the user. Personally.

Right now, most energy companies hand AI to the IT department or to a "digital innovation" team two reporting layers removed from anyone who touches a capital decision. That's how you get chatbots for the help desk and summarizers for the weekly all-hands. Useful. Worth maybe \$15M a year. And completely disconnected from the \$3B development program where the real leverage sits. Everyone writes a LinkedIn post about digital transformation. Meanwhile the analysts are still manually building variance tables for a billion-dollar credit facility.

The CFO who reviews four scenarios that someone else built is a different person than the CFO who directs a system that runs forty and flags the three that change the decision. Same title. Completely different leverage on the capital budget.

This is why the Concho loss is structural, not anecdotal. The team that selected 230-foot spacing had the domain expertise. What they lacked was scenario coverage – and scenario coverage was rationed because the analytical cost came out of G&A, while the decision it informed controlled billions in capital. The budget owner for the analysis and the budget owner for the decision were in different buildings. AI doesn't fix that misalignment. Organizational design does.

The practical version: energy companies that succeed with AI need someone in the room where capital gets allocated who also controls the AI deployment that feeds that decision. Call it what you want – Chief Decision Officer, VP of Decision Intelligence, the CFO who actually uses the tools. The title doesn't matter. What matters is that the person directing AI and the person signing the AFE are the same person, or sit next to each other. When they're three org layers apart, AI gets aimed at the G&A line. When they're the same person, it gets aimed at the capital budget.

Anyway, the 40:1 gap is just an org chart problem wearing a technology costume. The models are ready. The question is whether the person who benefits from better decisions is the same person who controls the AI budget. At most energy companies today, they're not. That probably changes at some point.

#### DRI CASE · TREASURY AND LENDER READINESS

A mid-market E&P with a \$1.2B reserve-based lending facility runs two borrowing-base redeterminations per year. Each cycle: 3 analysts, 3 weeks, rebuilding the same variance tables and lender Q&A packets from scratch. The lender asks the same fifteen questions every cycle. Most teams haven't structured last cycle's work product for reuse.

**One owner. One cycle. One budget line.** Direct labor: ~\$35K/year. AI inference: ~\$2K/year at API pricing (~\$20K with orchestration overhead). The direct savings are modest. The real value: freed analysts run 15 additional sensitivity scenarios on the next development program –

scenarios the team has been skipping because each one cost three engineer-weeks. On a \$500M–2B facility, a 1% improvement in borrowing-base utilization from better scenario coverage = \$5–20M.

**What stays human:** negotiation posture, lender relationship, representation of downside scenarios, signoff.

**What the system composes:** source comparison, variance tables, covenant extraction, Q&A draft generation, package assembly – delivered before the cycle starts, not rebuilt from scratch. [Full beachhead breakdown →](#)


## The job that's easiest to automate is almost never the job where automation matters most.

Score by technical exposure and you get one ranking. Score by dollars and you get a different one. That's the whole problem.

### EXHIBIT 5 · VALUE-RISK FRONTIER · 30 ROLES

**In energy, the biggest AI upside lives next to real downside.**

Everybody wants the lower-right corner of this chart: high value, low risk. Look at it. It's empty. Every role where AI creates massive value also carries real downside. That's not a bug – that's energy. Value and consequence have always traveled together here. The deploy-first corridor isn't where risk is zero. It's where the math works anyway.



X = value creation potential. Y = asymmetric downside risk. Size = automation exposure. The empty lower-right is the chart's most important feature – you can't deploy AI where it's most valuable without accepting real downside. That makes this a governance problem, not a technology problem.

THREE DIMENSIONS OF AI IMPACT · 30 ROLES

**One score hides the answer.  
Three scores reveal it.**

A reservoir engineer barely moves on automation but lights up on value creation. A roughneck scores low on automation but the downside risk is catastrophic. Score on one axis and you automate the wrong things. Score on three and you see which roles are actually worth touching first. These are the 30 we modeled in depth.

Search roles — try 'reservoir', 'landman', 'board', 'trader'...

30 of 30 roles

Advisory

Capital markets

Corporate

Governance

Physical

Technical

[↓ EXPORT 30 ROLES AS CSV](#)

## DEPLOYMENT

**Start where AI already pays for itself.**

This is not a vision deck. Three workflows you can start next quarter. Inputs named. Gates identified. Compression you can measure. You own the document assembly. The external parties (lenders, counterparties, regulators) own the outcome – but you don't need their permission to start the analytical pilot. Just someone willing to try it.

The gap between "models can do this" and "your company actually does this" is where all the early-mover advantage is. It's closing. These three are where to start – not because they're exciting, but because they can pay for themselves in a single cycle.

BEACHHEAD 01 · SEMI-CONTROLLED

**Start with the packet every lender asks for**

Borrowing-base support, covenant monitoring, lender Q&A, amendment packages

**WHAT FEEDS IN**

Reserve reports, production exports, price decks, hedge schedules, covenant definitions, prior lender memos

**WHAT COMPRESSES**

Source comparison, variance tables, covenant extraction, Q&A draft generation, package assembly

**WHAT STAYS HUMAN**

Negotiation posture, lender relationship management, representation of downside scenarios, signoff

**SUCCESS METRICS**

Package cycle time, analyst hours per cycle, questions answered inside 24 hours, error escapes

**DECISION CONTEXT** **Borrowing-base readiness · \$500M–2B credit facility**

**The first-principles question:** The borrowing-base package exists because lenders can't continuously monitor collateral. If AI makes continuous monitoring cheap, the semi-annual cycle disappears – the lender gets a live dashboard and the "package" becomes unnecessary. You can make the cycle 5x faster. Or you can ask why the cycle exists at all.

**Back-of-envelope economics (not audited):** 3 analysts × 3 weeks × 2 cycles/year ≈ \$35K in allocated labor. AI inference cost ≈ \$2K–20K/year. Direct savings: modest. The real value isn't the labor delta – it's the 15 additional scenarios the freed analysts can run, and what that does to the quality of the credit decision on a \$500M-2B facility. A 1% improvement in borrowing-base utilization on \$1B = \$10M.

See the model process a synthetic borrowing-base packet end-to-end. Every step inspectable.

**Operator memo**

Full workflow breakdown, data requirements, implementation path, signoff boundaries

**Eval specification**

How to measure success — metrics, baselines, pass/fail criteria, test structure

**Deal readiness memo**

Commercial positioning, deployment economics, buyer qualification criteria

BEACHHEAD 02 · SEMI-CONTROLLED

**Attack the paper trail that keeps cash and decisions stuck**

JOAs, AFEs, JIB review, division orders, curative, title chains

**WHAT FEEDS IN**

JOAs and amendments, lease and deed chains, title opinions, curative notes, division orders, AFEs, election notices

**WHAT COMPRESSES**

Clause extraction, document comparison, exception queue triage, owner-response drafts, curative clustering

**WHAT STAYS HUMAN**

Legal interpretation on edge cases, negotiation with counterparties, escalation on title risk

**SUCCESS METRICS**

Suspense resolution time, title queue aging, JIB dispute cycle time, % auto-triaged

**DECISION CONTEXT** Net revenue interest accuracy · \$50–500M working-interest exposure

**What would prove this wrong:** If AI can't reliably parse the clause structure of a 40-page JOA – specifically conditional consent provisions and pooling elections – the comparison step doesn't compress. Current models handle this for standardized JOAs but struggle with pre-2010 custom agreements. Test it on your worst JOA before committing.

BEACHHEAD 03 · SEMI-CONTROLLED

**Shorten the prep stack before the next rate case**

IRPs, rate-case exhibits, testimony support, discovery responses

**WHAT FEEDS IN**

Forecast workbooks, generation scenarios, depreciation schedules, historical filings, commission discovery requests

**WHAT COMPRESSES**

Exhibit assembly, testimony draft support with citations, discovery routing, consistency checks

**WHAT STAYS HUMAN**

Policy judgment, regulatory strategy, witness preparation, external advocacy

**SUCCESS METRICS**

Turnaround on data requests, revision count, witness-prep load, consistency escapes

**DECISION CONTEXT** Allowed return on equity · \$1–10B+ regulated rate base

**What would prove this wrong:** If a PUC rejects AI-assisted testimony or discovery responses – not because they're wrong, but because the commission won't accept the process. Regulatory conservatism is the binding constraint. Test with a non-controversial filing first, not the general rate case.

**Six things AI won't fix. No matter what the vendor**

tells you.

### External clocks

AI cannot speed up the regulator's calendar

You can assemble the rate-case exhibits in two days instead of two months. The PUC still takes fourteen months to issue an order. AI compresses time to prepare. Time to decide (agency review, public comment, commission deliberation) runs on a calendar you don't control.

### Organizational friction

It will not stop people from waiting on each other

The borrowing-base takes three weeks because the VP is traveling, the bank wants a different format, and the geologist is arguing with the reservoir engineer about type curves. Maybe 30% of cycle time is work product. The rest is waiting. AI compresses the 30%.

### Data quality

It will not rescue bad source data

Well logs from the 1970s were hand-transcribed onto paper by a guy named Earl. (Seriously. His name was usually Earl.) AI on bad data produces bad results faster. Garbage in, garbage out – now at the speed of light.

### Independence

The machine is not the signer

Reserve auditors exist because lenders require independent attestation. The bank doesn't care how smart your model is. They care that a human with a PE license signed the page.

### Never alone

Some calls stay human because the downside is physical

Final legal opinions, auditor conclusions, safety-critical field calls. These stay human. Not because the industry is slow to change. Because a wellhead blowout doesn't have an undo button.

### The downcycle

In a downturn, the value proposition flips

This analysis is implicitly mid-cycle. In a downturn, the prep stack gets gutted by layoffs before AI touches it. The value proposition flips: not "compress the work" but "maintain analytical capability after the RIF." The company that cut 40% of finance in 2020 and has AI can still run the analysis. The one that cut 40% without it can't.

WHAT ACTUALLY CHANGES INSIDE A ROLE

### Inside every role, the same split.

Some tasks shrink. Some disappear. Some become more valuable because the prep bottleneck is gone. The table below decomposes five roles into their task layers – then shows how time and value restructure when the assembly work gets cheap.

ROLE	TASKS AI COMPRESSES ■	TASKS AI AMPLIFIES ■	TASKS THAT STAY HUMAN ■	TIME SHIFT
<b>Reservoir engineer</b> Score: 7.1	Decline curve fitting, type-curve generation, reserve report assembly, data gathering from production databases, variance commentary drafts	Scenario comparison (can now run 40 instead of 4), sensitivity analysis across price/decline/spacing, pattern recognition across analogue wells	Subsurface judgment calls, well spacing decisions, reserve certification signoff, risk framing for the board	<del>60/15/25</del> <b>10/45/45</b>
<b>Treasury analyst</b> Score: 8.2	Borrowing-base packet assembly, covenant compliance checks, lender Q&A drafts, amendment redlining, data reconciliation across systems	Exception detection (catches covenant breaches earlier), cross-cycle comparison (persistent memory across redeterminations), scenario stress testing on covenants	Lender negotiation posture, downside framing, final representations, signoff authority, relationship management	<del>70/40/20</del> <b>15/40/45</b>
<b>Land / title analyst</b> Score: 7.8	Division order calculation, lease abstraction, JIB exception screening, curative document comparison, ownership chain reconciliation	Cross-asset title pattern matching (flag similar defects across properties), historical exception memory (recalls curative outcomes from prior cycles)	Curative negotiation, title opinion judgment calls, counterparty relationship management, legal liability decisions	<del>65/40/25</del> <b>10/40/50</b>
<b>Regulatory analyst (utility)</b> Score: 6.9	Exhibit assembly, discovery response drafting, data request compilation, testimony support document preparation, precedent citation lookup	Cross-docket pattern analysis (identifies commissioner tendencies), IRP scenario modeling (more alternatives evaluated), consistency checks across multi-year filing history	Regulatory strategy, testimony delivery, commissioner relationship management, settlement negotiation, policy judgment	<del>55/45/30</del> <b>10/40/50</b>
<b>Lineman</b> Score: 2.4	Paperwork: daily job briefing forms, time entry, incident reporting templates, outage documentation	Predictive routing (AI optimizes storm response dispatching), outage pattern recognition (learns from historical restoration sequences)	Climbing, switching, grounding, live-line work, safety assessment, crew leadership, storm response decisions	<del>40/5/85</del> <b>3/10/87</b>

Time-split estimates (compress / amplify / stays-human) are modeled from FERC procedure timelines, E&P workflow audits, and utility rate-case filing breakdowns. Pre-AI splits represent the approximate share of a full-time role's hours. Post-AI splits assume production deployment, not pilot.

### The role does not vanish. The economics of the role change.

The time-shift column tells a story that aggregate compressibility scores miss. When a reservoir engineer's time allocation moves from 60/15/25 to 10/45/45, the role doesn't disappear – it restructures from an assembly job into a judgment job. The economic consequence: the value per hour of that person's work increases because each hour now

involves scenario evaluation or decision-making rather than data compilation. A role that was 60% commodity work and 25% judgment work becomes 45% judgment work. The judgment didn't get cheaper — it became a larger share of the output. (This is the part the headcount-reduction crowd misses entirely.)

The amplification paradox, at the role level. A treasury analyst who spent 70% of their hours assembling borrowing-base packets was being paid mostly for document work. Post-compression, that same person spends 45% of their time on lender negotiation and downside framing. Same hourly rate. Different economic content per hour. The role's strategic value to the organization goes up even as headcount stays flat.

The lineman row is the control case. When only 10% of a role's time is in the compress column, AI changes the paperwork and the dispatch routing but not the job. The time shift (10/5/85 to 3/10/87) is real but marginal. This is why physical-layer roles cluster below 3.0 in the benchmark — there's not enough compressible time to restructure the role's economics.

#### FAILURE MODE TO MONITOR

The "stays human" column is only durable if organizations actually invest the freed-up time in judgment and analysis rather than simply reducing headcount. The risk: a company compresses its five-person treasury team to two people, but keeps the same workflow volume — so each remaining person rubber-stamps AI outputs at 3× speed instead of reviewing them at appropriate depth. The amplification column only works if people have the time budget to fill it. If compression translates to headcount cuts rather than role restructuring, the centaur model degrades into an automation model with a human-shaped rubber stamp at the end. This is an organizational design choice, not a technology constraint.

The pattern across all five: the **compress** column is document assembly and data reconciliation — it's where the hours are today and where AI makes the most immediate difference. The **amplify** column is scenario depth and pattern recognition — things the human could always do but couldn't afford to because prep consumed their capacity. The **stays human** column is judgment, liability, relationships, and physical execution. The time shifts quantify the centaur trade: organizations buy judgment capacity by spending less time on assembly. The question is whether they use that capacity or simply eliminate it.

WORKFORCE · SCENARIO

THREE-PART FRAMEWORK

### Work doesn't disappear. It migrates.

"Which jobs does AI replace?" is the wrong question. The right one: what part of the work compresses, what human capability gets amplified, and what new work shows up because AI now exists? The labor evidence so far points toward redesign, not extinction. The functions show up before anyone invents a title for them.

#### COMPRESSES

##### The assembly layer inside the job

Drafting, reconciliation, packaging, routing, repetitive evidence assembly. Not "the whole job." The prep work underneath the judgment. Where four analysts assembled borrowing-base

#### AMPLIFIES

##### Judgment seats get more leverage

Judgment, negotiation, relationship management, signoff, testimony, exception handling, field decision-making, political and regulatory sense-making. These people don't disappear. Their leverage increases because the

#### EMERGES

##### A new control layer around the machine

Every workflow AI automates creates a new control problem. Someone has to teach, audit, route, and sign for machine output. The first new jobs are

packets, two review AI-assembled packets and spend the freed hours on scenario analysis.

---

Analysts, packagers, coordinators, memo-builders, context routers, reconciliation specialists

prep work around them gets cheaper. As analysis gets cheap, judgment gets expensive.

---

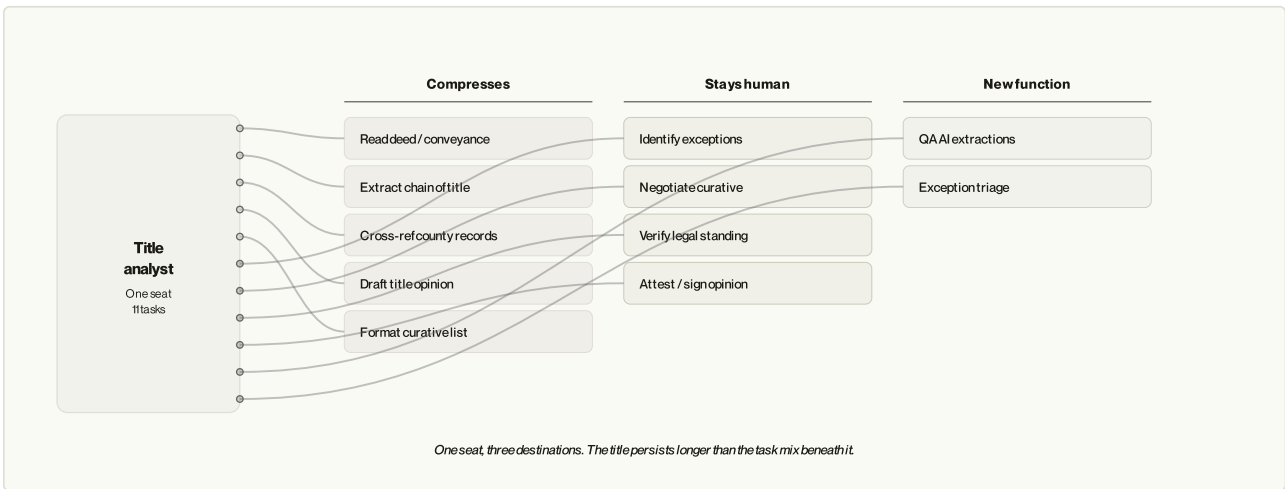
Approvers, negotiators, operators, witnesses, relationship owners, field decision-makers

not sci-fi. They are the people who ensure the machine's work is trustworthy enough to act on.

---

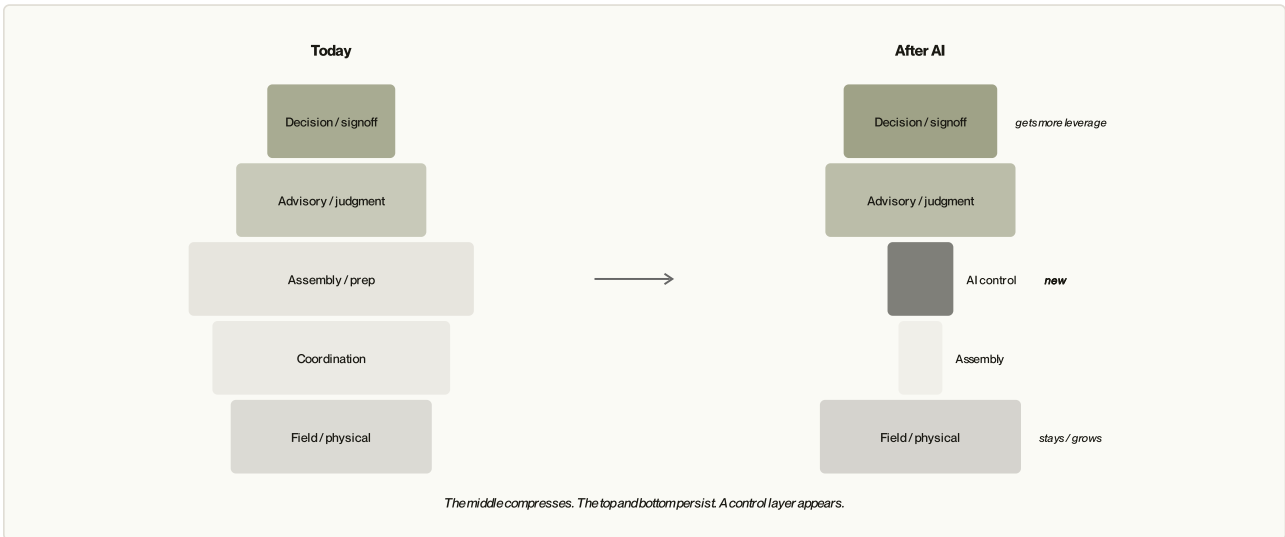
Evidence architects, exception managers, eval/QA leads, workflow owners, provenance leads, apprenticeship stewards

**DON'T ASK WHETHER AI REPLACES THE JOB**  
 Ask which part of the seat *was* the job.



Each tile is a task inside one role. Some flow to AI. Some stay with the human. A few recombine into new functions. The seat is being unbundled – not eliminated.

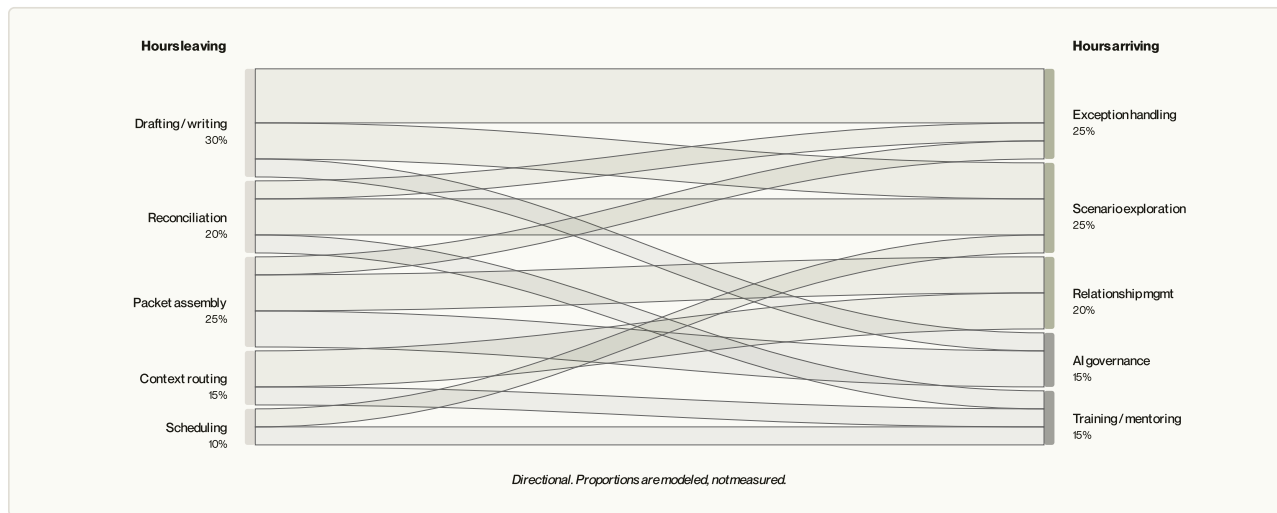
**AI DOESN'T DELETE THE ORG CHART**  
 It redraws the center.



The old pyramid thins in the middle. The decision layer stays. The field layer stays or grows. A new thin control layer appears around the machine.

**THE HOURS DON'T DISAPPEAR**

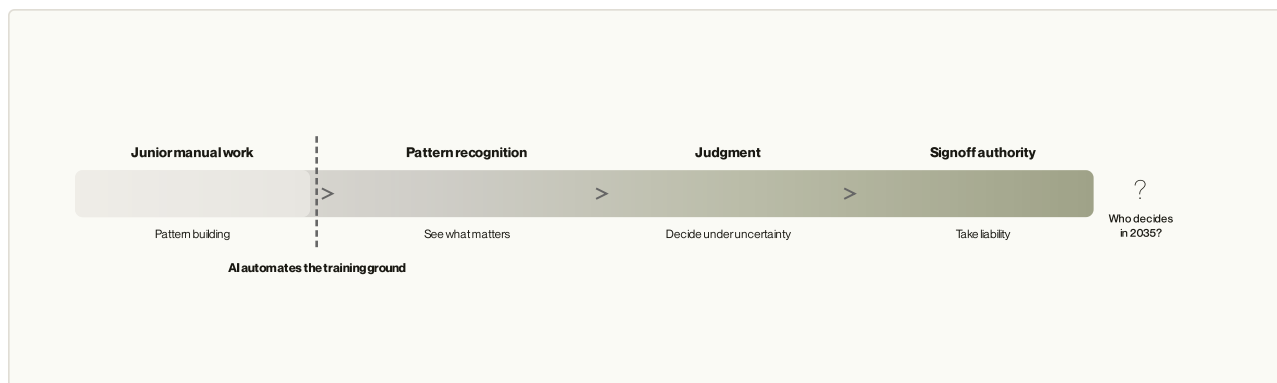
They move to higher-consequence work.



Hours leave drafting, reconciliation, and packet assembly. They arrive in exception handling, scenario exploration, and AI governance. The hours don't vanish – they migrate upward.

**YOU CAN AUTOMATE THE TRAINING GROUND OUT OF EXISTENCE**

Then who makes the decisions in 2035?



Junior manual work builds pattern recognition, which builds judgment, which earns signoff authority. AI slices out the junior layer. The firm now has a problem: how does it produce future decision-makers?

**NEW WORK FUNCTIONS, NOT JOB FORECASTS**

We estimate new *work*, not new *jobs*. Functions appear before titles do. The demand model:

**Scenario expansion** Work generated by running 40 scenarios instead of 4. In treasury, this creates an evidence architect who makes each redetermination cycle cheaper than the last.

<b>Exception adjudication</b>	The 5% of cases that create 95% of delays. In title/non-op, this is the exception manager who handles curative failures that AI flags but cannot resolve.
<b>Eval / governance</b>	Trust infrastructure. In regulatory, this is the provenance/QA lead who prevents citation escapes and witness-prep chaos before a rate case.
<b>Knowledge maintenance</b>	Keeping the institutional memory alive as senior staff retire. The workflow owner who ensures AI-assisted processes encode domain knowledge, not just patterns.
<b>Training / apprenticeship</b>	Across the company, the apprenticeship steward who keeps junior staff learning judgment, not just becoming prompt operators.

These functions deliver business value, not soft benefits. An evidence architect brings compounding memory. An exception manager brings downside protection. An eval lead brings the trust that lets you actually deploy. A workflow owner brings adoption and budgeted ROI. An apprenticeship steward brings future judgment supply.

#### THE NET EMPLOYMENT MATH

A mid-size operator creates 3–7 FTEs across emerging functions while restructuring 20–50 above-field positions. Net headcount likely declines. Per-person value rises. Total payroll may stay roughly flat. The IEA's 1.7-to-1 retirement-to-entrant ratio means in many cases AI-driven compression doesn't cause layoffs – it prevents the capability gap from widening. A 50-person technical team that loses 8 to retirement and replaces 5 (with AI assistance) maintains roughly the same output. HR calls this "transformation." Everyone else calls it Tuesday.

**2025–2026:** Pilot phase – emerging functions handled as side responsibilities. **2027–2028:** Production phase – workflow architect and eval roles become distinct positions. **2029+:** Regulatory pressure formalizes governance and provenance roles. The firms that benefit most treat this as workforce restructuring, not simple automation.

#### CONFIDENCE: SCENARIO

These work functions are projected from benchmark structure and cross-industry labor evidence, not observed in operating companies today. The pattern – technology shifts create control work around the seams of the technology – shows up every time (database administrators before databases; DevOps before cloud; compliance officers before Sarbanes-Oxley). We are confident the functions will exist. We are not confident about timing or headcount.

*The winners will not be the firms with "fewer people everywhere." They will be the firms with fewer people rebuilding the same packet, more people governing faster decisions, and a deliberate system for producing future judgment.*

#### DECISION LEVERAGE

**Your first workflow depends on what kind of company you run.**

Eight company types. Eight different entry points. The shale E&P starts in treasury. The utility starts in regulatory. The PE fund starts in portfolio monitoring. Start in the wrong place and you burn six months building a demo that gets polite applause and zero budget.


Each one is a DRI-shaped problem: a single owner, a recurring cycle, a budget line, and company-controlled prep work. No regulator, lender, or board permission needed to start the analytical pilot. Click any row.

COMPANY TYPE	HIGHEST-TORQUE APPLICATION	DECISION SCOPE	WHY IT'S ASYMMETRIC
<b>Shale E&amp;P</b> ▶	Development planning scenario acceleration	\$1-5B capex	5% better capital allocation = \$50-100M
<b>Integrated major</b> ▶	Cross-BU portfolio optimization	\$20-30B capex	Portfolio misallocation compounds over decades
<b>Utility</b> ▶	IRP scenario acceleration	\$10-50B fleet	Shapes the proceeding vs. being shaped by it
<b>PE fund</b> ▶	Portfolio monitoring and IC memo acceleration	\$2-10B AUM	Faster pattern recognition across 10-30 portcos
<b>Investment bank</b> ▶	CIM and diligence acceleration	\$500M-5B deals	3-week CIM → 1-week CIM = more mandates, faster close
<b>Midstream operator</b> ▶	Commercial contract management	\$500M-3B rev	Missed escalators and volume commitments leak margin
<b>LNG exporter</b> ▶	Cargo optimization and destination arb	\$100M-1B / yr	\$5-20M margin per cargo diversion decision
<b>Power developer</b> ▶	Interconnection queue management	\$500M-5B pipeline	80% of queue projects never built — faster screening wins

EXHIBIT 6 · BENEFIT BY COMPANY TYPE

**Most companies save 15%.  
One type of company  
gains 10x.**


Most consultants sell "cost savings from AI." The savings are real and nearly identical across company types: maybe 10-20% on analytical labor. The strategic value (better decisions, competitive moats, new capabilities) varies by 10x. It concentrates in the company types your consultant isn't modeling. (Your consultant is modeling headcount reduction. The money is somewhere else entirely.)



Gray = efficiency gains (cost savings, time compression). Olive = strategic value (decision quality, competitive moat, new capabilities). The gray bars are roughly equal. The olive bars aren't. Scores are modeled estimates – the ranking is the claim, not the precision.

## If you're building an AI platform for energy, cost savings aren't your market. Decision leverage is.


The next three exhibits shift the denominator from headcount to token demand, recurring revenue, and capture sequence.



X = annual token demand. Y = decision value per 1M tokens. Size = addressable companies. Olive = company-controlled; amber = semi-controlled (external party governs outcome); gray = externally governed. Upper-right quadrant is where recurring model spend compounds. That's the business.



WORKFLOW	COMPANY TYPE	BUDGET OWNER	CASES / YR	ANNUAL TOKENS (M)	FRESH %	VALUE AT STAKE	EXPANSION
<b>NRC licensing + filing</b>	Nuclear / SMR	VP Regulatory	1	<b>250M</b>	70%	\$5-20B program	Safety → o
<b>IRP scenario + testimony</b>	Utility	VP Planning	1	<b>320M</b>	70%	\$10-50B fleet	Resource grid mode
<b>Data center power procurement</b>	Hyperscaler	VP Energy	10	<b>170M</b>	65%	\$1B+ campus	PPA → que
<b>Rate-case exhibits + discovery</b>	Utility	VP Regulatory	1	<b>280M</b>	65%	\$1-10B rate base	IRP → testi compliance
<b>Development plan scenarios</b>	E&P	VP Development	2	<b>150M</b>	60%	\$1-5B capex	Full-cycle portfolio o
<b>LNG cargo optimization</b>	LNG trader	Commercial Head	200	<b>340M</b>	40%	\$5-20M / cargo	Netback → → fleet
<b>CIM + buyer list + diligence</b>	Inv. bank	MD / Group Head	8	<b>160M</b>	60%	\$500M-5B deal	Fairness o integration
<b>IC memo + portco monitoring</b>	PE fund	Managing Director	24	<b>190M</b>	55%	\$2-10B AUM	LP reporti exit prep
<b>Contract abstraction + diligence</b>	Advisory / Law	Partner / GC	40	<b>140M</b>	45%	\$50M-5B deal	Issue-spot
<b>Reserve report + audit support</b>	E&P	VP Reserves	1	<b>90M</b>	50%	\$1-10B PV-10	Type curv decline an
<b>Borrowing-base / RBL package</b>	E&P	Treasurer	2	<b>120M</b>	45%	\$500M-2B facility	Covenant lender das
<b>Interconnection queue mgmt</b>	Developer	VP Development	50	<b>130M</b>	55%	\$500M-5B pipeline	Permitting
<b>Refinery plan + blend optimization</b>	Refiner	VP Supply	12	<b>110M</b>	55%	\$500M-3B throughput	Crude sch
<b>Non-op JIB / AFE / division order</b>	E&P / Non-op	Controller	36	<b>210M</b>	35%	\$50-500M WI	Curative →
<b>Trade capture / confirm / settle</b>	Trading	COO / Back-office	5,000	<b>260M</b>	25%	\$100M-1B book	PnL → cre
<b>Environmental permitting package</b>	Developer	VP Permitting	10	<b>100M</b>	60%	\$200M-2B project	NEPA → st
<b>Title chain + curative</b>	Land / Title	Land Manager	500	<b>180M</b>	30%	\$100M property	Lease → R
<b>Earnings / disclosure prep</b>	All public	CFO	4	<b>80M</b>	50%	Market cap risk	SEC → prc
<b>Midstream commercial contracts</b>	Midstream	VP Commercial	50	<b>85M</b>	40%	\$500M-3B rev	Nominatio
<b>Production accounting close</b>	E&P	Controller	12	<b>95M</b>	30%	Revenue accuracy	JIB → susp




Panel A: API cost vs. decision value (log-log), with 10:1 / 100:1 / 1000:1 leverage lines. Panel B: cumulative revenue waterfall by capture priority. Panel C: price sensitivity at \$4, \$8, \$12/M blended token rates. All figures modeled. Color scheme matches Exhibit 7.

## Energy is the bottleneck on AI's next wave.

AI compute demand needs chips. Chips need data centers. Data centers need power. Power needs permitting: scenario analysis, regulatory filings, engineering studies, contract negotiation. Twenty workflows stand between a signed lease and a spinning turbine.

A major near-term bottleneck on the AI revolution is the interconnection, permitting, and contracting stack — much of it document-heavy. As of end-2024, nearly 2,300 GW sat in U.S. interconnection queues. Only ~13% of capacity entering queues from 2000–2019 reached commercial operation by end-2024. An estimated 60–70% of queue time is document work — a Sunya modeled estimate (medium confidence), not an observed measurement — that frontier models can already accelerate. NVIDIA uses AI to design chips. Google uses AI to train AI. Anthropic committed \$50B to U.S. computing infrastructure, separately signed a multi-year partnership with the U.S. Department of Energy, and just expanded into multiple gigawatts of next-generation TPU capacity with Google and Broadcom — suggesting the frontier labs view physical infrastructure and energy access as a binding constraint on their growth. Energy — the sector that determines how much AI the world gets — is still assembling rate-case exhibits in Word.

The economic loop is visible but undeveloped: faster permitting expands compute supply, which accelerates AI deployment.



Accelerate permitting by one quarter per project. Multiply across 2,300 GW of queued capacity. Freed-up TWh become data center build-outs. Data centers become AI compute. The flywheel spins.

But only if someone closes the loop at the energy layer. OpenAI's latest round names compute as the strategic advantage. Compute needs power. Power needs permitting. Permitting is document work.

## The most advanced technology on earth is waiting on a filing cabinet.

Energy companies sit on some of the most honest signal in any industry. Every transaction, every filing, every interconnection study, every redet cycle – these aren't survey responses or ad clicks. They're facts about how capital moves through the physical world. The quality of any intelligence layer is only as good as the signal feeding it. Energy has the signal. It just doesn't have the model yet.

*A dollar spent on AI for energy workflows doesn't just improve energy decisions. It expands the compute supply AI runs on. Energy isn't another vertical. It's the physical constraint that sizes every other AI market – including one where two frontier labs already claim a combined \$39B+ run rate, company-reported, not audited.*

### WHAT COULD BREAK THIS

Faster analysis doesn't guarantee faster approvals. Lender committees, regulator calendars, and board decisions operate on their own timelines. The 60–70% document-work estimate is a modeled decomposition, not observed from project-level data. Physical bottlenecks – site execution, procurement, interconnection queue position – remain binding regardless of analytical speed. The loop is real, but the compression rate depends on which phase is actually the bottleneck for a given project.

DOCKET NO. 2024-00187-EL

### Staff Data Request Set 3, Item 14

PUC - TX

**Request:** Provide the Company's actual and projected plant additions, retirements, and transfers for each functional category for the test year and each of the five preceding calendar years. Include explanations for any year-over-year variance exceeding 10%.

**Supplemental:** For each variance explanation, identify the specific capital project(s), their FERC account classification(s), the date placed in service, and whether the investment was included in the Company's most recent depreciation study filed in Docket No. 2021-00042-EL.

**Format:** Provide in Excel format with supporting workpapers. Cross-reference to the Company's response to Staff DR Set 1, Item 7 (rate base roll-forward) and OPC DR Set 2, Item 22 (depreciation schedules).

*Response deadline: 10 business days from date of service. Objections due within 5 business days.*

AI ANALYSIS

Processing document...

STRUCTURED OUTPUT · 4.2 SECONDS

WHAT THE ANALYST WOULD DO NEXT

1

**MANUAL**

Spreadsheets and Word. Documents emailed between teams.



2

**POINT TOOLS**

ChatGPT/Claude used ad-hoc. No integration. No governance.



3

**INTEGRATED PILOTS**

AI connected to internal data. 1-3 formal workflows. Measuring time savings.



**4**  
**SYSTEMATIC**  
 AI across 5+ departments. Feedback loops. New roles.  
 ● ●

**5**  
**NATIVE-AI**  
 Continuous decision systems. Real-time scenario refresh. The Bridgewater model.  
 ← *Nobody is here yet*

<p style="text-align: center;"><b>TIER 1</b>  <b>Autonomous</b></p> <p style="text-align: center;"> <span>Payroll</span> <span>Trade confirms</span> <span>AP / AR</span> <span>Scheduling</span> </p> <p style="text-align: center;">Model tier (as of April 2026): Haiku 4.5 / GPT-5.4 Mini · ~\$1–3/M tokens · High volume, low stakes</p>	<p style="text-align: center;"><b>TIER 2</b>  <b>Draft &amp; review</b></p> <p style="text-align: center;"> <span>Lender packages</span> <span>CIMs</span> <span>Board packs</span> <span>Contract summaries</span> </p> <p style="text-align: center;">Model tier (as of April 2026): Sonnet 4.6 / GPT-5.4 · ~\$3–9/M tokens · Human reviews every output</p>
<p style="text-align: center;"><b>TIER 3</b>  <b>Assist &amp; decide</b></p> <p style="text-align: center;"> <span>Reserve analysis</span> <span>Scenario modeling</span>  <span>Portfolio optimization</span> <span>IRP planning</span> </p> <p style="text-align: center;">Model tier (as of April 2026): Opus 4.6 / GPT-5.4 Pro · ~\$15–30/M tokens · Frontier reasoning, low volume</p>	<p style="text-align: center;"><b>TIER 4</b>  <b>Human only</b></p> <p style="text-align: center;"> <span>Safety-critical ops</span> <span>Legal opinions</span>  <span>Auditor attestation</span> <span>Regulatory testimony</span> </p> <p style="text-align: center;">No model. Human judgment, human liability, human signature.</p>

Model tier mapping is directional, not prescriptive. Current frontier pricing (April 2026): Opus 4.6 at \$5/\$25 per million tokens, Sonnet 4.6 at \$3/\$15, Haiku 4.5 at \$1/\$5. OpenAI GPT-5.4 at \$2.50/\$15, Mini at ~\$0.75/\$4.50, Pro at \$30/\$180. Prompt caching and batch APIs can reduce effective cost by up to 95%.

The \$/M blended rate used in Exhibit 9 assumes a weighted mix across tiers — most token volume flows through Tier 1-2 (cheap), most decision value flows through Tier 3 (expensive). The governance tier determines the model tier. Get that mapping right and the unit economics follow.

---

EVERYTHING ABOVE: WHAT TO DEPLOY NOW

---


**EVERYTHING BELOW: WHERE THE INDUSTRY GOES**

---

If you've been through a failed "digital transformation" — and most energy executives have — here's why this is different. The last wave tried to change the workflow. New platform, new data architecture, new operating model. Eighteen months of IT integration before anyone saw a result. This doesn't change the workflow. The borrowing-base cycle still exists. The treasury team still runs it. The lender still asks the same fifteen questions. The only thing that changes is that the assembly takes two days instead of three weeks. Nobody adopts a new system. They review AI output instead of building from scratch. Correction, not creation.

That's the structural difference: digital transformation was a platform play — replace the system, retrain the team, migrate the data, hope it works. This is a workflow play — same system, same team, same data, different starting point. Instead of a blank screen, you start with a draft. The analyst who spent three weeks building the packet now spends two days reviewing it. The skills don't change. The starting point does.

The companies that survive this still have to do three things the others don't.

**SURVIVAL TACTIC 01**

Start with the workflow the team already hates. Borrowing-base assembly, not reservoir analysis. Few will fight to keep a hated process. The treasury team that dreads the three-week redet cycle is your first adopter – they'll champion anything that ends it.

**SURVIVAL TACTIC 02**

Let the AI draft. Let the human edit and take credit. Measure time saved, not people replaced. The metric is "hours back" not "heads out." The analyst who used to spend three weeks building the borrowing-base packet now spends two days reviewing the AI's draft and a week running scenarios the team never had time for.

**SURVIVAL TACTIC 03**

Instead of waiting for the reservoir engineer's spreadsheet, send them an AI-drafted variance table built from production data. They'll spend twenty minutes correcting it instead of three days building it. Correction is faster than creation. Use that everywhere.


SCENARIO: THE NUMBER NOT WIDELY REPORTED

# \$1.0<sub>T</sub>

in capital reallocated over a decade

from 40 scenarios to 4

Cumulative capital reallocated



The math is not complicated. Global energy investment reached \$3.3 trillion in 2025. A 3% improvement in capital allocation – the difference between 4 scenarios and 40 – redirects ~\$100 billion per year to higher-return projects. Over a decade: ~\$1 trillion. Not saved. Reallocated. From worse wells to better ones, from overbuilt projects to ones that should have been built instead. Every year you run 4 scenarios instead of 40, you're not losing that money to a competitor. You're losing it to a version of yourself that bothered to do the analysis.

**01 · THE DECISION CLOCK****Development plans become continuous, not annual.**

The dev plan is annual because it takes six months to build. The borrowing-base is semi-annual because each cycle takes three weeks. When reasoning is cheap, these become continuous. New KPIs follow: not "how many wells did we drill?" but "how many scenarios did we run before choosing?"

**02 · INFORMATION ASYMMETRY FLATTENS****The buyer assembles its own reserve estimate before entering the data room.**

The seller traditionally knew the asset better than the buyer. AI flattens that. A buyer can now assemble its own reserve estimate from public data. Your information moat gets thinner. Same logic applies to lenders, regulators, and counterparties.

**03 · SPEED RISK****Fast analysis without quality architecture is dangerous.**

When the analysis cycle was three weeks, errors got caught in review. At three hours, one bad assumption propagates through 40 scenarios before anyone looks. The companies that skip quality architecture will move fast and break expensive things. Governance is the constraint, not speed.

**What remains human-controlled:** Field safety. Regulatory judgment. Relationship capital. Physical execution. Liability. The native-AI company removes humans from the assembly line, not from the decisions. Each carries external accountability or physical consequence that no model absorbs. **The deeper question:** the prep stack exists because humans need information in narrative form. A decision system that operates directly on the data needs no translation layer — no memo, no package, no slide deck. The bottleneck was never the decision. It was the preparation. That distinction reshapes where the ROI model points.

**THE QUESTION UNDERNEATH EVERYTHING****What does your company understand that is genuinely hard to understand — and is that understanding getting deeper every day?**


If the answer is nothing, AI is cost optimization. Cut headcount, improve margins for a few quarters, get absorbed. If the answer is deep — interconnection queue dynamics, reservoir behavior across basins, regulatory filing patterns, lender decision logic — then AI doesn't just augment your company. It reveals what your company actually is.

This benchmark is either a one-time report or the seed of a world model for energy decision-making. The 404 positions are capabilities. The compressibility scores are a first-pass intelligence layer. The deal flow, the filing history, the cycle data — that's honest signal waiting for a model. The question is whether anyone builds it.

## EXHIBIT 10 · THE POWER BOTTLENECK

**2,300 GW are waiting  
in line. Most will never  
make it through.**

As of end-2024, the U.S. interconnection queue holds more capacity than the entire installed fleet (LBNL Queued Up 2025). The median project waits 4+ years. A major share of that wait is document-heavy – interconnection studies, environmental review, permitting, contract negotiation. That's the bottleneck AI was built to compress.



Data: LBNL Queued Up 2025 Edition. Completion rate (13% of capacity) from projects entering queues 2000–2019 as of end-2024. Funnel widths are approximate. Document share is a Sunya estimate – modeled, not observed.

A 1 GW AI data center throws off roughly \$800M–\$1B a month. Every month it sits in queue, that revenue waits with it. The paperwork is the holdup. And the paperwork is exactly what AI compresses fastest.

## \$10–12B

### ANNUAL REVENUE PER GW OF AI DATA CENTER

A 1 GW AI data center throws off \$10-12B a year. Every month it sits in queue, that's roughly \$1B waiting. Compress the paperwork and the project comes online a year earlier. Multiply that across 100+ GW of planned capacity and you're looking at tens of billions in revenue sitting in a queue, waiting on documents.

Revenue per GW: SemiAnalysis estimate

## 2 paths

### GRID INTERCONNECTION AND ONSITE GENERATION

The hyperscalers aren't waiting for the grid. Oracle/OCI collaborated with VoltaGrid for 2.3 GW of onsite generation in Texas. xAI reportedly ordered 1.9 GW from Doosan (per industry press – not confirmed by primary source). Both grid-connected and behind-the-meter paths have document bottlenecks – interconnection studies on one side, air quality permits and gas supply contracts on the other. Different paperwork, same compression opportunity.

## 3–5 yr

### NRC LICENSING FOR NEW REACTOR DESIGNS

Constellation's 20-year PPA with Microsoft for the Crane Clean Energy Center restart (announced Sept 2024). Google's agreement with Kairos Power for SMRs (announced Oct 2024). Amazon acquiring nuclear-adjacent capacity across multiple sites. NRC licensing is 30–40% document-intensive – safety analysis reports, environmental review, inspection documentation. Compress that by 12 months and you've accelerated the nuclear buildout that every hyperscaler is betting on for baseload.

### WHAT THE NEXT VERSION MODELS

What hyperscaler power teams actually need to know.

#### State-level breakdown

In 2023, data centers consumed about 26% of Virginia's electricity supply. Texas gets tens of GW in monthly requests. The queue, the regulatory stack, and the power mix are completely different in each market.

#### Efficiency crossover

Blackwell is ~4x more efficient per token than Hopper. If chip efficiency improves 4x every 2 years but demand grows 2x, when does efficiency stop outrunning demand? That crossover point is the most important number in the industry.

#### Bottleneck migration

2023: CoWoS packaging. 2024-25: data center power. 2026+: semiconductor fabs. Bottlenecks move. The analysis needs to show the sequence, not just the current constraint.

#### Turbine supply chain

If every hyperscaler goes onsite gas, who makes the turbines? GE Vernova, Siemens, Doosan, Wärtsilä, Bloom, Caterpillar. The turbine order book is the leading indicator of data center capacity.

### IF SUPERINTELLIGENCE IS AROUND THE CORNER

Energy isn't a sector AI impacts.  
It's the sector that determines  
how much AI the world gets.

This report covers wave one – the filing cabinet, the prep stack, the back office. But if AGI is close, the binding constraint shifts fast. Away from analyst-hours and toward megawatts, interconnection time, and political permission. The question flips from "which function compresses?" to "which region delivers power fastest?" Here are the four waves, from now to the horizon.

### Wave 1

#### Workflow compression

Filings, packages, contracts, reports. What this page scores. Deployable now.

### Wave 2

#### Continuous capital allocation

Dev plans refresh weekly. Every asset ranked against every alternative. The ~\$1T reallocation prize.

### Wave 3

#### Energy-system control

AI as the operating system for the grid. The IEA (2025) estimates AI could unlock up to 175 GW of additional transmission capacity from existing lines and save up to \$110B/yr in the electricity sector if widely adopted.


### Wave 4

#### Physics

Materials science, fusion, novel reactor design, storage chemistry. What changes when intelligence is no longer the bottleneck – and the physical world still is.

#### GLOBAL DATA CENTER ELECTRICITY DEMAND

How much power will AI  
actually need?



A data center goes up in two to three years. The power plant it needs takes longer. The transmission line takes longer still. The permit to build the transmission line takes longest of all. Energy becomes a speed problem before it becomes a cost problem.

IEA published scenarios: base = 945 TWh by 2030, lift-off = 1,700 TWh by 2035. "Base ext." extrapolates the IEA 2030 base case to ~1,200 TWh by 2035 using internal trend analysis – not an IEA-published 2035 figure. "Headwinds" is a Sunya downside scenario. Year-by-year values are interpolations, not exact data points. US share from Berkeley Lab 2025: 325–580 TWh by 2028 (7–12% of national electricity).

**Compresses**

Analysts, packagers, coordinators, memo-builders. The assembly layer inside every role.

**Amplifies**

Approvers, negotiators, operators, field decision-makers. As analysis gets cheap, judgment gets expensive.

**Emerges**

35 new roles. Builders, bridgers, orchestrators. [See Exhibit 11 ↓](#)

*You can automate the training ground out of existence. If entry-level analytical work disappears, where do future operators, traders, engineers, and executives learn judgment?*


In nuclear roles, 1.7 workers nearing retirement for every young worker entering. For grid roles: 1.4x (IEA 2025).

EXHIBIT 11 · MODELED NEW — 35 EMERGING ROLES

35 roles that didn't exist two years ago. The talent war is already on.

Previous exhibits map where existing value compresses. This one maps where new value gets created, and who captures it. Big Tech energy hiring jumped 34% year-over-year. Amazon alone added 605 energy hires. AI-energy hybrid roles stay open 90+ days with 30–40% salary premiums. And AI-skilled workers in energy are 40% less concentrated than in tech or finance. The talent gap is structural.

Three origin categories. Two scoring axes. The roles that matter sit where leverage is highest and talent is scarcest.



X = leverage (decision value flowing through the role). Y = scarcity (time-to-fill, salary premium, talent pool size). Bubble size = estimated industry-wide headcount demand. Color: ■ Builders (AI-for-Energy) · ■ Bridgers (Energy-for-AI & AI-lab GTM) · ■ Orchestrators (hybrid mutations of existing roles). Upper-right quadrant = the talent war zone: high leverage, high scarcity, where every PE firm, utility, and hyperscaler is competing for the same 200 people. All scores are modeled estimates – the clustering pattern is the claim, not individual precision.

#### **TALENT WAR ZONE · UPPER-RIGHT QUADRANT**

Seven roles cluster where leverage exceeds 7 and scarcity exceeds 7: Data Center Energy Lead, Energy Strategy Director (AI Lab), Power Procurement Specialist, Interconnection Queue Manager, AI Safety Engineer (Grid), Enterprise Sales – Energy (AI Lab), and AI-Augmented Reservoir Engineer. Microsoft poached GE's former CFO to run energy strategy. Google hired Duke's Tyler Norris for energy market innovation. Amazon staffed 605 energy positions in a single year. They're not hiring AI people. They're hiring energy people who understand AI. There are maybe 200 of them on earth.

#### **WHAT COULD BREAK THIS**

These roles assume sustained AI investment and continued energy-AI convergence. If frontier model costs collapse faster than expected, some builder roles commoditize. If hyperscaler power strategies shift to long-term utility contracts (removing the custom procurement layer), bridger demand softens. If AI tools become genuinely autonomous, orchestrator roles shrink rather than grow. The scarcity scores also assume current training pipelines – if universities and bootcamps spin up energy-AI programs at scale, premiums compress within 3–5 years. These are 2026 scores, not permanent conditions.

This report is wave one. The binding constraints shift from analyst-hours to megawatts, interconnection time, and political permission. The beachheads work whether the destination is 3x better or 300x different. Here's what would prove us wrong.

Most objections to this analysis make one of four moves. Naming them is not a dismissal — it's an invitation to check whether your own objection clears the bar.

#### Question drift

We claim "these tasks compress." The objection answers "AI won't replace people." Different question. We agree — the roles persist, the prep work inside them shrinks.

#### Perfect-proxy fallacy

Compressibility scores are directional proxies, not deployment blueprints. Attacking the proxy for not being a perfect predictor misses the purpose — it ranks where to look first, not what to buy.

#### Autonomy strawman

We never argue for autonomous AI in critical workflows. The thesis is augmented prep: AI drafts, humans verify and decide. Attacking "unsupervised AI" attacks a claim we don't make.

#### Edge-case universalizing

Finding one workflow where AI fails and generalizing it to the whole stack. Some workflows won't compress. That's already in the data — the "[Six things AI won't fix](#)" section maps exactly where.

**The single strongest objection that survives all four filters:** compressibility as a single axis is too blunt. A high score tells you the task is compressible — it doesn't tell you whether the organization can adopt it, whether the data exists to train on, or whether the regulatory environment allows it.

We agree. That's why 304 of 404 positions carry three scores, not one: compressibility (can the prep work shrink?), criticality (what breaks if AI gets it wrong?), and reasoning demand (does this need frontier-class models or commodity inference?). If you only look at compressibility, you'll automate the wrong things. The three-axis view is how you avoid that.

#### What we know, what we believe, and what we're guessing

**The scores are hypotheses.** Scored by one analyst ([Raj Mistry](https://rajmistry.com)). No inter-rater reliability, no blinded scoring, no adjudication log. Employment counts are formulaic (layer base × tier multiplier), not survey-sourced. Compensation uses six layer-level proxies, not role-specific wages. The directional layer pattern is the claim; individual role numbers are structural estimates. Multi-rater validation is planned for a future edition. **CONFIDENCE:** medium. Layer rankings are high confidence; individual scores are low.

**The ~\$1T number is a scenario, not a forecast.** \$3.3T annual energy investment × 3% decision-quality improvement × 10 years. At 1%: ~\$330B. At 5%: ~\$1.65T. The 3% approximates the difference between running 4 scenarios and 40 on a typical development program. The entire

gap between "nice efficiency play" and "industry transformation" rests on whether cheaper analysis leads to more analysis (Jevons) or just faster versions of the same 4 scenarios. We believe Jevons is directionally true. The magnitude is speculative. CONFIDENCE: low. The range (\$330B–\$1.65T) is the honest answer.

**60-70% of queue time as document work is a modeled estimate.** Derived from FERC Order 2023 procedure timelines – system impact study, facilities study, environmental review, permitting, PPA negotiation – vs. physical construction. A developer or interconnection consultant could confirm or refute it with project-level data. We haven't found one willing to share yet. CONFIDENCE: medium. Procedure-derived, not project-validated.

**The native-AI energy company is a hypothesis, not a finding.** We've identified no venture-backed company that has demonstrated a fully AI-native operating model in energy workflows as of early 2026. The analogies (Snowflake, Uber) are retrospective – we know they worked. Whether the pattern transfers to energy is an open question. CONFIDENCE: speculative.

**Known gaps.** No competitive landscape mapping (Microsoft Copilot, Palantir, C3.ai, and dozens of startups target overlapping workflows). US-centric regulatory framing (FERC, PUC, NRC, RRC) – international frameworks deserve equal depth. Uneven coverage depth: upstream E&P and utilities are granular, refining and petrochemicals less so, the trading floor alone could be 30+ workflows. And the IOCs are already spending hundreds of millions on AI – this research is most relevant to the companies that haven't started. CONFIDENCE: N/A. Scope disclosures.

SPECULATION · NO PRECEDENT EXISTS

---

*A venture-backed, AI-native energy operator is plausible – built not to sell tools to energy companies, but to be one. Capital allocation powered by continuous scenario analysis instead of quarterly human assembly. No venture-backed company has demonstrated this yet. But cheap reasoning, expensive expertise, and recurring decisions are three ingredients that have built category-defining companies in every other industry they've appeared in.*

---

*This report is an illustration of its own thesis.*

---

If you're reading this and thinking "my company should be producing this kind of analysis internally" — that thought is the thesis. You're living in it.

~\$800

MODEL COST TO PRODUCE THIS PAGE

1

ANALYST WITH DOMAIN EXPERTISE

~\$450K

EQUIVALENT RESEARCH TEAM COST

90–560x

MODEL COST LEVERAGE

373 roles, 24 workflows, 7 artifacts. 11 interactive exhibits. Full dataset, claims ledger, and scoring methodology published. Built with Claude Opus 4.6 and GPT-5.4 Pro — millions of tokens of research, drafting, and iteration. Directed, scored, and verified by Raj Mistry (https://rajmistry.com). The models assembled; Raj made the calls. A six-person team working six months would have produced comparable breadth at far higher cost. That's the thesis illustrated.

Methodology note: the \$800 figure reflects direct API billing for the final production sessions only (Claude Opus 4.6 and GPT-5.4 Pro); earlier research, iteration, and discarded drafts are not included. Total model spend including iteration is estimated at \$800–\$5K, giving the 90–560x range shown above. Labor comparator is a benchmark estimate for a comparable industry research team at market rates — not a quote. Neither figure is audited. Token logs and model-mix breakdown will be published in a future edition.

We didn't set out to prove the thesis while writing about it. But one analyst directing two frontier models just produced what would have taken six people six months. The models did the prep work. The analyst made every judgment call. That's the thesis — illustrated, not argued.

GO DEEPER

Discuss this research with AI

Opens a new conversation with a prompt that links back to this report. The model can read the page and discuss the findings, methodology, or implications for your company.

Chat with Claude (https://claude.ai/new?q=Please%20read%20the%20full%20report%20at%20https%3A%2F%2Fresearch.sunya.ai%20%20E2%80%94%20it%20scores%20404%20energy%20positions%20by%20AI%20compressibility%20and%20finds%2090%25%20of%20AI-exposed%20wages%20bill%20its%20above%20the%20field%20in%20treasury%20C%20title%20regulatory%20C%20and%20commercial%20work.%20Read%20the%20page%20first%20C%20then%20let%27s%20discuss%20the%20findings.)

Chat with ChatGPT (https://chatgpt.com/?q=Please%20read%20the%20full%20report%20at%20https%3A%2F%2Fresearch.sunya.ai%20%20E2%80%94%20it%20scores%20404%20energy%20positions%20by%20AI%20compressibility%20and%20finds%2090%25%20of%20AI-exposed%20wages%20bill%20its%20above%20the%20field%20in%20treasury%20C%20title%20regulatory%20C%20and%20commercial%20work.%20Read%20the%20page%20first%20C%20then%20let%27s%20discuss%20the%20findings.)

**A note on the data.** 60 of 404 positions have BLS-sourced employment counts. The remaining 344 use a formulaic estimate: layer base × tier multiplier (three bands: 0.75×, 1.0×, 1.25×). These are structural scaffolding for directional wage-pool analysis, not census-quality headcounts. Compensation uses six layer-level proxies (\$58K–\$180K), not role-specific wages. The directional finding — that roughly nine of every ten AI-exposed wage dollars sit above the field — holds across a wide range of employment assumptions. Individual role numbers should not be cited as precise. Full methodology: [methodology FAQ](#).

**How to reproduce every number in this benchmark**

The core benchmark — scores, employment, compensation, and the -90% headline — can be reconstructed from the published CSV and the three formulas below. Some interactive exhibits (role explorer, workflow economics, demand scenarios) use additional embedded data structures visible in the page source.

**FORMULA 1 — EXPOSURE-WEIGHTED WAGE BILL**

$$EWB = \text{estimated\_employment} \times \text{layer\_median\_comp\_proxy} \times \text{compressibility\_score} / 10$$

This is the single metric that drives ranking, the slope chart, the heatmap area, and the -90% headline. The core exhibits trace back to this formula applied row-by-row across the 404-position CSV.

**FORMULA 2 — EMPLOYMENT ESTIMATION (344 FORMULAIC ROWS)**

$$\text{employment} = \text{layer\_base} \times \text{tier\_multiplier}$$

Where:

**layer\_base** Physical: 22,000 · Technical: 8,000 · Corporate: 10,000 · Advisory: 7,000 · Capital mkts: 5,000 · Governance: 3,000

**tier\_multiplier** Three explicit bands: Low (0.75×), Mid (1.0×), High (1.25×). Each role is assigned to a tier by a stable hash of its name: `h=0`; for each char `c`: `h=((h<<5)-h+charCode(c))|0`; `tier=|h|%3`. Tier 0→0.75×, 1→1.0×, 2→1.25×. This produces three discrete employment levels per layer rather than pseudo-random integers that would imply false precision. The tier assignment is deterministic and reproducible from role name alone.

*Methodology note (v32.1):* Employment estimation is decoupled from compressibility scoring. An earlier version used `max(0.15, 1.3 - score/8)` as a score factor, which gave low-scoring roles more people by construction. The continuous jitter `0.7 + (hash mod 600)/1000` was also replaced with three explicit tier bands to eliminate false precision. All employment figures for formulaic rows are displayed with rounding and a "-" prefix to signal estimation uncertainty.

The 60 anchor rows override this formula entirely — their employment figures come from BLS Occupational Employment and Wage Statistics or industry-derived analogs, flagged in the `employment_source` column. Note: two of these 60 rows are not people-roles (Board pack production, A&D diligence support) — they are workflow/artifact rows with industry-estimated volume, not BLS occupational matches.

**FORMULA 3 — COMPENSATION PROXIES**

Physical	Technical	Corporate	Advisory	Capital Mkts	Governance
\$58K	\$95K	\$72K	\$105K	\$130K	\$180K

These are BLS OEWS-derived median annual wages for the layer’s representative occupations, not role-specific compensation. Every row in the same layer uses the same proxy. Figures reflect wages only; employer-paid benefits are not included.

**COMPRESSIBILITY SCORE — WHAT IT MEASURES**

Each position’s compressibility score (1–10) represents how plausibly current AI can compress the preparation stack around that role or workflow. It is built from five sub-dimensions:

**Task exposure** What fraction of the role’s tasks involve document-heavy, pattern-matching, or analytical work that current models handle well

**Adoption feasibility** How quickly the workflow could realistically adopt AI given regulatory, organizational, and technical constraints

- Economic compression** Whether compressing this workflow changes the cost structure enough for a budget owner to notice
- Workflow standardization** How repeatable and templated the work is across cycles and counterparties
- Template density** How much of the output follows a known structure (forms, packets, exhibits, standard reports)

The score is a calibrated composite, not a simple average. Roles with high task exposure but low adoption feasibility (e.g., safety-critical field roles) score lower than roles where both align (e.g., document assembly and reconciliation). The full scoring rationale is published per-row in the CSV's [rationale](#) column and in the [scoring method note](#).

**Augmented dataset**

404 positions with compressibility scores, employment estimates, and augmented axes. 60 BLS-anchored rows, 344 formulaic. The raw material behind every exhibit.  
[Download 404-position CSV](#) →

**Claims ledger**

Every major claim tagged by type: measured, modeled, scenario, hypothesis, cited, or self-critique.  
[Download claims ledger CSV](#) →

**Scenario table**

Base, aggressive, and extreme annual token scenarios across six workflow families.  
[Download token-demand scenarios](#) →

**Headline metrics**

Registry of key metrics with definitions, denominators, and derivation paths.  
[Download metrics registry](#) →

**Scoring method**

The five sub-dimensions, calibration approach, and heuristics behind every compressibility score.  
[Read scoring methodology](#) →

**Sensitivity analysis**

Monte Carlo stress test: the 90% finding survives ±1 and ±2 point scoring errors across all 404 rows. 1,000 iterations.  
[Read Monte Carlo stress test](#) →

**External sources**

Validation table for every external data point: URL, date, exact/approximated status, and confirmation flags.  
[Read source validation table](#) →

**Counterforces**

Where the thesis bends, slows, or fails – and how to test each assumption.  
[Read counterforces](#) →

Every major claim on this page is sourced or tagged as modeled. Every assumption is attackable. If we're wrong, the data is right here.

If we're right, you're running four scenarios right now. Your competitor is about to run forty. That's not a technology gap. It's a calendar problem – and the calendar is already moving.

**IF YOU RUN AN ENERGY OPERATOR**

Name the five decision loops where your team still rations analysis. Start where the cycle recurs quarterly, the team

**IF YOU INVEST IN OR LEND TO ENERGY**

Ask your portfolio companies how many scenarios they ran before their last billion-dollar

**IF YOU BUILD AI TOOLS FOR ENERGY**

The workflows in Exhibits 7–8 are ranked by recurring reasoning demand and decision

dreads the assembly, and you own the prep — no regulator permission needed.

[See the three beachheads](#)

commitment. If the answer is four, ask why. The gap between 4 and 40 is where the next write-down hides.

[Download the dataset](#)

value per token. That's the revenue signal. Start where the budget owner can feel the gain within one operating cycle.

[Open claims ledger](#)

*Four scenarios on a billion-dollar decision. Or forty. That's the only question this page asks.*

---

This research is for informational purposes only and does not constitute investment, financial, legal, or professional advice. Compressibility scores, cost estimates, and scenario outputs are modeled approximations — not audited figures, not forecasts, and not recommendations to act. AI-assisted content was human-directed and editorially reviewed but has not been independently verified. Use your own judgment and consult qualified advisors before making decisions based on anything presented here.

© 2026 Sunya Research · [Raj Mistry](https://rajmistry.com) (<https://rajmistry.com>) · [Methodology](#) · [Privacy](#) · [Terms](#)

Built with Claude Opus 4.6 + GPT-5.4 Pro · Scored and directed by a human