

Fortune 500 Insurance x NODES

Fortune 500 Insurance | 215+ Locations | Company-Wide Deployment | January 2025 to Present

NODES deployed inside the carrier's VPC with zero data egress, scoring **730,000+ candidates** across 28+ behavioral dimensions. Fit Scores appear natively inside Avature with a plain-English rationale for every recommendation. Legal approved in 17 days after rejecting 6 vendors over 18 months.

80+% top performer accuracy over 4,000+ hires.

730K+

Candidates Processed

28+

Behavioral Dimensions
Scored Per Candidate

2.47x

Top Performer Rate (High-
Scored vs Low-Scored Hires)

17

Days to Legal Approval

Reading this from outside insurance? This case study uses insurance terminology. SNA (first production milestone) and RPA (sustained production award) are the industry's standardized measures of whether a hire actually produces. Every industry has equivalents: quota attainment in sales, ramp completion in engineering, certification in skilled trades, retention and productivity milestones in manufacturing. The methodology applies universally. The data happens to come from insurance.

The Challenge

A Fortune 500 insurance holding company operating across 215 locations. They receive 1.5 million applications per year through their Avature ATS. Their recruiters could manually review approximately 1.5% of total applicants, meaning the best candidates were buried in the 98.5% never reviewed.

Hiring decisions relied on keyword filters: insurance experience, sales background, degree. These filters had never been validated against actual production outcomes. Recruiters had no visibility into which hires would actually produce. The pattern recognition their best hiring managers built over decades was disappearing through retirements.

What We Found When We Validated the Filters

After deployment, NODES enabled something no previous system could: connecting ATS screening criteria to actual production milestones at the individual agent level. We analyzed 747 NODES-scored agents contracted January through March 2025, with 9+ months of tenure. Of these, 535 had parseable ATS skills text. We measured against two standardized industry milestones: SNA (Sale New Agent, a sales threshold) and RPA (Rookie Production Award, sustained production).

The results challenged every assumption about what makes a good hire.

Your #1 Filter Eliminates 80% of Your Award Winners

We tested each of the 6 standard ATS screening keywords against RPA achievement:

Screening Keyword	RPA With	RPA Without	Ratio	p-value	% RPA Eliminated
Insurance experience	10.9%	9.0%	1.20x	0.56	80%
Sales experience	13.8%	6.0%	2.16x	0.007	34%
Customer service	9.3%	9.4%	0.99x	1.00	42%
Leadership	9.2%	9.6%	0.96x	0.87	30%
Communication	10.2%	8.8%	1.17x	0.65	58%
Teamwork	8.3%	9.6%	0.87x	0.85	84%

Insurance experience, the filter most employers apply first, has zero predictive power for RPA (p=0.56) but eliminates 80% of RPA achievers in a single step. **Sales experience** is the only traditional signal that significantly predicts RPA (2.16x, p=0.007). But by the time it is applied, the insurance filter has already eliminated 80% of winners.

Cumulative Impact: 98% of Award Winners Eliminated

Among 50 RPA achievers with ATS skills text, applying each filter cumulatively:

Filter Applied (cumulative)	RPA Achievers Remaining	% Surviving
No filter	50	100%
+ Insurance Experience	10	20%
+ Sales Experience	9	18%
+ Customer Service	6	12%
+ Leadership	4	8%
+ Communication	3	6%
+ Teamwork	1	2%

If an employer screens for the "ideal hire" using all 6 standard criteria, **98% of their Rookie Production Award winners would never have been interviewed.** The same pattern holds on SNA: 96% eliminated.

Predictive Accuracy: Screening vs Coin Flip vs NODES

Predictor	SNA Prediction	RPA Prediction	RPA p-value
Traditional Fit (6 keywords)	AUC 0.512	AUC 0.548	0.25 (null)
Random (coin flip)	AUC 0.500	AUC 0.500	--
NODES Fit Score	AUC 0.567	AUC 0.618	0.006
Sales + NODES combined	AUC 0.576	AUC 0.644	--

The 6-keyword composite (AUC=0.512 SNA, AUC=0.548 RPA) performs identically to a coin flip. NODES behavioral scoring (AUC=0.618 RPA, p=0.006) is the strongest single predictor. Combined with the one traditional signal that works (sales experience), it reaches AUC=0.644.

The Combined Signal: 5.37x Production Gap

Segment	n	SNA Rate	RPA Rate
Sales keyword + NODES ≥ 72	189	25.9%	15.3%
Sales keyword + NODES < 72	50	16.0%	8.0%
No sales + NODES ≥ 72	163	19.0%	8.6%
No sales + NODES < 72	105	14.3%	2.9%

Sales + High NODES vs No Sales + Low NODES: **RPA rate 15.3% vs 2.9% = 5.37x (p=0.0007)**. This is the strongest statistical signal in the dataset. It requires combining ATS keyword parsing with NODES behavioral scoring with HRIS production milestones. No single system can produce this insight.

Methodology note: All findings use binary production milestones (SNA achieved yes/no, RPA achieved yes/no), not cumulative revenue metrics. All agents had 9+ months of tenure at time of measurement. Traditional Fit Score, NODES score, and all screening keywords are orthogonal to contract month (all $p > 0.10$). These findings are immune to temporal confounds by design.

Speed to Production: The Hidden Yield Equation

Across 10,362 hires (2022–2025), we tested whether speed to first production milestone predicts annual output. The answer depends entirely on whether scoring is present.

Without scoring (2022–2024, n=1,011): $r = +0.045$, $p = 0.15$. No meaningful correlation between speed and production. The fastest hires did not outproduce the slowest. Speed was noise.

With scoring (2025, n=679): Monotone decrease. Every bucket step slower = lower production. Fastest 30-day hires produce 1.8x the slowest 121+ day hires. Scoring creates the relationship between speed and output that did not exist before.

The Constant: \$54.35 Per Day Per Person

Because scoring predicts both production and speed, a constant falls out of the data: each day faster to first production milestone = \$54.35 more in annual production per person. This is derivable from the linear regression: $APC = \$13,964 - \$54.35 \times \text{days_to_milestone}$.

Observed acceleration: 47 days (median 62 days with scoring vs 109 days without). At 2,000 hires per year: $\$54 \times 47 \times 2,000 = \5.11M in additional annual production. At 500 hires per year: \$1.28M. The math scales linearly with hiring volume.

Every quarter without scoring is a quarter of hires ramping at the unscored baseline. A company that deploys scoring in Q1 2026 versus Q1 2027 has a \$5.11M production gap. The gap compounds forward: each quarter of scored hiring produces people who produce faster, retain longer, and generate outcomes that feed the next quarter's pipeline.

The General Case

Your ATS has its own screening fields. Your HRIS has production data. If your screen does not predict your production, the equation still holds — you just need the right signal. When you find it, the constant falls out. The fields change by industry. The math does not.

Why Every AI Vendor Gets Blocked

Legal's Question	Vendor's Answer	Result
Where does the data go?	Our cloud, then OpenAI	Rejected
Can we validate the models?	They are proprietary	Rejected
Liability exposure?	It's complicated, section 14	Rejected

The Solution

NODES deployed inside the carrier's VPC in 34 days. Legal approved the architecture in 17 days, the fastest AI vendor approval in company history. The system integrates natively with Avature with zero workflow change.

Every candidate profile now shows:

Fit Score (0-100): calibrated against the carrier's top-performer personas by location and role, computed across 28+ behavioral, skill, and cultural dimensions.

Plain-English rationale: explaining why the candidate scored the way they did. Strengths, skill gaps, and improvement areas listed. No black boxes.

Decision traces: that capture the reasoning behind every recommendation, creating a queryable library of hiring intelligence that compounds over time.

AI Interview (optional): to gather structured answers mid-funnel and auto-update the Fit Score. Full transcripts available to hiring managers.

Persona-based sourcing: that identifies high-fit prospects who haven't applied and pushes them into Avature.

Results

"We are screening thousands in hours instead of weeks, and the quality is measurably better. Very excited for this partnership to evolve even more."

Director of Field Sourcing

Metric	Result
Candidates Processed	730,000+ across 215 locations
Contracted Hires (2025)	2,335
Top Performer Lift (High-Scored Hires)	2.47x (p=0.006)
Resume Screening Reduction	40% fewer manual screens
Interview Reduction	2 fewer interviews per hire
Time-to-Hire	127 days to 38 days (70% reduction)
Customer Savings	\$1.58M in year one
Deployment Time	34 days from contract to production
Legal Approval	17 days (fastest AI vendor in company history)
Prior Vendor Rejections	6 vendors blocked over 18 months
Data Egress	Zero.

Production Milestone Validation

Insurance agent success is measured by standardized milestones. SNA (Sale New Agent) is achieved when an agent hits a specific sales threshold within a defined timeframe. RPA (Rookie Production Award) indicates sustained production above a threshold. Both are standard industry metrics. Every industry has equivalents: quota attainment, ramp completion, certification milestones, or retention benchmarks.

NODES Score Predicts Both Milestones

NODES Quintile	n	SNA Rate	RPA Rate	Trend
Q1 (lowest)	157	15.9%	5.1%	
Q2	145	14.5%	8.3%	
Q3	163	19.6%	8.0%	
Q4	144	18.1%	9.7%	
Q5 (highest)	138	22.5%	13.0%	2.55x Q1→Q5

The trend is monotonic: top performer rate increases from 5.1% (lowest-scored quintile) to 13.0% (highest-scored). Hires with NODES score 72 or above achieve top performer milestones at **2.47x the rate of those below 72 (p=0.006)**.

What the Fit Score Analyzes: 28+ Dimensions

Unlike keyword matching or resume parsing, NODES computes a multi-dimensional behavioral profile for every candidate. Each Fit Score is the weighted output of 28+ individual assessments across four categories, all calibrated against the carrier's validated top-performer persona by location and role.

Category	Dimensions Scored	What It Captures
Behavioral (6 scores)	Teamwork, Leadership, Adaptability, Communication, Problem-Solving, Resilience	Behavioral patterns derived from career history: how someone navigates transitions, handles responsibility, and operates under pressure.
Skills (16+ scores)	Sales Acumen, Closing Skills, Consultative Selling, Persuasive Communication, Objection Handling, Client Relationship Building, Self-Motivation, Technical Aptitude, Ethical Conduct, and more	Scored against the carrier's persona. A warehouse worker's "equipment repair" maps to technical aptitude. A restaurant server's "17 years customer-facing" maps to relationship building.
Cultural Fit (6 scores)	Integrity, Collaboration, Accountability, Drive for Results, Continuous Learning, Customer Centricity	Values alignment derived from career history and role patterns. Predicts cultural integration, not just performance.
Career Analysis (3 metrics)	Promotion Rate, Complexity Growth, Achievement Density	Career trajectory patterns. Are they growing in responsibility? Are their roles getting more complex?

No ATS, HRIS, or BI dashboard provides this. Avature stores keywords. NODES computes behavioral intelligence. That is the difference between knowing someone listed "sales experience" and knowing they demonstrate high adaptability, strong self-motivation, and a career trajectory that matches your top performers.

What NODES Sees That Resumes Don't

Traditional ATS screening uses keyword filters: insurance experience, sales background, degree. NODES analyzes behavioral patterns that keywords cannot capture. The result: candidates that traditional screening would overlook are identified as high-potential, and candidates with "perfect" resumes who lack the right behavioral signals are flagged.

Three Hires Your ATS Would Have Missed

Agent	Background	Resume Verdict	NODES Score	Production Result	Key Behavioral Signals
Agent A	Store worker, disc golf shop. No insurance exp. No sales exp.	Skip	86 (Strong Fit)	SNA + RPA achieved.	Adaptability 0.8, Technical Aptitude 0.7, Self-Motivation 0.8
Agent B	Restaurant server, 17 years. No insurance exp.	Skip	92 (Spotlight)	SNA achieved. Top producer.	Client Relationship 0.9, Resilience 0.8, Communication 0.8
Agent C	Insurance + sales exp. "Perfect" resume.	Hire	39 (Below Core)	No SNA. Did not produce.	Resilience 0.2, Adaptability 0.3, Relationship Building 0.2

The behavioral signals that separate top performers from underperformers are invisible to keyword matching. Adaptability, resilience, self-motivation, and relationship building cannot be detected by searching for industry experience or job-specific keywords. They require the multi-dimensional analysis NODES provides.

Why It Works: Fine-Tuned on Your Data

A small model fine-tuned on your actual hiring outcomes, your performance data, and your culture signals outperforms GPT-5, Claude 4.5 Sonnet, and Gemini 2.5 Pro on every benchmark that matters for talent prediction. Generic frontier models have never seen your hiring data. They do not know what success looks like at your company. See Technical Appendix for the full benchmark comparison.

Addressing the Three Risks

Every enterprise buyer evaluates three risks before deploying AI in hiring.

Risk	The Concern	Carrier's Evidence
Implementation	What if deployment takes 6 months and disrupts our workflows?	34 days to production. 17 days legal. Zero workflow change. Native Avature integration. Recruiters never left Avature.
Adoption	What if recruiters ignore it or it adds steps?	Sustained usage across 215 locations for 8+ months. Adoption driven by less noise in shortlists, not more tools.
Accuracy	What if predictions are wrong and we face bias claims?	Validated against SNA and RPA milestones (p=0.006). Full audit trails. Explainable scores. Bias monitoring. EEOC/OFCCP compliant.

Why It Landed at the Carrier

Less noise. The first slate is already strong. Managers stop wading through low-yield resumes. Every candidate has a score and a reason.

Faster cycles. Prioritized shortlists + fewer interviews to decision. 40% reduction in manual screening. 2 fewer interviews per hire.

Higher quality. Hires scored 72+ by NODES achieve top performer milestones at 2.47x the rate of those below 72 (p=0.006).

Explainable. Each recommendation includes the "why." Strengths, gaps, improvement areas per candidate. No black-box behavior.

Secure. Deployed inside the carrier's VPC. Zero data egress. SOC 2 compliant. Legal approved in 17 days.

The Compounding Advantage

Most enterprise software delivers the same value on day 365 as day 1. NODES gets better.

Timeline	What Compounds
Months 1-3	Success profiles calibrated. Screening begins. First decision traces captured. 28+ dimensional scoring active. Time-to-hire starts dropping.
Months 4-8	Outcome data validates predictions. Models retrain on your performance data inside your VPC. Accuracy improves without any data leaving your infrastructure. Production milestone data feeds back into scoring.
Year 2+	Queryable library of hiring reasoning. Success profiles extend to workforce development. Persona-based headhunting scales. The intelligence gap between you and competitors without this data becomes insurmountable.

The cost of waiting: A company running NODES for 12 months has validated success profiles, thousands of decision traces, and outcome data that exists nowhere else. A competitor starting in 12 months has nothing. This data cannot be bought, cannot be scraped, and cannot be built overnight.

Where the Intelligence Comes From

Most AI hiring tools read resumes and stop there. NODES connects to the three systems of record that contain the full picture of talent decisions, all processed inside your VPC:

ATS (Candidate Data): Who applied, who advanced, who got hired. Resume text parsed into 28+ dimensional behavioral profiles, not just keywords.

HRIS (Performance and Outcome Data): Production milestones, ramp timelines, retention rates, performance reviews. This is where you learn who actually succeeded post-hire. It closes the feedback loop no resume screen can replicate.

CRM (Decision Context): Call transcripts, recruiter notes, hiring manager debriefs. Where the reasoning lives. Why did the hiring manager override the score? What did the top performer say that signaled they would ramp fast?

Separately, each system tells you a fragment. Connected through NODES, they tell you the full story of why your best people succeed. The ATS shows who you hired. The HRIS shows who performed. The CRM shows why the decision was made. NODES weaves these into success profiles that get sharper with every hire.

At the carrier, Phase 1 connected to Avature (ATS) and HRIS performance data. The filter validation study above is from that integration alone. Phase 2, now in deployment, adds CRM call transcript ingestion to further refine success profiles.

What Becomes Possible

Hiring is where NODES starts. It is not where it stops. The same success profiles and decision traces that identify top performers before they are hired extend into every stage of the talent lifecycle.

Phase 1: Hiring Intelligence. "Show me every exception we granted for candidates without a degree, and how they performed." "What sourcing channels actually produced top performers?" "Why did this internal move work?"

Phase 2: Workforce Development. "What did our fastest-ramping sales reps in the Southeast have in common during their first 90 days?" "Which onboarding sequences correlate with 12-month retention above 85%?"

Phase 3: Full Talent Intelligence Platform. When ATS, HRIS, and CRM data are connected inside a single context graph, the system becomes a queryable institutional memory spanning the entire talent lifecycle.

Why Now

1. Open-source models reached enterprise capability. Models in the 3B-20B parameter range can now be fine-tuned for domain-specific tasks and deployed on standard enterprise hardware without GPU dependencies. This was not possible 18 months ago.

2. Regulatory clarity arrived. The EEOC issued AI hiring guidance. NYC Local Law 144 requires bias audits. The EU AI Act classified hiring as high-risk. The Workday \$280M settlement created board-level awareness. For the first time, Fortune 500 legal teams have explicit guidance on what they cannot approve. That is not a headwind. It is our entire go-to-market.

3. Institutional knowledge is disappearing. Your best hiring managers, the ones who can look at a resume and just know, are retiring. The judgment they built over decades is not written down. Candidate volume is 10-14x what it was a decade ago. Without a system to capture that reasoning, it vanishes.

NODES

AI Talent Infrastructure for the Enterprise

Saad Bin Shafiq | saad@nodes.inc | nodes.inc

Native ATS integration | Enterprise security | SOC 2 Type II