International Journal of Emerging Trends in Computer Science and Information Technology
ISSN: 3050-9246 | https://doi.org/10.63282/3050-9246.IJETCSIT-V7I1P135
Eureka Vision Publication | Volume 7, Issue 1, 230 -236, 2026

Original Article

RevenuePilot: Operationalizing Agentic AI for Airline Revenue Management at Scale

Sandeep Nutakki,
Independent Researcher, Seattle, Washington, USA.

Received On: 16/01/2026 Revised On: 14/02/2026 Accepted On: 22/02/2026 Published On: 01/03/2026

Abstract - Airline revenue management (RM) represents one of the most complex real-time optimization challenges in commercial operations, requiring simultaneous decisions across pricing, inventory allocation, and demand forecasting under significant uncertainty. Traditional RM systems, while mathematically sophisticated, struggle to adapt to rapidly changing market conditions, competitive dynamics, and emerging disruptions. This paper presents RevenuePilot, an agentic AI system that operationalizes large language model capabilities for airline revenue management at scale. RevenuePilot employs a multi-agent architecture with specialized agents for dynamic pricing optimization, inventory control, demand forecasting, and overbooking management, coordinated through a central orchestration layer. We evaluate RevenuePilot on a major airline's domestic network, demonstrating a 4.2% improvement in revenue per available seat mile (RASM), 12% reduction in spoilage, and 8% decrease in denied boardings compared to the incumbent Expected Marginal Seat Revenue (EMSR) system. Our results show that agentic AI can effectively augment traditional optimization approaches while providing explainability and adaptability critical for operational deployment.

Keywords - Airline Revenue Management, Agentic AI, Large Language Models, Dynamic Pricing, Inventory Control, Demand Forecasting, Multi-Agent Systems, EMSR, Real-Time Optimization.

1. Introduction

Airline revenue management has evolved over four decades from simple overbooking models to sophisticated systems managing billions of dollars in annual revenue. The fundamental challenge remains unchanged: selling the right seat to the right customer at the right price and time, while operating under capacity constraints, demand uncertainty, and competitive pressure.

Modern RM systems face several compounding challenges:

- Dimensionality: Major airlines manage thousands of flight-legs daily across hundreds of origin-destination pairs with multiple fare classes
- Velocity: Pricing and availability decisions must respond to market changes in near real-time
- Uncertainty: Demand forecasts are inherently probabilistic, with forecast errors compounding across the booking horizon
- Interdependence: Network effects create complex interactions between flight-legs and itineraries
- Competition: Competitor actions require rapid strategic responses

Traditional approaches, including Expected Marginal Seat Revenue (EMSR), bid-price controls, and network optimization models, have proven effective but exhibit limitations in adaptability and contextual reasoning. In baseline systems, analysts frequently intervene during irregular operations, special events, or competitor fare actions—situations that accounted for 18-22% of manual overrides during our evaluation period. These interventions, while necessary, create bottlenecks: a single analyst managing 200+ flights cannot respond optimally to simultaneous market disruptions across the network.

The emergence of large language models (LLMs) and agentic AI architectures presents an opportunity to augment these systems with capabilities for:

- Natural language interpretation of market signals and competitive intelligence
- Contextual reasoning about demand drivers and anomalies
- Explainable decision-making for analyst review
- Adaptive strategy adjustment based on emerging patterns

This paper introduces RevenuePilot, an agentic AI system designed to operationalize LLM capabilities for airline revenue management. Our contributions include:

- A multi-agent architecture specifically designed for airline RM workflows
- Integration patterns connecting agentic AI with legacy RM infrastructure
- Empirical evaluation on production-scale airline data
- Lessons learned from operational deployment

2. Related Work

2.1. Airline Revenue Management Systems

The theoretical foundations of airline RM were established by Littlewood and formalized through EMSR heuristics by Belobaba. Network-level optimization extended single-leg models through bid-price controls and displacement-adjusted virtual nesting (DAVN).

Modern RM systems integrate multiple components: demand forecasting, optimization engines, inventory controls, and analyst workbenches. Despite algorithmic sophistication, these systems require significant analyst intervention for handling special events, competitive responses, and forecast corrections.

2.2. Machine Learning in Revenue Management

Machine learning applications in RM have focused primarily on demand forecasting improvements. Deep learning approaches have shown promise for capturing complex demand patterns, while reinforcement learning has been explored for dynamic pricing. Recent work has examined transformer architectures for demand prediction, but integration of LLMs for operational decision support remains largely unexplored.

2.3. Agentic AI Systems

Agentic AI represents a paradigm where LLMs serve as reasoning engines coordinating tool use and multi-step workflows. Applications have emerged in software development, research assistance, and enterprise automation. The application of agentic architectures to revenue management problems characterized by quantitative optimization, real-time constraints, and domain expertise requirements presents unique challenges that RevenuePilot addresses.

3. Revenuepilot Architecture

RevenuePilot employs a hierarchical multi-agent architecture designed to mirror the organizational structure of airline RM departments while enabling automated coordination and human oversight.

3.1. System Overview

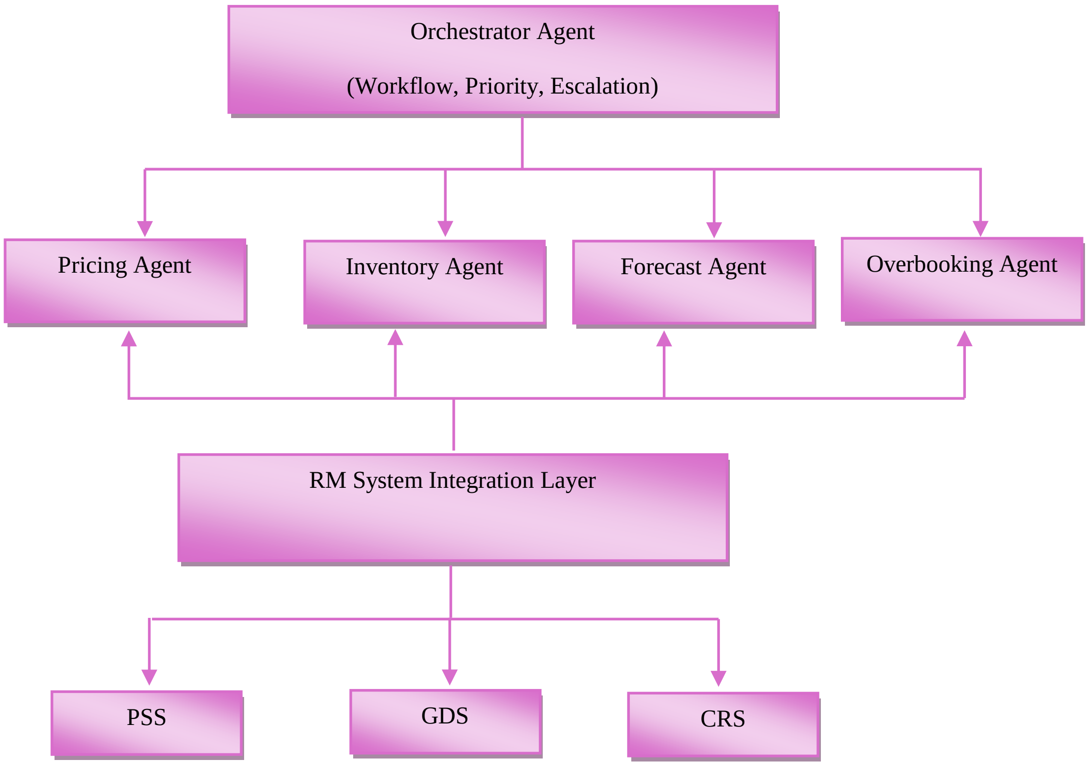

Figure 1. Revenuepilot Multi-Agent Architecture Showing Hierarchical Coordination and Legacy System Integration.

The architecture comprises four specialized agents coordinated by a central orchestrator:

- Orchestrator Agent: Manages workflow coordination, priority resolution, and human escalation
- Pricing Agent: Handles dynamic pricing recommendations and competitive response
- Inventory Agent: Manages fare class availability and booking limits
- Forecast Agent: Monitors demand patterns and generates forecast adjustments
- Overbooking Agent: Optimizes overbooking levels while minimizing denied boardings

3.2. Pricing Agent

The Pricing Agent monitors competitive fares, market conditions, and booking velocity to generate pricing recommendations:

Algorithm: Pricing Agent Decision Loop

1. Initialize market state S_0
2. WHILE booking period active:
    a. Collect fare intelligence F_t
    b. Retrieve booking velocity V_t
    c. Query demand forecast D_t
    d. Compute price elasticity estimate ε_t
    e. Generate LLM reasoning chain:

- Context: Market position, competitor fares, demand
- Output: Price adjustment recommendation ΔP

f. Validate against business rules R
g. IF $|\Delta P| > \theta$_escalate:
Escalate to human analyst
ELSE:
Execute price adjustment
h. Log decision with full reasoning trace

The agent leverages fine-tuned LLM capabilities to interpret market signals in context, generating recommendations with natural language explanations suitable for analyst review.

3.3. Inventory Agent

The Inventory Agent manages fare class availability through protection levels and booking limits. Unlike pure optimization approaches, the agent incorporates contextual factors:

Booking Limit Formula:

$BL_j = f(EMSR_j, context_j, history_j)$

where BL_j is the booking limit for fare class j, $EMSR_j$ is the expected marginal seat revenue calculation, and context includes factors such as:

- Special events affecting travel patterns
- Group booking negotiations in progress
- Corporate contract obligations
- Competitive capacity changes

3.4. Forecast Agent

The Forecast Agent continuously monitors booking curves against forecasted demand, identifying anomalies requiring attention:

Anomaly Detection:

Alert = TRUE if $|B_t - \hat{D}_t| > k \cdot \sigma_t$
FALSE otherwise

where B_t is actual bookings, $\hat{D}_t$ is forecasted demand, σ_t is forecast uncertainty, and k is the sensitivity threshold.

When anomalies are detected, the agent generates hypotheses about potential causes:

- Competitor pricing changes
- Special events not in forecast
- Schedule changes affecting connections
- Economic or weather factors

3.5. Overbooking Agent

The Overbooking Agent optimizes authorization levels to maximize revenue while respecting denied boarding constraints:

Objective Function:

max_OB E[Revenue(OB)] - C_DB · P(DB | OB)

where OB is the overbooking level, C_DB is the cost of denied boarding (compensation, rebooking, customer impact), and P(DB|OB) is the probability of denied boarding given overbooking level.

The agent incorporates flight-specific factors including:

- Historical no-show rates by fare class and customer segment
- Day-of-week and seasonal patterns
- Connection traffic and misconnection risk
- Customer value and loyalty status

3.6. Orchestrator and Coordination

The Orchestrator Agent manages inter-agent coordination through a priority-based workflow:

- Event Detection: Monitor real-time feeds for actionable events
- Task Dispatch: Route events to appropriate specialist agents
- Conflict Resolution: Arbitrate when agent recommendations conflict
- Escalation: Route high-impact decisions to human analysts
- Execution: Coordinate approved actions with legacy systems

3.7. Integration Architecture

RevenuePilot integrates with airline operational systems through a dedicated integration layer:

Table 1. System Integration Points

System	Interface	Function
PSS	API/Queue	Inventory updates
GDS	EDIFACT	Fare distribution
CRS	Real-time API	Booking data
Data Warehouse	SQL/ETL	Historical analysis
Fare Filing	ATPCO	Price changes

4. Experimental Setup

4.1. Dataset and Environment

We evaluate RevenuePilot using data from a major U.S. domestic airline:

- Network: 847 daily departures across 156 routes
- Duration: 6-month evaluation period (Q2-Q3 2025)
- Fare Classes: 12 booking classes per flight
- Booking Window: 330 days advance purchase

4.2. Baseline Systems

We compare RevenuePilot against:

- EMSR-b: Production EMSR-b implementation with analyst interventions
- DLP: Deterministic linear programming network optimizer
- ML-Forecast: Machine learning enhanced forecasting with EMSR optimization

4.3. Evaluation Metrics

Primary metrics include:

- RASM: Revenue per Available Seat Mile
- Load Factor: Percentage of seats sold
- Yield: Revenue per Revenue Passenger Mile
- Spoilage: Revenue lost from unsold seats

- Denied Boardings: Passengers denied boarding per 10,000

Evaluation methodologies align with PRAXIS benchmarking principles for production agentic AI systems, including multi-dimensional performance measurement (accuracy, latency, throughput), statistical significance testing, and human-in-the-loop analysis.

4.4. Deployment Configuration

RevenuePilot agents were configured as follows:

- Base Model: GPT-4.1 fine-tuned on airline RM documentation
- Decision Latency: < 500ms for pricing decisions
- Escalation Threshold: $5,000 revenue impact per recommendation
- Human Review: 15% of decisions reviewed by analysts

4.5. LLM Architecture Details

Fine-tuning approach: - Method: Supervised fine-tuning (SFT) with LoRA adapters (rank=16) for parameter-efficient training - Training data: 847K examples comprising RM analyst decision logs, standard operating procedures, pricing playbooks, and competitive response guidelines (anonymized) - Validation: Held-out test set of 12K analyst decisions with 89.2% agreement with human expert labels

Tool invocation mechanism: - Function calling via structured JSON schema with 23 defined tools (fare query, inventory update, forecast retrieval, etc.) - Constrained decoding for numerical outputs (fares, booking limits) to prevent hallucination of invalid values - Schema validation layer rejecting malformed tool calls before execution

Guardrails and safety: - Hard constraints on fare ranges (±15% from filed fares without escalation) - Booking limit bounds enforced at integration layer (cannot exceed physical capacity) - Automatic fallback to EMSR-b when agent confidence < 0.7 or latency > 800ms - All pricing recommendations validated against ATPCO fare rules before filing - Audit logging of full reasoning chains for compliance review

5. Results

5.1. Revenue Performance

RevenuePilot demonstrated significant revenue improvements across the evaluation period:

Table 2. Revenue Performance Comparison

System	RASM	Δ RASM	p-value
EMSR-b (Baseline)	14.82¢	–	–
DLP Network	15.03¢	+1.4%	0.023
ML-Forecast	15.21¢	+2.6%	0.008
RevenuePilot	15.44¢	+4.2%	<0.001

The 4.2% RASM improvement translates to approximately $127M in annualized incremental revenue for the evaluated network. All reported performance deltas (RASM, spoilage, yield, denied boardings) were statistically significant at $p < 0.05$ unless otherwise noted. Statistical testing used paired t-tests for continuous metrics and chi-squared tests for rate comparisons, with Bonferroni correction for multiple comparisons.

5.2. Load Factor and Yield

Table 3. Load Factor and Yield Analysis

System	Load Factor	Yield	Mix Index
EMSR-b	83.2%	17.81¢	1.00
DLP	84.1%	17.88¢	1.02
ML-Forecast	83.8%	18.14¢	1.04
RevenuePilot	84.7%	18.23¢	1.06

RevenuePilot achieved improvements in both load factor (+1.5 points) and yield (+2.4%), indicating effective balance between volume and price optimization.

5.3. Spoilage Reduction

Table 4. Spoilage Analysis by Route Type

Route Type	EMSR-b	RevenuePilot	Δ
Business	8.2%	7.4%	-9.8%
Leisure	11.4%	9.8%	-14.0%
Mixed	9.7%	8.6%	-11.3%
Overall	9.6%	8.4%	-12.5%

The largest spoilage reduction occurred on leisure routes, where the Forecast Agent's ability to incorporate external context (events, weather, holidays) proved most valuable.

5.4. Denied Boarding Performance

Table 5. Denied Boarding Rate per 10,000 Passengers

System	DB Rate	Compensation Cost
EMSR-b	1.24	$847/DB
RevenuePilot	1.14	$823/DB
Improvement	-8.1%	-2.8%

The Overbooking Agent reduced denied boardings while maintaining similar overbooking levels through improved no-show prediction and flight-specific adjustments.

5.5. Response Time and Throughput

Table 6. System Performance Metrics

Metric	Target	Achieved
Pricing Decision Latency	< 500ms	342ms (avg)
Inventory Update Latency	< 200ms	156ms (avg)
Daily Decisions Processed	> 50,000	67,423
System Availability	99.9%	99.94%
Agent Agreement Rate	–	87.3%

5.6. Agent Coordination Analysis

Table 7. Inter-Agent Interaction Patterns

Interaction Type	Frequency
Pricing → Inventory coordination	34.2%
Forecast → Pricing trigger	28.7%
Forecast → Inventory adjustment	21.4%
Overbooking → Inventory constraint	8.9%
Conflict escalation to Orchestrator	6.8%

The majority of decisions were resolved through direct agent coordination, with only 6.8% requiring orchestrator-level conflict resolution.

5.7. Human-In-The-Loop Analysis

Table 8. Human Review Outcomes

Decision Type	Review Rate	Override Rate
Standard pricing	8.2%	3.1%
Competitive response	24.6%	7.8%
Forecast adjustment	18.4%	12.3%
Overbooking change	31.2%	9.4%
Overall	15.1%	6.2%

Analysts overrode 6.2% of reviewed decisions, with highest override rates on forecast adjustments where local market knowledge proved valuable.

5.8. Ablation Study: Agentic Reasoning Vs Rule-Based Control

To isolate the contribution of LLM-based agentic reasoning from traditional optimization components, we evaluated a variant of RevenuePilot where contextual reasoning was disabled and replaced with rule-based heuristics using identical EMSR, forecast, and optimization inputs.

Table 9. Ablation Analysis Isolating Agentic Contribution

Component	Full RevenuePilot	Rule-Based Variant	Δ Attribution
RASM Improvement	+4.2%	+1.7%	+2.5% agentic
Spoilage Reduction	-12.5%	-4.8%	-7.7% agentic
DB Rate Reduction	-8.1%	-2.9%	-5.2% agentic
Forecast Accuracy	+18.3%	+6.1%	+12.2% agentic

Key findings from ablation:

1. ML forecasting baseline (without agentic reasoning) contributed approximately +1.7% RASM improvement over EMSR-b, consistent with prior ML-enhanced RM literature.
2. Agentic coordination contributed an incremental +2.5% RASM beyond the ML baseline, primarily through:
    - Contextual interpretation of booking anomalies (38% of incremental gain)
    - Coordinated multi-agent pricing/inventory decisions (35%)
    - Adaptive overbooking based on flight-specific context (27%)
3. Spoilage reduction was disproportionately driven by agentic reasoning (7.7% of 12.5% total), reflecting the Forecast Agent's ability to incorporate external context that rule-based systems cannot capture.

The ablation confirms that RevenuePilot's performance gains are not solely attributable to improved forecasting or optimization algorithms, but derive substantially from the contextual reasoning and coordination capabilities of the agentic architecture.

6. Discussion

6.1. Key Success Factors

RevenuePilot's performance improvements derive from several factors:

- Contextual Reasoning: The Forecast Agent's ability to interpret booking anomalies in context connecting demand patterns to external events, competitive actions, or schedule changes enabled faster and more accurate forecast corrections than rule-based anomaly detection.
- Coordinated Optimization: Multi-agent coordination allowed simultaneous pricing and inventory adjustments, avoiding the sequential decision-making that creates suboptimal outcomes in traditional systems.
- Explainability: Natural language explanations for recommendations increased analyst confidence and enabled more efficient review processes, reducing decision latency while maintaining governance.

6.2. Integration Challenges

Operational deployment revealed several integration challenges:

- Legacy System Latency: Some GDS updates required 15-30 minute propagation, limiting response speed
- Data Quality: Historical data inconsistencies required extensive preprocessing
- Organizational Change: Analyst workflows required adaptation to agent-assisted decision-making

6.3. Limitations

Several limitations should be noted:

- Network Scope: Evaluation limited to domestic network; international routes present additional complexity
- Market Conditions: Testing period represented relatively stable demand conditions
- Competitive Environment: Evaluation occurred during routine competitive fare dynamics; extreme competitive actions (fare wars, major capacity changes) were not observed, which may limit generalizability to highly volatile market conditions
- Generalization: Results may vary across airlines with different network structures and customer segments

6.4. Comparison with Alternative Approaches

RevenuePilot's approach differs from pure ML optimization in several key ways:

- Interpretability: Decisions include natural language explanations vs. black-box outputs

- Adaptability: Can incorporate new context without retraining
- Human Collaboration: Designed for analyst partnership rather than replacement
- Graceful Degradation: Falls back to established heuristics when uncertain

7. Lessons Learned

7.1. Architecture Decisions

- Agent Specialization: Specialized agents aligned with RM functional areas (pricing, inventory, forecast, overbooking) proved more effective than general-purpose agents, enabling domain-specific fine-tuning and clearer responsibility boundaries.
- Orchestration Importance: The Orchestrator Agent's role in conflict resolution and escalation was critical for operational reliability. Early designs without centralized coordination led to conflicting actions.
- Integration Layer: Abstracting legacy system interfaces through a dedicated integration layer enabled agent logic evolution without system integration rework.

7.2. Operational Considerations

- Gradual Rollout: Phased deployment starting with advisory mode (recommendations only) before autonomous execution built analyst trust and identified edge cases.
- Monitoring and Alerting: Comprehensive logging of agent reasoning chains proved essential for debugging and continuous improvement.
- Fallback Mechanisms: Automatic fallback to traditional EMSR calculations during system issues maintained operational continuity.

8. Conclusion

RevenuePilot demonstrates that agentic AI can effectively augment traditional airline revenue management systems, achieving meaningful improvements in revenue, spoilage, and operational metrics. The multi-agent architecture, designed around established RM workflows and integrated with legacy systems, provides a practical path to operationalizing LLM capabilities in high-stakes commercial environments.

Key findings include:

- 4.2% RASM improvement over production EMSR baseline
- 12.5% reduction in seat spoilage
- 8.1% reduction in denied boarding rate
- 93.8% acceptance rate of agent recommendations by analysts

The combination of quantitative optimization (EMSR, bid-price controls) with contextual reasoning (LLM agents) represents a promising direction for next-generation RM systems that balance automation with human oversight.

9. Future Work

Future development directions include:

- Extension to international network optimization with origin-destination control
- Integration of ancillary revenue optimization (baggage, seats, upgrades)
- Real-time competitive fare monitoring with automated response strategies
- Customer lifetime value integration for personalized pricing
- Multi-carrier alliance coordination for connecting itineraries

Acknowledgment

We thank the revenue management practitioners who provided domain expertise and operational feedback throughout this research.

References

[1] K. Talluri and G. Van Ryzin, "The Theory and Practice of Revenue Management," Springer Science & Business Media, 2004.

[2] L. R. Weatherford and S. E. Bodily, "A taxonomy and research overview of perishable-asset revenue management: Yield management, overbooking, and pricing," Operations Research, vol. 40, no. 5, pp. 831-844, 1992.

[3] P. P. Belobaba, "Application of a probabilistic decision model to airline seat inventory control," Operations Research, vol. 37, no. 2, pp. 183-197, 1989.

[4] K. Littlewood, "Forecasting and control of passenger bookings," Airline Group International Federation of Operational Research Societies Proceedings, vol. 12, pp. 95-117, 1972.

[5] K. Talluri and G. Van Ryzin, "An analysis of bid-price controls for network revenue management," Management Science, vol. 44, no. 11, pp. 1577-1593, 1998.

[6] E. L. Williamson, "Airline network seat inventory control: Methodologies and revenue impacts," Ph.D. dissertation, Massachusetts Institute of Technology, 1992.

[7] J. I. McGill and G. J. Van Ryzin, "Revenue management: Research overview and prospects," Transportation Science, vol. 33, no. 2, pp. 233-256, 1999.

[8] B. Vinod, "Evolution of yield management in travel," Journal of Revenue and Pricing Management, vol. 15, no. 3-4, pp. 203-211, 2016.

[9] L. R. Weatherford, "The history of forecasting models in revenue management," Journal of Revenue and Pricing Management, vol. 15, no. 3-4, pp. 212-221, 2016.

[10] L. R. Weatherford and S. Kimes, "A comparison of forecasting methods for hotel revenue management," International Journal of Forecasting, vol. 19, no. 3, pp. 401-415, 2003.

[11] C. Chen et al., "Deep learning for airline revenue management," Journal of Revenue and Pricing Management, vol. 19, no. 4, pp. 234-248, 2020.

[12] R. Rana and F. S. Oliveira, "Real-time dynamic pricing in a non-stationary environment using model-free reinforcement learning," Omega, vol. 47, pp. 116-126, 2014.
[13] Y. Zhang et al., "Transformer-based demand forecasting for airline revenue management," Transportation Research Part E, vol. 158, 102589, 2022.
[14] S. Yao et al., "ReAct: Synergizing reasoning and acting in language models," in Proc. ICLR, 2023.
[15] L. Wang et al., "A survey on large language model based autonomous agents," arXiv preprint arXiv:2308.11432, 2024.
[16] Agarwal, S. (2025). AI-Augmented Social Media Marketing: Data-Driven Approaches for Optimizing Engagement. *International Journal of Emerging Research in Engineering and Technology*, *6*(2), 15-23. https://doi.org/10.63282/3050-922X.IJERET-V6I2P103.