Toward an Autonomous Business Operating System for the U.S. Private Security Industry: A Multi-Agent AI Framework for End-to-End Workflow Automation

K M FAZLE RABBI
Department of Computer Science and Engineering
St. Cloud State University

***Abstract*—Private security providers coordinate scheduling, emergency coverage, timekeeping, reporting, compliance, payroll, billing, and client communication through workflows that are often distributed across separate systems and manual processes. This paper proposes ASBOS, an Autonomous Security Business Operating System that uses an event-driven, multi-agent architecture to coordinate these functions while preserving human approval for consequential actions. Six specialized agents support workforce scheduling, emergency voice outreach, timesheet and payroll reconciliation, report generation, compliance and human resources, and billing and client relationship management. A policy-aware orchestration layer records event provenance, limits agent permissions, prevents duplicate actions, and routes exceptions to supervisors. The design is grounded in public workforce data, current security-workforce platforms, prior scheduling research, and documented voice-AI capabilities. A prospective evaluation protocol is presented for measuring administrative workload, shift-fill performance, payroll exceptions, report timeliness, credential compliance, human escalation, and operational reliability. An economic sensitivity model illustrates how assumptions about labor savings, exception reduction, and annual system cost affect expected value. ASBOS is positioned as a research framework for auditable cross-system coordination rather than as a replacement for guards, managers, or authoritative business systems.**

***Index Terms*—AI agents, multi-agent systems, private security, workforce management, voice AI, scheduling, timekeeping, compliance, human-in-the-loop automation, event-driven architecture.**

I. Introduction

Private security companies manage a continuous flow of operational decisions: assigning qualified guards, responding to call-offs, reconciling time records, delivering client reports, maintaining licenses, and converting approved service hours into invoices. These tasks are tightly connected, yet many organizations still coordinate them through a mixture of workforce software, spreadsheets, messages, telephone calls, and manager judgment. A missed clock-in, for example, can become a staffing problem, a client-communication issue, a payroll exception, and a billing adjustment within minutes.

The U.S. Bureau of Labor Statistics reports 1,272,400 workers in the combined security-guard and gambling-surveillance occupation in 2024 and projects approximately 162,300 openings per year during 2024–2034, primarily because workers leave the occupation or labor force [1]. A Center for American Progress analysis of Census Quarterly Workforce Indicators reports 50.8% annual turnover in the security industry in 2023, compared with 38.4% across the private sector [2]. These figures indicate that staffing continuity and replacement demand are persistent operating concerns.

Commercial platforms already provide substantial combinations of scheduling, time and attendance, guard touring, payroll, billing, compliance, reporting, and analytics [9], [10], [11]. Prior research has also addressed security-guard scheduling through robust optimization, genetic algorithms, and reinforcement learning [12], [13], [14]. The remaining systems question is not whether individual functions can be digitized. It is whether a security-specific orchestration layer can coordinate them across systems, act within explicit limits, preserve an interpretable record of each decision, and improve outcomes without transferring unacceptable risk to automated agents.

This paper contributes: (1) a six-agent ASBOS architecture; (2) a typed event model for cross-system coordination; (3) a tiered human-oversight model; (4) an emergency shift-coverage workflow using bounded voice automation; and (5) a prospective evaluation and economic sensitivity framework.

II. Industry and Technology Context

A. Workforce Operations

Security operations are unusually sensitive to time, location, qualifications, and continuity. A schedule is not merely a list of names and hours: each assignment may depend on a guard license, training status, site-specific authorization, overtime constraints, travel time, client rules, and post orders. Operational systems must therefore distinguish between a feasible assignment and an assignment that is legally, contractually, and operationally acceptable.

Timekeeping is similarly connected to downstream processes. A clock event may influence payroll, client billing, attendance history, staffing metrics, and an investigation. Research on employee time theft shows that the behavior is multidimensional and cannot be reduced to a single universal payroll-loss percentage [3]. Historical biometric timekeeping studies reported measurable reductions in buddy punching and payroll error, but those findings are

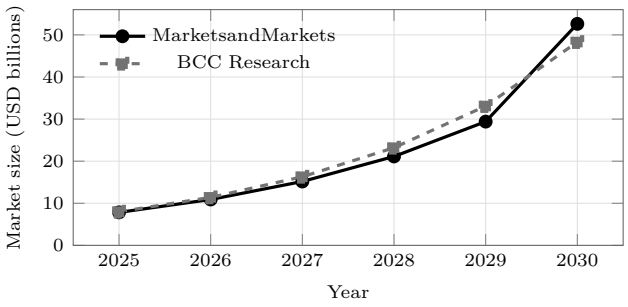

Fig. 1. Two published AI-agent market forecasts. Intermediate values are interpolated only to visualize the reported 2025 and 2030 endpoints [6], [7].

not specific to private security and should not be treated as expected ASBOS outcomes [4], [5]. A field deployment should establish its own baseline exception rate before estimating savings.

B. Market Direction

Analyst forecasts indicate rapid expansion in agent-based software, although market estimates vary by scope and methodology. MarketsandMarkets projects the AI-agents market from $7.84 billion in 2025 to $52.62 billion in 2030, while BCC Research projects growth from approximately $8.0 billion to $48.3 billion over the same period [6], [7]. Market.us projects the voice-AI-agent market from $2.4 billion in 2024 to $47.5 billion by 2034 [8]. These forecasts do not establish operational effectiveness, but they indicate that the underlying model, speech, telephony, and orchestration infrastructure is becoming commercially accessible.

C. Related Systems and Research Gap

Trackforce describes integrated scheduling, time and labor, guard touring, reporting, payroll, billing, compliance, and analytics [9]. Belfry describes scheduling, geofenced timecards, payroll, invoicing, field operations, and incident reporting [10]. Celayix describes connected scheduling, time and attendance, guard touring, payroll, billing, and third-party integrations [11]. These products demonstrate that broad workflow coverage is already possible.

ASBOS therefore focuses on coordinated autonomous execution across functions. Its research question is: *Can an event-driven multi-agent layer reduce administrative delay and exception handling while maintaining clear authority boundaries, reversible actions, and human control?* The design extends scheduling research by connecting an assignment decision to communication, compliance, timekeeping, reporting, and billing events rather than treating scheduling as an isolated optimization problem.

III. ASBOS Architecture

ASBOS is an orchestration layer above authoritative systems of record. Agents receive only the data and tools required for their assigned function. They do not receive unrestricted access to payroll, personnel, licensing, client, and billing data. Every event includes a unique identifier, timestamp, source system, organization and site scope, data classification, correlation identifier, and processing status. Agent actions are designed to be idempotent so replayed events do not create duplicate shifts, messages, reports, or invoice lines.

A. Event and Policy Layer

The orchestration service receives events for shift creation, missed clock-ins, expiring credentials, overdue reports, opened incidents, and invoices coming due. For each event, it evaluates policy, selects an agent, records tool calls and outputs, applies retry and duplicate-prevention rules, and routes exceptions to a supervisor. A shared schema supports interoperability without forcing every application into a single database.

Policy is separated from model behavior. Deterministic rules handle hard constraints such as expired credentials, maximum approved pay premiums, prohibited assignments, and required approvals. Models may rank options, summarize context, or draft communication, but they may not override a hard rule. This separation also enables versioned policy review and reproducible testing.

B. Workforce Scheduling Agent

The Scheduling Agent proposes assignments using availability, qualifications, overtime risk, site rules, travel constraints, historical attendance, and employee preferences. It produces a ranked recommendation with constraint explanations and alternatives. During early deployment, a manager approves the schedule before publication. Later automation can be limited to low-risk assignments that satisfy all required rules.

C. Emergency Shift Coverage Voice Agent

The Voice Agent receives an approved ranked list of eligible replacements and a bounded call script. It identifies the organization and automated nature of the call, communicates the location, start time, expected duration, approved compensation, and response deadline, then records acceptance or decline. Ambiguous requests, negotiations outside the approved range, emergencies, and complaints transfer immediately to a person.

Bland documents outbound calls, batch calling, API actions during calls, disclosure controls, guardrails, and transfers [15]. Retell reports end-to-end latency as low as 600 ms under suitable conditions [16]; Vapi describes approximately 500–700 ms as an optimized voice-to-voice range [17]. These capabilities support implementation feasibility, but the relevant operational outcome is the time required to obtain suitable coverage, not model latency alone.

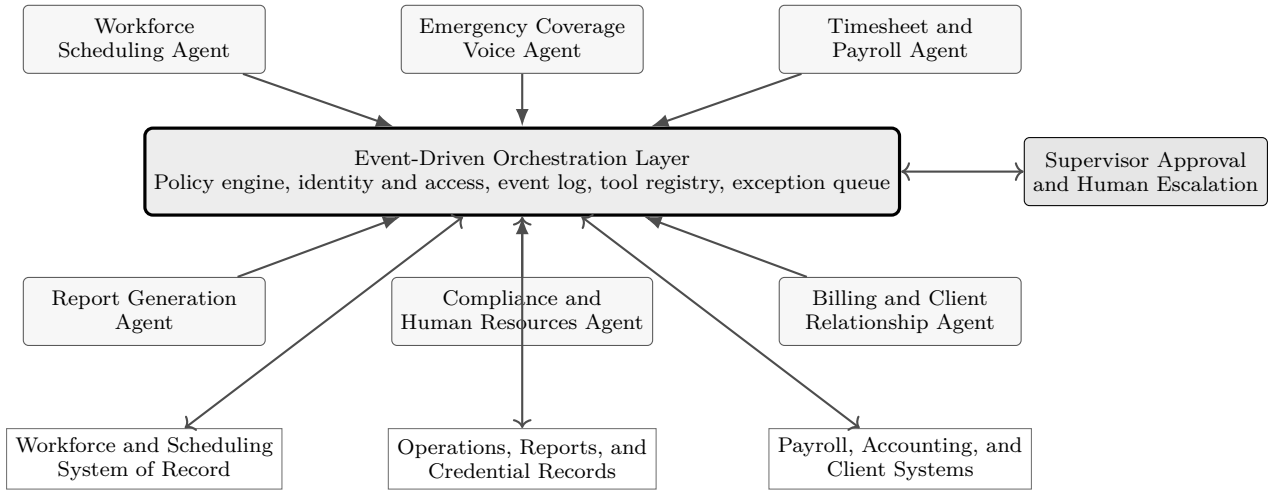

Fig. 2. ASBOS multi-agent architecture. Specialized agents act through a shared orchestration and policy layer; authoritative systems and supervisor approval remain explicit.

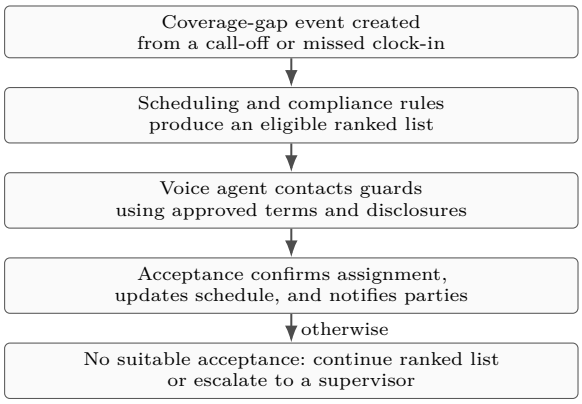

Fig. 3. Bounded emergency shift-coverage workflow. Eligibility and compensation limits are established before automated outreach.

D. Timesheet and Payroll Agent

This agent reconciles scheduled shifts, approved clock events, location evidence, pay rules, and exception codes. It creates a review queue for missing punches, overlaps, off-site clock-ins, hours materially beyond the schedule, and unapproved rate changes. It does not label an employee as fraudulent. Payroll submission requires explicit authorization and a reversible export or API transaction.

E. Report Generation Agent

The Report Agent assembles structured activity data and drafts daily activity reports, incident narratives, and client summaries. Each statement must link to a source event or be clearly identified as an inference requiring review. High-severity incident communications require human approval. The system stores the source record set, draft, revisions, approver, and final report.

F. Compliance and Human Resources Agent

The Compliance Agent tracks credential type, jurisdiction, issue and expiration dates, document status, required training, and site eligibility. Deterministic rules can block an assignment when a mandatory credential has expired. Incomplete or disputed records are escalated. Labor-rule logic must be versioned by jurisdiction and maintained by qualified organizational personnel.

G. Billing and Client Relationship Agent

The Billing Agent prepares invoice lines from reconciled hours, contract rates, approved premiums, and service exceptions. It may issue routine reminders and flag aging balances, but write-offs, contract changes, and disputed charges require human action. Client-health indicators should be interpretable and must not incorporate protected characteristics.

IV. Human Oversight, Privacy, and Security

ASBOS uses three action tiers. Tier 1 actions are low-risk and reversible, such as drafting a schedule, preparing a report, or sending an approved reminder. Tier 2 actions may execute automatically while notifying a supervisor, such as confirming a replacement within a pre-approved compensation range. Tier 3 actions require approval, including unbudgeted premiums, adverse personnel actions, payroll changes, credential exceptions, and high-severity incident communications.

Voice deployment requires jurisdiction-specific review of recording consent, AI disclosure, call authorization, data retention, and opt-out handling. Federal Trade Commission guidance describes disclosure, consent, caller-identification, and opt-out obligations for covered telemarketing and prerecorded-call practices, while state recording laws may impose additional requirements [18]. Internal workforce calls are not automatically telemarketing, but organizations must review the actual use case and jurisdictions involved.

TABLE I
Recommended Evaluation Outcomes

Outcome	Operational definition
Workload	Manager minutes per 100 scheduled guard-hours, separated by workflow
Coverage	Median and 90th-percentile fill time; percentage filled before start
Timekeeping	Exceptions per 1,000 worked hours, correction time, and false-positive rate
Reporting	On-time delivery, revision count, and unsupported-statement rate
Compliance	Assignment blocks, overrides, credential incidents, and correction time
Oversight	Escalations, override rate, approval latency, and reason codes
Reliability	Availability, duplicate actions, failed integrations, incidents, and recovery time

Illinois's Biometric Information Privacy Act includes voiceprints within its definition of biometric identifiers [19]. Ordinary audio recording is not necessarily voiceprint processing. If the system performs speaker recognition, voice authentication, or creates a voice template, it must apply relevant notice, consent, retention, security, and deletion controls. The default architecture does not require biometric voice templates.

Security controls include least-privilege service identities, encryption in transit and at rest, tenant and site isolation, scoped tool permissions, immutable event history, secrets management, rate limits, data-retention policies, and emergency suspension of an agent or integration. Every consequential action must identify the initiating event, policy version, model and tool versions, inputs, output, and approving person when approval is required.

V. Prospective Evaluation

A research deployment should use a pre-specified protocol rather than relying only on before-and-after impressions. A 90-day design can include a baseline period, staged activation of lower-risk agents, supervised voice outreach, and a final observation period. A stepped-wedge or matched-site design is preferable when multiple sites are available because it reduces the effect of seasonal demand and site-specific differences.

Primary outcomes should include administrative minutes per scheduled guard-hour, median and 90th-percentile coverage time, the percentage of shifts filled before start, payroll exceptions per 1,000 hours, report delivery compliance, credential-related blocks, human escalation, false positives, and system availability. Secondary outcomes may include guard acceptance, supervisor workload, client satisfaction, and perceived fairness.

The study should retain raw event logs, outcome definitions, policy versions, model and tool versions, human overrides, and data-exclusion rules. Results should report confidence intervals and distinguish statistical significance from operational significance. Failures and near misses should be included rather than removed from the analysis.

TABLE II
Illustrative Annual Sensitivity Model for 100 Guards

Scenario	Benefit	Cost	Net ROI	B/C
Low	\$30,399	\$60,000	−49%	0.51
Base	\$60,798	\$45,000	35%	1.35
High	\$100,797	\$30,000	236%	3.36

VI. Economic Sensitivity Analysis

The economic model is intentionally limited to two categories: administrative labor saved and payroll exceptions prevented or corrected. It excludes retention, client renewal, days-sales-outstanding, liability reduction, and new revenue to reduce double counting. For illustration, a 100-guard organization is modeled using 2,080 annual hours per guard and the BLS May 2024 median wage of \$18.46 per hour [1], producing approximately \$3.84 million in annual gross guard payroll.

Low, base, and high scenarios assume 10, 20, and 30 administrative hours saved per week at a \$40 loaded hourly cost, together with payroll improvement equal to 0.25%, 0.50%, and 1.00% of gross payroll. These percentages are scenario assumptions, not universal loss rates. Annual system costs are modeled at \$60,000, \$45,000, and \$30,000.

The model shows why evaluation quality matters. Under conservative assumptions the system does not recover its cost; under stronger but still plausible assumptions the benefit-cost ratio improves substantially. A pilot should report gross benefit, recurring cost, implementation cost, net benefit, benefit-cost ratio, and uncertainty separately.

VII. Discussion

The central contribution of ASBOS is coordination. Current platforms digitize broad portions of security workforce management, and prior studies optimize schedules. ASBOS adds an event and policy layer that links decisions across operational domains while preserving explicit authority boundaries. A coverage gap can trigger eligibility checks, ranked outreach, assignment confirmation, schedule updates, client notification, timekeeping preparation, and an event record without allowing a language model to bypass hard constraints.

The emergency voice agent is the highest-risk component because it combines automated outreach, possible recording, personal data, and negotiation-like interaction. Its scope should remain narrow: approved recipients, approved pay bands, standard shift information, explicit disclosure, limited retention, and immediate human handoff. The second major risk is silent propagation of incorrect data. A wrong credential date, site rule, or pay rate can affect several systems. Event provenance, policy versioning, reversible writes, and reconciliation are therefore foundational controls rather than optional implementation details.

ASBOS should not be evaluated only by how many tasks it automates. A mature implementation may deliberately escalate a meaningful share of cases because the objective

is reliable operations, not maximum autonomy. Measures such as override quality, false-positive exceptions, recovery time, and employee acceptance are as important as speed.

VIII. Limitations and Future Work

This paper presents an architecture and evaluation plan, not results from a completed production deployment. Market figures are third-party forecasts and differ by publisher and category definition. Public product documentation may omit capabilities, constraints, implementation details, and roadmap items. Voice performance varies with accents, noise, network conditions, prompts, and telephony providers. Scheduling and compliance rules differ across organizations and jurisdictions.

Future work should include a registered pilot protocol, de-identified event dataset, baseline work-sampling study, failure-mode and effects analysis, fairness testing of replacement ranking, usability testing across accents and noisy environments, independent security and privacy review, and comparison with both manual operations and existing workforce-management automation.

IX. Conclusion

ASBOS proposes a security-specific, event-driven multi-agent layer for coordinating scheduling, emergency coverage, timekeeping, reporting, compliance, payroll, billing, and client communication. Its architecture keeps authoritative systems, deterministic rules, and human approval explicit while using agents for ranking, drafting, communication, reconciliation, and exception handling. The framework is technically feasible with current workforce, integration, and voice-AI infrastructure, but its value must be established through a reproducible field study. The next step is not broader claims of autonomy; it is evidence showing which workflows can be coordinated safely, where human intervention remains necessary, and whether measured operational benefits exceed total system cost.

References

[1] U.S. Bureau of Labor Statistics, "Security Guards and Gambling Surveillance Officers," *Occupational Outlook Handbook*, updated Aug. 28, 2025. [Online]. Available: https://www.bls.gov/ooh/protective-service/security-guards.htm

[2] A. Glass, "Low Standards Hurt Security Officers' Ability To Make Ends Meet," Center for American Progress, Oct. 1, 2025. [Online]. Available: https://www.americanprogress.org/article/low-standards-hurt-security-officers-ability-to-make-ends-meet/

[3] C. M. Harold, B. Hu, and J. Koopman, "Employee time theft: Conceptualization, measure development, and validation," *Personnel Psychology*, vol. 75, no. 2, pp. 347–382, 2022, doi: 10.1111/peps.12477.

[4] Nucleus Research, "ROI Evaluation Report: Kronos 4500 Touch ID Terminal," Research F24, Apr. 28, 2005. [Online]. Available: https://nucleusresearch.com/research/single/roi-evaluation-report-kronos-4500-touch-id-terminal/

[5] Nucleus Research, "Automating Time and Attendance with Biometrics Reduces Payroll Error and Boosts Productivity," Apr. 30, 2009. [Online]. Available: https://nucleusresearch.com/news/automating-time-and-attendance-with-biometrics-reduces-payroll-error-and-boosts-productivity/

[6] MarketsandMarkets, "AI Agents Market—Global Forecast to 2030," Report TC 9168, Apr. 2025. [Online]. Available: https://www.marketsandmarkets.com/Market-Reports/ai-agents-market-15761548.html

[7] BCC Research, "AI Agents: Technologies, Applications and Global Markets," Report AIT012A, Oct. 2025. [Online]. Available: https://www.bccresearch.com/market-research/artificial-intelligence-technology/ai-agent-market.html

[8] Market.us, "Voice AI Agents Market Size, Share, Statistics Analysis," Apr. 2025. [Online]. Available: https://market.us/report/voice-ai-agents-market/

[9] Trackforce, "Security Workforce Management Optimized," product overview, 2026. [Online]. Available: https://www.trackforce.com/

[10] Belfry, "Security Guard Management System," product overview, 2026. [Online]. Available: https://www.belfrysoftware.com/

[11] Celayix, "Security Guard Management Software," product overview, 2026. [Online]. Available: https://www.celayix.com/industries/security-guard-officer-scheduling-software/

[12] T. Okimoto, "Robust Security Guard Scheduling Problem in Dynamic Environment," *Transactions of the Japanese Society for Artificial Intelligence*, vol. 39, no. 5, 2024, doi: 10.1527/tjsai.39-5_D-O14.

[13] I. F. Ashari *et al.*, "Security Guard Scheduling System Using Genetic Algorithm and Tournament Selection," *Journal of Applied Informatics and Computing*, vol. 5, no. 2, pp. 202–207, 2021, doi: 10.30871/jaic.v5i2.3464.

[14] C. Jiang, W. Kusakunniran, N. Pomprasatpol, C. Limsuwankesorn, and Y. Li, "Smart Security Guard Scheduling System Based on Reinforcement Learning," in *Proc. Int. Computer Science and Engineering Conf.*, pp. 214–218, 2017, doi: 10.1109/ICSEC.2017.8443946.

[15] Bland AI, "Welcome to Bland" and "Guard Rails," official documentation, 2026. [Online]. Available: https://docs.bland.ai/welcome-to-bland; https://docs.bland.ai/tutorials/guard-rails

[16] Retell AI, "Check Estimated Latency," official documentation, 2026. [Online]. Available: https://docs.retellai.com/reliability/check-estimated-latency

[17] Vapi, "Core Models" and "Definitions," official documentation, 2026. [Online]. Available: https://docs.vapi.ai/quickstart; https://docs.vapi.ai/glossary

[18] Federal Trade Commission, "Complying with the Telemarketing Sales Rule," updated 2025. [Online]. Available: https://www.ftc.gov/business-guidance/resources/complying-telemarketing-sales-rule

[19] Illinois General Assembly, "Biometric Information Privacy Act," 740 ILCS 14, current text. [Online]. Available: https://www.ilga.gov/Legislation/ILCS/Articles?ActID=3004&ChapterID=57

[20] M. Wooldridge and N. R. Jennings, "Intelligent agents: Theory and practice," *The Knowledge Engineering Review*, vol. 10, no. 2, pp. 115–152, 1995.

[21] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed. Pearson, 2020.