

C A S E S T U D Y

How a CLU Agent Took Over 87% of a Help Desk in 37 Days

6,775 real cases. 37 days in production. Not a pilot.

81% Cost Reduction	88.9% SLA Compliance	87.4% Workload Takeover	6,775 Cases Resolved
------------------------------	--------------------------------	-----------------------------------	--------------------------------

Powered by **CLU** | cluagents.com

Industry: BPO / Identity Verification | Period: February – March 2026

Executive Summary

A BPO services provider operating an identity verification help desk for a national organization faced a critical operational challenge: fluctuating ticket volumes, strict 4-hour SLAs, and a 3-person team costing \$2,700 USD/month that could only operate during business hours.

Using CLU's AI agent orchestration platform, the organization deployed an autonomous agent designed to handle document verification cases end-to-end. Deployed on February 18, 2026, the agent processed 6,775 real cases in just 37 days—becoming the primary operator handling 87.4% of all volume by March.

KEY RESULTS

81% reduction in monthly operating cost (USD 2,700 → USD 500)

88.9% SLA compliance vs. 83.1% for human agents (+5.8 percentage points)

Near-instant median response (0.00 business hours vs. 1.68 hours for humans)

82.2% approval rate — identical to human agents (82.3%), confirming quality parity

The Challenge

The organization's help desk operations were built on a traditional model: human agents manually reviewing identity documents against verification criteria. This created compounding inefficiencies that impacted both cost structure and service quality.

Operational Reality

- Monthly operating cost of \$2,700 USD for a 3-person team handling variable volumes
- Processing limited to business hours, leaving overnight ticket queues unaddressed
- Volume fluctuations from 1,586 cases (January) to 6,000+ (peak months) made staffing unpredictable
- SLA compliance inconsistent across agents, ranging from 57.1% to 96.1%

The core problem was not a lack of skilled agents—the team was competent. The problem was that the manual, business-hours-only model could not scale to meet demand without proportionally increasing headcount and cost.

The Solution

CLU deployed an autonomous AI agent designed to process identity verification cases end-to-end. Unlike a rule-based bot, the CLU agent operates within CLU's orchestration framework, coordinating document analysis, verification logic, and resolution decisions autonomously.

How It Works

The agent was built on CLU's GRID framework and deployed alongside the existing human team—no disruption, no big-bang migration. Key capabilities include:

- Batch processing of verification cases overnight, clearing the queue before human agents arrive
- Real-time document analysis applying the same approval/rejection criteria as human reviewers
- Autonomous SLA management, prioritizing cases approaching the 4-hour threshold
- Seamless escalation to human agents for edge cases requiring manual review

The deployment moved from integration to production in under one week, with the agent handling real cases from day one.

Results in Detail

All metrics below are from 37 days of sustained production operations (February 18 – March 26, 2026), measured against the existing human team's baseline. These are not pilot results or projections. Data is derived from 60,346 real operational records.

1. Cost Efficiency: 81% Reduction

Metric	Before	After
Monthly operating cost	USD 2,700	USD 500
Cost per case	—	USD 0.07
Monthly savings	—	USD 2,200
Projected annual savings	—	USD 26,400

At \$500/month, the CLU agent costs 81% less than the previous 3-person team while handling a larger share of the total volume.

2. Speed: Near-Instant Resolution

Metric	Human Agents	CLU Agent
Median response time	1.68 hours	0.00 hours
Average response time	2.46 hours	1.39 hours
Speed improvement	—	44% faster (avg)

Response time measured in business hours. CLU achieves near-zero median by processing cases in overnight batch runs.

3. Scale: 6,775 Cases in 37 Days

Over the 37-day measurement period, the CLU agent processed 6,775 cases—averaging 183 cases per day. Volume grew from 2,137 cases in February (partial month, deployed Feb 18) to 4,631 in March, a 117% increase demonstrating both reliability and organic trust from the operations team.

January 2026	February 2026	March 2026
0%	33%	87.4%
CLU not deployed 1,586 cases — all human	CLU deployed Feb 18 2,137 of 6,453 cases	CLU primary operator 4,631 of 5,301 cases

4. Quality: Identical Decision Accuracy

The CLU agent maintained an approval rate of 82.2%—virtually identical to the human team’s 82.3%. This is the strongest indicator that the agent applies the same verification criteria as trained human reviewers: no rubber-stamping, no over-rejection, no bias drift.

SLA compliance reached 88.9% for the CLU agent versus 83.1% for human agents—a 5.8 percentage point improvement measured against the 4-hour business-hours threshold.

What Made the Difference

This was not a chatbot deployment. Three architectural decisions drove the results:

Overnight Batch Processing

83% of the CLU agent’s processing happens between midnight and 7 AM. Human agents arrive each morning to a cleared queue rather than a backlog. This turns a business-hours-only operation into a 24/7 capability without adding shifts.

Human-in-the-Loop by Design

The agent was deployed alongside the existing team, not as a replacement. Human agents handle edge cases, quality assurance, and escalations. The 12.6% of March cases still handled by humans represent genuine complexity—not agent failures.

Production-First Approach

CLU deployed the agent into production handling real cases from day one. Every metric in this case study reflects actual operational performance with real documents, real SLAs, and real business processes. The ramp from 0% to 87% was organic, driven by demonstrated reliability.

Looking Ahead

With the CLU agent now handling the majority of verification cases, the client's operations team has shifted from volume processing to quality oversight and exception management. The next phase focuses on expanding the agent's scope to additional document types and verification workflows, further reducing the need for manual intervention while maintaining the same quality standards.

Ready to orchestrate your AI agents?

CLU helps enterprises deploy autonomous AI agents that deliver measurable results in production—not just in demos.

[Book a live demo at cluagents.com](https://cluagents.com)

Disclaimer: Results are based on actual production data from a specific deployment (February 18 – March 26, 2026). Individual outcomes may vary depending on operational complexity, data quality, and integration scope. All metrics derived from 60,346 operational records. Cost comparisons use the client's pre-deployment team economics as baseline. Client identity anonymized at their request.