Published by :
https://www.ijert.org/
An International Peer-Reviewed Journal

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 12, UCBO - 2026

AI-Enabled Service Desk Process Improvement: A Framework for Ticket Resolution Automation and SAP MM Procurement Integration

Tanishka Jeevan Wagh
MBA, School of Business,
Indira University,
Pune, India 411033
tanishkawagh7188@gmail.com

Dr. Nitin Joshi
Associate Professor and Area Chair - MBA, School of Business,
Indira University,
Pune, India 411033
nitinjoshi.sob@indirauniversity.edu.in

Dr. Hetal Bhinde
Associate Dean Ph.D. Programs,
School of Business,
Indira University,
Pune, India-411033
hetalb19@gmail.com
ORCID: https://orcid.org/0000-0002-6435-5113

Abstract - **In industrial equipment manufacturing, after-sales responsiveness depends on how quickly customer enquiries can be interpreted, matched with order and material data held in enterprise systems, and answered consistently. Contemporary literature positions artificial intelligence (AI) as a means of automating routine, rules-bound service work while keeping people responsible for exceptions. This applied, exploratory study examines the customer service ticket-resolution process at Atlas Copco India, where Customer Service Officers (CSOs) resolve enquiries by moving between the Halo service desk platform, the SAP/HUB information layer and Standard Operating Procedures (SOPs). The objectives were to (1) document the current ticket-resolution workflow, (2) analyze ticket volume and lead-time trends, (3) design a human-in-the-loop AI orchestration layer with decision rules, risk controls and key performance indicators (KPIs), and (4) recommend a phased implementation roadmap. The methodology combined process observation and mentor-supported walk-throughs during an on-the-job training placement that began in November 2025, analysis of aggregate Halo dashboard data for October 2025 to July 2026, a structured literature review and conceptual solution design. Exposure to warehouse operations and SAP Materials Management (SAP MM) supplied supporting context.**

The analysis showed that 3,924 tickets were received between October 2025 and June 2026, about 436 a month, while 3,567 were closed, leaving a cumulative gap of about 357 tickets. Average lead time to close rose from 3 in October 2025 to 22 in April and May 2026 and to 38 in July 2026, a more than twelve-fold increase, while monthly intake stayed between roughly 400 and 510 after the first month. Process mapping located the effort in system-switching, manual SOP look-up, ad hoc routing and manual drafting of recurring responses. A five-layer architecture is proposed in which an AI agent classifies tickets, extracts entities, performs read-only SAP/HUB lookups and applies explicit rules, with CSO review for exceptions and an audit log. The design is a proposal, not a validated system, and its benefits must be tested in a pilot. It is recommended that management standardize master data and ticket categories, digitize SOPs, fix KPI baselines, pilot a human-in-the-loop assistant on ETA and order-status enquiries, and expand automation only after validated results.

Keywords— AI Service Desk, Human-in-the-Loop, Process Orchestration, SAP/HUB, Halo, Customer Service Automation, Industrial Equipment

I. INTRODUCTION

Industrial equipment manufacturers compete not only on product performance but on the speed and reliability of the service that surrounds the product. Compressors, vacuum systems and tools stay in operation for many years, and each year of operation generates spare-parts orders, delivery enquiries and technical requests that reach a service desk. The India compressor market alone has been estimated at about USD 2.30 billion in 2025 and is projected to reach USD 3.65 billion by 2032, a compound annual growth rate of roughly 6.8% (MarkNtel Advisors, 2026). As installed bases and order volumes grow, service desks face a rising stream of enquiries, many of which are routine questions about expected delivery dates, order status and material availability whose answers already sit in enterprise systems.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 12, UCBO - 2026

Atlas Copco India, the Pune-headquartered Indian subsidiary of the Atlas Copco Group, operates within this environment (The Machine Maker, 2025). Its customer service function logs enquiries in the Halo service desk platform and relies on SAP/HUB, an order and material information layer, together with SOP documents to resolve them. During an on-the-job training placement that began in November 2025, the first author observed that a CSO resolving a typical ticket reads and interprets the enquiry, searches for the order or material in SAP/HUB, consults the relevant SOP, decides the action, drafts a reply and updates Halo. Each system switch adds time, and the outcome depends on individual familiarity with navigation, SOP content and classification judgement. The same placement included exposure to warehouse operations and SAP MM procurement, which showed how closely service quality depends on the accuracy of inventory and master data.

To address these friction points, this study presents an applied investigation of the Halo and SAP/HUB ticket-resolution process and proposes an AI orchestration layer that works around the existing systems rather than replacing them. Using process mapping, analysis of aggregate Halo dashboard data and a lean-waste lens, the paper documents the current workflow, quantifies recent volume and lead-time trends, designs a rules-constrained and human-supervised agent, and sets out a KPI framework, a risk matrix and a seven-phase roadmap. The contribution is a practical design and measurement framework for adding AI to a legacy ticketing and ERP stack in an industrial after-sales setting, with the limits of the evidence stated openly.

II. LITERATURE REVIEW AND THEORETICAL BACKGROUND

A. Literature Review

Huang and Rust (2018) explain that different forms of artificial intelligence support diverse types of service activities. Their framework suggests that repetitive tasks involving clearly defined procedures are more suitable for early automation, whereas activities requiring contextual judgement and emotional understanding continue to depend more heavily on human involvement. and empathy remain with people for longer. Davenport and Ronanki (2018) identify several practical uses of artificial intelligence within organizations, including the automation of operational activities, the analysis of information and technology-supported interactions. Their findings also emphasize the value of introducing AI through manageable projects that strengthen existing employee capabilities rather than immediately replacing established workflows. Wirtz et al. (2018) examined service robots in the frontline and stressed that technology acceptance and the service encounter design shape whether automation improves service.

Evidence on customer-facing AI points to the importance of task complexity. Xu et al. (2020) found that usage intention for AI customer service depended on the complexity of the task and on the problem-solving ability users attributed to the system, which supports routing simple enquiries to automation and complex ones to people. Adam et al. (2021) studied how the design of AI-based chatbots affects user compliance in customer service. Evidence from Brynjolfsson et al. (2023) indicates that generative AI assistance can improve productivity in customer support operations. The study reported an average productivity improvement of approximately 14%, with substantially greater gains among less experienced employees. These findings suggest that AI tools may help reduce differences in performance associated with employee experience. which suggests that assistance can also narrow the gap between experienced and inexperienced agents. These studies concern chat-based consumer support rather than industrial after-sales desks that depend on ERP data.

Technical advances explain why such assistance is now feasible. Pre-trained language models such as BERT (Devlin et al., 2019) and large generative models that can be prompted with few examples (Brown et al., 2020) make ticket classification and information extraction possible without large labelled datasets, and retrieval-augmented generation grounds answers in a defined body of documents such as SOPs (Lewis et al., 2020). Amershi et al. (2019) emphasize the importance of designing AI systems that communicate their capabilities and limitations to users. In the proposed service desk environment, these principles support the inclusion of confidence indicators, human review mechanisms and the ability of Customer Service Officers to modify or reject AI-generated recommendations. (Amershi et al., 2019). Dellermann et al. (2019) frame this as hybrid intelligence, in which human and machine capabilities are combined rather than substituted. Broader reviews highlight data quality, governance and organisational readiness as persistent barriers to AI adoption (Dwivedi et al., 2021).

Davenport (1998) highlights that enterprise information systems influence organizational processes in addition to connecting information across departments. For the present study, this suggests that AI integration should respect established SAP workflows and operational controls rather than introduce changes that could disrupt the existing system of record. For a company whose service staff already depend on SAP data, an AI layer that reads from, and does not rewrite, the system of record is therefore a lower-risk starting point. In the sources reviewed, few studies address how to add AI orchestration to a legacy ticketing and ERP stack in industrial after-sales service while preserving business rules, auditability and the CSO role; this is the gap the present study addresses.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 12, UCBO - 2026

B. *Theoretical Background*

This study is grounded in service desk management, in which tickets move through intake, classification, prioritisation, routing, resolution and closure, and performance is judged by Service Level Agreements (SLAs), response and resolution time, first-contact resolution, reopen rate and customer satisfaction (Cannon, 2011). Escalation procedures transfer tickets that are complex, ambiguous or close to breaching an SLA to more experienced staff. These categories of metric underpin the KPI framework proposed in this paper.

The study applies lean management principles to examine activities that contribute little to the customer's service experience. Drawing on Womack and Jones (2003), tasks involving unnecessary waiting, repeated information retrieval and avoidable processing are examined as potential sources of operational inefficiency. (Womack & Jones, 2003). Although the waste taxonomy was developed for manufacturing, categories such as waiting, motion, over-processing, defects and unused talent translate directly to information work and are used here to classify the non-value-adding activities in the ticket workflow. The analysis also draws on supply chain visibility, the ability to see inventory, orders and shipments in near real time, which is a precondition for responsive customer service (Chopra & Meindl, 2016).

The proposed solution rests on three further ideas. First, an AI agent perceives input, reasons against rules and knowledge, uses tools such as read-only look-ups and acts, and in enterprises its autonomy is deliberately constrained (Russell & Norvig, 2020). Second, human-in-the-loop design places review points where confidence is low, risk is high or a case falls outside defined rules (Amershi et al., 2019). The governance approach adopted in this study emphasizes traceable decisions, restricted access to enterprise information and clearly defined responsibility for AI-supported actions. These safeguards reflect broader principles of responsible AI implementation discussed by the World Economic Forum (2022). (World Economic Forum, 2022). Finally, because technology adoption is more successful when introduced in stages with staff involvement, the roadmap follows the change-management logic of Kotter (1996). In SAP terms, the supporting context is the Materials Management module and its master data (material master, vendor master, purchasing information record and source list), which determine the quality of the data on which both CSOs and any agent rely (SAP SE, n.d.).

III. MATERIALS AND METHODS

A. *Research Objectives*

This study examines the customer service ticket-resolution process at Atlas Copco India to achieve four objectives: (1) to document the current Halo and SAP/HUB ticket-resolution workflow and identify its manual, repetitive and non-value-adding activities; (2) to analyse ticket volume, closure and lead-time trends from the Halo dashboard; (3) to design a human-in-the-loop AI service desk agent with decision rules, a risk matrix and a KPI framework for future validation; and (4) to recommend a phased, governed implementation roadmap.

B. *Research Design and Data Collection*

An applied, exploratory, primarily qualitative case study of a single organisational process was carried out. The case is the ticket-handling process of the customer service function at Atlas Copco India, which is incorporated as a private limited company, is headquartered in Pune and was delisted from the Indian stock exchanges in 2011 (Atlas Copco Group, 2011). The study period began with the placement start on 3 November 2025, and the dashboard data used covers October 2025 to August 2026.

Four sources of data were used. (1) Direct observation of the service desk workflow and structured walk-throughs with the industry mentor and the customer service team, generalised for confidentiality. (2) Aggregate monthly figures read from the labels of the Halo Evolution Tickets dashboard: tickets received (October 2025 to June 2026), tickets closed (October 2025 to August 2026) and average lead time to close (October 2025 to July 2026). (3) Academic and professional literature on service desk management, lean thinking, SAP MM and AI. (4) Public information about the company and sector. Supporting exposure to warehouse operations, data preparation for a supply planning dashboard and SAP MM transactions (ME01 for source lists, ME11 for purchasing information records and ME31K for contracts) informed the understanding of the data that service staff rely on. No ticket-level content, customer records or proprietary system specifications were collected or are reproduced.

C. *Analytical Framework and Quality Evaluation*

The current state was mapped as an AS-IS process from observation. Ticket data was analysed through monthly counts, the closed-to-received ratio, the cumulative gap between receipts and closures and the trend in average lead time. Non-value-adding activities were classified using the lean-waste categories of waiting, motion, over-processing, defects and unused talent, and each classification is labelled by its basis: observed in the workflow, indicated by dashboard data, or hypothesised and awaiting validation.

The TO-BE design followed three principles: non-invasive integration, meaning the agent works as an orchestration layer around Halo and SAP/HUB; rules-constrained decision-making, meaning actions follow explicit business rules and SOPs rather than unconstrained model output; and human-in-the-loop oversight for exceptions, low-confidence cases and high-risk actions. A risk matrix assessed impact and probability qualitatively, and a KPI framework defined future measurement with formulas, frequency and ownership. Because no

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 12, UCBO - 2026

performance data for an AI agent exists, every proposed target is stated as pilot-dependent and no benefit is asserted. Validation is proposed through a phased, CSO-reviewed pilot (Amershi et al., 2019). Dashboard values were read from chart labels, so the unit of lead time and the treatment of incomplete months were treated as assumptions to be confirmed with the process owner.

IV. DATA ANALYSIS AND DISCUSSION

Document the Existing Ticket-Resolution Process (AS-IS). The observed workflow begins when a customer submits an enquiry and a ticket is created in Halo. The CSO reads the free-text enquiry, identifies the requirement, searches for the relevant order or material, opens SAP/HUB to retrieve the information, checks the applicable SOP, decides the action, drafts a reply and updates the ticket before closing or escalating it. The workflow, with the system involved and the friction observed at each stage, is summarised in Table I.

TABLE I.
AS-IS SERVICE DESK WORKFLOW AND POINTS OF FRICTION

Stage	System / source	Observed friction
1. Ticket created	Halo	Free-text enquiry; no automatic classification
2. Read and interpret	Halo	Interpretation depends on CSO experience
3. Search order or material	SAP/HUB	Manual navigation and switch between systems
4. Retrieve information	SAP/HUB	Repeated look-ups for routine ETA and status enquiries
5. Check SOP	SOP documents	Manual reference; risk of inconsistent application
6. Decide action and route	CSO judgement	Ad hoc routing; rules not explicit or documented
7. Draft response	Halo / e-mail	Typed manually, even for recurring enquiries
8. Update and close or escalate	Halo	Manual update; limited audit trail

Table I shows that three of the eight stages (3 to 5) involve searching another system or document, and that two further stages, decision and routing and response drafting, rely on individual judgement or typing for enquiries that are often routine. The workflow also holds knowledge that is largely tacit: the rules a CSO applies when deciding whether a case is straightforward are learned through experience and are not written as testable conditions. This matters for automation, because an agent can only apply rules that have been made explicit.

Analyse Ticket Volume, Closure and Lead-Time Trends. The Halo dashboard provides monthly counts and average lead time, reproduced in Table II and plotted in Fig. 1. The first month, October 2025, may be a partial month, and the July and August closure counts are low and probably reflect incomplete periods, so they are shown but not used to draw conclusions.

TABLE II.
ONTHLY TICKETS RECEIVED, CLOSED AND AVERAGE LEAD TIME TO CLOSE (HALO DASHBOARD)

Month	Received	Closed	Closed / received	Avg lead time
Oct 2025	272	188	0.69	3
Nov 2025	412	388	0.94	6
Dec 2025	435	376	0.86	7
Jan 2026	448	405 A	0.90	12
Feb 2026	495	461	0.93	15
Mar 2026	444	407	0.92	15
Apr 2026	509	408	0.80	22
May 2026	397	550	1.39	22
Jun 2026	512	384	0.75	19
Jul 2026	n/a	208	n/a	38
Aug 2026	n/a	74	n/a	n/a

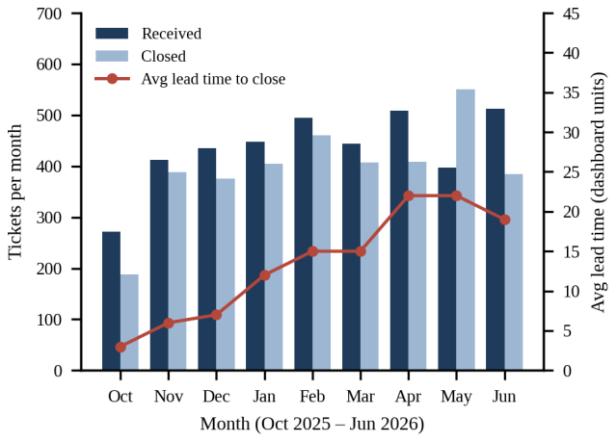


Fig. 1. Tickets received and closed per month and average lead time to close, October 2025 to June 2026 (plotted from Table II)

Between October 2025 and June 2026 the service desk received 3,924 tickets, an average of about 436 per month, and closed 3,567, so closures trailed receipts by about 357 tickets. The closed-to-received ratio was below 1.0 in eight of the nine months, and only May 2026 exceeded it, which suggests a partial clearing of earlier backlog. Because closures in a month can include tickets received earlier, the gap is an indicator of a growing open-ticket backlog rather than an exact count.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 12, UCBO - 2026

Average lead time to close rose from 3 in October 2025 to 22 in April and May 2026 and to 38 in July 2026, a more than twelve-fold increase over the period. Volume alone does not explain this. From November to April monthly intake grew by about 24%, from 412 to 509 tickets, while lead time rose from 6 to 22, more than threefold. A plausible reading is that the manual process has limited slack: once intake approaches the capacity of the team, waiting time accumulates faster than volume. The dashboard does not show ticket complexity, staffing or seasonality, so this remains an interpretation to be tested rather than a causal finding. The unit of lead time should also be confirmed with the process owner, although it is most likely days.

Identify Non-Value-Adding Activities and Classify Waste. Applying the lean-waste lens to the AS-IS workflow gives the classification in Table III. The basis column distinguishes what was observed directly, what the dashboard indicates and what is hypothesised.

TABLE III.
EAN WASTE CLASSIFICATION OF THE TICKET-RESOLUTION WORKFLOW

Waste	Manifestation at the service desk	Basis
Waiting	Tickets queue before a CSO can act; lead time and backlog gap rise	Indicated by dashboard
Motion	Switching between Halo, SAP/HUB and SOP documents for each ticket	Observed
Over-processing	Repeated manual look-ups and drafting for recurring ETA and status enquiries	Observed
Defects	Misclassification, wrong routing and inconsistent SOP application between CSOs	Hypothesised; to be validated
Unused talent	Experienced CSOs spend time on retrieval instead of exception handling	Hypothesised; to be validated

Table III indicates that the largest share of non-value-adding effort is motion and over-processing, both of which are rules-based information retrieval and are therefore the most suitable for automation. Waiting is the consequence visible in the data, while defects and unused talent cannot be confirmed without ticket-level analysis. This ordering shapes the design: the agent is aimed first at retrieval and drafting for routine enquiries and leaves judgement-heavy cases to CSOs, consistent with the task-complexity evidence reviewed earlier (Xu et al., 2020).

Relate the Findings to the Warehouse and SAP MM Context. The placement also covered warehouse operations and SAP MM, and these areas bear on service desk quality because the answers a CSO or an agent gives depend on the accuracy of the records behind them. Table IV summarises the link.

TABLE IV.
UPPORTING PLACEMENT EXPOSURE AND ITS RELEVANCE TO THE SERVICE DESK

Area	Activity during placement	Relevance to service desk
Warehouse operations	Receiving, storage, material handling, dispatch, layout	Stock and dispatch records determine availability and ETA answers
Data preparation and dashboard	Verification, cleaning and consolidation of inventory data; supply planning dashboard	Clean data and KPI dashboards are the basis for baselines and monitoring
SAP MM master data	Material master, Business Partner, source list (ME01), info record (ME11)	Quality of material and supplier data retrieved through SAP/HUB
Procurement cycle	MRP, requisition, RFQ, purchase order, contract (ME31K)	Delivery dates and planner details depend on purchasing and planning records

Design the Human-in-the-Loop AI Service Desk Agent (TO-BE). The proposed agent is an orchestration layer around Halo and SAP/HUB, shown in Fig. 2. An AI/NLP layer classifies the intent of an incoming ticket and extracts entities such as the order or material number. A read-only integration retrieves order status, delivery date, material information, MD04-related information and planner identity from SAP/HUB, and a knowledge base returns the applicable SOP. A decision engine applies explicit business rules and confidence thresholds to decide whether to draft a response, route the ticket, create a child ticket or escalate, and a human-in-the-loop layer sends exceptions and low-confidence cases to a CSO. Every decision is written to an audit log, and an SLA monitor tracks at-risk tickets.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 12, UCBO - 2026

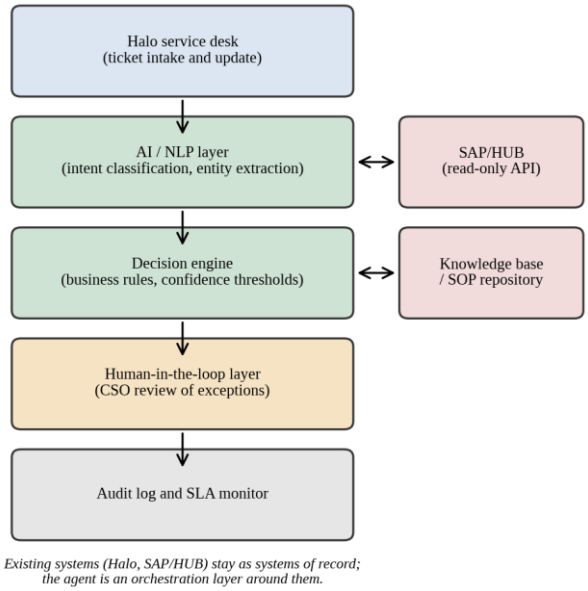


Fig. 2. Proposed architecture of the enterprise AI service desk agent around Halo and SAP/HUB

Table V compares the AS-IS and TO-BE workflows step by step. The expected benefits are design expectations drawn from the literature and the process map and are not measured results.

TABLE V.
S-IS VERSUS TO-BE WORKFLOW COMPARISON

Step	AS-IS (current)	TO-BE (proposed)
Interpretation	CSO reads and classifies	NLP intent classification
Data look-up	Manual SAP/HUB navigation	Read-only API look-up
SOP reference	Manual document search	Knowledge-base retrieval
Routing	Judgement-based	Rule-based, with reasons recorded
Response	Typed by CSO	AI-drafted; CSO review where required
Escalation	Ad hoc, CSO-initiated	Rules-triggered with clear criteria
Ticket update	Manual in Halo	Automatic, with audit-log entry
SLA monitoring	Periodic manual review	Continuous automated monitoring

Table VI gives illustrative decision rules. They are examples of the form the rules would take and are not drawn from actual tickets. In the ETA scenario, the agent classifies the intent, extracts the order number, retrieves the delivery date and, if the order carries no exception flag, drafts a response for CSO approval during the pilot. If the order shows a delay or the number cannot be extracted with confidence, the case goes to a CSO with the reason flagged. In every case where entity extraction fails, the agent does not attempt a look-up, which avoids a false match.

TABLE VI.
LLUSTRATIVE DECISION RULES FOR THE AGENT

Condition	Resulting action
ETA enquiry classified with confidence; order found; no exception flag	Draft response with retrieved delivery date (CSO approves in pilot)
Order or material number not confidently extracted	No look-up; route to CSO for clarification
Order shows delay or exception status in SAP/HUB	Route to CSO with supporting data
Intent outside defined SOP categories	Escalate to CSO or specialised team
Ticket marked urgent or close to SLA breach	Immediate CSO escalation with priority flag

A prototype data model was also defined as a class diagram, approved for external sharing. Its main entities are summarised in Table VII. The Ticket entity carries the extracted fields needed for look-ups and routing, the Employee entity records skills, availability and capacity so that tickets can be matched to staff, and the Assignment entity stores a score, a reason and a step trace for every decision, which provides the audit trail required for governance (World Economic Forum, 2022).

TABLE VII.
AIN ENTITIES OF THE PROTOTYPE DATA MODEL

Entity / component	Key fields or operations	Purpose
Ticket	Department, topic, customer message, HUB order and material number, delivery date, planner, recommended action, routed team, status, reason	Holds the enquiry and extracted fields
Employee	Skills, availability, active tickets, maximum capacity	Matches tickets to suitable staff
Assignment	Ticket, employee, score, reason, step trace	Audit record of each decision
Assignment and Resolution engines	Skill and workload matching; value normalisation, validation steps, decision	Apply rules to assign and resolve

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 12, UCBO - 2026

Assess Risks and Define the KPI Framework. Introducing an AI layer into a customer-facing process creates risks that must be controlled before any automation is allowed. Table VIII lists the principal risks with their assessed impact and probability and the mitigation designed into the architecture. The assessments are qualitative judgements by the authors and should be reviewed by the host organisation's IT and risk functions.

TABLE VIII.

I RISK MATRIX (QUALITATIVE ASSESSMENT)

Risk	Impact / probability	Mitigation
Hallucinated or incorrect response	High / medium	Constrain answers to retrieved SAP/HUB and SOP data; CSO review of auto-resolved cases in pilot
Incorrect order look-up	High / low to medium	Validate number format; confidence thresholds; no look-up on unresolved reference
Data privacy exposure	High / low	Role-based access; data minimisation; limited retention
Unauthorised system access	High / low	Read-only API scope; strong authentication
Wrong routing	Medium / low to medium	Explicit, testable rules; periodic review
System integration failure	Medium / medium	Fallback to manual CSO process; monitoring and alerts
Knowledge-base inconsistency	Medium / medium	SOP review and version control
Human over-reliance on AI	Medium / medium	Training, audits and retained review requirements

Because no AI performance data exists, the KPI framework in Table IX uses the dashboard to fix baselines where possible and otherwise states what must be measured. The framework follows the efficiency, quality and customer-experience categories of service desk measurement (Cannon, 2011), so that gains in speed are not obtained at the expense of resolution quality.

TABLE IX.

ROPOSED KPI FRAMEWORK WITH AVAILABLE BASELINES

KPI	Baseline	Measurement
Monthly tickets received	About 436 (Oct 2025 to Jun 2026)	Halo dashboard
Average lead time to close	3 (Oct 2025) rising to 38 (Jul 2026)	Halo timestamps
Closed-to-received ratio	Below 1.0 in 8 of 9 months	Halo dashboard
Average response and resolution time	To be measured A	Halo timestamps
Manual handling time per ticket	To be measured	CSO activity study
Auto-resolution and escalation rate	Not applicable before pilot	AI and escalation logs
Classification and extraction accuracy	Not applicable before pilot	CSO validation of AI output
SLA compliance, reopen rate, satisfaction	To be measured	Halo SLA reports, surveys

Recommend a Phased and Governed Implementation Roadmap. The roadmap in Table X moves from preparation to controlled automation in seven phases, each with an exit criterion that must be met before the next begins. The sequence reflects the finding that data quality and process standardisation are prerequisites for AI, and the change-management principle that adoption should be incremental and involve the staff affected (Kotter, 1996).

TABLE X.

HASED IMPLEMENTATION ROADMAP AND EXIT CRITERIA

Phase	Activity	Exit criterion
1	Standardise master data and ticket categories; record baseline KPIs	Agreed categories and cleaned data
2	Digitise SOPs into a searchable knowledge base	Each SOP versioned with an owner
3	Prototype classification and extraction on anonymised tickets	Accuracy measured and reviewed
4	Establish read-only SAP/HUB integration	IT and security approval
5	Human-in-the-loop pilot on ETA enquiries	CSO-approved results against KPIs
6	Expand to validated, low-risk categories P	Pilot targets met
7	Institutionalise KPI monitoring and rule review	Regular review cycle running

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 12, UCBO - 2026

Taken together, the analysis supports a cautious reading. The dashboard shows a service desk under strain, the process map shows that much of the effort is rules-based retrieval and typing, and the literature indicates that such tasks are the most suitable for assistance. The data does not show how much time an agent would save or how accurate it would be, and the paper therefore stops at a design, a measurement framework and a staged plan.

V. FINDINGS, CONCLUSIONS AND SUGGESTIONS

A. Findings

The investigation of the Halo and SAP/HUB ticket-resolution process at Atlas Copco India produced the following findings. First, ticket handling depends on manual interpretation, repeated system-switching and SOP look-up, and routing decisions that are not documented as explicit rules; three of eight workflow stages involve searching another system or document and two more depend on individual judgement or manual typing. Second, the Halo dashboard shows 3,924 tickets received and 3,567 closed between October 2025 and June 2026, a cumulative gap of about 357, with closures below receipts in eight or nine months. Third, average lead time to close rose from 3 to 38 between October 2025 and July 2026, and its rise outpaced the growth in intake, which indicates limited capacity slack in a manual process. Fourth, the lean-waste classification identified motion and over-processing as the largest observed non-value-adding activities, both of which are rules-based retrieval tasks. Fifth, supporting exposure to warehouse operations and SAP MM showed that accurate master data and inventory records are the foundation of reliable lookups, so that data quality is a prerequisite for any automation. Sixth, a human-in-the-loop agent that classifies tickets, performs read-only look-ups, applies explicit rules and escalates exceptions is feasible in design, and its principal risks, including incorrect responses and incorrect lookups, can be reduced by confidence thresholds, read-only access and audit logging.

B. Conclusions

This study shows that the ticket-resolution process at Atlas Copco India is manually intensive and that recent dashboard data are consistent with a process operating near the limit of its capacity. The workflow assessment indicates that ETA and order-status enquiries provide practical starting points for AI assistance because they frequently involve identifiable information requirements and structured SAP/HUB lookups. However, their suitability for automation must be established through ticket-level validation and controlled pilot testing. and that the agent should be introduced in a governed, phased way with CSO review for ambiguous, exception and high-risk cases. The proposed architecture, decision rules, risk matrix, KPI framework and roadmap are design proposals and not validated implementation. Their benefits, including any reduction in handling time and lead time, remain to be established through a pilot using real ticket data. The study also confirms, in an industrial after-sales setting, the wider observation in the literature that data quality and process standardization, rather than model sophistication, are the main conditions for success.

C. Suggestions

To obtain tangible improvement while controlling risk, the following suggestions are directed at the main stakeholders.

For Management and Process Owners: Fix the KPI baseline now, using Halo timestamps, so that a pilot can be judged on evidence. Review the rise in lead time against staffing, ticket complexity and seasonality before attributing it to process design, and approve a limited human-in-the-loop pilot restricted to ETA and order-status enquiries.

For Customer Service Officers and Team Leads: Write the tacit rules used to classify and route tickets as explicit, testable conditions, and take part in validating AI outputs during the pilot, since CSO feedback is the main source of accuracy measurement and of acceptance.

For SAP, IT and Security Teams: Assess the availability of a read-only API to SAP/HUB, define role-based access and data-retention rules, and clean and standardise master data and ticket categories across SAP, Halo and SAP/HUB before integration.

For Knowledge and SOP Owners: Digitise SOPs into a structured, versioned and searchable knowledge base, with a named owner for each document, so that both CSOs and any future agent apply procedures consistently.

D. Implications of the study

The implications of the study extend to several groups. For industrial after-sales organisations, the paper offers a template for adding AI to a legacy Halo-and-ERP stack without replacing core systems, with explicit rules, confidence thresholds and audit trails as the controls that make adoption defensible.

Management and Customer Service Leaders: A measurement framework with baselines drawn from existing dashboards allows investment in automation to be judged on evidence rather than expectation, and the ordering of automation from routine to complex cases limits the risk of service failure during transition.

CSOs and Service Staff: Removing repetitive look-ups and typing for routine enquiries can free experienced staff for exceptions and complex cases, while human review preserves their accountability for customer-facing decisions.

IT, Security and Governance Functions: The read-only, audit-logged design gives a concrete basis for assessing privacy, access-control and explainability requirements before any system is built.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14, Issue 12, UCBO - 2026

Researchers and Educators: The study illustrates how lean-waste categories can be applied to information work and how a qualitative, dashboard-based case study can ground an AI design where ticket-level data is not available.

E. Future Scope of the study

The design and baselines established in this study provide several avenues for future inquiry:

Pilot Validation: A controlled human-in-the-loop pilot can measure handling time, auto-resolution rate, classification and extraction accuracy, escalation rate and customer satisfaction against the baselines reported here.

Ticket-Level Analysis: Analysis of ticket content, complexity and staffing, under appropriate confidentiality arrangements, can separate routine from complex demand and test the capacity explanation proposed for the rise in lead time.

Retrieval-Augmented Knowledge Support: The knowledge base can be extended to procurement and warehouse queries, and retrieval-augmented responses can be tested for accuracy and traceability.

Integrated Analytics: Linking the Halo dashboard with inventory, delivery and supplier-performance data can identify the root causes of delay and support predictive demand and inventory planning.

Economic Evaluation: A cost-benefit analysis using actual cost and productivity data can establish whether the investment in development, integration, training and maintenance is justified.

ACKNOWLEDGMENT

The authors express their sincere gratitude to Atlas Copco India for the opportunity to undertake this study, to the industry mentor, Mr. Omkar Phansalkar, for guidance on warehouse, SAP and service desk processes, and to the customer service, SAP MM and warehouse teams for their patient support. The authors also thank the School of Business, Indira University for including the on-the-job training component in the MBA curriculum. All confidential organisational information has been withheld or generalised in line with the host organisation's expectations.

REFERENCES

[1] Adam, M., Wessel, M., & Benlian, A. (2021). AI-based chatbots in customer service and their effects on user compliance. *Electronic Markets*, *31*(2), 427–445. https://doi.org/10.1007/s12525-020-00414-7

[2] Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human-AI interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1–13. https://doi.org/10.1145/3290605.3300233

[3] Atlas Copco Group. (2011, March 18). *Atlas Copco to increase ownership in Indian subsidiary and delist shares* [Press release]. https://www.atlascopcogroup.com/en/media/press-releases/2011/110318-atlas-copco-to-delist-in-india

[4] Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., et al. (2020). Language models are few-shot learners. *Advances in Neural Information Processing Systems*, *33*, 1877–1901.

[5] Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). *Generative AI at work* (NBER Working Paper No. 31161). National Bureau of Economic Research. https://doi.org/10.3386/w31161

[6] Cannon, D. (2011). *ITIL service operation* (2011 ed.). The Stationery Office.

[7] Chopra, S., & Meindl, P. (2016). *Supply chain management: Strategy, planning, and operation* (6th ed.). Pearson.

[8] Davenport, T. H. (1998). Putting the enterprise into the enterprise system. *Harvard Business Review*, *76*(4), 121–131.

[9] Davenport, T. H., & Ronanki, R. (2018). Artificial intelligence for the real world. *Harvard Business Review*, *96*(1), 108–116.

[10] Dellermann, D., Ebel, P., Söllner, M., & Leimeister, J. M. (2019). Hybrid intelligence. *Business & Information Systems Engineering*, *61*(5), 637–643. https://doi.org/10.1007/s12599-019-00595-2

[11] Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. *Proceedings of NAACL-HLT 2019*, 4171–4186. https://doi.org/10.18653/v1/N19-1423

[12] Dwivedi, Y. K., Hughes, L., Ismagilova, E., Aarts, G., Coombs, C., Crick, T., et al. (2021). Artificial Intelligence (AI): Multidisciplinary perspectives on emerging challenges, opportunities, and agenda for research, practice and policy. *International Journal of Information Management*, *57*, 101994. https://doi.org/10.1016/j.ijinfomgt.2019.08.002

[13] Huang, M.-H., & Rust, R. T. (2018). Artificial intelligence in service. *Journal of Service Research*, *21*(2), 155–172. https://doi.org/10.1177/1094670517752459

[14] Kotter, J. P. (1996). *Leading change*. Harvard Business School Press.

[15] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. *Advances in Neural Information Processing Systems*, *33*, 9459–9474.

[16] MarkNtel Advisors. (2026). *India compressor market to reach USD 3.65 billion by 2032, growing at a CAGR of 6.8% during 2026–2032*. https://www.marknteladvisors.com/press-release/india-compressor-market-share

[17] Russell, S., & Norvig, P. (2020). *Artificial intelligence: A modern approach* (4th ed.). Pearson.

[18] SAP SE. (n.d.). *SAP Help Portal: Materials Management (MM)*. https://help.sap.com

[19] The Machine Maker. (2025). *Atlas Copco Group opens cutting-edge manufacturing facility in Pune*. https://themachinemaker.com/news/atlas-copco-group-opens-cutting-edge-manufacturing-facility-in-pune/

[20] Wirtz, J., Patterson, P. G., Kunz, W. H., Gruber, T., Lu, V. N., Paluch, S., & Martins, A. (2018). Brave new world: Service robots in the frontline. *Journal of Service Management*, *29*(5), 907–931. https://doi.org/10.1108/JOSM-04-2018-0119

[21] Womack, J. P., & Jones, D. T. (2003). *Lean thinking: Banish waste and create wealth in your corporation* (2nd ed.). Free Press.

[22] World Economic Forum. (2022). *AI governance: A comprehensive approach to implement ethics into AI*. https://www.weforum.org

[23] Xu, Y., Shieh, C.-H., van Esch, P., & Ling, I.-L. (2020). AI customer service: Task complexity, problem-solving ability, and usage intention. *Australasian Marketing Journal*, *28*(4), 189–199.

(This work is licensed under a Creative Commons Attribution 4.0 International License.)