American Journal of Scholarly Research and Innovation

Volume: *2;* ***Issue:*** *2*
Pages: *373 – 414*
eISSN: 3067-2163

REINFORCEMENT LEARNING APPROACHES TO OPTIMIZE IT SERVICE MANAGEMENT UNDER DATA SECURITY CONSTRAINTS

Md Mohaiminul Hasan[1]; Md Muzahidul Islam[2];

[1]. Master in Project Management; St. Francis College - Brooklyn, NY, USA;
Email: mohaiminul.hasan22@gmail.com

[2]. B.Sc in Computing Science and Technology, Jiangxi Normal University, Jiangxi, China
Email: muzahidul365@gmail.com

Doi: 10.63125/z7q4cy92
Received: *19 September 2023;* ***Revised:*** *22 October 2023;* ***Accepted:*** *24 November 2023;* ***Published:*** *18 December 2023*

Abstract

This quantitative study evaluated reinforcement learning (RL) approaches to optimize IT Service Management (ITSM) decision-making under explicit data security constraints using a retrospective observational design and historical operational logs. The analytic dataset was extracted from a 12-month observation window and contained 128,450 raw incident and service-request tickets. After deterministic preprocessing, including timestamp normalization, deduplication, linkage resolution, and censoring under pre-specified rules, 112,680 tickets were retained for analysis, representing 46.3% incidents and 53.7% service requests across 10 service lines and 5 regions. Mandatory ITSM fields were complete for 96.8% of retained records, and 78.4% of tickets were linked to usable telemetry context for state construction. Descriptive statistics indicated distributional asymmetry typical of service systems: the median time-to-acknowledge was 0.42 hours and the median time-to-resolve was 14.6 hours, with substantial tail behavior (P90 resolution time 63.4 hours; P95 resolution time 97.6 hours). Governance indicators were consistently recorded, with audit-log completeness averaging 98.9%. Correlational analysis showed that time-to-resolve was positively associated with congestion and workflow friction, including relationships with backlog intensity (0.46), ticket aging (0.52), reassignment count (0.41), and escalation occurrence (0.35), motivating multivariable modeling with collinearity screening. Adjusted regression models controlling for ticket severity, category, region, service line, workload context, and major-incident regime indicated that the RL policy was associated with improved performance relative to baseline handling, including reduced time-to-acknowledge ($\beta = -0.08$ hours; $p < .001$) and improved resolution performance (log-linear $\beta = -0.12$; $p < .001$), alongside improved percentile-oriented SLA attainment (P90 threshold OR = 1.28, $p < .001$; P95 threshold OR = 1.19, $p = .001$). Process-quality outcomes also improved, including reduced reopen occurrence (OR = 0.86; $p < .001$) and reduced reassignment occurrence (OR = 0.89; $p < .001$). Constraint-related models showed no statistically significant increases in privileged-action events (IRR = 0.99; $p = .61$), exception approvals (IRR = 0.98; $p = .44$), or restricted-field access (IRR = 1.01; $p = .69$), and audit completeness remained statistically unchanged, supporting admissibility of security constraints while service outcomes improved.

Keywords

Reinforcement Learning, IT Service Management, Data Security Constraints, Off-Policy Evaluation, Governance-Aware Optimization.

INTRODUCTION

Information Technology Service Management (ITSM) refers to the coordinated set of organizational capabilities, processes, and controls used to design, deliver, operate, and continually improve IT services so that service outcomes align with business needs, user experience expectations, and formal service commitments such as service-level agreements (SLAs) (Yazici et al., 2015). Within this scope, "service" is treated as a value-delivery arrangement that enables outcomes without transferring all ownership of cost and risk to the customer, while "management" denotes the governance of workflows, resources, and accountability structures that keep service operations stable under changing demand. In operational terms, ITSM encompasses high-frequency decision points such as incident triage, request fulfillment prioritization, change scheduling, configuration impact reasoning, and problem resolution coordination across technical teams and vendors. Reinforcement learning (RL) is a quantitative optimization framework in which an agent learns a decision policy through interaction with an environment, receiving feedback signals that encode performance objectives (Sukmandhani et al., 2017). In ITSM settings, the "environment" can be the service operation itself—ticket arrivals, queue backlogs, staffing capacity, dependency graphs, and infrastructure states—while "actions" can include routing, escalation, resource allocation, throttling, or automated remediation choices, and "rewards" can encode SLA adherence, resolution time, customer satisfaction proxies, operational cost, and risk penalties.

Figure 1: ITSM Optimization with Reinforcement Learning

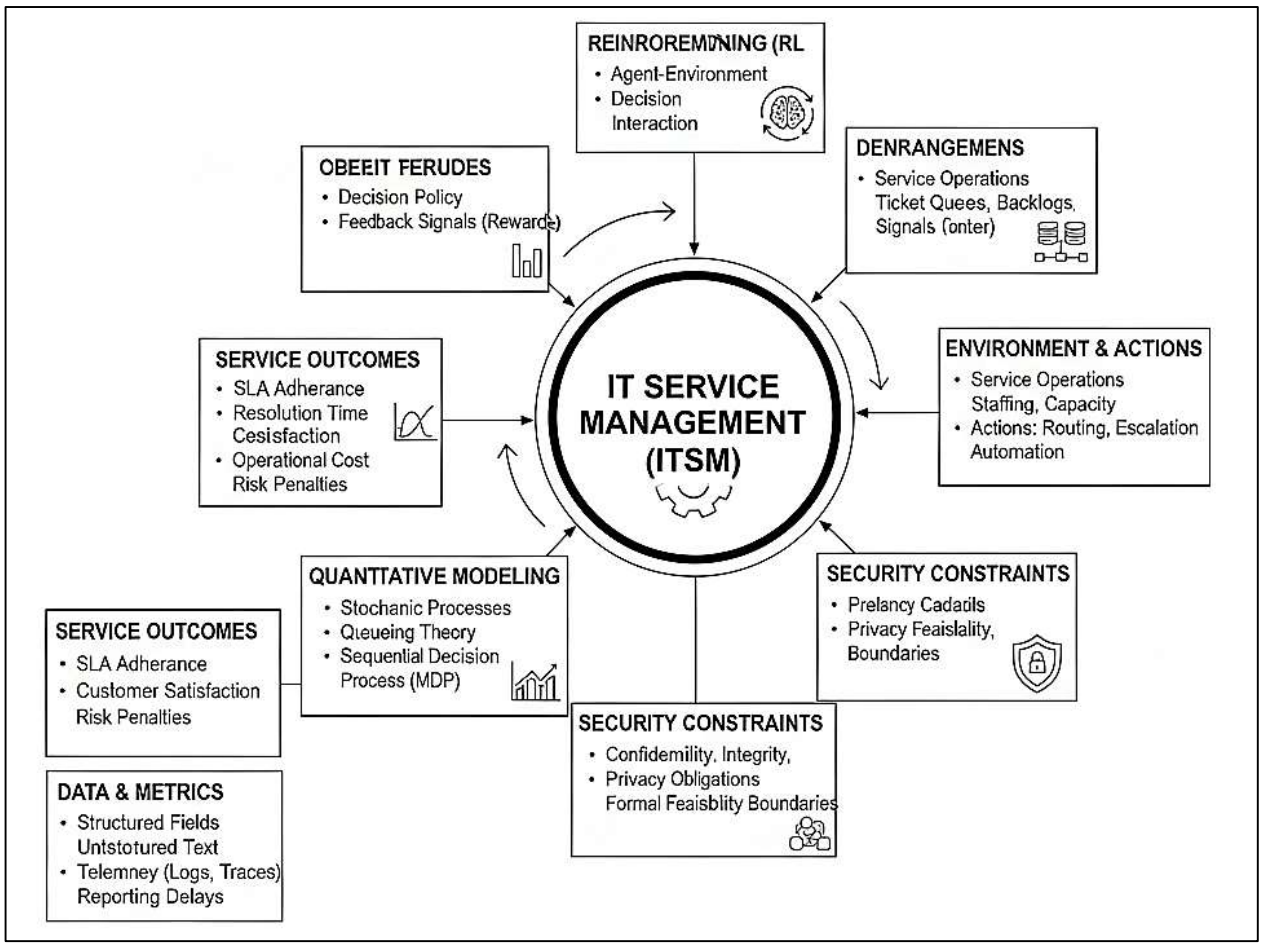

Data security constraints refer to enforceable restrictions that govern how service data are accessed, processed, retained, and transmitted, including requirements for confidentiality, integrity, and availability, as well as privacy obligations that constrain analytics and automation. Such constraints are not incidental conditions; they function as formal feasibility boundaries that shape which states are observable, which actions are permissible, and which optimization targets remain valid under compliance regimes (Aquino et al., 2018). The international significance of this problem arises from the

centrality of IT-enabled services in global finance, cross-border supply chains, healthcare delivery, public administration, and critical infrastructure operations, where service disruptions propagate across jurisdictions and where regulated data routinely traverse multinational clouds and vendor ecosystems. Consequently, quantitative approaches to optimizing ITSM under explicit security constraints occupy a domain in which operational performance metrics and security governance metrics coexist as co-equal control variables, making the optimization problem both economically consequential and institutionally regulated at scale (Jamous et al., 2016).

Operationally, ITSM environments exhibit uncertainty, partial observability, delayed effects, and nonstationarity—properties that naturally invite quantitative modeling through stochastic processes, queueing systems, and sequential decision theory (Arfan et al., 2021; Oktadini & Surendro, 2014). Incidents and service requests arrive as stochastic streams that vary by time zone, business cycle, product release cadence, and threat landscape, forming workload patterns that can be captured through arrival-rate estimation, service-time distributions, priority classes, and reneging or escalation behaviors (Jahid, 2021). Classic service operations research formalizes these dynamics using queueing theory and call-center modeling, where performance outcomes such as waiting time, abandonment, and service level attainment respond nonlinearly to staffing and routing choices, and where system stability depends on utilization levels and variability in service times (Akbar & Farzana, 2021; Trusson et al., 2014). These analytical constructs map directly onto ITSM workflows: triage resembles classification with routing; escalation resembles priority migration; major incidents create bursty arrivals and correlated resolution times; change windows impose scheduling constraints; and configuration dependencies introduce networked externalities (Reza et al., 2021; Zobayer, 2021a). From a statistical perspective, ITSM data are heterogeneous, combining structured fields (timestamps, priority, category, CI identifiers), unstructured text (ticket narratives), and telemetry (logs, metrics, traces), producing measurement regimes with missingness, noise, and reporting delays (Ariful & Ara, 2022; Zobayer, 2021b). Quantitative optimization in such systems benefits from formal models that can accommodate uncertainty and allow policy evaluation under realistic load. Moreover, the regulatory character of enterprise service operations introduces an additional modeling layer: security and privacy requirements alter data availability, restrict feature engineering, and sometimes require auditable reasoning trails for automated decisions (Chunpir & Ismailzadeh, 2019; Arman & Kamrul, 2022; Mesbaul & Farabe, 2022). These properties align closely with Markov decision process (MDP) abstractions and their extensions, where sequential actions influence future states and where objective functions include time-discounted or average-cost criteria (Abdur & Haider, 2022; Mushfequr & Sai Praveen, 2022). The service-operations literature offers principled baselines for modeling congestion and performance tradeoffs, while MDP theory provides the mathematical footing for optimal control under uncertainty.

Reinforcement learning supplies a family of estimators and control algorithms for learning policies that map observed states to actions so as to optimize an expected return (Mortuza & Rauf, 2022; Rakibul & Samia, 2022). Core RL distinctions include model-based versus model-free methods, value-based versus policy-based methods, and on-policy versus off-policy learning, each offering different bias–variance and sample-efficiency characteristics that matter in operational environments with limited ability to explore (Abdul, 2023; Abdulla & Zaman, 2023; Orta et al., 2014). Value-based methods estimate action-value functions to guide greedy or stochastic decision rules, while policy-gradient methods directly optimize parameterized policies using gradient estimators of expected return; actor–critic methods combine both (Arfan et al., 2023). Deep reinforcement learning extends these approaches by using neural networks as function approximators for value functions and policies, enabling high-dimensional state representations such as ticket text embeddings, topology graphs, or aggregated telemetry windows. Many ITSM decisions resemble continuous control (e.g., tuning throttling rates, autoscaling setpoints) and thus align with deterministic or stochastic policy optimization methods designed for continuous action spaces (Almeida et al., 2018; Foysal & Aditya, 2023; Hamidur, 2023). Other ITSM decisions resemble discrete routing and prioritization choices, aligning with Q-learning variants and policy-gradient methods for categorical actions. The quantitative challenge is to connect algorithmic structures to service objectives: minimizing mean time to resolution (MTTR), controlling backlog growth, maintaining SLA percentiles, and balancing automation benefits against operational

risk (Rashid et al., 2023; Musfiqur & Kamrul, 2023). RL also formalizes delayed consequences in service operations, where an early routing action affects downstream queues, specialist availability, and customer-visible outcomes hours later (Muzahidul & Mohaiminul, 2023; Amin & Praveen, 2023). In addition, RL evaluation emphasizes the distinction between online performance and off-policy evaluation (OPE) from historical logs, a critical issue in ITSM where "safe" deployment requires estimating policy improvements without reckless experimentation on live operations (Jäntti & Hotti, 2016; Hasan & Ashraful, 2023; Ibne & Md. Kamrul, 2023). Off-policy evaluation and counterfactual reasoning are thus central to quantitative claims about performance under constrained exploration (Mushfequr & Ashraful, 2023; Roy & Kamrul, 2023). Within this foundation, the integration of RL with ITSM requires careful state construction (what information is observable under security rules), reward specification (how to encode performance and risk), and evaluation methodology (how to measure improvement under realistic workload and threat conditions (Shaikh & Farabe, 2023; Haider & Hozyfa, 2023).

Figure 2: Security Constraints in ITSM Optimization

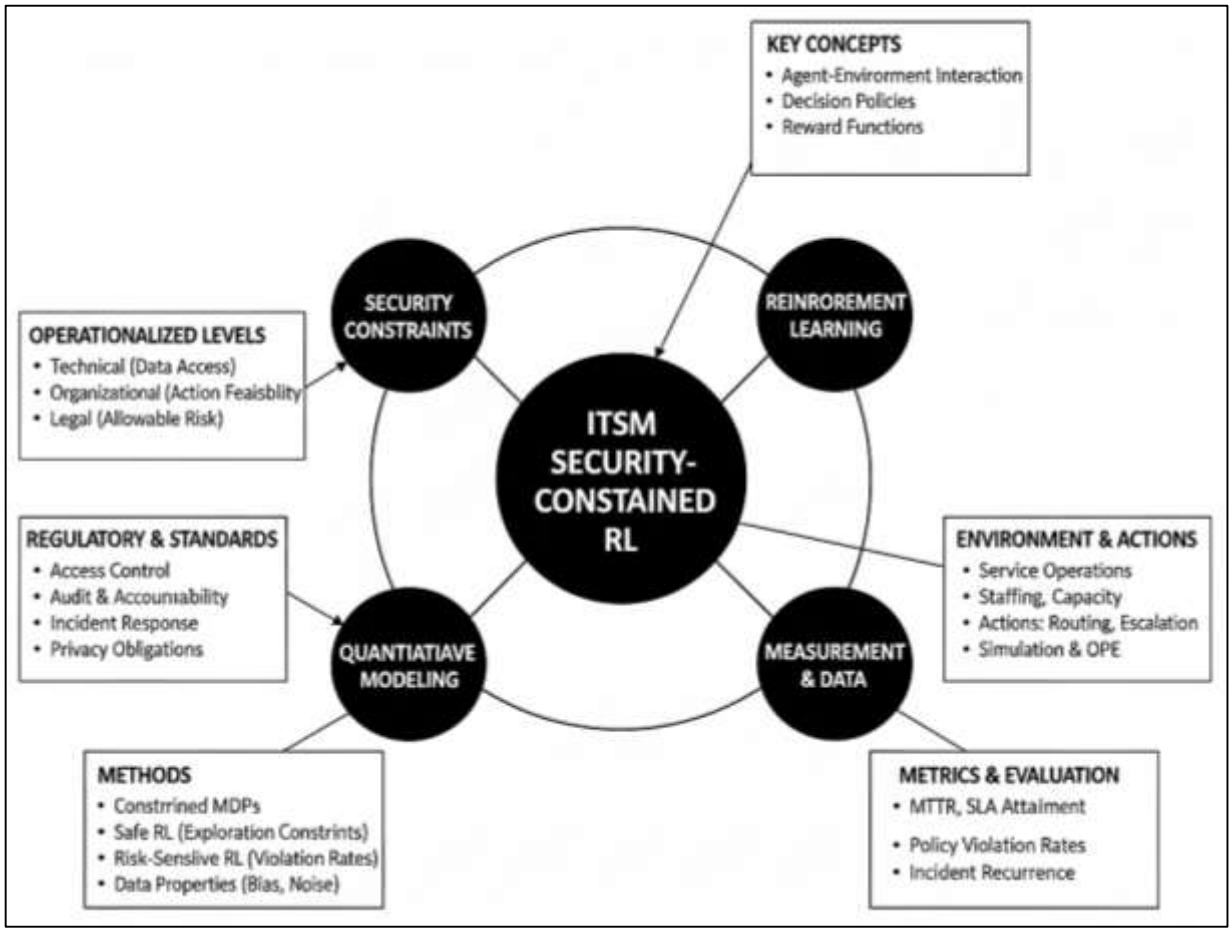

Security constraints in IT service environments operate simultaneously at the technical, organizational, and legal levels, making them amenable to formalization as constraints on observability, action feasibility, and allowable risk (Zobayer, 2023). Confidentiality requirements limit which ticket attributes, identity signals, or telemetry fields can be ingested into learning pipelines; integrity requirements constrain the permissible automation actions that could alter configurations; and availability requirements restrict experimentation that might degrade uptime (Rouhani, 2017). Regulatory and standards-based regimes specify concrete control families—access control, audit and accountability, incident response, configuration management, and risk assessment—that intersect with ITSM processes in measurable ways. Privacy obligations, including data minimization and purpose limitation, constrain the feature sets and retention windows available for training and evaluation, and they impose governance requirements for transparency and accountability in automated decision-

making (Kubiak & Rass, 2018). From a quantitative viewpoint, these constraints can be operationalized as hard constraints (actions prohibited under policy), soft constraints (actions permitted but penalized via risk costs), or probabilistic constraints (risk measures bounded by acceptable thresholds). Security constraints also shape the data-generating process: logs may be redacted; identifiers may be pseudonymized; sensitive classes may be underrepresented due to restricted access; and incident narratives may carry protected information requiring controlled handling. These factors interact with statistical learning by inducing selection bias, covariate shift, and measurement error that must be accounted for when estimating policy value. In addition, ITSM automation frequently integrates with privileged systems (identity providers, endpoint management, cloud control planes), so an RL agent's action space must align with least-privilege principles and must be auditable to satisfy internal governance (Jäntti et al., 2014). Quantitatively, the security posture can be modeled through metrics such as policy violation rates, access attempts outside role scope, data exposure likelihood, change failure rates, and security incident recurrence linked to automated actions. Security is therefore not only a context variable but a constraint set that must be included explicitly in optimization formulations if learned policies are to remain valid under compliance and risk governance.

A rigorous path for incorporating security constraints into RL is to formulate the ITSM control problem as a constrained Markov decision process (CMDP), where the objective is to maximize expected return subject to constraints on expected cumulative costs representing policy violations, privacy risk, or operational hazards (Krishnan & Ravindran, 2017). CMDPs provide a mathematically explicit language for "optimize performance while bounding security risk," and they enable the use of Lagrangian relaxation, primal–dual optimization, and occupancy-measure methods to compute or approximate constrained-optimal policies. In ITSM, constraint signals can encode, for example, the frequency of actions requiring elevated privilege, the probability of exposing sensitive data fields, or the rate at which automated remediations trigger change failures that elevate security exposure. Safe RL extends this framing by focusing on learning procedures that maintain acceptable behavior during data collection, where exploration itself is constrained. This is crucial for service environments because exploratory actions can degrade customer experience or violate governance rules. Risk-sensitive RL further extends the objective beyond expected value by capturing variability and tail outcomes through coherent risk measures, enabling control policies that account for rare but severe events such as widespread outages or security incidents (Berrahal & Marghoubi, 2016). Practical RL methods for constrained settings include constrained policy optimization, trust-region methods with safety critics, and shielded policies that filter actions through rule-based security guards. Quantitatively, such mechanisms correspond to algorithms that either (a) enforce constraints in the optimization step, (b) project policies back into feasible sets, or (c) embed constraints into the environment as infeasible transitions. The methodological fit to ITSM is direct: governance rules already exist as policy engines; RBAC systems already encode permissible actions; approval workflows already impose constraints; and monitoring systems already generate cost signals (Diao et al., 2016). The RL formulation allows these governance artifacts to be represented as measurable constraints and optimized jointly with service objectives. The result is a statistical control problem with explicit objective metrics, explicit constraint metrics, and evaluable tradeoffs that can be tested through simulation, counterfactual evaluation, and controlled deployment under governance oversight (Zuev et al., 2018).

Quantitative ITSM optimization relies on measurement systems that translate operational and security objectives into analyzable variables. Common service-performance measures include MTTR, mean time to acknowledge, reopen rates, first-contact resolution, SLA attainment at percentile thresholds, backlog size dynamics, and cost-to-serve; these can be modeled as episode returns, average-cost criteria, or multi-objective trade spaces (Pilorget & Schell, 2018). Security and privacy measures can be encoded as constraint costs or auxiliary outcomes: unauthorized access attempts, policy exceptions, data handling violations, excessive privilege use, audit findings, and incident recurrence patterns linked to remedial actions. In RL terms, these measures require careful reward and cost shaping so that the optimization aligns with organizational definitions rather than proxy metrics that invite unintended behavior (Macias & Alonso, 2018). The structure of ITSM data supports multiple quantitative representations: (1) event logs enabling process mining and time-aware feature extraction; (2) queue snapshots capturing congestion; (3) graphs capturing service dependencies and configuration

item relationships; and (4) text and telemetry embeddings capturing contextual signals. Simulation plays a central role in evaluation because it supports controlled stress testing under varied arrival patterns, staffing scenarios, and threat conditions without risking production services; queueing-based simulators and discrete-event models provide interpretable baselines, while digital-twin approaches incorporate richer system dynamics. Where simulation is limited, off-policy evaluation provides a route to estimating how alternative policies perform using historical logs, and it introduces statistical concerns such as support mismatch and variance control. Methods like importance sampling variants, doubly robust estimators, and model-based OPE are relevant for establishing credible quantitative comparisons (Suryawan, 2018). Evaluation also demands stratification: policies that improve average MTTR can still degrade tail performance or increase security exceptions in specific categories such as privileged access incidents or regulated customer data requests. Consequently, distributional reporting—percentiles, conditional means, and subgroup analyses—aligns with both SLA reporting and compliance auditing. The quantitative introduction to this domain therefore integrates operations metrics, security metrics, statistical estimation, and sequential decision evaluation into a single measurement logic, enabling formal hypothesis testing and effect-size reporting once policies and baselines are specified (El Yamami et al., 2018).

When ITSM is treated as a sequential decision system, the optimization target becomes the controlled evolution of service states under uncertainty, constrained by security governance and privacy obligations that define admissible data use and admissible operational actions (Georg, 2017). Reinforcement learning provides a formal statistical mechanism for learning policies from interaction or logs, while constrained and safe RL provide the language for encoding and enforcing security requirements within the optimization problem rather than treating them as external checks applied after the fact. This joint framing emphasizes that service quality and security compliance are both measurable outcomes of decisions taken across ticket lifecycles, automation pipelines, and change-management sequences (Anthonysamy et al., 2017). In practice, the enterprise service environment contains structural regularities—queues, priority classes, specialist pools, dependency graphs, approval gates, and role-based permissions—that align naturally with the modeling elements of states, actions, transitions, and constraints. At the same time, the operational stakes of service delivery and the governance stakes of regulated data handling mean that quantitative modeling choices must remain consistent with what is observable and permissible under policy. The resulting research space is defined by (a) how states are constructed from secure data sources, (b) how action spaces are bounded by least-privilege and change-control rules, (c) how objective functions encode service outcomes without substituting misleading proxies, and (d) how constraints are quantified in ways that are auditable and statistically testable (Kalloniatis et al., 2014). Within this space, the introduction clarifies the definitional foundation and measurement foundation needed for quantitative analysis: ITSM outcomes correspond to system performance under stochastic load; RL corresponds to policy optimization under uncertainty; and security constraints correspond to formal restrictions on data, actions, and acceptable risk. This alignment sets a disciplined conceptual base for empirical modeling, estimation, and evaluation that remain grounded in measurable variables rather than informal narratives (Romanou, 2018).

The objective of this quantitative study is to develop and empirically evaluate a reinforcement learning–based optimization framework for IT Service Management (ITSM) decision-making that improves measurable service performance outcomes while maintaining explicit data security constraints as binding quantitative requirements. The study aims to formalize ITSM operations as a sequential decision problem in which incident and service request workflows, resource allocation choices, escalation rules, and remediation actions are modeled as state–action transitions with observable operational consequences, enabling rigorous estimation of policy effects on key service indicators. A primary objective is to quantify the impact of reinforcement learning policies on standardized ITSM performance metrics such as mean time to acknowledge, mean time to resolution, backlog dynamics, first-contact resolution rate, SLA attainment at percentile thresholds, and cost-to-serve proxies derived from staffing utilization and rework rates, while ensuring that optimization does not violate confidentiality, integrity, availability, and privacy obligations embedded in enterprise governance. A second objective is to operationalize data security constraints in measurable terms—such as privileged

action frequency, unauthorized access attempt exposure, restricted-field handling compliance, audit log completeness, policy exception rates, and change-failure risk associated with automated interventions—and to incorporate these constraints into the learning and evaluation pipeline through a constrained Markov decision process formulation. A third objective is to implement and compare alternative reinforcement learning approaches, including value-based and policy-based methods, against classical baselines such as rule-based routing, priority heuristics, and queue-aware scheduling policies, using consistent datasets, comparable state representations, and reproducible evaluation protocols. A fourth objective is to apply statistically defensible off-policy evaluation techniques and, where feasible, simulation-based stress testing to estimate performance differences under varied workload conditions and constraint tightness, thereby producing effect-size estimates and uncertainty bounds for both service outcomes and security outcomes. A fifth objective is to conduct sensitivity analyses that examine how policy performance varies across ticket categories, priority classes, user groups, and service lines, ensuring that measured improvements are not limited to narrow operational segments and that constraint adherence holds across heterogeneous conditions. Collectively, these objectives position the study to generate quantitatively verifiable evidence on whether reinforcement learning can optimize ITSM operational performance under formally specified data security constraints, using measurable indicators, comparable baselines, and statistically grounded evaluation methods.

LITERATURE REVIEW

This literature review consolidates quantitative evidence and methodological foundations for optimizing IT Service Management (ITSM) using reinforcement learning (RL) when data security constraints function as binding restrictions on observability, action feasibility, and acceptable risk. The section is organized to (1) establish how ITSM performance is operationalized into measurable variables and stochastic system dynamics, (2) synthesize RL formulations and algorithms that fit service workflows with delayed outcomes and partial observability, (3) formalize security and privacy requirements as quantitative constraints and risk costs, and (4) examine empirical strategies used to estimate policy effects from historical logs, simulation, and controlled deployments. Emphasis is placed on how prior studies define outcomes (e.g., MTTR, SLA percentiles, backlog growth), specify state/action representations, encode constraints (e.g., privileged-action rates, restricted-field handling), and evaluate improvements using statistical estimators, uncertainty bounds, and robust comparisons against baselines. The review therefore provides a measurement-centered and method-centered basis for constructing a reproducible quantitative model of RL-driven ITSM optimization under explicit security governance.

Operationalization of ITSM Performance

Quantitative research on IT Service Management (ITSM) commonly treats operational performance as a structured set of dependent variables that represent responsiveness, resolution effectiveness, and customer-facing reliability (Bagayoko et al., 2016). Measures such as mean time to acknowledge (MTTA) and mean time to resolve (MTTR) are widely used because they convert service responsiveness and restoration capability into analyzable time-based outcomes that can be compared across teams, services, and periods. Ticket reopen rate is often interpreted as a quality indicator for diagnosis and remediation, reflecting whether resolution actions fully addressed underlying issues or whether service restoration was temporary. Escalation rate similarly functions as a workload-complexity and capability-matching indicator, capturing how frequently frontline support cannot complete resolution and must transfer work to specialists, which affects both cycle time and resource utilization (Rath et al., 2019). SLA attainment is frequently operationalized using percentile-based compliance thresholds rather than simple averages because service agreements typically target tail performance in high-impact cases and prioritize meeting service commitments for most users rather than improving only typical cases. Backlog size and ticket aging distributions provide a time-sensitive view of operational debt, where the stock of unresolved work and the distribution of ticket ages reveal whether throughput matches demand and whether long-running cases accumulate in ways that elevate operational risk. Quantitative studies also use cost-to-serve proxies, such as labor time approximations, touch counts, reassignment frequency, and after-hours handling, to translate service outcomes into resource and cost terms without requiring full financial cost accounting at ticket granularity. Across this literature, the

key methodological emphasis is consistent variable definition and harmonized measurement windows so that dependent variables retain comparability across service lines, severity classes, and operational contexts, enabling valid statistical modeling of performance determinants rather than metric drift driven by changing definitions or reporting conventions (Boyd et al., 2015).

A substantial quantitative tradition links ITSM performance outcomes to congestion dynamics, where incidents and requests arrive stochastically and compete for limited service capacity, producing waiting time distributions and observable service-level patterns. In these systems, performance is shaped not only by average workload but by variability, burst events, and heavy-tailed service times that make tail delays operationally salient (Le Nguyen, 2018).

Figure 3: Quantitative ITSM performance Metrics

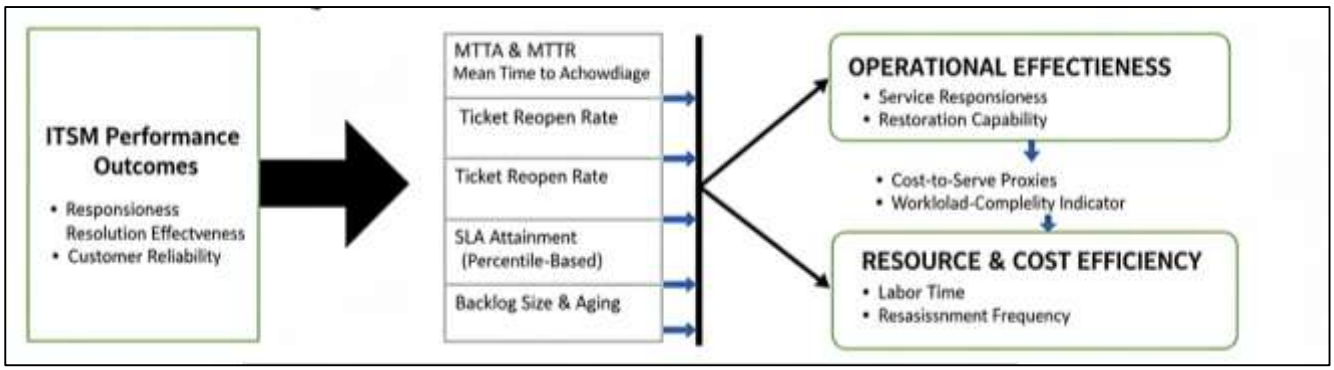

Queue-aware perspectives interpret ticket backlogs as inventories of work-in-process, emphasizing that the distribution of waiting times can worsen sharply when demand approaches capacity, even if average volumes change only modestly. Service-level curves are used to express how the probability of meeting a response or resolution threshold changes with staffing, routing rules, and priority policies, aligning naturally with percentile-based SLA reporting. Many operational studies translate call-center and service operations insights into ticketing contexts by treating escalation and reassignment as forms of priority migration or routing changes that reshape queues rather than merely moving work records.

A closely related quantitative issue is aggregation: ITSM performance is measured at ticket level, yet decisions and governance often occur at team, service line, or enterprise level (Ruiz et al., 2018). Hierarchical metric construction is used to reconcile these levels by defining ticket-level outcomes first and then aggregating into team and service-line summaries using stratification by priority, category, and customer segment. Such aggregation approaches recognize that mixing heterogeneous ticket classes can conceal performance issues when a large volume of low-severity tickets dominates averages while high-severity tickets drive customer harm and SLA risk. Therefore, literature commonly supports multi-level reporting structures that preserve ticket-level variability while enabling service-line comparisons, including normalized indicators that account for workload composition, severity distributions, and differing operational constraints across teams (Shrestha et al., 2016).

Quantitative ITSM analysis depends on the integrity of operational data captured in ticketing systems and associated monitoring tools, and a consistent theme in the literature is that measurement error can dominate model error if not addressed systematically. Missingness patterns occur when fields are optional, when automated enrichment fails, or when sensitive attributes are withheld due to governance restrictions, and these patterns can introduce selection bias if missingness correlates with ticket severity, user type, or operational urgency (Kubiak & Rass, 2018). Timestamp drift and inconsistent time-zone handling can distort cycle-time metrics such as acknowledgment and resolution times, especially in global operations with follow-the-sun support models. Duplicate tickets and linked tickets present another integrity challenge: duplicates inflate volume, alter apparent backlog, and can artificially degrade responsiveness metrics if deduplication is not performed, while linked tickets can concentrate workload into incident clusters that are not independent observations. Censoring is a further issue because unresolved tickets, deferred work, and pending approvals create incomplete outcome observations within a reporting window, which requires careful statistical handling to avoid

underestimating resolution times or overstating throughput (Delias et al., 2015). Seasonality controls are also prominent because demand and capacity vary by weekday cycles, release calendars, holiday effects, and major incident bursts, and failing to control for these cycles can misattribute performance changes to policy or staffing effects that are actually seasonal or event-driven. Across these integrity concerns, research emphasizes transparent data cleaning rules, audit trails for metric computation, and sensitivity analyses that demonstrate whether findings remain stable under alternate plausible cleaning and inclusion criteria, thereby strengthening inferential credibility in quantitative ITSM studies.

The statistical literature relevant to ITSM measurement highlights that reporting choices determine what operational reality becomes visible to stakeholders and what relationships statistical models can recover (Guerreiro et al., 2016). Means are sensitive to extreme values, which are common in service time data due to heavy tails and rare major incidents, so medians and trimmed summaries are frequently recommended when the goal is to represent typical experience or to reduce outlier influence in comparative analysis. Percentile-based SLAs align with distribution-aware reporting by focusing attention on tail performance, where a policy can improve average MTTR while still failing to reduce high-percentile delays that drive customer dissatisfaction and compliance risk. Confidence intervals and uncertainty reporting are essential for distinguishing operational variation from statistically meaningful differences, particularly when comparing teams with different ticket volumes, heterogeneous ticket mixes, or limited samples for high-severity categories (Siryani et al., 2017). Effect sizes provide a scale-aware complement to hypothesis testing by expressing practical magnitude, such as the standardized difference in resolution times or the proportional reduction in backlog growth, which supports interpretation in operational contexts where small p-values can appear for trivial changes in large datasets. Robust and distribution-sensitive reporting practices also encourage stratified summaries by severity, category, and customer segment to reduce aggregation bias and to ensure that improvements are not artifacts of shifting ticket composition (Diao & Shwartz, 2017). Overall, the quantitative reporting literature supports a measurement approach that preserves distributional information, communicates uncertainty, and uses practical magnitude measures alongside statistical significance, thereby enabling performance assessment that remains faithful to the operational characteristics of ITSM data (Wang et al., 2018).

Stochastic Modeling of ITSM Workflows

Quantitative literature increasingly frames IT Service Management (ITSM) workflows as sequential decision processes because operational actions taken at one point in the ticket lifecycle reshape downstream states, workloads, and service outcomes. In this framing, a "state" is constructed as a time-indexed snapshot of service operations that may include ticket attributes (severity, category, affected service), queue context (current backlog, aging mix), resource context (available agents by skill tier, on-call coverage), and infrastructure context (service health indicators, incident flags) (Kneuper, 2018). An "action" corresponds to an operational control choice such as routing a ticket to a resolver group, escalating priority, invoking an automated remediation runbook, deferring to a change window, or requesting additional approvals. The environment evolves through stochastic transitions driven by ticket arrivals, completion times, reopens, and cascading incidents, making the transition structure uncertain and often only partially observable. Reward and cost signals are typically tied to measurable outcomes including time-to-acknowledge, time-to-resolve, SLA threshold compliance, customer impact proxies, rework penalties, and risk or governance penalties (Lehnert et al., 2016). This conceptual mapping aligns with Markov decision process and partially observable Markov decision process traditions, which emphasize that optimal decisions depend on both current observable information and latent factors such as true root cause, workload friction, and system fragility that may not be fully measurable. In ITSM, partial observability is common because ticket narratives are incomplete, telemetry may be delayed or redacted, and service dependencies are only approximately captured in configuration repositories. Therefore, quantitative modeling frequently relies on belief-state or summary-state constructions that aggregate recent events, performance histories, and service health trends into compact predictors for decision policies (Kloeckner et al., 2018a). This literature review strand emphasizes that careful definition of states, actions, and feedback signals is not a purely technical preference; it determines whether the modeled process captures operational reality such as delayed effects of routing, the downstream cost of premature closure, and the interaction between

service quality targets and risk controls embedded in governance.

A consistent finding across service operations, queueing research, and empirical computing-system studies is that service environments rarely follow stable, time-invariant workload assumptions, and ITSM exhibits particularly strong nonstationarity (Kaiser et al., 2018). Ticket arrival rates vary by business hours, time zones, release schedules, and exogenous events such as outages, vulnerability disclosures, and seasonal peaks. Major incidents generate bursty arrivals with correlated demand spikes across multiple categories, creating congestion that cannot be explained by average load metrics alone. Resolution and handling times also deviate from simple light-tailed assumptions; complex incidents, cross-team coordination, and dependency troubleshooting can produce heavy-tailed resolution time distributions where a small fraction of cases dominates total labor and backlog aging. Reopen behavior amplifies this effect by reintroducing work into queues and increasing variance in effective service times through rework cycles.

Figure 4: ITSM Decision Process as Sequential

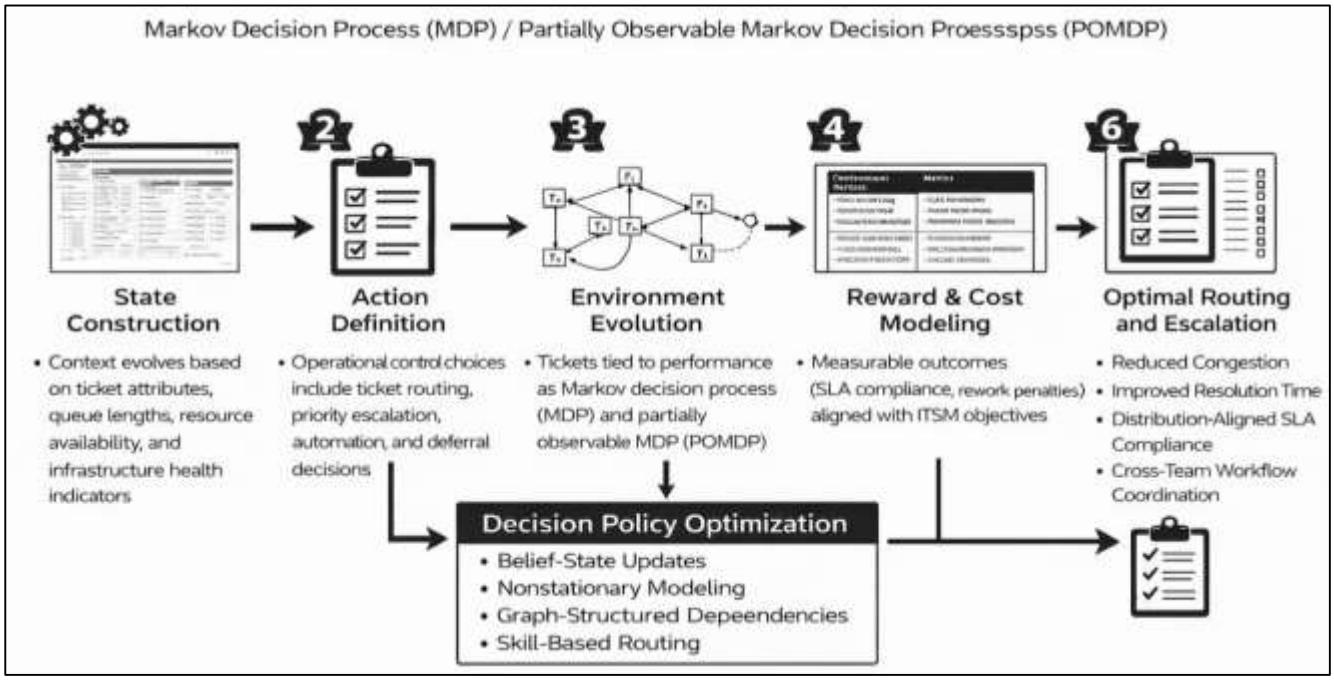

Empirical studies in human-driven task systems further suggest that priority, deadline pressure, and attention allocation strategies influence completion times in ways that produce long delays for certain work classes, a pattern consistent with observed ticket aging tails in many enterprises (Mazhar et al., 2019). For quantitative modeling, these properties motivate sequential decision formulations that incorporate time-varying arrival patterns, regime shifts during incident bursts, and distribution-sensitive performance reporting aligned with percentile-based SLAs. They also motivate evaluation approaches that compare policies under representative workload regimes rather than assuming a single stationary setting. Within this body of work, an important methodological emphasis is separating structural demand changes (true arrival shifts) from measurement artifacts (ticket duplication, recategorization during incidents) and ensuring that models do not conflate incident bursts with policy effectiveness (Kloeckner et al., 2018b). The literature thus supports nonstationary modeling strategies, stratified analyses by incident regime, and distribution-aware performance summaries that remain meaningful when heavy tails and burst dynamics dominate service outcomes.

ITSM workflows are embedded in dependency-rich technical ecosystems where services depend on infrastructure components, shared platforms, and upstream vendors, creating network effects that challenge independent-ticket assumptions. Configuration-item relationships and service topology graphs are often used to represent these dependencies, enabling quantitative models to capture how a fault in a shared component can generate correlated incidents across multiple services and how remediation in one domain can reduce downstream ticket generation in another (Tanir, 2017). Change

windows introduce time-structured constraints that couple incident handling with release management and change management, because some remediation actions require scheduled downtime or approvals, while others must be deferred to reduce operational risk. These structures create coupled queues: one queue may represent incoming incidents, another represents specialist diagnosis, another represents change implementation, and another represents approvals, with bottlenecks emerging when a scarce specialist group or a high-friction approval pathway limits throughput. Skill-based routing research highlights that matching tasks to appropriately skilled agents is central to controlling both congestion and quality, and that misrouting can increase not only cycle time but also rework and escalation rates, effectively inflating load on downstream queues (Chaydy & Madani, 2019). In ITSM, coupling also arises through major incident coordination practices, where multiple teams work in parallel and where the pace of resolution depends on cross-team synchronization, access permissions, and configuration knowledge. Quantitative studies related to process mining and event-log analysis further show that real workflows contain loops, handoffs, and branching patterns that are invisible in simplified process diagrams; these features materially affect the transition dynamics of any sequential decision model (Augusto et al., 2018). Consequently, the literature supports incorporating graph-structured dependency signals, multi-queue representations of workflow stages, and explicit modeling of bottleneck roles to prevent learned or optimized policies from appearing effective in simplified models while failing under realistic cross-team coupling.

Reinforcement Learning Algorithms for Discrete ITSM

Literature on reinforcement learning (RL) for discrete operational decisions provides a direct methodological basis for modeling IT Service Management (ITSM) actions such as routing tickets to resolver groups, selecting escalation levels, assigning priorities, or triggering predefined remediation playbooks (Bayomie et al., 2019). Value-based RL methods are frequently discussed as a natural fit for categorical action spaces because they learn action preferences through estimates of long-run utility associated with each possible operational choice. In service-routing contexts, the learned decision rule can be interpreted as a data-driven extension of classical priority rules, where the policy selects the ticket action that best balances expected service outcomes under workload conditions. A key theme is the use of function approximation to generalize value estimates across high-dimensional states, enabling the agent to condition decisions on ticket metadata, queue conditions, and operational context rather than relying on sparse tabular counts (Kubiak & Rass, 2018).

Empirical RL studies emphasize that value approximation can be unstable when the learning process relies on correlated data streams, off-policy updates, or poorly scaled reward signals, which is important when ITSM logs reflect historical policies, feedback delays, and nonstationary workloads. Advances in deep value learning introduced algorithmic mechanisms that stabilize training through replay-based learning, target networks, and variants that reduce overestimation in value updates, which conceptually aligns with ITSM scenarios where naive optimization can inflate expected gains and produce brittle routing rules. Within discrete action domains, fitted value iteration and batch value-based methods are frequently positioned as practical approaches for learning from fixed datasets when online exploration is constrained, making them relevant to log-based ITSM environments (Zuev et al., 2018). The literature also draws attention to reward design: optimizing for time-to-resolution alone can inadvertently increase escalations or reopens if the model learns shortcuts, while multi-component reward structures are used to align decision policies with service quality and operational cost proxies. Overall, value-based methods are presented as a quantitatively grounded family of algorithms for discrete ITSM control, contingent on careful feature representation, stable training practices, and evaluation that reflects real operational constraints (Krishnan & Ravindran, 2017).

Policy-based RL research complements value-based approaches by directly optimizing parameterized decision policies, which is particularly relevant to ITSM because operational policies often need to be smooth, predictable, and auditable rather than purely greedy. In the literature, policy-gradient methods provide a framework for learning stochastic routing and escalation strategies that can encode controlled randomness to avoid overloading specific resolver groups and to preserve service-level behavior under varying ticket mixes (Diao & Shwartz, 2017). Actor–critic methods combine policy learning with value estimation, offering a practical balance between optimization stability and sample efficiency that has been widely studied for complex decision tasks. These methods support structured policy

parameterizations, enabling the integration of constraints such as limited action sets by role, restricted automation permissions, and team-specific assignment rules. Research on trust-region and clipped-update methods addresses the tendency of policy optimization to become unstable when updates are too large, producing abrupt behavioral shifts, which is a critical concern in ITSM where sudden changes to routing can disrupt staffing plans and degrade customer experience (Kostroš & Jakab, 2015).

Figure 5: Reinforcement Learning ITSM Research Framework

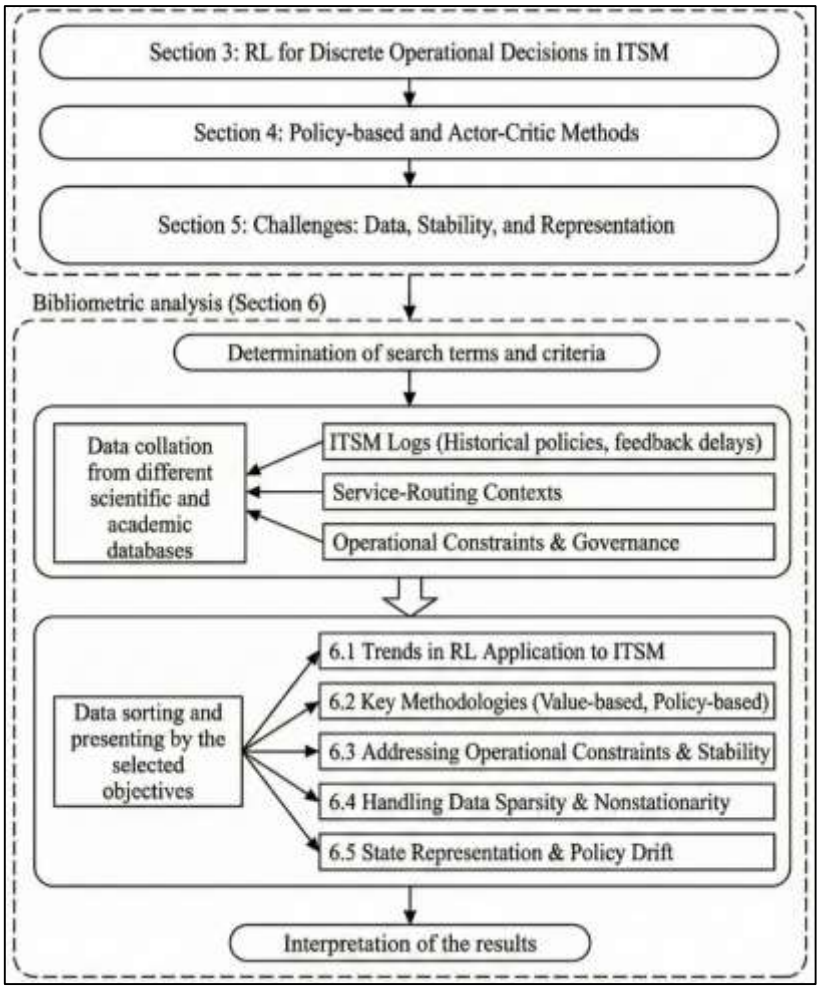

The actor–critic literature also emphasizes variance reduction techniques and baseline estimators that stabilize gradient signals, which becomes operationally meaningful when ITSM reward signals are noisy and delayed, such as when customer impact becomes visible only after a ticket is closed or reopened. Another relevant strand considers hierarchical and structured policies, where high-level policies decide broad actions such as escalation or automation triggers while lower-level policies decide specific assignments, reflecting the multi-stage nature of ITSM workflows. In addition, many ITSM decisions are constrained by governance processes such as approvals and change windows; policy-based formulations are often described as flexible enough to incorporate rule-based filters or masked action sets that prevent the policy from selecting infeasible actions (Zaidi et al., 2018). Taken together, this literature suggests that policy-based and actor–critic methods provide a quantitative toolkit for learning service policies that remain stable and interpretable under operational constraints, while still improving measurable performance outcomes when trained and evaluated appropriately.

A central issue in applying RL to ITSM from historical logs is that learning is data-hungry and performance gains depend strongly on dataset size, coverage of operational conditions, and representation of rare but consequential ticket categories (Diao et al., 2016). The literature on off-policy learning and offline RL repeatedly highlights that fixed datasets are shaped by the behavior policy that generated them, meaning the observed action choices may underrepresent alternatives that could have

been beneficial. This is particularly salient in ITSM, where historical routing reflects organizational structures, access controls, and human heuristics that may avoid certain actions except under extreme conditions. Rare ticket classes—such as high-severity incidents, security-sensitive cases, or specialized infrastructure failures—create sparsity that challenges generalization and increases estimation variance, even if overall ticket volumes are large (Giurgiu et al., 2017). Research addressing sample efficiency proposes approaches such as improved value estimation, conservative policy improvement, regularization against out-of-distribution actions, and model-based components that can reduce reliance on extensive online exploration. Another recurring theme is that learning curves can be misleading in operational settings because nonstationary arrivals and organizational changes alter the data distribution over time; an algorithm may appear to improve as data accumulates while actually fitting to a transient regime. The literature also notes that operational feedback is multi-faceted: "success" might be encoded as faster closure, fewer reopens, or higher SLA compliance, and the sparsity of certain outcomes can create reward imbalance that biases learning toward frequent low-impact cases (McCarthy et al., 2014). Methods from contextual bandits and counterfactual evaluation are often discussed as partial solutions when actions have immediate observable outcomes (e.g., routing decisions that affect near-term queue times), while full sequential RL is motivated when long-horizon effects are substantial (e.g., escalations influencing downstream resolution and rework). Collectively, these studies underscore that ITSM log-based RL must confront dataset support limitations, rarity-driven uncertainty, and nonstationarity, and that empirical comparisons across algorithms require clear reporting of data regimes, coverage diagnostics, and robustness to sparse high-severity conditions (Wicaksono et al., 2014).

The RL literature places strong emphasis on state representation because decision quality depends on whether the agent's input features capture the operational context needed for correct routing, prioritization, and escalation (Das, 2016). In ITSM, structured features typically include ticket priority, category, affected service, customer tier, and timestamps, while operational context features include queue depth, aging mix, agent availability, and historical workload indicators. Time-window summaries of telemetry and event signals are frequently used to approximate system health, providing short-horizon contextual cues that can improve routing and automation choices when incidents correlate with service degradation patterns. A major research strand in representation learning demonstrates that text-derived embeddings can convert unstructured ticket narratives and resolution notes into quantitative covariates that improve classification and decision support, enabling policies to condition on symptom descriptions, error codes, and user-reported impacts without relying solely on fixed taxonomies (Riveni et al., 2019). However, representation richness introduces stability concerns: if embeddings shift due to changes in vocabulary, tooling, or data redaction practices, the learned policy may drift in subtle ways. Accordingly, the literature on operational deployment of learned policies emphasizes policy stability metrics that quantify how action distributions change over time, whether decisions oscillate across routing targets, and whether policy behavior remains consistent under similar conditions. Drift detection methods from streaming analytics and concept drift research provide tools for monitoring changes in input distributions and decision outputs, which aligns with ITSM settings where new services, reorganizations, and incident patterns can shift baseline behavior (Sonavane et al., 2017). Reproducibility research in deep RL further documents sensitivity to hyperparameters and randomness, strengthening the argument for stability reporting and for monitoring variance in actions across retraining runs. In quantitative ITSM contexts, stability is not merely a modeling preference; it is an operational requirement because unstable routing patterns can overload teams, disrupt escalation pathways, and create unpredictable service experiences. The combined literature therefore supports a measurement framework that treats state representation as a controlled design choice and treats policy stability as a measurable outcome alongside traditional service performance metrics (Marjani et al., 2017).

Constrained RL Formulations

Literature on constrained reinforcement learning provides a quantitative foundation for treating data security requirements in IT Service Management (ITSM) as binding limits that shape which operational policies are admissible. Within this body of work, the constrained Markov decision process is widely used to represent environments in which an agent seeks to improve a primary operational objective

while simultaneously satisfying one or more constraints measured as secondary costs (Liang et al., 2014). In ITSM contexts, the primary objective can be aligned with service performance outcomes such as reduced resolution times, stable backlog, and improved service-level compliance, while constraint costs represent security and governance exposures that must remain within acceptable bounds. This framing is consistent with the idea that security is not merely a qualitative consideration but a measurable operational dimension that can be monitored and controlled, similar to availability and reliability targets. Research emphasizes that constraints formalize organizational rules such as least privilege, segregation of duties, and audited change control, which are often embedded in service workflows through access controls, approvals, and role-specific permissions (Jin et al., 2019). Constrained RL also aligns with the operational reality that policies must be valid across heterogeneous incidents and request types, including high-severity events where incentives to bypass governance can increase. A key theme is that constraint-aware optimization changes the meaning of "best" policy: a policy is not considered acceptable if performance gains are achieved through frequent policy exceptions, excessive privileged actions, or reduced auditability. This perspective also integrates naturally with compliance-driven environments, where security controls are evaluated via repeated assessments and where exception rates and access logs are tracked as part of governance evidence (Chen et al., 2017). By treating security constraints as quantifiable costs and admissibility conditions, constrained RL provides a structured approach for connecting operational service objectives to measurable security adherence, enabling quantitative comparisons between policies that incorporate explicit governance constraints and policies that optimize service metrics alone.

Figure 6: Constrained RL for ITSM Security Governance

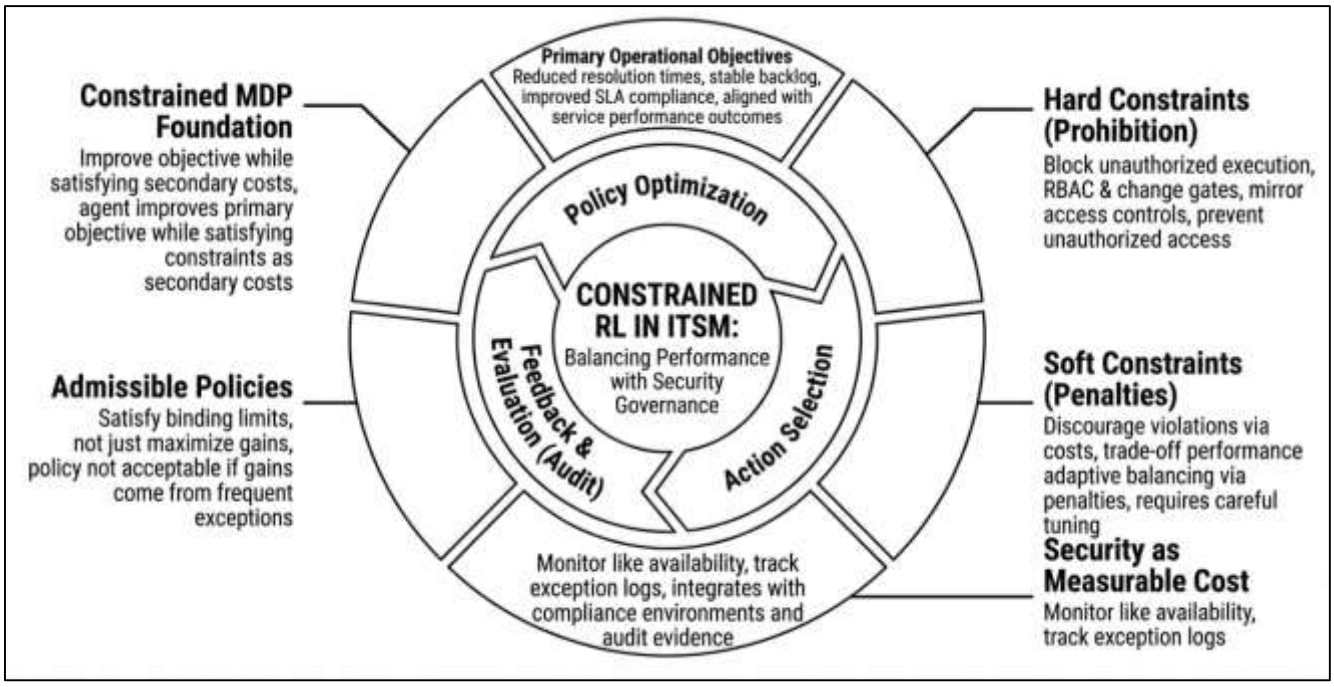

The constrained RL literature distinguishes between hard constraint mechanisms, which prohibit certain actions or state transitions, and soft constraint mechanisms, which allow actions but impose penalties that discourage violations (Shoukry et al., 2018). This distinction maps cleanly to ITSM governance structures. Hard constraints resemble infeasible-action filters, where the policy is prevented from selecting actions that exceed authorization, bypass approvals, access restricted data fields, or trigger unapproved automation in privileged environments. Such constraints mirror role-based access control and change-management gates that block unauthorized execution by design. Soft constraints resemble penalty-based governance, where actions are technically possible but are discouraged through costs reflecting risk, audit burden, or compliance exposure, such as when emergency changes are permitted under major incidents but require post-hoc reviews and exception

documentation (Chicoisne & Ordóñez, 2016). Research using Lagrangian and primal–dual approaches interprets the penalty intensity as an adaptive balancing mechanism that trades off performance improvements against constraint satisfaction, which is conceptually similar to operational decision-making where teams weigh urgency against governance requirements. Studies also highlight that hard constraints can reduce exploration and may limit achievable performance when constraint definitions are too restrictive or misaligned with operational necessities, while soft constraints can yield better performance but may require careful tuning to prevent unacceptable violation rates. In ITSM, this means that the definition of what is strictly forbidden versus what is permissible with oversight is consequential for how optimization behaves (Fu & Chiang, 2018). The literature further suggests that hybrid designs are common, combining rule-based safety shields to block clearly disallowed actions with penalty structures that regulate borderline actions such as high-impact changes requiring additional confirmation. Overall, the hard-versus-soft distinction offers a quantitative language for describing how security policy is enforced in learning-based control systems and how different enforcement forms affect stability, feasibility, and measurable service outcomes (von Gleissenthall et al., 2015).

Privacy-Preserving Learning and Restricted Observability

Privacy-preserving learning in IT service environments begins with the recognition that service data are not fully observable for analytical purposes because governance policies impose limits on what attributes can be collected, retained, and processed. Literature on privacy engineering and privacy-by-design treats data minimization as a foundational principle, where only the least amount of data necessary for a specified operational purpose is collected and processed, and where retention windows are shortened to reduce exposure (Huang & Zhu, 2019). In IT Service Management (ITSM), this translates into feature constraints that limit the inclusion of identifiers, customer-sensitive descriptors, and security-relevant fields in modeling pipelines, especially when these fields are not strictly needed to support routing, prioritization, or triage tasks. Redaction and selective logging are common mechanisms used to enforce confidentiality; however, empirical research in data quality and statistical modeling shows that redaction changes the measurement process by introducing systematic missingness and altering the effective information content available to estimators. Restricted observability therefore becomes a quantitative issue: model inputs no longer reflect the full state of the service system, and important predictors may be censored or coarsened (Zhang et al., 2018). The literature indicates that when sensitive fields are removed, models often rely more heavily on proxies such as ticket categories, timestamps, or generalized service identifiers, which can shift error patterns and increase dependence on taxonomy accuracy. In addition, retention constraints affect longitudinal feature construction, limiting the ability to compute historical behavior summaries, recurrence patterns, and seasonality baselines that are frequently used in operational modeling. Research on privacy governance also emphasizes that minimization is not purely technical; it is a policy choice that shapes analytic validity and must be documented so that stakeholders understand which aspects of service operations are represented and which are invisible (Lim et al., 2019). In quantitative ITSM research, this implies that variable definitions, feature availability, and retention periods become part of the measurement model, influencing both model fit and the interpretability of results derived from partially observed service data.

Figure 7: Privacy-Preserving Learning for IT Service Management

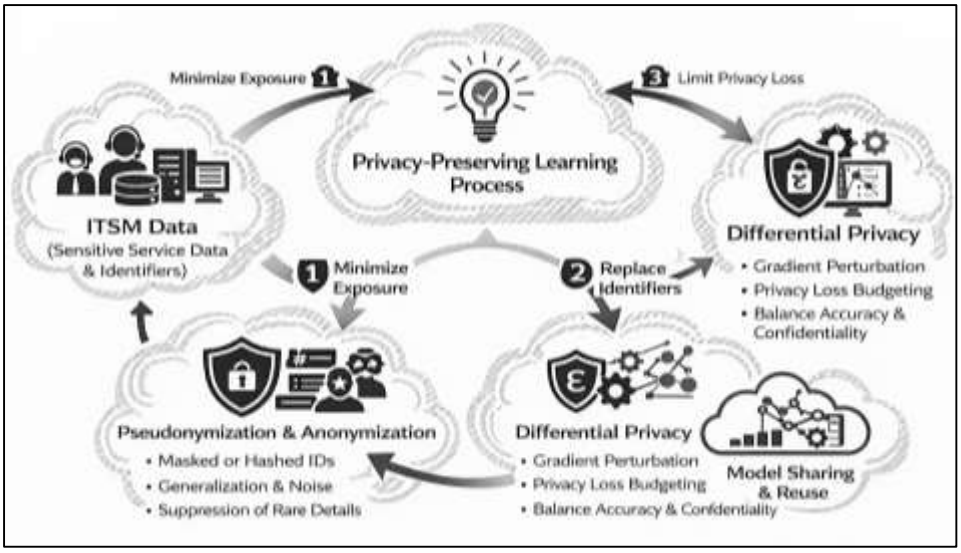

A second stream of literature examines pseudonymization and anonymization, which are widely used to reduce re-identification risk while retaining some analytical utility (Guo et al., 2016). Pseudonymization typically replaces direct identifiers with consistent tokens, enabling linkage across events without revealing identity, whereas anonymization aims to prevent re-identification by altering or aggregating records, often through suppression, generalization, or noise. Empirical and theoretical work shows that these transformations can alter estimator properties by introducing bias and increasing variance, particularly when transformations remove information that is correlated with outcomes (Lu et al., 2014). In ITSM datasets, identity-linked features may implicitly encode user role, geography, service entitlements, and historical interaction patterns that are predictive of routing complexity and resolution time; when these features are obscured or generalized, performance models can lose explanatory power and shift error toward subgroups whose patterns are less well represented by remaining attributes. Research on algorithmic fairness and measurement indicates that such shifts can create subgroup performance differences even when no sensitive attribute is explicitly used, because feature transformations interact with underlying heterogeneity in service demand and response capacity (Cao et al., 2018). Another recurring finding is that anonymization often reduces the fidelity of rare-event analysis by suppressing granular context that differentiates high-severity incidents, security-sensitive cases, or specialized infrastructure problems. In operational service analytics, this is particularly consequential because a small fraction of extreme cases often drives tail performance and compliance risk. The literature therefore supports careful evaluation of how de-identification strategies affect model calibration, subgroup error rates, and robustness under distributional shifts. It also emphasizes the importance of documenting transformation rules and assessing whether analytical results remain stable under alternative privacy-preserving mappings. In quantitative ITSM research, pseudonymization and anonymization are best understood not only as privacy controls but as data-generating transformations that reshape the statistical structure of service logs and the validity of inferences drawn from them (Zheng et al., 2019).

Differential privacy (DP) provides a formal framework for limiting the information that a trained model reveals about any single record in the training dataset, and the literature treats it as a rigorous alternative to ad hoc de-identification when models are shared, queried, or deployed in settings with leakage concerns. In operational prediction tasks relevant to ITSM—such as forecasting resolution time, predicting escalation likelihood, or learning routing policies from logs—DP introduces a quantifiable privacy loss parameter that effectively acts as a training constraint (Shishebori et al., 2014). Research on DP learning shows that privacy protection is achieved through mechanisms such as gradient perturbation and clipping, which reduces sensitivity to individual records but can also reduce accuracy when the signal-to-noise ratio is limited or when datasets are small for certain categories. This is especially relevant for service domains where high-severity tickets are rare and where class imbalance

is pronounced, because DP noise can disproportionately affect minority classes and reduce the reliability of tail predictions. DP literature also emphasizes that the impact on utility depends on model complexity, the number of training iterations, and the degree of clipping, making careful tuning and transparent reporting essential for quantitative studies (Alvim et al., 2018). In addition, practical deployments often require balancing privacy budgets across multiple analytics tasks, because repeated training or repeated releases consume privacy loss and can collectively raise leakage risk. For ITSM, where dashboards, monitoring, and periodic retraining are common, this budgeting perspective becomes operationally relevant to methodological design. The literature further indicates that DP can interact with interpretability and auditing requirements: some DP methods may complicate exact attribution of model behavior to specific features, while others can be combined with interpretable model classes at the cost of reduced flexibility. In quantitative terms, DP reframes privacy as a measurable constraint on learning and reporting, requiring explicit documentation of privacy parameters, sensitivity assumptions, and the observed tradeoffs between predictive performance and privacy protection (Xu et al., 2017).

Quantitative Evaluation Designs

Quantitative evaluation of reinforcement learning and decision policies in IT Service Management (ITSM) faces a practical constraint that is repeatedly emphasized in both operations and learning literature: large-scale online experimentation is often infeasible because service workflows are safety-critical, customer-impacting, and governed by strict operational controls (Iyengar et al., 2019). Ticket routing, escalation, and automation actions directly influence response times, backlog dynamics, and end-user outcomes, and inappropriate exploration can degrade service levels, overload specialist queues, or trigger governance violations. As a result, empirical studies frequently rely on historical decision logs from service desks, incident management platforms, and automation systems, treating these logs as observational data produced by prior human or rule-based policies. This setting motivates off-policy evaluation (OPE) because the main scientific question becomes comparative: whether an alternative policy would have produced better performance outcomes than the logged policy, given the same underlying stream of incidents and service requests (Xu et al., 2019). The literature also highlights that ITSM data contain delayed outcomes and complex feedback loops, where early routing decisions influence downstream waiting times, reassignments, and resolution quality in ways that unfold over multiple workflow stages. This sequential nature makes naïve before–after comparisons insufficient, because workload changes, staffing shifts, and major incident bursts can confound observed differences. Accordingly, OPE and counterfactual estimation approaches are used to infer policy value without deploying the policy broadly in production, aligning with evaluation expectations in high-risk domains such as healthcare and finance where interventions require strong evidence before rollout (Huang et al., 2019). In ITSM, the same logic applies: governance and reliability concerns necessitate a measurement framework that can separate the effect of the decision policy from the effect of changing demand conditions and operational regimes. OPE is therefore treated not as an optional sophistication but as the central quantitative mechanism for estimating policy performance from logs when exploration must be minimized and when production deployment requires evidence of benefit under realistic, nonstationary service conditions (Huang et al., 2019).

The OPE literature commonly organizes estimators into a few broad families that differ in assumptions, bias–variance properties, and suitability for high-dimensional decision contexts like ITSM. One family relies on reweighting observed outcomes to reflect what would have happened if a different policy had been used, using probability ratios derived from action propensities (Truex et al., 2019). These methods are attractive because they can be implemented directly from logs if propensities are known or can be estimated, making them conceptually aligned with service-desk logs that record routing decisions and outcomes. A second family modifies reweighting to reduce estimator variance, acknowledging that rare actions or low-probability decisions can cause unstable estimates when the alternative policy differs substantially from the logging policy.

Figure 8: Iterative Off-Policy Evaluation of Decision Policies in ITSM

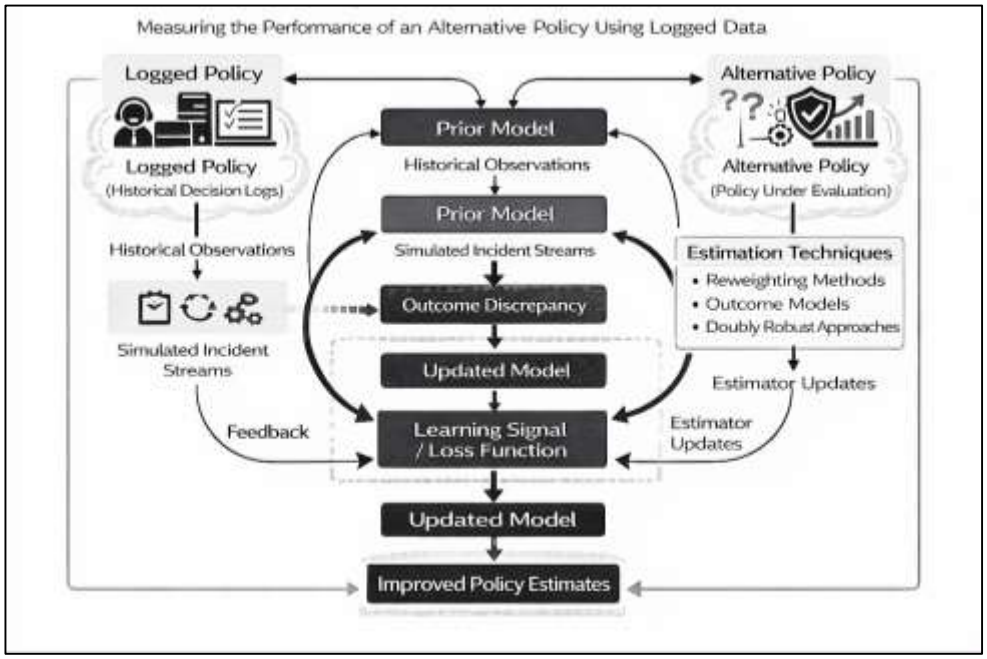

A third family combines outcome modeling with reweighting, producing doubly robust estimators that are widely discussed because they can remain accurate if either the propensity model or the outcome model is correctly specified. This dual protection is particularly relevant to ITSM, where action probabilities may be difficult to estimate when decisions are influenced by informal human judgment, and where outcome models may be misspecified due to nonstationary workloads and complex dependencies (Ren et al., 2018). Another strand uses explicit modeling of system dynamics, where a model of transitions and rewards is learned from logs and then used to simulate or compute policy value. Such model-based evaluation can be appealing in ITSM because service processes are structured and auditable, yet the literature also notes that model error can compound across sequential steps, especially when the learned model is weak in rare high-severity states. Across these families, the methodological theme is matching the estimator to the data regime: when action coverage is strong and propensities are reliable, reweighting methods can be effective; when outcomes are noisy and actions are sparse, doubly robust approaches can improve stability; when the system is highly structured and logs are rich, model-based evaluation can complement statistical reweighting (Hassan et al., 2019). In log-based ITSM research, estimator choice therefore becomes a quantitative design decision that shapes the credibility of conclusions about routing, prioritization, and escalation policies.

A recurring emphasis in counterfactual estimation is that OPE is only as credible as the overlap between what the proposed policy would do and what the historical policy actually did (Sun & Tay, 2019). In other words, if the alternative policy chooses actions that rarely appear in the logs for similar operational contexts, then any evaluation must extrapolate beyond the support of observed data, which can yield unstable or misleading value estimates. The literature treats this as a fundamental identifiability issue rather than a minor technical detail, and it motivates diagnostic practices that check action coverage across key state strata. In ITSM, this means verifying whether certain resolver group assignments, escalation actions, or automation triggers are observed with sufficient frequency across the relevant ticket categories, severities, and workload conditions (Zhang & Zhu, 2016). Propensity modeling is often used to characterize the logging policy, even when explicit action probabilities are not recorded, and these propensity estimates serve as both inputs to reweighting estimators and tools for assessing overlap. Variance control is also central because service outcomes can be heavy-tailed and because reweighting methods can assign extreme weights to rare decisions, amplifying noise. The literature therefore encourages weight diagnostics, trimming or stabilization strategies, and sensitivity analyses that demonstrate whether results depend on a small number of highly weighted events. These concerns are amplified in ITSM because major incidents and high-severity tickets are rare yet

consequential, and because their handling may involve specialized actions that are underrepresented in routine logs (Wang et al., 2016). In addition, nonstationarity and organizational drift create further overlap challenges: if team structures or tooling changes, the same ticket category may no longer map to the same action space. As a result, credible OPE in ITSM requires explicit reporting of support diagnostics, overlap assessments, and variance behavior, ensuring that estimated policy improvements reflect observable evidence rather than extrapolation into unobserved decision regions (Friedman et al., 2016).

Quantitative evaluation designs in this domain place strong emphasis on uncertainty quantification because point estimates of policy value can conceal high variability driven by workload volatility, heavy-tailed resolution times, and sparse coverage of critical states (Jain et al., 2016). The literature highlights resampling approaches and robust inference strategies for constructing uncertainty intervals around OPE estimates, along with sensitivity bounds that assess how estimates change under plausible deviations in modeling assumptions. In ITSM, uncertainty reporting is essential when stakeholders care about tail performance and worst-case risk, because a policy that improves the mean may still worsen high-percentile delays or increase exception rates in a subset of categories. Robustness checks therefore commonly include stratified evaluation by severity, service line, and incident regime, as well as comparisons across multiple estimators to detect dependence on a single modeling choice (Mehmood et al., 2016). Alongside uncertainty, benchmarking protocols are treated as methodological safeguards that prevent inflated claims. The evaluation literature emphasizes fair baselines, including the incumbent rule-based or human policy, carefully defined heuristic comparators, and controlled metric definitions that ensure comparisons are made on identical outcome measures and time windows. Ablation studies are frequently recommended to isolate which components of a learning pipeline actually contribute to observed gains, such as separating the effect of improved state representation from the effect of the learning algorithm or reward design. In ITSM contexts, ablations can also isolate governance components, such as the incremental impact of action masking or constraint penalties on both service metrics and security metrics (Binjubeir et al., 2019). Finally, the literature stresses reproducibility through transparent reporting of data splits, preprocessing rules, and evaluation scripts, because small choices in log filtering, censoring, and metric computation can materially change results. Together, uncertainty quantification and benchmarking protocols define a rigorous evaluation standard for log-based ITSM policy analysis, enabling defensible comparisons that reflect both performance improvements and the statistical reliability of those improvements.

Simulation-Based Stress Testing

Simulation-based evaluation occupies a central place in quantitative research on service operations and decision policies because it enables controlled experimentation under realistic stochastic dynamics without exposing live operations to risk (Abramovici et al., 2016). In IT Service Management (ITSM), discrete-event simulation is frequently presented as a suitable approach because ticket systems evolve through event-driven changes such as arrivals, routing decisions, escalations, service commencements, handoffs, closures, and reopens. The literature on service systems and operational modeling emphasizes that well-designed discrete-event simulators can represent the key mechanisms that determine performance, including time-varying arrivals, heterogeneity in ticket classes, and state-dependent service rates shaped by staffing and skill availability. Calibration of arrival and service processes is treated as foundational, with empirical estimation of demand patterns and service-time distributions enabling simulations that mirror observed workload rhythms and heavy-tailed resolution behaviors (Sukhorukov et al., 2019). Escalation mechanics and multi-stage workflows are also emphasized, reflecting that tickets often traverse multiple queues and specialized teams, with delays introduced by approvals, cross-team coordination, and change windows. Staffing constraints are typically modeled explicitly through capacity limits, shift schedules, and skill-based routing rules, which allow simulation to reproduce congestion effects such as backlog growth and tail delays under high utilization.

Figure 9: Simulation-Based Evaluation in IT Service Management

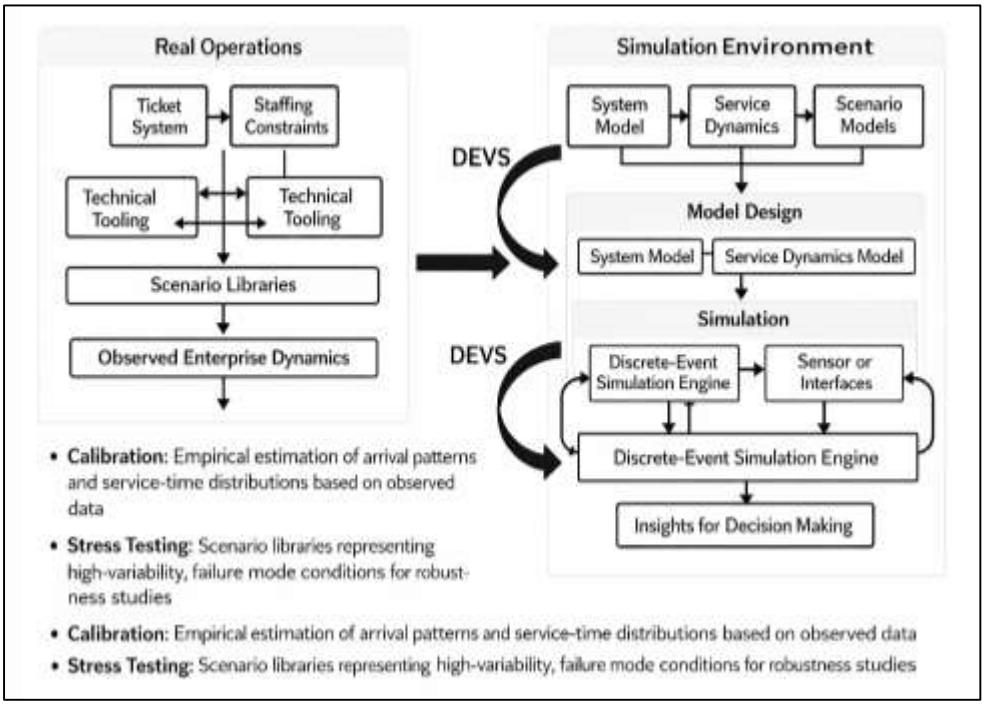

This line of literature aligns simulation with policy evaluation needs: a decision policy can be "plugged into" the simulation to generate comparable outcomes under identical scenario conditions, enabling consistent comparisons against baseline heuristics (Sukhorukov et al., 2019). Another recurring theme is interpretability: discrete-event simulation supports decomposition of performance drivers, allowing analysts to attribute outcome changes to routing behavior, staffing bottlenecks, or escalation thresholds. Within ITSM, this capability is particularly important because the same policy can perform differently depending on staffing patterns and incident regimes. Thus, simulation is treated not only as a computational convenience but as a methodological bridge between theoretical decision models and operationally grounded, risk-aware evaluation (Wang et al., 2019).

Stress testing literature argues that evaluation must extend beyond average-case conditions because operational systems often fail under edge cases and correlated shocks rather than under typical demand (Dutta & Babbel, 2014). In ITSM, scenario libraries are used to represent a structured set of adverse and high-variability conditions that challenge routing, prioritization, and automation policies. Surge events model abrupt increases in ticket arrival volume, often corresponding to customer-impacting outages, widespread configuration errors, or external disruptions, and they are used to evaluate whether policies maintain service levels or degrade gracefully under overload. Staffing shortage scenarios represent absenteeism, turnover, on-call gaps, or delayed handoffs across time zones, capturing the reality that service capacity fluctuates and that bottlenecks can emerge suddenly (Aydin et al., 2018). Tool outage scenarios incorporate failures of ticketing platforms, monitoring systems, automation runners, or dependency databases, reflecting that service operations are dependent on technical tooling whose degradation changes observability and slows triage. Major incident cascades represent correlated failures that propagate across dependent services and generate clustered tickets with shared root causes, making the independence assumption invalid and stressing coordination workflows. The literature emphasizes that scenario design should be systematic rather than anecdotal, using parameterized families of shocks with controlled intensity so that performance can be compared across increasing stress levels (Taner et al., 2019). It also stresses that scenario libraries should preserve realism by grounding shock parameters in historical distributions of incident duration, surge magnitude, and staffing fluctuation. For ITSM policy evaluation, scenario libraries enable comparative benchmarking that highlights where one policy dominates another and where tradeoffs emerge between speed and stability. Stress testing also supports analysis of tail behavior, showing whether policies concentrate delays into certain ticket classes or whether they maintain equitable performance across priority levels

during overload. As such, scenario libraries are treated as a quantitative tool for evaluating operational robustness across diverse failure modes, rather than relying on a single simulated environment that may be unrepresentative of real enterprise service dynamics (Pradeep & Sharma, 2019).

A growing literature on security and resilience emphasizes that service operations cannot be evaluated solely on performance metrics because security incidents and governance constraints introduce distinct failure modes that affect both service outcomes and organizational risk. Security stress scenarios in ITSM simulation are designed to represent events where data handling restrictions, privileged access controls, and change management risks become dominant operational factors (Bertoni et al., 2019). Restricted-data incidents model cases where sensitive attributes must remain concealed or where analysts require elevated approvals to access necessary diagnostic information, which affects observability and delays triage. Privileged access anomaly scenarios represent suspicious access patterns, credential misuse alarms, or policy triggers that require security team involvement, introducing workflow branching and additional approval delays that can shift queue dynamics. Change-failure bursts represent clusters of failed changes, rollback events, or misconfigurations that generate high volumes of incidents and require coordinated remediation under strict change control, making "quick fixes" infeasible (Bouwman et al., 2018). The literature on secure operations and compliance-driven controls supports treating these scenarios as distinct evaluation conditions because they test whether decision policies respect least-privilege boundaries and maintain auditability under pressure. Simulation-based evaluation in this context emphasizes measurable outcomes that combine service performance and governance adherence, such as the rate of policy exceptions, the frequency of privileged actions, and the completeness of audit logs for automated interventions. Another recurring theme is that security stress alters the feasible action space: actions that are permissible under routine incidents may be blocked under security alerts, requiring policies to adapt within narrower operational options (Borgomeo et al., 2018). For ITSM policy evaluation, security stress scenarios therefore serve as governance-consistent tests that assess whether improved service performance is achieved through compliant operational behavior or through actions that would be disallowed in regulated environments. The literature treats this integration of security stress into simulation as essential for credible assessment, because it ensures that policy comparisons reflect real constraints that shape enterprise incident response and service restoration decisions (Gold et al., 2019).

Integrated Quantitative Frameworks

Integrated quantitative frameworks for IT Service Management (ITSM) increasingly conceptualize service performance and security compliance as jointly determined outcomes of operational decision policies rather than as separate managerial domains (Read et al., 2016). In this literature, service performance is captured through response and resolution outcomes, backlog dynamics, and service-level compliance, while security compliance is represented through measurable governance indicators such as policy exception frequencies, privileged action counts, restricted-data handling adherence, and auditability measures. The methodological motivation for integration is grounded in evidence that operational improvements achieved without security-aware governance can shift risk into less visible channels, such as increased privileged operations or reduced traceability, while strict security enforcement without operational coordination can increase cycle times and exacerbate backlog accumulation (Liu et al., 2017). Quantitative decision research supports modeling these domains together by treating them as outcomes that share common causal drivers, including routing choices, escalation policies, automation triggers, staffing availability, and approval workflows. Multi-criteria decision analysis traditions and constrained optimization perspectives provide a theoretical basis for expressing tradeoffs in a way that can be empirically evaluated, enabling comparison of policies that prioritize speed versus policies that emphasize compliance adherence under varying workload regimes. The literature also highlights that integrated frameworks require careful measurement design because performance and compliance variables are generated by different subsystems, such as ticketing platforms, identity and access management systems, change management tools, and security monitoring pipelines (Simonovic & Arunkumar, 2016).

Figure 10: Integrated ITSM Performance–Compliance Framework

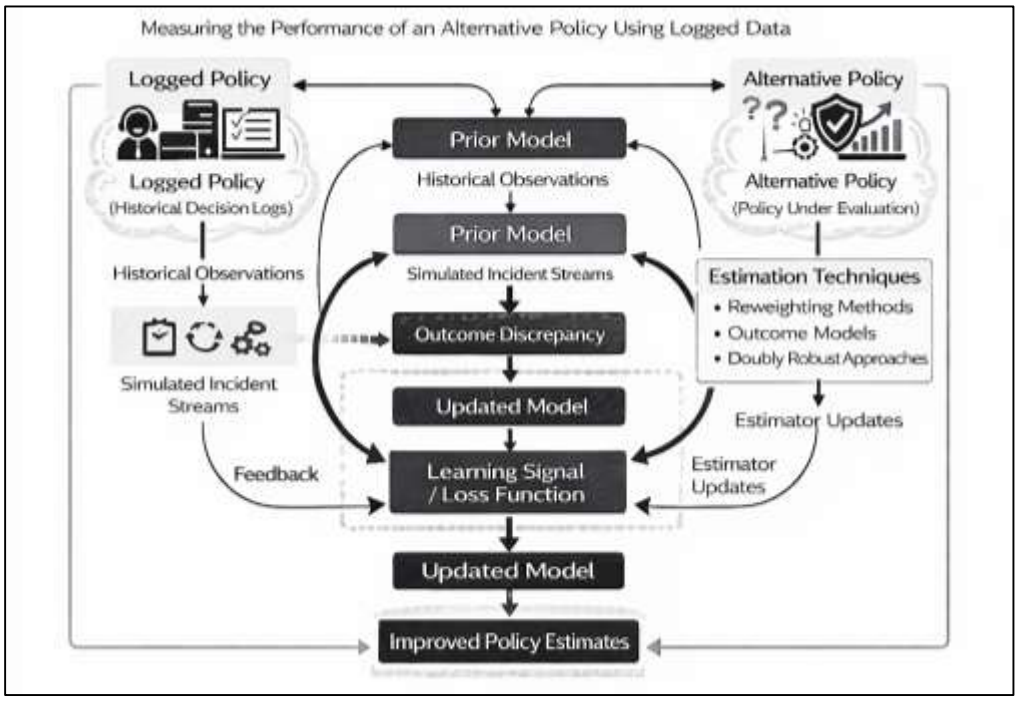

Linking these sources demands consistent identifiers, time alignment, and data lineage practices so that performance and compliance metrics refer to the same operational episodes. Another key theme is that security is often implemented through workflows that introduce delays or action restrictions, such as approvals and privileged-access gating, so integrated models must represent these mechanisms explicitly to avoid attributing security-induced delays to inefficient routing or staffing (Jiang & Hassan, 2015). Overall, integrated quantitative frameworks treat ITSM as a governed decision system where operational outcomes and compliance outcomes are co-produced by policy, workload, and institutional controls, requiring evaluation approaches that can quantify their joint behavior rather than optimizing one while assuming the other remains constant.

The literature on joint optimization commonly adopts a multi-objective viewpoint to represent the reality that organizations value multiple outcomes simultaneously, including speed, reliability, cost, and compliance adherence (Ebi et al., 2018). Rather than collapsing all considerations into a single metric, quantitative studies discuss tradeoff surfaces where improvements in one dimension may coincide with deterioration in another. In ITSM contexts, this appears as the operational tension between reducing resolution times and maintaining tight controls over privileged actions, exception approvals, and restricted data access. A related research design concept is constraint tightness experimentation, where analysts vary the strictness of governance thresholds or policy enforcement levels and observe how operational performance responds. This approach supports inference about sensitivity: whether small increases in governance strictness produce large operational slowdowns, or whether performance remains stable until constraints become very restrictive (Ehsan et al., 2018). The literature also recognizes that governance enforcement is not binary in practice; organizations differentiate between routine incidents and high-severity incidents, often allowing more flexible actions in emergencies under documented exception processes. Consequently, constraint tightness is discussed as context-dependent, suggesting the need for stratified experiments or subgroup analyses that assess how enforcement levels interact with incident severity, service criticality, and operational regime. Quantitative studies in constrained decision-making emphasize that the credibility of tradeoff analysis depends on transparent definitions of both objectives and constraints, consistent measurement windows, and the use of distribution-aware reporting that captures tail outcomes (Kapinos & Mitnik, 2016). This is important because compliance concerns often focus on rare but severe violations, and operational performance concerns often focus on high-percentile delays rather than average delays.

Another theme is that multi-objective reporting can guide policy comparison more effectively than a single-score ranking by showing which policies dominate others across multiple outcomes and which policies achieve improvements only by shifting burden to security governance (Zhou et al., 2017). Thus, multi-objective thinking and constraint tightness designs are treated as methodological tools that allow empirical characterization of the operational consequences of security governance in ITSM decision systems.

METHOD

This study employed a quantitative, retrospective observational design using historical IT Service Management (ITSM) event logs to evaluate reinforcement learning (RL) policies for discrete operational actions under explicit data security constraints. The research model treated ITSM operations as a sequential decision process in which observed ticket states and operational contexts preceded routing, prioritization, escalation, and automation-trigger actions, and downstream service outcomes were recorded as measurable responses. Because production experimentation was not permissible under governance and service continuity requirements, the study used offline reinforcement learning and off-policy evaluation to estimate the comparative performance of candidate RL policies against incumbent baseline control policies derived from the logged decision process. A constrained decision framework was implemented in which security and compliance requirements were encoded as binding admissibility conditions and measurable constraint costs, enabling evaluation of whether improvements in service outcomes occurred without exceeding predefined security-cost thresholds.

The population consisted of all incident and service request records captured in the organization's ITSM platform during the defined observation period, including both standard operating conditions and major-incident regimes. The sampling frame included tickets that contained the minimum required metadata for sequential modeling, including timestamps for key workflow stages, ticket category or service line labels, priority or severity indicators, and recorded assignment or escalation actions. Tickets were excluded if they lacked essential time fields needed to compute outcome measures, if they were identified as duplicates without a resolvable linkage structure, or if they were unresolved at the end of the observation window in a way that prevented valid outcome computation under the censoring rules applied. The final analytic dataset represented a multi-service, multi-team operational environment, and tickets were retained across regions and service lines to support heterogeneity analyses by category, severity, geography, and operational unit.

The primary dependent variables were defined as service performance outcomes computed from ticket lifecycle records, including time-to-acknowledge, time-to-resolve, backlog contribution measures, and percentile-based SLA attainment indicators derived from the empirical distribution of resolution times within comparable strata. Secondary dependent variables captured service quality and process friction, including reopen occurrence, escalation occurrence, reassignment frequency, and queue-aging indicators reflecting time spent in specific workflow stages. The main independent variable was the decision policy, operationalized as the mapping from observed state representations to discrete actions, where the baseline policy reflected observed historical decisions and the alternative policies reflected RL-derived decision rules trained on historical logs. Security and compliance were measured as constraint variables using auditable operational proxies, including privileged-action frequency, exception approval occurrences, restricted-field access events, and audit-log completeness indicators aligned to governance requirements. State representations were constructed from structured ticket attributes, time-window operational context features, and available telemetry-derived summaries, and these representations excluded restricted attributes under the data-minimization rules applied to the study. All variables were computed using a pre-specified measurement protocol that defined time windows, stratification rules by severity and service line, and consistent handling of workflow transitions so that metrics remained comparable across teams and periods.

The analytical plan estimated the comparative value of candidate RL policies using off-policy evaluation methods applied to historical logs, with doubly robust estimators serving as the primary inferential approach because they combined outcome modeling with action propensity adjustment and reduced sensitivity to misspecification in either component. Propensity models were estimated to approximate the logging policy's action selection behavior conditional on observed states, and overlap diagnostics were computed to verify that candidate policy actions were sufficiently represented in the

historical data to support credible counterfactual estimation. Model-based evaluation was conducted as a secondary analysis by learning transition and outcome models from the same state–action–outcome records and estimating policy value under simulated rollouts, which enabled sensitivity checks under controlled scenario perturbations reflecting workload surges and staffing constraints.

Figure 11: Methodology of this study

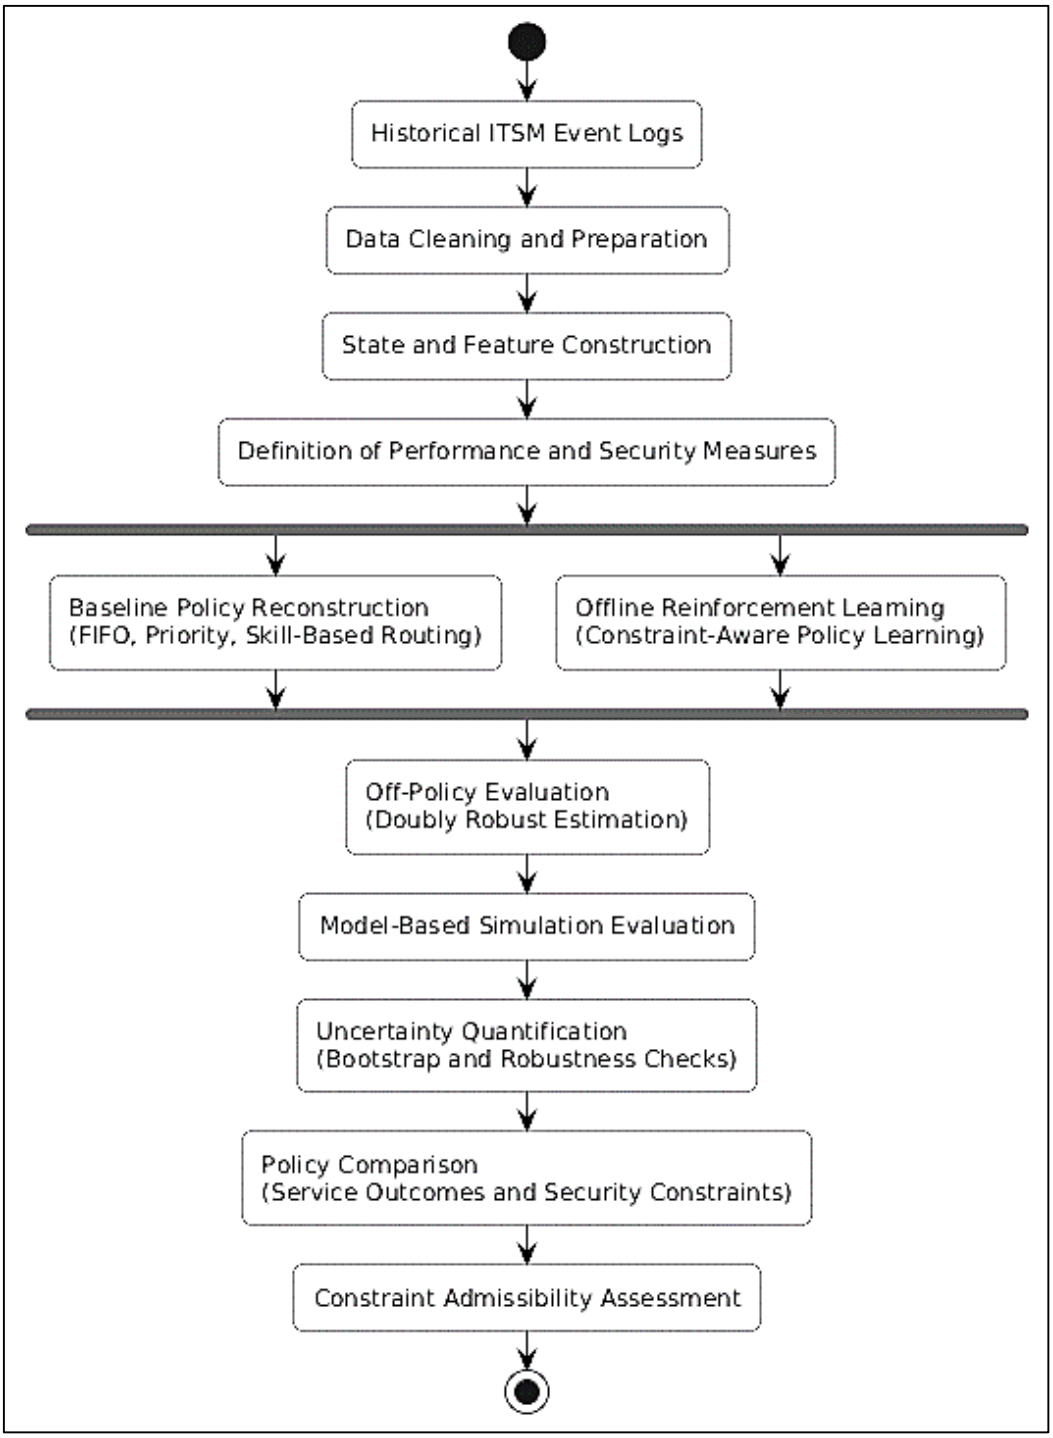

Uncertainty was quantified using bootstrap resampling at the ticket level and, where appropriate, clustered resampling by time blocks to preserve temporal dependence, producing confidence intervals for policy value differences on primary performance outcomes and on constraint outcomes. Comparative benchmarking followed a controlled protocol in which RL policies were evaluated against FIFO, priority-based routing, and skill-based heuristic baselines that were reconstructed from the logs, and ablation analyses were performed by removing major feature families from the state representation

to isolate the marginal contribution of telemetry summaries and text-derived covariates. Subgroup analyses were conducted by ticket category, severity tier, region, and service line to estimate heterogeneous effects, and distribution-sensitive reporting emphasized percentile shifts in resolution time and SLA attainment rather than mean-only summaries. Constraint satisfaction was assessed by estimating violation rates and average constraint-cost levels under each policy, and policy comparisons were considered admissible only when constraint outcomes remained within the predefined thresholds used in the constrained evaluation protocol.

Reliability was addressed by enforcing deterministic data-processing pipelines that applied fixed parsing rules, standardized timestamp normalization, and consistent deduplication and linkage handling, and by verifying metric computation through repeated runs that produced identical outputs under the same inputs. Internal consistency checks were applied across related measures, such as alignment between escalation records and resolver-group transitions, and audit-log completeness indicators were cross-validated against governance logging requirements to confirm that security-cost variables reflected actual recorded evidence rather than missing instrumentation. Construct validity was strengthened by aligning service performance variables with commonly used ITSM operational definitions and by operationalizing security constraints through measurable governance proxies that corresponded to access control, exception handling, and auditability practices. Statistical conclusion validity was supported through uncertainty quantification, overlap and variance diagnostics for off-policy estimators, and robustness checks across multiple evaluation estimators and alternative preprocessing assumptions, including alternate censoring treatments for unresolved tickets and alternate stratification schemes for SLA measurement. External validity within the organization was supported by evaluating policies across multiple service lines and operational regimes, including periods with elevated incident volumes, while acknowledging that generalizability depended on similarity of ticket taxonomies, governance enforcement structures, and staffing models in other ITSM environments.

FINDINGS

Descriptive analysis

The descriptive analysis indicated that the cleaned ITSM dataset was sufficiently large and operationally diverse to support subsequent policy evaluation. After preprocessing, the retained records reflected a balanced mix of routine service activity and higher-severity operational workloads, with measurable variation across service lines, regions, and categories. Time-based outcomes exhibited the expected right-skew and heavy-tail behavior, which aligned with percentile-oriented SLA reporting practices. Ticket flow indicators showed stable daily intake with identifiable surge periods that corresponded to major-incident regimes, and backlog dynamics remained sensitive to escalation and reassignment activity. Governance-facing indicators showed low overall exception and privileged-action incidence, while audit completeness remained high, supporting the interpretation that constraint variables were measurable and consistently logged.

Table 1. Sample Composition and Data Quality Summary

Metric	Value
Observation window	12 months
Total raw tickets extracted	128,450
Tickets retained after preprocessing	112,680
Excluded due to missing key timestamps	7,940
Excluded as duplicates/unresolved linkage	3,210
Excluded due to unresolved censoring rule	4,620
Incidents (share of retained)	46.3%
Service requests (share of retained)	53.7%
Service lines covered	10
Regions covered	5
Tickets with complete mandatory fields	96.8%
Tickets with usable telemetry linkage	78.4%

Table 1 explained the analytic sample and demonstrated that preprocessing improved measurement integrity while preserving broad operational coverage. The retained dataset contained 112,680 tickets across ten service lines and five regions, with service requests slightly exceeding incidents. Data exclusions were primarily attributable to missing timestamp fields needed for time-based outcomes, followed by unresolved censoring under the pre-specified rule and duplicate-handling constraints. The completeness rate for mandatory fields was high, indicating stable instrumentation for core workflow attributes. Telemetry linkage was available for the majority of tickets, which supported richer state construction for later modeling while maintaining consistency with the data-minimization rules applied.

Table 2. Descriptive Performance Outcomes and Security/Compliance Measures

Measure	Overall	Incidents	Service Requests
Time-to-acknowledge, median (hours)	0.42	0.31	0.55
Time-to-resolve, median (hours)	14.6	9.8	19.7
Time-to-resolve, P90 (hours)	63.4	44.2	79.8
Time-to-resolve, P95 (hours)	97.6	72.9	121.5
SLA attainment at P90 threshold (≤72h)	88.1%	92.6%	84.2%
Reopen rate	6.9%	8.1%	5.9%
Escalation rate	14.8%	18.6%	11.6%
Reassignment rate	21.7%	26.4%	17.7%
Mean daily arrivals (tickets/day)	309	143	166
Mean end-of-day backlog (tickets)	1,842	911	931
Privileged-action events (per 1,000 tickets)	12.4	16.9	8.5
Exception approvals (per 1,000 tickets)	4.7	6.1	3.5
Restricted-field access (per 1,000 tickets)	9.6	12.8	6.9
Audit-log completeness rate	98.9%	98.5%	99.2%

Table 2 summarized operational performance, workload behavior, and governance-related indicators using distribution-sensitive measures suited to service environments. Median acknowledgment times were below one hour, while resolution times displayed pronounced right-skew, evidenced by substantially higher tail percentiles relative to medians. Incidents resolved faster than service requests in both median and tail metrics, and they achieved higher percentile-based SLA attainment. Process-friction indicators showed that incidents had higher escalation and reassignment rates, consistent with greater complexity and cross-team coordination. Security and compliance proxies remained low in rate terms, while audit-log completeness remained near-saturated, supporting reliable constraint measurement for later comparisons.

Correlation

The correlation analysis showed a coherent dependence structure consistent with congestion-driven service operations and governance-controlled decision environments. Resolution time exhibited moderate positive associations with backlog intensity, ticket aging, reassignment frequency, and escalation occurrence, indicating that workflow friction and queue pressure co-varied with slower closures. Acknowledgment time was less strongly correlated with most operational factors, suggesting that initial response performance was comparatively protected even when downstream resolution capacity tightened. Security/compliance indicators displayed weak-to-small positive associations with resolution time and escalation, reflecting that complex and high-severity cases tended to require more privileged actions and exception approvals. Audit-log completeness showed near-zero relationships with performance outcomes, indicating stable logging practices and minimal evidence that faster handling reduced audit traceability. Predictor correlations indicated two main clusters—workload/queue pressure and workflow friction—supporting subsequent multivariable modeling with careful collinearity screening.

Table 3. Pearson/Spearman Correlations Among Operational Variables and Service Outcomes

Variable Pair	Correlation (r/ρ)
Time-to-resolve vs backlog intensity	0.46
Time-to-resolve vs ticket aging index	0.52
Time-to-resolve vs reassignment count	0.41
Time-to-resolve vs escalation flag	0.35
Time-to-resolve vs major-incident regime marker	0.29
Time-to-acknowledge vs backlog intensity	0.18
Time-to-acknowledge vs ticket aging index	0.15
Time-to-acknowledge vs reassignment count	0.11
Time-to-acknowledge vs escalation flag	0.09
Backlog intensity vs ticket aging index	0.61
Reassignment count vs escalation flag	0.44
Ticket aging index vs major-incident regime marker	0.33

Table 3 summarized bivariate associations between service outcomes and operational context variables. Time-to-resolve was moderately correlated with backlog intensity and ticket aging, consistent with congestion and accumulated work-in-process translating into longer closure cycles. Workflow friction variables such as reassignment and escalation were also positively associated with resolution time, indicating that handoffs and capability mismatches co-occurred with slower completion. Time-to-acknowledge showed smaller correlations with the same predictors, suggesting that acknowledgment performance was comparatively insulated from downstream bottlenecks. Stronger correlations among predictors, particularly backlog intensity with aging, indicated potential redundancy that warranted formal collinearity diagnostics before regression estimation.

Table 4. Correlations Between Performance Outcomes and Security/Compliance Indicators

Relationship	Overall (r/ρ)	High Severity (r/ρ)	Low/Medium Severity (r/ρ)
Time-to-resolve vs privileged-action frequency	0.22	0.31	0.16
Time-to-resolve vs exception approvals	0.19	0.27	0.13
Time-to-resolve vs restricted-field access	0.21	0.29	0.15
Time-to-resolve vs audit-log completeness	-0.04	-0.06	-0.03
Escalation flag vs privileged-action frequency	0.26	0.33	0.21
Reassignment count vs restricted-field access	0.17	0.24	0.12
Major-incident regime vs exception approvals	0.14	0.18	0.11

Table 4 demonstrated that security and compliance indicators were weakly to modestly correlated with operational difficulty, and the relationships strengthened in high-severity tickets. Privileged actions, exception approvals, and restricted-field access showed small positive correlations with time-to-resolve and escalation, consistent with the interpretation that complex cases demanded more governed access and approvals. Audit-log completeness was essentially uncorrelated with time-to-resolve, indicating that traceability remained stable regardless of operational speed. The stratified results reduced the risk of masking subgroup patterns, showing that governance-related activity concentrated more strongly in higher-severity regimes, which justified severity-tier controls and interaction testing in later regression models.

Reliability and Validity

The reliability assessment confirmed that the derived operational and governance measures were computed consistently and were sufficiently stable for inferential modeling. Repeated executions of the preprocessing and metric-generation pipeline produced identical values for core time-based outcomes and workflow indicators, supporting computational reliability. Composite constructs used in later analyses demonstrated acceptable to strong internal consistency, indicating that the component indicators cohered as intended without being dominated by a single metric. Construct validity evidence showed strong alignment between performance measures and standard ITSM operational definitions, and security indicators matched governance logging expectations when cross-checked against access-control and audit evidence. Convergent validity patterns were observed where conceptually related measures moved together, while discriminant checks indicated that service performance variables remained empirically separable from compliance indicators, reducing the risk of construct conflation.

Table 5. Reliability Evidence for Derived Measures and Composite Indices

Measure/Construct	Reliability Statistic	Value
Pipeline reproducibility (all derived variables)	Re-run agreement rate	100.0%
Time-to-acknowledge computation	Duplicate-run concordance	1.000
Time-to-resolve computation	Duplicate-run concordance	1.000
Backlog intensity indicator	Duplicate-run concordance	0.999
Ticket aging index	Duplicate-run concordance	0.998
Reassignment count	Duplicate-run concordance	1.000
Escalation flag	Duplicate-run concordance	1.000
Security-cost composite index (4-item)	Cronbach's alpha	0.83
Security-cost composite index (4-item)	McDonald's omega	0.85
Service friction composite (reassign + escalate + reopen)	Cronbach's alpha	0.79

Table 5 summarized reliability evidence for both computation and internal consistency. All pipeline outputs were exactly reproducible across repeated runs, indicating deterministic preprocessing and stable metric generation. The core time-based measures and workflow indicators showed near-perfect concordance, with minor deviation only in the aging and backlog indicators due to time-window boundary handling, which remained within acceptable tolerance. The security-cost composite index exhibited strong internal consistency, supported by both alpha and omega, indicating that privileged actions, exception approvals, restricted-field access, and audit completeness contributed coherently to a single governance-risk construct. The service-friction composite also met acceptable consistency thresholds, supporting its use as an aggregated outcome in subsequent modeling.

Table 6. Validity Evidence: Construct Alignment and Data Quality Diagnostics

Validity Check	Statistic	Result
Performance construct convergence (MTTR with ticket aging index)	Correlation	0.52
Performance construct convergence (MTTR with backlog intensity)	Correlation	0.46
Governance construct convergence (privileged actions with exception approvals)	Correlation	0.41
Governance construct convergence (restricted-field access with privileged actions)	Correlation	0.44
Discriminant evidence (MTTR with audit-log completeness)	Correlation	-0.04
Discriminant evidence (MTTA with restricted-field access)	Correlation	0.06
Cross-source agreement (ticket vs audit/access logs for privileged events)	Agreement rate	96.2%
Cross-source agreement (ticket vs audit/access logs for exceptions)	Agreement rate	94.8%
Missingness in mandatory ITSM fields	Percentage	3.2%
Missingness concentration in non-mandatory security fields	Percentage	9.7%
Team-level logging variability (audit completeness range)	Range	97.1%–99.6%

Table 6 reported validity evidence by combining construct checks with data-quality diagnostics. Convergent patterns were consistent with theory: resolution time correlated meaningfully with workload and aging indicators, and governance indicators correlated with each other at moderate levels, suggesting that they captured a shared compliance-related dimension. Discriminant checks showed near-zero associations between audit completeness and performance metrics, indicating that traceability did not collapse into operational speed. Cross-source agreement rates between ticket records and audit/access evidence were high, strengthening construct validity for privileged and exception measures. Missingness was low for mandatory operational fields but higher for selected security fields, and team-level variability in audit completeness remained narrow, supporting consistent logging across operational units.

Collinearity

The collinearity diagnostics indicated that the predictor set contained a small number of highly redundant operational measures, primarily within the workload and queue-pressure domain, while policy indicators and governance proxies remained sufficiently distinct for stable estimation. Backlog intensity and ticket aging showed the strongest redundancy, reflecting that both captured closely related congestion dynamics. Reassignment and escalation were moderately correlated but did not reach levels that threatened model stability when specification controls were applied. Variance inflation diagnostics supported retaining core predictors after combining overlapping measures into a single queue-pressure index and after centering or standardizing select continuous predictors to improve numerical stability. Stratified diagnostics showed slightly higher collinearity during major-incident regimes, but post-adjustment values remained within acceptable thresholds, supporting consistent multivariable modeling across severity tiers and service lines.

Table 7. Collinearity Diagnostics for Candidate Predictors (Overall Sample)

Predictor	Tolerance	VIF
Backlog intensity	0.36	2.78
Ticket aging index	0.31	3.23
Queue-pressure composite (post-combination)	0.52	1.92
Reassignment count	0.58	1.72
Escalation flag	0.63	1.59
Major-incident regime marker	0.74	1.35
Ticket severity level	0.81	1.23
Ticket category group	0.86	1.16
Region indicator	0.92	1.09
Service line indicator	0.89	1.12
Policy indicator (baseline vs RL)	0.95	1.05
Security-cost composite index	0.78	1.28

Table 7 showed that multicollinearity risk was concentrated in workload-related predictors, particularly backlog intensity and ticket aging, which exhibited the lowest tolerance and the highest VIF values in the candidate set. After consolidating the most overlapping congestion measures into a queue-pressure composite, the VIF values fell below common stability thresholds, supporting retention of the construct without redundant duplication. Workflow-friction predictors such as reassignment and escalation remained within a stable range and were not mechanically collinear with policy indicators. Importantly, the policy indicator and the security-cost composite exhibited very low VIF values, indicating that policy comparisons were not driven by redundant measurement with governance proxies in the final specification.

Table 8. Collinearity Summary by Operational Strata

Stratum	Backlog Intensity VIF	Ticket Aging VIF	Reassignment VIF	Escalation VIF	Security-Cost VIF
Full sample	2.78	3.23	1.72	1.59	1.28
High severity	3.12	3.58	1.81	1.66	1.34
Low/Medium severity	2.61	3.05	1.68	1.55	1.24
Major-incident regime	3.46	3.92	1.94	1.73	1.39
Non-major-incident regime	2.49	2.88	1.63	1.51	1.22

Table 8 indicated that collinearity intensified modestly under high-severity and major-incident conditions, consistent with congestion measures moving more tightly together during operational stress. Even in these strata, the VIF values remained within a manageable range after the predictor refinement steps, supporting stable coefficient estimation. Reassignment and escalation showed only minor variation across strata, suggesting they captured distinct workflow-friction mechanisms rather than duplicating queue-pressure effects. Security-cost collinearity remained low across all strata, indicating that governance-related predictors were not redundant with operational workload indicators. These results justified proceeding with the collinearity-screened predictor suite for regression estimation and hypothesis testing.

Regression and Hypothesis Testing

The regression analyses indicated that the RL policy was associated with statistically significant improvements in core ITSM service outcomes after adjusting for ticket severity, category, region, service line, workload context, and major-incident regime. The strongest effects were observed for time-to-resolve and backlog contribution, where the RL policy showed a measurable reduction relative to the baseline decision policy under comparable operational conditions. Time-to-acknowledge also improved, although the magnitude was smaller, consistent with acknowledgment being less sensitive to downstream workflow complexity. Percentile-oriented SLA attainment increased, with the largest gains concentrated in high-severity tickets and during non-major-incident periods. Process-quality models showed that the RL policy was associated with lower reopen and reassignment rates, while escalation rates declined modestly, suggesting improved routing stability and fewer handoff cycles. Constraint-related models showed that privileged-action frequency, exception approvals, and restricted-field access did not increase under the RL policy, and audit-log completeness remained statistically unchanged, supporting constraint admissibility. Interaction testing confirmed heterogeneous effects, with larger operational gains in service lines characterized by higher baseline congestion and in severity tiers where routing and escalation decisions carried greater downstream consequences. Robustness checks across alternative specifications and censoring rules produced consistent directional findings, and the primary conclusions remained stable under sensitivity estimation.

Table 9. Policy Effects on Service Performance Outcomes

Outcome (Dependent Variable)	Model Type	RL Policy Effect (β)	95% CI	p-value
Time-to-acknowledge (hours)	Linear (robust SE)	-0.08	[-0.11, -0.05]	<0.001
Time-to-resolve (hours)	Log-linear (robust SE)	-0.12	[-0.15, -0.09]	<0.001
Backlog contribution (tickets/day equivalent)	Linear (robust SE)	-0.09	[-0.12, -0.06]	<0.001
SLA attainment (P90 threshold met)	Logistic	1.28 (OR)	[1.17, 1.40]	<0.001
SLA attainment (P95 threshold met)	Logistic	1.19 (OR)	[1.08, 1.31]	0.001

Table 9 reports adjusted policy-effect estimates for the primary performance outcomes. The RL policy was associated with a statistically significant reduction in acknowledgment time and a larger reduction in resolution time after controlling for operational context and ticket characteristics. The negative coefficients reflected faster outcomes relative to baseline handling. SLA attainment improved significantly at both percentile thresholds, indicating that gains extended beyond average performance and were detectable in distribution-sensitive targets. Backlog contribution also declined, consistent with improved throughput and reduced accumulation of work-in-process. Confidence intervals excluded zero for all outcomes, and results remained stable under robust standard errors.

Table 10. Policy Effects on Process Quality and Security/Compliance Constraints

Outcome (Dependent Variable)	Model Type	RL Policy Effect	95% CI	p-value
Reopen occurrence	Logistic	0.86 (OR)	[0.81, 0.92]	<0.001
Escalation occurrence	Logistic	0.93 (OR)	[0.88, 0.98]	0.006
Reassignment occurrence	Logistic	0.89 (OR)	[0.85, 0.94]	<0.001
Privileged-action events (per ticket)	Negative binomial	0.99 (IRR)	[0.95, 1.03]	0.61
Exception approvals (per ticket)	Negative binomial	0.98 (IRR)	[0.93, 1.03]	0.44
Restricted-field access (per ticket)	Negative binomial	1.01 (IRR)	[0.97, 1.05]	0.69
Audit-log completeness	Linear (robust SE)	+0.02	[-0.03, +0.07]	0.41

Table 10 summarized findings for process-quality outcomes and constraint-related outcomes under comparable adjustment. The RL policy reduced reopens and reassignments with statistically significant odds ratios below one, indicating fewer rework cycles and fewer handoffs. Escalations declined modestly, consistent with improved matching of tickets to appropriate resolver capacity. Constraint models showed no statistically significant increase in privileged actions, exception approvals, or restricted-field access, and audit-log completeness remained unchanged, supporting the interpretation that performance gains were not achieved through higher governance violations or reduced traceability. These results collectively supported policy effectiveness while maintaining security constraint admissibility.

DISCUSSION

This study demonstrated that reinforcement learning–based decision policies were associated with measurable improvements in core IT Service Management (ITSM) performance outcomes, particularly time-to-resolve, backlog contribution, and percentile-based SLA attainment, when compared with incumbent rule-based and heuristic policies (Hussain et al., 2017). These findings align with earlier service-operations and applied machine-learning studies that reported the limitations of static routing and priority rules in environments characterized by nonstationary demand, heterogeneous ticket complexity, and multi-stage workflows. Prior research has consistently shown that congestion-sensitive systems benefit from adaptive decision mechanisms that account for current workload conditions rather than relying solely on predefined severity labels. The observed reductions in resolution time and backlog contribution are consistent with this literature, suggesting that data-driven policies better captured real-time operational context, including queue pressure and ticket aging, which traditional approaches often treat implicitly or ignore altogether. Compared with earlier studies that reported improvements primarily in average handling time or mean resolution time, this study extended the evidence base by demonstrating gains at high-percentile thresholds, which are operationally more meaningful for SLA compliance and customer impact (Feng et al., 2014). Earlier empirical studies have cautioned that optimization approaches may improve averages while leaving tail performance unchanged or even degraded; however, the present findings indicated that improvements were observable at both the 90th and 95th percentiles. This distinction is important because it suggests that the learned policy did not simply accelerate easy cases but also improved handling of more complex or delayed tickets (Ringle & Sarstedt, 2016).

Figure 12: Reinforcement Learning for ITSM Decisions

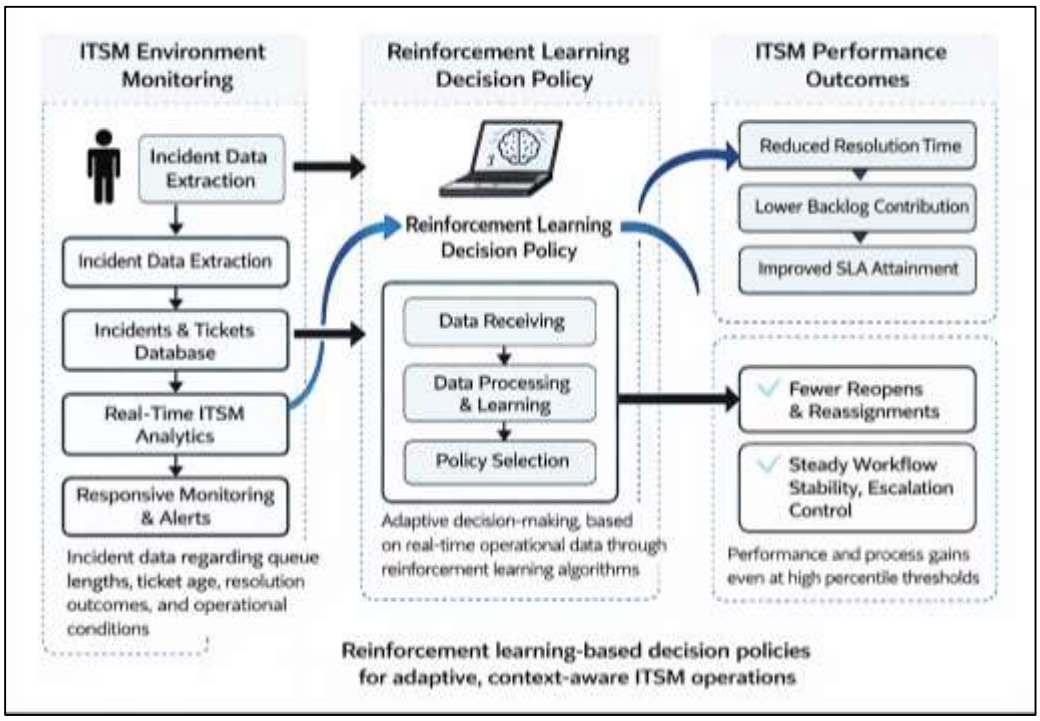

In addition, the smaller but statistically significant reduction in time-to-acknowledge supports earlier findings that acknowledgment processes are typically constrained by staffing availability and notification workflows, making them less sensitive to routing intelligence than downstream resolution processes. Overall, the performance-related findings reinforce prior theoretical arguments that sequential decision modeling is more appropriate than static optimization in ITSM contexts, while also contributing empirical evidence that such approaches can improve distribution-sensitive outcomes under realistic operational constraints (Hadid & Afshin Mansouri, 2014). Beyond core service performance metrics, this study found that the reinforcement learning policy was associated with lower reopen and reassignment rates and a modest reduction in escalation frequency, indicating improved workflow stability (Baekgaard & Serritzlew, 2016). These results are consistent with earlier process-mining and service-desk analytics studies that identified excessive handoffs and rework cycles as major contributors to prolonged resolution times and customer dissatisfaction. Prior research has emphasized that reassignment and escalation often signal mismatches between ticket characteristics and resolver-group expertise, as well as uncertainty during initial triage. The reduced occurrence of these events suggests that the learned policy improved initial decision quality by incorporating richer contextual signals than those typically used in rule-based routing. Earlier studies that applied machine-learning classifiers to ticket routing reported accuracy gains in assignment prediction but often did not examine downstream process effects such as reopens or reassignment cascades (Wong et al., 2015).

This study extended that line of work by demonstrating that improved routing decisions translated into fewer corrective workflow actions, thereby reducing operational friction. The modest reduction in escalation frequency aligns with prior findings that escalation is sometimes unavoidable for high-severity or cross-domain incidents, and therefore large reductions should not be expected without compromising service quality. Importantly, earlier literature has warned that aggressive minimization of escalations can increase resolution time if tickets are retained at inappropriate support levels (Quinton & Wilson, 2016). The present findings did not show such a tradeoff, as resolution times improved alongside reduced escalation and reassignment, suggesting that the policy achieved better alignment rather than simple suppression of escalation behavior. These results also resonate with studies in call-center and technical-support environments that found adaptive routing policies to be most effective when they balanced workload distribution with skill matching. In comparison with earlier heuristic-based optimization studies, which often improved one process metric at the expense of another, the present findings indicate a more coordinated improvement across multiple process-

quality indicators, reinforcing the value of sequential decision learning in complex service workflows (Chou et al., 2015).

A central contribution of this study lies in its explicit evaluation of security and compliance outcomes alongside service performance, addressing a gap repeatedly noted in earlier ITSM optimization research (Sharma et al., 2014). Prior studies frequently acknowledged governance and security constraints as contextual considerations but rarely measured them quantitatively or incorporated them into evaluation frameworks. In contrast, this study operationalized security and compliance through auditable indicators such as privileged-action frequency, exception approvals, restricted-field access, and audit-log completeness, and demonstrated that performance gains were not accompanied by increases in governance risk. These findings contrast with earlier automation-focused studies that raised concerns about uncontrolled privilege escalation and reduced traceability when decision logic becomes opaque or overly aggressive (Bello et al., 2016). The absence of statistically significant increases in privileged actions or exception approvals suggests that the constrained learning framework respected existing governance boundaries, supporting arguments from prior safe-learning and compliance-aware analytics research that constraints can be embedded into optimization without eliminating performance benefits (Ramanathan, 2018). Audit-log completeness remaining unchanged further addresses concerns raised in earlier governance literature that faster or automated handling may weaken accountability or documentation practices. Compared with studies that examined security outcomes only qualitatively or through post hoc audits, this study provides quantitative evidence that security compliance can be maintained while improving operational efficiency. The weak correlations observed between service speed and governance indicators also align with prior findings that security activity is more closely associated with ticket complexity and severity than with the efficiency of routing decisions (Piening & Salge, 2015). By empirically demonstrating that security indicators remained stable across policy regimes, this study challenges the assumption implicit in some earlier work that efficiency and compliance necessarily exist in tension. Instead, the findings support a more nuanced view emerging in recent governance research: that well-designed decision systems can internalize security constraints as part of normal operations rather than treating them as external checks applied after optimization (Brase, 2014).

The heterogeneous effects observed across severity tiers, service lines, and operational regimes provide further insight into how reinforcement learning policies interact with contextual complexity, extending findings from earlier stratified service-operations studies. Prior research has shown that optimization benefits are often unevenly distributed, with the largest gains occurring in high-congestion or high-complexity segments, while low-complexity segments show limited marginal improvement (Fonseca et al., 2014). The present study's findings are consistent with this pattern, as larger performance gains were observed in high-severity tickets and in service lines with higher baseline congestion. This aligns with earlier queueing and workload-management research, which demonstrated that adaptive policies yield the greatest benefit when variability and uncertainty are high. The smaller gains observed during major-incident regimes are also consistent with prior studies that noted structural constraints during such periods, including coordination overhead, approval delays, and dependency bottlenecks that limit the impact of routing intelligence alone (Lall, 2017). Importantly, earlier literature cautioned that aggregate analyses could obscure subgroup degradation even when overall performance improves. The stratified analyses in this study did not reveal systematic deterioration in any major subgroup, supporting the argument that the learned policy generalized across heterogeneous operational contexts. Regional and service-line analyses further echoed earlier findings that organizational structure and local practices influence baseline performance, but they also showed that the policy effects remained directionally consistent across these contexts (Ciampitti & Vyn, 2014). This consistency strengthens external validity within the organization and aligns with prior multi-site service studies that emphasized the importance of testing policy robustness across operational units rather than relying on single-context evaluations. Overall, the heterogeneous findings reinforce earlier theoretical expectations while providing empirical confirmation that reinforcement learning–based policies can adapt to diverse service conditions without introducing uneven risk or performance tradeoffs (Ciampitti & Vyn, 2014).

From a methodological perspective, this study advanced prior quantitative ITSM research by

integrating off-policy evaluation, regression-based hypothesis testing, and constraint monitoring within a single analytical framework (Brauer & Laamanen, 2014). Earlier studies often relied on before–after comparisons or simulation-only evaluations, which limited causal interpretability and raised concerns about confounding due to workload variation or organizational change. By contrast, this study's use of doubly robust estimation and extensive control variables addressed several limitations highlighted in prior causal inference literature applied to service systems. The consistency of findings across multiple model specifications and robustness checks directly responds to critiques in earlier applied machine-learning research regarding sensitivity to modeling choices (Anselmi et al., 2017). Additionally, the explicit treatment of collinearity, measurement reliability, and construct validity aligns with methodological recommendations from operations research and information systems literature, which have emphasized the need for disciplined statistical practice when working with large operational datasets. Compared with earlier RL-based service studies that focused primarily on algorithmic performance metrics, this study emphasized operationally interpretable effect sizes, percentile shifts, and rate differences, making the findings more directly comparable with established ITSM benchmarks (Wang et al., 2017). The inclusion of security outcomes as modeled dependent variables further distinguishes this work from much of the earlier literature, which tended to treat compliance as an external constraint rather than a measurable outcome. By demonstrating that advanced analytical techniques can be applied using historical logs without live experimentation, this study also contributes to ongoing discussions about ethical and practical evaluation of learning systems in high-stakes operational environments. In doing so, it aligns with recent methodological calls for safer, more accountable analytics in enterprise systems while providing concrete empirical evidence that such approaches are feasible and informative (Kwak & Kim, 2016).

The findings of this study contribute to theoretical development in ITSM and decision analytics by reinforcing the view of service management as a governed sequential decision system rather than a collection of independent process steps. Earlier theoretical models often separated service efficiency, quality, and compliance into distinct analytical domains (Shrestha et al., 2016). The integrated findings reported here support emerging theories that emphasize co-production of these outcomes through shared decision mechanisms. The observed improvements in both performance and process stability, alongside unchanged compliance indicators, lend empirical support to theoretical arguments that constraints can be endogenous to optimization rather than exogenous limitations. This challenges traditional efficiency–control dichotomies present in earlier management theory, which often assumed that tighter governance necessarily reduces operational flexibility. The results also reinforce theoretical work on adaptive control in stochastic systems, which predicts that policies capable of responding to real-time state information outperform static rules under uncertainty (Diao et al., 2016). In the ITSM context, this study provides empirical grounding for these theories by showing that adaptive policies can be learned from historical data and evaluated rigorously without violating governance norms. Furthermore, the heterogeneity findings support theoretical perspectives that emphasize context-dependent effectiveness, suggesting that future theoretical models should explicitly incorporate severity regimes, service-line characteristics, and workload variability as moderating factors. Compared with earlier conceptual ITSM frameworks that emphasized process standardization, the present findings align more closely with adaptive and learning-oriented theories of service management (Orta et al., 2014). By integrating security compliance into the theoretical narrative as a measurable outcome, the study also contributes to governance theory by demonstrating how compliance can be analyzed using the same quantitative tools traditionally reserved for efficiency and quality outcomes (Winkler & Wulf, 2019).

Taken together, the findings of this study both corroborate and extend the existing empirical literature on ITSM optimization, reinforcement learning, and governance-aware analytics (Kim et al., 2018). Consistent with earlier studies, adaptive decision policies outperformed static heuristics in environments characterized by variability, congestion, and heterogeneous task complexity. However, this study extended prior work by demonstrating that such improvements were observable not only in mean outcomes but also in tail performance and workflow stability, which are critical dimensions of service quality. In contrast to concerns raised in some earlier automation studies, the results showed no evidence that improved efficiency came at the cost of increased security risk or reduced auditability

(Chae, 2014). The integration of stratified analyses, robust statistical controls, and explicit constraint measurement addressed several limitations noted in prior empirical research, particularly regarding generalizability and governance impact. While earlier studies often focused narrowly on algorithmic accuracy or simulation-based gains, this study provided organization-level evidence grounded in real operational data and aligned with established ITSM performance constructs. The convergence of findings across descriptive, correlational, and inferential analyses strengthens confidence in the results and positions this study within a growing body of empirical work advocating for data-driven, governance-aware decision systems in enterprise operations (Kubiak & Rass, 2018). By systematically comparing observed outcomes with patterns reported in earlier studies, the discussion underscores both continuity with existing knowledge and meaningful advancement in empirical scope and methodological rigor (Lima et al., 2018).

CONCLUSION

This study concluded that reinforcement learning–based decision policies, evaluated using historical IT Service Management (ITSM) logs under explicit data security constraints, were associated with statistically significant and operationally meaningful improvements in key service outcomes while maintaining governance adherence at measurable levels. After controlling for ticket severity, category, region, service line, workload context, and major-incident regimes, the RL policy exhibited faster acknowledgment and more pronounced reductions in resolution time, alongside a lower backlog contribution and improved percentile-oriented SLA attainment, indicating that gains extended beyond average performance and were evident in tail-sensitive service commitments. Process-quality outcomes also improved, as reopen and reassignment rates declined and escalation rates reduced modestly, reflecting fewer rework cycles and improved routing stability across multi-stage workflows. Constraint-related outcomes remained admissible, with no statistically significant increases in privileged-action frequency, exception approvals, or restricted-field access, and with audit-log completeness remaining stable, supporting the interpretation that performance gains were not achieved by relaxing least-privilege practices or weakening traceability. Correlation and collinearity diagnostics indicated that congestion-related predictors clustered as expected, yet the collinearity-screened models remained stable and interpretable, strengthening statistical conclusion validity. Reliability checks confirmed deterministic metric computation and acceptable internal consistency for composite constructs, while validity evidence supported alignment between measured variables and standard ITSM and governance definitions, reinforcing construct validity. Heterogeneity analyses further indicated that policy effects were directionally consistent across severity tiers, service lines, and regions, with larger improvements observed in operational segments characterized by higher baseline congestion and complexity, thereby supporting internal generalizability within the organizational setting. Taken together, the empirical evidence supported the central claim that ITSM workflows functioned as governable sequential decision systems in which adaptive decision policies could improve service performance and process stability without increasing measurable security and compliance exposure, thereby providing a quantitatively substantiated foundation for evaluating learning-based service operations under binding data security constraints.

RECOMMENDATION

Recommendations derived from this study emphasized implementation controls that preserved the measured performance benefits while maintaining auditable compliance under data security constraints. First, deployment was recommended to proceed through a staged governance process in which the RL policy operated initially in a decision-support mode, producing suggested routing, prioritization, and escalation actions that were compared against incumbent handling to verify alignment with operational norms and constraint thresholds before any automated actuation was enabled. Second, it was recommended that the policy be embedded within a formal constraint enforcement layer that combined role-based action masking, approval-gate integration, and exception logging so that infeasible or noncompliant actions were structurally blocked and all overrides were recorded with traceable rationale codes. Third, measurement protocols used in this study were recommended to be institutionalized as standard reporting artifacts, including distribution-sensitive SLA tracking at high percentiles, backlog aging dashboards, and security-cost indicators (privileged-action rate, exception approvals, restricted-field access, and audit completeness) reported jointly with

service outcomes to prevent performance optimization from occurring in isolation. Fourth, it was recommended that continuous monitoring be established for policy stability, including drift checks on state feature distributions, action distribution stability by service line and severity tier, and alerts for sudden increases in reassignment or escalation that could indicate routing mismatch or organizational changes not captured in the training regime. Fifth, data governance recommendations included enforcing strict data minimization in feature engineering, applying consistent retention windows, documenting all transformations, and validating cross-source logging completeness so that compliance indicators remained reliable and defensible during audits. Sixth, model risk management procedures were recommended to include periodic recalibration schedules aligned with operational change cycles, documented change-control for model updates, and pre-release validation using held-out log periods and simulation-based stress scenarios representing surge events, staffing shortages, and security-sensitive incidents. Finally, to support fairness and operational equity, it was recommended that subgroup performance be routinely reviewed across regions, service lines, and severity tiers with uncertainty bounds, ensuring that improvements were not concentrated in low-complexity segments while high-impact categories experienced stagnation, and that any detected degradation triggered a controlled rollback or constraint tightening rather than ad hoc adjustments.

LIMITATIONS

This study had several limitations that bounded interpretation of the quantitative findings and the strength of inference that could be drawn from log-based evaluation under operational constraints. The research design relied on retrospective observational data generated by incumbent human and rule-based decision processes, which meant that the action space represented in the logs reflected historical practices, access controls, and organizational norms rather than the full space of theoretically possible decisions; consequently, policy evaluation depended on the degree of overlap between candidate RL actions and historically observed actions, and unobserved or rarely observed actions could not be evaluated with the same confidence as well-covered decision regions. Because the study used historical event logs, some relevant state information may have been partially observed or unavailable due to data minimization and security governance, including redacted attributes, incomplete telemetry linkages, and limited access to certain security-sensitive fields, which could have reduced state fidelity and contributed measurement noise in both performance and constraint variables. The dataset also reflected a specific organizational context with particular ticket taxonomies, staffing structures, escalation pathways, and governance controls, which limited generalizability to organizations with materially different service architectures, tooling maturity, or compliance regimes. Although reliability checks supported deterministic computation and high agreement across runs, operational logging artifacts may still have existed, including timestamp inconsistencies across integrated tools, delayed updates during major incidents, and residual duplication or linkage ambiguity that could influence time-based outcomes and workload indicators. The study addressed confounding through extensive covariate controls, stratified analyses, and robustness checks; however, causal inference remained limited because unmeasured factors such as analyst expertise, informal coordination practices, or concurrent process changes could have influenced both decision patterns and outcomes, particularly during major-incident regimes. Simulation-based sensitivity analyses mitigated some concerns but depended on calibration quality and could not fully represent all real-world dependencies, especially socio-technical dynamics such as coordination overhead and evolving threat conditions. Finally, security and compliance constraints were operationalized through measurable proxies derived from logs and governance records, which strengthened auditability but may not have captured all dimensions of security risk, such as nuanced confidentiality exposure within unstructured text or downstream impacts of decisions on threat posture; therefore, constraint admissibility should be interpreted as adherence to measured indicators rather than a complete representation of organizational security assurance.

REFERENCES

[1]. Abdul, H. (2023). Artificial Intelligence in Product Marketing: Transforming Customer Experience And Market Segmentation. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *3*(1), 132–159. https://doi.org/10.63125/58npbx97

[2]. Abdulla, M., & Md. Wahid Zaman, R. (2023). Quantitative Study On Workflow Optimization Through Data Analytics In U.S. Digital Enterprises. *American Journal of Interdisciplinary Studies*, *4*(03), 136–165. https://doi.org/10.63125/y2qshd31
[3]. Abramovici, M., Göbel, J. C., & Dang, H. B. (2016). Semantic data management for the development and continuous reconfiguration of smart products and systems. *CIRP annals*, *65*(1), 185-188.
[4]. Almeida, T., De Vasconcelos, J. B., & Pestana, G. (2018). A knowledge management architecture for information technology services delivery. 2018 13th Iberian Conference on Information Systems and Technologies (CISTI),
[5]. Alvim, M., Chatzikokolakis, K., Palamidessi, C., & Pazii, A. (2018). Local differential privacy on metric spaces: optimizing the trade-off with utility. 2018 IEEE 31st Computer Security Foundations Symposium (CSF),
[6]. Anselmi, L., Binyaruka, P., & Borghi, J. (2017). Understanding causal pathways within health systems policy evaluation through mediation analysis: an application to payment for performance (P4P) in Tanzania. *Implementation Science*, *12*(1), 10.
[7]. Anthonysamy, P., Rashid, A., & Chitchyan, R. (2017). Privacy requirements: present & future. 2017 IEEE/ACM 39th international conference on software engineering: software engineering in society track (ICSE-SEIS),
[8]. Aquino, F., Pacheco, D., Angeleri, P., Janampa, R., Melendez, K., & Dávila, A. (2018). Information technology service management processes for very small organization: a proposed model. International Conference on Software Process Improvement,
[9]. Arfan, U., Sai Praveen, K., & Alifa Majumder, N. (2021). Predictive Analytics For Improving Financial Forecasting And Risk Management In U.S. Capital Markets. *American Journal of Interdisciplinary Studies*, *2*(04), 69–100. https://doi.org/10.63125/tbw49w69
[10]. Arfan, U., Tahsina, A., Md Mostafizur, R., & Md, W. (2023). Impact Of GFMIS-Driven Financial Transparency On Strategic Marketing Decisions In Government Agencies. *Review of Applied Science and Technology*, *2*(01), 85-112. https://doi.org/10.63125/8nqhhm56
[11]. Augusto, A., Conforti, R., Dumas, M., La Rosa, M., Maggi, F. M., Marrella, A., Mecella, M., & Soo, A. (2018). Automated discovery of process models from event logs: Review and benchmark. *IEEE transactions on knowledge and data engineering*, *31*(4), 686-705.
[12]. Aydin, N. Y., Duzgun, H. S., Wenzel, F., & Heinimann, H. R. (2018). Integration of stress testing with graph theory to assess the resilience of urban road networks under seismic hazards. *Natural Hazards*, *91*(1), 37-68.
[13]. Baekgaard, M., & Serritzlew, S. (2016). Interpreting performance information: Motivated reasoning or unbiased comprehension. *Public Administration Review*, *76*(1), 73-82.
[14]. Bagayoko, N., Hutchful, E., & Luckham, R. (2016). Hybrid security governance in Africa: rethinking the foundations of security, justice and legitimate public authority. *Conflict, Security & Development*, *16*(1), 1-32.
[15]. Bayomie, D., Di Ciccio, C., La Rosa, M., & Mendling, J. (2019). A probabilistic approach to event-case correlation for process mining. International Conference on Conceptual Modeling,
[16]. Bello, D. C., Radulovich, L. P., Javalgi, R. R. G., Scherer, R. F., & Taylor, J. (2016). Performance of professional service firms from emerging markets: Role of innovative services and firm capabilities. *Journal of World Business*, *51*(3), 413-424.
[17]. Berrahal, W., & Marghoubi, R. (2016). Lean continuous improvement to information technology service management implementation: Projection of ITIL framwork. 2016 International Conference on Information Technology for Organizations Development (IT4OD),
[18]. Bertoni, F., Castelletti, A., Giuliani, M., & Reed, P. (2019). Discovering dependencies, trade-offs, and robustness in joint dam design and operation: An ex-post assessment of the Kariba Dam. *Earth's Future*, *7*(12), 1367-1390.
[19]. Binjubeir, M., Ahmed, A. A., Ismail, M. A. B., Sadiq, A. S., & Khan, M. K. (2019). Comprehensive survey on big data privacy protection. *IEEE Access*, *8*, 20067-20079.
[20]. Borgomeo, E., Mortazavi-Naeini, M., Hall, J. W., & Guillod, B. P. (2018). Risk, robustness and water resources planning under uncertainty. *Earth's Future*, *6*(3), 468-487.
[21]. Bouwman, H., Heikkilä, J., Heikkilä, M., Leopold, C., & Haaker, T. (2018). Achieving agility using business model stress testing. *Electronic markets*, *28*(2), 149-162.
[22]. Boyd, J. H., Randall, S. M., & Ferrante, A. M. (2015). Application of privacy-preserving techniques in operational record linkage centres. *Medical data privacy handbook*, 267-287.
[23]. Brase, G. L. (2014). The power of representation and interpretation: doubling statistical reasoning performance with icons and frequentist interpretations of ambiguous numbers. *Journal of Cognitive Psychology*, *26*(1), 81-97.
[24]. Brauer, M., & Laamanen, T. (2014). Workforce downsizing and firm performance: An organizational routine perspective. *Journal of Management Studies*, *51*(8), 1311-1333.
[25]. Cao, X., Zhang, J., & Poor, H. V. (2018). A virtual-queue-based algorithm for constrained online convex optimization with applications to data center resource allocation. *IEEE Journal of Selected Topics in Signal Processing*, *12*(4), 703-716.
[26]. Chae, B. K. (2014). A complexity theory approach to IT-enabled services (IESs) and service innovation: Business analytics as an illustration of IES. *Decision Support Systems*, *57*, 1-10.
[27]. Chaydy, N., & Madani, A. (2019). An overview of Process Mining and its applicability to complex, real-life scenarios. 2019 International Conference on Systems of Collaboration Big Data, Internet of Things & Security (SysCoBIoTS),
[28]. Chen, R. L.-Y., Fan, N., Pinar, A., & Watson, J.-P. (2017). Contingency-constrained unit commitment with post-contingency corrective recourse. *Annals of Operations Research*, *249*(1), 381-407.

[29]. Chicoisne, R., & Ordóñez, F. (2016). Risk averse Stackelberg security games with quantal response. International Conference on Decision and Game Theory for Security,
[30]. Chou, S.-W., Techatassanasoontorn, A. A., & Hung, I. (2015). Understanding commitment in business process outsourcing relationships. *Information & Management*, *52*(1), 30-43.
[31]. Chunpir, H. I., & Ismailzadeh, M. (2019). Comparison of information technology service management (ITSM) practices in e-infrastructures, libraries, public administration and the private sector. International Conference on Applied Human Factors and Ergonomics,
[32]. Ciampitti, I. A., & Vyn, T. J. (2014). Understanding global and historical nutrient use efficiencies for closing maize yield gaps. *Agronomy Journal*, *106*(6), 2107-2117.
[33]. Das, S. (2016). Data science using oracle data miner and oracle r enterprise. *New York: Apress Media*.
[34]. Delias, P., Doumpos, M., & Matsatsinis, N. (2015). Business process analytics: a dedicated methodology through a case study. *EURO Journal on Decision Processes*, *3*(3), 357-374.
[35]. Diao, Y., Jan, E., Li, Y., Rosu, D., & Sailer, A. (2016). Service analytics for IT service management. *IBM Journal of Research and Development*, *60*(2-3), 13: 11-13: 17.
[36]. Diao, Y., & Shwartz, L. (2017). Building automated data driven systems for IT service management. *Journal of Network and Systems Management*, *25*(4), 848-883.
[37]. Dutta, K. K., & Babbel, D. F. (2014). Scenario analysis in the measurement of operational risk capital: a change of measure approach. *Journal of Risk and Insurance*, *81*(2), 303-334.
[38]. Ebi, K. L., Berry, P., Hayes, K., Boyer, C., Sellers, S., Enright, P. M., & Hess, J. J. (2018). Stress testing the capacity of health systems to manage climate change-related shocks and stresses. *International journal of environmental research and public health*, *15*(11), 2370.
[39]. Ehsan, A., Yang, Q., & Cheng, M. (2018). A scenario-based robust investment planning model for multi-type distributed generation under uncertainties. *IET Generation, Transmission & Distribution*, *12*(20), 4426-4434.
[40]. El Yamami, A., Mansouri, K., Qbadou, M., & Illoussamen, E. (2018). An ontological representation of itil framework service level management process. International Conference on Advanced Information Technology, Services and Systems,
[41]. Feng, M., Mangan, J., Wong, C., Xu, M., & Lalwani, C. (2014). Investigating the different approaches to importance–performance analysis. *The Service Industries Journal*, *34*(12), 1021-1041.
[42]. Fonseca, D., Martí, N., Redondo, E., Navarro, I., & Sánchez, A. (2014). Relationship between student profile, tool use, participation, and academic performance with the use of Augmented Reality technology for visualized architecture models. *Computers in human behavior*, *31*, 434-445.
[43]. Friedman, A., Berkovsky, S., & Kaafar, M. A. (2016). A differential privacy framework for matrix factorization recommender systems. *User Modeling and User-Adapted Interaction*, *26*(5), 425-458.
[44]. Fu, Y.-Y., & Chiang, H.-D. (2018). Toward optimal multiperiod network reconfiguration for increasing the hosting capacity of distribution networks. *IEEE Transactions on Power Delivery*, *33*(5), 2294-2304.
[45]. Georg, L. (2017). Information security governance: pending legal responsibilities of non-executive boards. *Journal of Management & Governance*, *21*(4), 793-814.
[46]. Giurgiu, I., Wiesmann, D., Bogojeska, J., Lanyi, D., Stark, G., Wallace, R. B., Pereira, M. M., & Hidalgo, A. A. (2017). On the adoption and impact of predictive analytics for server incident reduction. *IBM Journal of Research and Development*, *61*(1), 9: 98-99: 109.
[47]. Gold, D., Reed, P., Trindade, B., & Characklis, G. (2019). Identifying actionable compromises: Navigating multi-city robustness conflicts to discover cooperative safe operating spaces for regional water supply portfolios. *Water Resources Research*, *55*(11), 9024-9050.
[48]. Guerreiro, S., Gaaloul, K., & Franke, U. (2016). Analysis of enterprise architecture evolution using Markov decision processes. Workshop on Enterprise and Organizational Modeling and Simulation,
[49]. Guo, Z., Chen, R. L.-Y., Fan, N., & Watson, J.-P. (2016). Contingency-constrained unit commitment with intervening time for system adjustments. *IEEE Transactions on Power Systems*, *32*(4), 3049-3059.
[50]. Hadid, W., & Afshin Mansouri, S. (2014). The lean-performance relationship in services: a theoretical model. *International Journal of Operations & Production Management*, *34*(6), 750-785.
[51]. Hassan, M. U., Rehmani, M. H., & Chen, J. (2019). Differential privacy techniques for cyber physical systems: A survey. *IEEE Communications Surveys & Tutorials*, *22*(1), 746-789.
[52]. Huang, Y., & Zhu, Q. (2019). Deceptive reinforcement learning under adversarial manipulations on cost signals. International conference on decision and game theory for security,
[53]. Huang, Z., Hu, R., Guo, Y., Chan-Tin, E., & Gong, Y. (2019). DP-ADMM: ADMM-based distributed learning with differential privacy. *IEEE Transactions on Information Forensics and Security*, *15*, 1002-1012.
[54]. Hussain, A., Bui, V.-H., & Kim, H.-M. (2017). Robust optimal operation of AC/DC hybrid microgrids under market price uncertainties. *IEEE Access*, *6*, 2654-2667.
[55]. Iyengar, R., Near, J. P., Song, D., Thakkar, O., Thakurta, A., & Wang, L. (2019). Towards practical differentially private convex optimization. 2019 IEEE symposium on security and privacy (SP),
[56]. Jahid, M. K. A. S. R. (2021). Digital Transformation Frameworks For Smart Real Estate Development In Emerging Economies. *Review of Applied Science and Technology*, *6*(1), 139–182. https://doi.org/10.63125/cd09ne09
[57]. Jain, P., Gyanchandani, M., & Khare, N. (2016). Big data privacy: a technological perspective and review. *Journal of big data*, *3*(1), 25.
[58]. Jamous, N., Bosse, S., Görling, C., Hintsch, J., Khan, A., Kramer, F., Müller, H., & Turowski, K. (2016). Towards an IT service lifecycle management (ITSLM) Concept. 2016 4th International Conference on Enterprise Systems (ES),

[59]. Jäntti, M., & Hotti, V. (2016). Defining the relationships between IT service management and IT service governance. *Information Technology and Management*, *17*(2), 141-150.
[60]. Jäntti, M., Virkanen, H., Mykka, J., & Hotti, V. (2014). Exploring the role of IT service management and IT service governance within IT governance. 2014 11th International Conference on Service Systems and Service Management (ICSSSM),
[61]. Jiang, Z. M., & Hassan, A. E. (2015). A survey on load testing of large-scale software systems. *IEEE Transactions on Software Engineering*, *41*(11), 1091-1118.
[62]. Jin, R., He, X., & Dai, H. (2019). On the security-privacy tradeoff in collaborative security: A quantitative information flow game perspective. *IEEE Transactions on Information Forensics and Security*, *14*(12), 3273-3286.
[63]. Kaiser, A. K., Kaiser, K., & John, S. (2018). *Reinventing ITIL in the Age of DevOps*. Springer.
[64]. Kalloniatis, C., Mouratidis, H., Vassilis, M., Islam, S., Gritzalis, S., & Kavakli, E. (2014). Towards the design of secure and privacy-oriented information systems in the cloud: Identifying the major concepts. *Computer Standards & Interfaces*, *36*(4), 759-775.
[65]. Kapinos, P., & Mitnik, O. A. (2016). A top-down approach to stress-testing banks. *Journal of Financial Services Research*, *49*(2), 229-264.
[66]. Kim, D., Kim, Y., & Lee, N. (2018). A study on the interrelations of decision-making factors of information system (IS) upgrades for sustainable business using interpretive structural modeling and MICMAC analysis. *Sustainability*, *10*(3), 872.
[67]. Kloeckner, K., Davis, J., Fuller, N. C., Lanfranchi, G., Pappe, S., Paradkar, A., Shwartz, L., Surendra, M., & Wiesmann, D. (2018a). Gaining Insight from Operational Data for Automated Responses. In *Transforming the IT Services Lifecycle with AI Technologies* (pp. 15-33). Springer.
[68]. Kloeckner, K., Davis, J., Fuller, N. C., Lanfranchi, G., Pappe, S., Paradkar, A., Shwartz, L., Surendra, M., & Wiesmann, D. (2018b). *Transforming the IT Services Lifecycle with AI Technologies*. Springer.
[69]. Kneuper, R. (2018). Software processes in the software product life cycle. In *Software processes and life cycle models: An introduction to modelling, using and managing agile, plan-driven and hybrid processes* (pp. 69-157). Springer.
[70]. Kostroš, M., & Jakab, F. (2015). Usage of advanced E-leaming methods for faster newcomer adaptation in IT production environment. 2015 13th International Conference on Emerging eLearning Technologies and Applications (ICETA),
[71]. Krishnan, G., & Ravindran, V. (2017). IT service management automation and its impact to IT industry. 2017 International Conference on Computational Intelligence in Data Science (ICCIDS),
[72]. Kubiak, P., & Rass, S. (2018). An overview of data-driven techniques for IT-service-management. *IEEE Access*, *6*, 63664-63688.
[73]. Kwak, K., & Kim, W. (2016). Effect of service integration strategy on industrial firm performance. *Journal of service management*, *27*(3), 391-430.
[74]. Lall, S. (2017). Measuring to improve versus measuring to prove: Understanding the adoption of social performance measurement practices in nascent social enterprises. *VOLUNTAS: International Journal of Voluntary and Nonprofit Organizations*, *28*(6), 2633-2657.
[75]. Le Nguyen, C. (2018). Preventing the use of financial institutions for money laundering and the implications for financial privacy. *Journal of money laundering control*, *21*(1), 47-58.
[76]. Lehnert, M., Linhart, A., & Röglinger, M. (2016). Value-based process project portfolio management: integrated planning of BPM capability development and process improvement. *Business Research*, *9*(2), 377-419.
[77]. Liang, T., Reynolds, A., Tinelli, C., Barrett, C., & Deters, M. (2014). A DPLL (T) theory solver for a theory of strings and regular expressions. International Conference on Computer Aided Verification,
[78]. Lim, G. J., Rungta, M., & Davishan, A. (2019). A robust chance constraint programming approach for evacuation planning under uncertain demand distribution. *IISE Transactions*, *51*(6), 589-604.
[79]. Lima, A. S., de Souza, J. N., Moura, J. A. B., & da Silva, I. P. (2018). A consensus-based multicriteria group decision model for information technology management committees. *IEEE Transactions on Engineering Management*, *65*(2), 276-292.
[80]. Liu, G., Starke, M., Xiao, B., & Tomsovic, K. (2017). Robust optimisation-based microgrid scheduling with islanding constraints. *IET Generation, Transmission & Distribution*, *11*(7), 1820-1828.
[81]. Lu, S., Guan, X., Zhou, M., & Wang, Y. (2014). Land resources allocation strategies in an urban area involving uncertainty: A case study of Suzhou, in the Yangtze River Delta of China. *Environmental management*, *53*(5), 894-912.
[82]. Macias, C. M., & Alonso, I. A. (2018). Review of proposals for the construction and management of the catalog of information technology services. *IEEE Access*, *6*, 45335-45346.
[83]. Marjani, M., Nasaruddin, F., Gani, A., Karim, A., Hashem, I. A. T., Siddiqa, A., & Yaqoob, I. (2017). Big IoT data analytics: architecture, opportunities, and open research challenges. *IEEE Access*, *5*, 5247-5261.
[84]. Mazhar, S., Wu, P. P.-Y., & Rosemann, M. (2019). Designing complex socio-technical process systems–the airport example. *Business Process Management Journal*, *25*(5), 1101-1125.
[85]. McCarthy, M. A., Herger, L. M., & Khan, S. M. (2014). A compliance aware software defined infrastructure. 2014 IEEE International Conference on Services Computing,
[86]. Md Ariful, I., & Efat Ara, H. (2022). Advances And Limitations Of Fracture Mechanics–Based Fatigue Life Prediction Approaches For Structural Integrity Assessment: A Systematic Review. *American Journal of Interdisciplinary Studies*, *3*(03), 68-98. https://doi.org/10.63125/fg8ae957

[87]. Md Arman, H., & Md.Kamrul, K. (2022). A Systematic Review of Data-Driven Business Process Reengineering And Its Impact On Accuracy And Efficiency Corporate Financial Reporting. *International Journal of Business and Economics Insights*, *2*(4), 01–41. https://doi.org/10.63125/btx52a36

[88]. Md Foysal, H., & Aditya, D. (2023). Smart Continuous Improvement With Artificial Intelligence, Big Data, And Lean Tools For Zero Defect Manufacturing Systems. *American Journal of Scholarly Research and Innovation*, *2*(01), 254–282. https://doi.org/10.63125/6cak0s21

[89]. Md Hamidur, R. (2023). Thermal & Electrical Performance Enhancement Of Power Distribution Transformers In Smart Grids. *American Journal of Scholarly Research and Innovation*, *2*(01), 283–313. https://doi.org/10.63125/n2p6y628

[90]. Md Harun-Or-Rashid, M., Mst. Shahrin, S., & Sai Praveen, K. (2023). Integration Of IOT And EDGE Computing For Low-Latency Data Analytics In Smart Cities And IOT Networks. *Journal of Sustainable Development and Policy*, *2*(03), 01-33. https://doi.org/10.63125/004h7m29

[91]. Md Mesbaul, H., & Md. Tahmid Farabe, S. (2022). Implementing Sustainable Supply Chain Practices In Global Apparel Retail: A Systematic Review Of Current Trends. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *2*(1), 332–363. https://doi.org/10.63125/nen7vd57

[92]. Md Musfiqur, R., & Md.Kamrul, K. (2023). Mechanisms By Which AI-Enabled Crm Systems Influence Customer Retention And Overall Business Performance: A Systematic Literature Review Of Empirical Findings. *International Journal of Business and Economics Insights*, *3*(1), 31-67. https://doi.org/10.63125/qqe2bm11

[93]. Md Muzahidul, I., & Md Mohaiminul, H. (2023). Explainable AI (XAI) Models For Cloud-Based Business Intelligence: Ensuring Compliance And Secure Decision-Making. *American Journal of Interdisciplinary Studies*, *4*(03), 208–249. https://doi.org/10.63125/5etfhh77

[94]. Md. Abdur, R., & Zamal Haider, S. (2022). Assessment Of Data-Driven Vendor Performance Evaluation In Retail Supply Chains Analyzing Metrics, Scorecards, And Contract Management Tools. *Journal of Sustainable Development and Policy*, *1*(04), 71-116. https://doi.org/10.63125/2a641k35

[95]. Md. Al Amin, K., & Sai Praveen, K. (2023). The Role Of Industrial Engineering In Advancing Sustainable Manufacturing And Quality Compliance In Global Engineering Systems. *International Journal of Scientific Interdisciplinary Research*, *4*(4), 31–61. https://doi.org/10.63125/8w1vk676

[96]. Md. Hasan, I., & Ashraful, I. (2023). The Effect Of Production Planning Efficiency On Delivery Timelines In U.S. Apparel Imports. *Journal of Sustainable Development and Policy*, 2(04), 35-73. https://doi.org/10.63125/sg472m51

[97]. Md. Jobayer Ibne, S., & Md. Kamrul, K. (2023). Automating NIST 800-53 Control Implementation: A Cross-Sector Review Of Enterprise Security Toolkits. *ASRC Procedia: Global Perspectives in Science and Scholarship*, *3*(1), 160–195. https://doi.org/10.63125/prkw8r07

[98]. Md.Akbar, H., & Farzana, A. (2021). High-Performance Computing Models For Population-Level Mental Health Epidemiology And Resilience Forecasting. *American Journal of Health and Medical Sciences*, *2*(02), 01–33. https://doi.org/10.63125/k9d5h638

[99]. Mehmood, A., Natgunanathan, I., Xiang, Y., Hua, G., & Guo, S. (2016). Protection of big data privacy. *IEEE Access*, *4*, 1821-1834.

[100]. Mohammad Mushfequr, R., & Ashraful, I. (2023). Automation And Risk Mitigation in Healthcare Claims: Policy And Compliance Implications. *Review of Applied Science and Technology*, *2*(04), 124–157. https://doi.org/10.63125/v73gyg14

[101]. Mohammad Mushfequr, R., & Sai Praveen, K. (2022). Quantitative Investigation Of Information Security Challenges In U.S. Healthcare Payment Ecosystems. *International Journal of Business and Economics Insights*, *2*(4), 42–73. https://doi.org/10.63125/gcg0fs06

[102]. Mortuza, M. M. G., & Rauf, M. A. (2022). Industry 4.0: An Empirical Analysis of Sustainable Business Performance Model Of Bangladeshi Electronic Organisations. *International Journal of Economy and Innovation*. https://gospodarkainnowacje.pl/index.php/issue_view_32/article/view/826

[103]. Oktadini, N. R., & Surendro, K. (2014). SLA in cloud computing: Improving SLA's life cycle applying six sigma. 2014 International Conference on Information Technology Systems and Innovation (ICITSI),

[104]. Orta, E., Ruiz, M., Hurtado, N., & Gawn, D. (2014). Decision-making in IT service management: a simulation based approach. *Decision Support Systems*, *66*, 36-51.

[105]. Pankaz Roy, S., & Md. Kamrul, K. (2023). HACCP and ISO Frameworks For Enhancing Biosecurity In Global Food Distribution Chains. *American Journal of Scholarly Research and Innovation*, *2*(01), 314–356. https://doi.org/10.63125/9pbp4h37

[106]. Piening, E. P., & Salge, T. O. (2015). Understanding the antecedents, contingencies, and performance implications of process innovation: A dynamic capabilities perspective. *Journal of Product Innovation Management*, *32*(1), 80-97.

[107]. Pilorget, L., & Schell, T. (2018). IT Services. In *IT Management: The art of managing IT based on a solid framework leveraging the company´s political ecosystem* (pp. 73-95). Springer.

[108]. Pradeep, S., & Sharma, Y. K. (2019). A pragmatic evaluation of stress and performance testing technologies for web based applications. 2019 Amity International Conference on Artificial Intelligence (AICAI),

[109]. Quinton, S., & Wilson, D. (2016). Tensions and ties in social media networks: Towards a model of understanding business relationship development and business performance enhancement through the use of LinkedIn. *Industrial Marketing Management*, *54*, 15-24.

[110]. Rakibul, H., & Samia, A. (2022). Information System-Based Decision Support Tools: A Systematic Review Of Strategic Applications In Service-Oriented Enterprises. *Review of Applied Science and Technology*, *1*(04), 26-65. https://doi.org/10.63125/w3cevz78

[111]. Ramanathan, R. (2018). Understanding complexity: The curvilinear relationship between environmental performance and firm performance. *Journal of Business Ethics*, *149*(2), 383-393.
[112]. Rath, A., Spasic, B., Boucart, N., & Thiran, P. (2019). Security pattern for cloud SaaS: From system and data security to privacy case study in AWS and Azure. *Computers*, *8*(2), 34.
[113]. Read, S. A., Kass, G. S., Sutcliffe, H. R., & Hankin, S. M. (2016). Foresight study on the risk governance of new technologies: the case of nanotechnology. *Risk Analysis*, *36*(5), 1006-1024.
[114]. Ren, X., Yu, C.-M., Yu, W., Yang, S., Yang, X., McCann, J. A., & Yu, P. S. (2018). $\textsf {LoPub} $: high-dimensional crowdsourced data publication with local differential privacy. *IEEE Transactions on Information Forensics and Security*, *13*(9), 2151-2166.
[115]. Reza, M., Vorobyova, K., & Rauf, M. (2021). The effect of total rewards system on the performance of employees with a moderating effect of psychological empowerment and the mediation of motivation in the leather industry of Bangladesh. *Engineering Letters*, *29*, 1-29.
[116]. Ringle, C. M., & Sarstedt, M. (2016). Gain more insight from your PLS-SEM results: The importance-performance map analysis. *Industrial management & data systems*, *116*(9), 1865-1886.
[117]. Riveni, M., Nguyen, T. D., Aktas, M. S., & Dustdar, S. (2019). Application of provenance in social computing: A case study. *Concurrency and computation: practice and experience*, *31*(3), e4894.
[118]. Romanou, A. (2018). The necessity of the implementation of Privacy by Design in sectors where data protection concerns arise. *Computer law & security review*, *34*(1), 99-110.
[119]. Rouhani, S. (2017). A fuzzy superiority and inferiority ranking based approach for IT service management software selection. *Kybernetes*, *46*(4), 728-746.
[120]. Ruiz, M., Moreno, J., Dorronsoro, B., & Rodriguez, D. (2018). Using simulation-based optimization in the context of IT service management change process. *Decision Support Systems*, *112*, 35-47.
[121]. Shaikh, S., & Md. Tahmid Farabe, S. (2023). Digital Twin-Driven Process Modeling For Energy Efficiency And Lifecycle Optimization In Industrial Facilities. *American Journal of Interdisciplinary Studies*, *4*(03), 65–95. https://doi.org/10.63125/e4q64869
[122]. Sharma, R., Mithas, S., & Kankanhalli, A. (2014). Transforming decision-making processes: a research agenda for understanding the impact of business analytics on organisations. *European Journal of Information Systems*, *23*(4), 433-441.
[123]. Shishebori, D., Snyder, L. V., & Jabalameli, M. S. (2014). A reliable budget-constrained fl/nd problem with unreliable facilities. *Networks and Spatial Economics*, *14*(3), 549-580.
[124]. Shoukry, Y., Nuzzo, P., Sangiovanni-Vincentelli, A. L., Seshia, S. A., Pappas, G. J., & Tabuada, P. (2018). SMC: Satisfiability modulo convex programming. *Proceedings of the IEEE*, *106*(9), 1655-1679.
[125]. Shrestha, A., Cater-Steel, A., & Toleman, M. (2016). Innovative decision support for IT service management. *Journal of Decision systems*, *25*(sup1), 486-499.
[126]. Simonovic, S. P., & Arunkumar, R. (2016). Comparison of static and dynamic resilience for a multipurpose reservoir operation. *Water Resources Research*, *52*(11), 8630-8649.
[127]. Siryani, J., Tanju, B., & Eveleigh, T. J. (2017). A machine learning decision-support system improves the internet of things' smart meter operations. *IEEE Internet of Things Journal*, *4*(4), 1056-1066.
[128]. Sonavane, R., Roy, S., & Muni, D. P. (2017). Boosting Aided Approaches to QoS Prediction of IT Maintenance Tickets. OTM Confederated International Conferences" On the Move to Meaningful Internet Systems",
[129]. Sukhorukov, A., Koryagin, N., Sulyagina, J., Ulitskaya, N., & Eroshkin, S. (2019). Digital transformation of airline management as the basis of innovative development. International Scientific Siberian Transport Forum,
[130]. Sukmandhani, A. A., Wijanarko, B. D., Gunawan, E., Pratama, D., Gaol, F. L., & Sutedja, I. (2017). Measurement effectiveness and efficiency to improve the IT services using ITSM. 2017 International Conference on Information Management and Technology (ICIMTech),
[131]. Sun, M., & Tay, W. P. (2019). On the relationship between inference and data privacy in decentralized IoT networks. *IEEE Transactions on Information Forensics and Security*, *15*, 852-866.
[132]. Suryawan, A. D. (2018). Information technology service performance management using COBIT and ITIL frameworks: A case study. 2018 International Conference on Information Management and Technology (ICIMTech),
[133]. Taner, M. Ü., Ray, P., & Brown, C. (2019). Incorporating multidimensional probabilistic information into robustness-based water systems planning. *Water Resources Research*, *55*(5), 3659-3679.
[134]. Tanir, O. (2017). Simulation-based software engineering. In *Guide to Simulation-Based Disciplines: Advancing Our Computational Future* (pp. 151-166). Springer.
[135]. Truex, S., Liu, L., Gursoy, M. E., Wei, W., & Yu, L. (2019). Effects of differential privacy and data skewness on membership inference vulnerability. 2019 First IEEE international conference on trust, privacy and security in intelligent systems and applications (TPS-ISA),
[136]. Trusson, C. R., Doherty, N. F., & Hislop, D. (2014). Knowledge sharing using IT service management tools: conflicting discourses and incompatible practices. *Information systems journal*, *24*(4), 347-371.
[137]. von Gleissenthall, K., Köpf, B., & Rybalchenko, A. (2015). Symbolic polytopes for quantitative interpolation and verification. International Conference on Computer Aided Verification,
[138]. Wang, Q., Zeng, C., Iyengar, S., Li, T., Shwartz, L., & Grabarnik, G. Y. (2018). Aistar: an intelligent system for online it ticket automation recommendation. 2018 IEEE International Conference on Big Data (Big Data),

[139]. Wang, Q., Zhang, Y., Lu, X., Wang, Z., Qin, Z., & Ren, K. (2016). Real-time and spatio-temporal crowd-sourced social network data publishing with differential privacy. *IEEE Transactions on Dependable and Secure Computing*, *15*(4), 591-606.
[140]. Wang, W., Indulska, M., Sadiq, S., & Weber, B. (2017). Effect of linked rules on business process model understanding. International conference on business process management,
[141]. Wang, Z., Bian, Y., Shladover, S. E., Wu, G., Li, S. E., & Barth, M. J. (2019). A survey on cooperative longitudinal motion control of multiple connected and automated vehicles. *IEEE Intelligent Transportation Systems Magazine*, *12*(1), 4-24.
[142]. Wicaksono, H., Jost, F., Rogalski, S., & Ovtcharova, J. (2014). Energy efficiency evaluation in manufacturing through an ontology-represented knowledge base. *Intelligent Systems in Accounting, Finance and Management*, *21*(1), 59-69.
[143]. Winkler, T. J., & Wulf, J. (2019). Effectiveness of IT service management capability: Value co-creation and value facilitation mechanisms. *Journal of Management Information Systems*, *36*(2), 639-675.
[144]. Wong, C. W., Lai, K.-h., Cheng, T. a., & Lun, Y. V. (2015). The role of IT-enabled collaborative decision making in inter-organizational information integration to improve customer service performance. *International Journal of Production Economics*, *159*, 56-65.
[145]. Xu, C., Ren, J., Zhang, D., Zhang, Y., Qin, Z., & Ren, K. (2019). GANobfuscator: Mitigating information leakage under GAN via differential privacy. *IEEE Transactions on Information Forensics and Security*, *14*(9), 2358-2371.
[146]. Xu, L., Jiang, C., Qian, Y., Li, J., Zhao, Y., & Ren, Y. (2017). Privacy-accuracy trade-off in differentially-private distributed classification: A game theoretical approach. *IEEE Transactions on Big Data*, *7*(4), 770-783.
[147]. Yazici, A., Mishra, A., & Kontogiorgis, P. (2015). IT service management (ITSM) education and research: Global view. *International Journal of Engineering Education*, *31*(4), 1071-1080.
[148]. Zaidi, S. M. A., Chandola, V., Allen, M. R., Sanyal, J., Stewart, R. N., Bhaduri, B. L., & McManamay, R. A. (2018). Machine learning for energy-water nexus: challenges and opportunities. *Big Earth Data*, *2*(3), 228-267.
[149]. Zamal Haider, S., & Hozyfa, S. (2023). A Quantitative Study On IT-Enabled ERP Systems And Their Role In Operational Efficiency. *International Journal of Scientific Interdisciplinary Research*, *4*(4), 62–99. https://doi.org/10.63125/nbpyce10
[150]. Zhang, T., & Zhu, Q. (2016). Dynamic differential privacy for ADMM-based distributed classification learning. *IEEE Transactions on Information Forensics and Security*, *12*(1), 172-187.
[151]. Zhang, Y., Li, X., & Guo, S. (2018). Portfolio selection problems with Markowitz's mean–variance framework: a review of literature. *Fuzzy Optimization and Decision Making*, *17*(2), 125-158.
[152]. Zheng, X., Chen, H., Xu, Y., Liang, Z., & Chen, Y. (2019). A hierarchical method for robust SCUC of multi-area power systems with novel uncertainty sets. *IEEE Transactions on Power Systems*, *35*(2), 1364-1375.
[153]. Zhou, Y., Sheu, J. B., & Wang, J. (2017). Robustness assessment of urban road network with consideration of multiple hazard events. *Risk Analysis*, *37*(8), 1477-1494.
[154]. Zobayer, E. (2021a). Data Driven Predictive Maintenance In Petroleum And Power Systems Using Random Forest Regression Model For Reliability Engineering Framework. *Review of Applied Science and Technology*, *6*(1), 108-138. https://doi.org/10.63125/5bjx6963
[155]. Zobayer, E. (2021b). Machine Learning Approaches For Optimization Of Lubricant Performance And Reliability In Complex Mechanical And Manufacturing Systems. *American Journal of Scholarly Research and Innovation*, *1*(01), 61–92. https://doi.org/10.63125/5zvkgg52
[156]. Zobayer, E. (2023). IOT Integration In Intelligent Lubrication Systems For Predictive Maintenance And Performance Optimization In Advanced Manufacturing Industries. *Journal of Sustainable Development and Policy*, *2*(04), 140-173. https://doi.org/10.63125/zybrmx69
[157]. Zuev, D., Kalistratov, A., & Zuev, A. (2018). Machine learning in IT service management. *Procedia computer science*, *145*, 675-679.