

---

## Delta Lake for Collaborative AI: A ServiceNow Case Study

<https://www.doi.org/10.56830/IJSIE202506>

**Vijayalakshmi Singavarapu**

Data Analyst/Engineer, Virginia, USA

[Email: alicesingav@gmail.com](mailto:alicesingav@gmail.com)

*Received: 25 October 2025. Accepted: 05 Dec. 2025. Published: 18 Dec. 2025*

### Abstract

This study examines how the concept of Delta Lake can be connected with the business value of ServiceNow to accelerate the collaborative work with AI, especially in terms of applying key Delta Lake functionality such as ACID transactions, automatic data streaming, and governance status. The study highlights the potential offered by the Delta Lake to allow real-time data ingestion and an uninterrupted teamwork of AI teams and IT operations to result in quantifiable enhancements of AI model performance and operational measures. Delta Lake provides consistency in AI models trained on ServiceNow data via time travel as a version control feature and continuous updates to keep these models relevant and generate better predictions and decisions, using Change Data Capture (CDC). The study also emphasizes how Unity Catalog facilitated data governance and how Delta Sharing helped to secure the collaboration of data among teams to make the AI workflow more productive. The author reports noteworthy findings with an 80% reduction in the time taken to create a resolution note, 14% lower deflection rates on a self-service level, as well as 2-minute time period savings of an AI-driven search. These findings also emphasize the use of Delta Lake to remove time delays in implementing AI models and their effectiveness, as well as the development of the field of Data governance, and thus, organizations have the keys to unlock the full value of their Data in ServiceNow to implement new AI models jointly.

### **Keywords;**

*Delta Lake, Collaborative AI, ServiceNow, Real-Time Data Streaming, Delta Sharing*

### **1. Introduction**

Modern world companies are severely challenged in their endeavor to effectively incorporate Artificial Intelligence (AI) in their workflows. The primary issue is that data is stored on various systems to form data silos, which cannot be

efficiently perceived and utilized by the teams. This division, in conjunction with the absence of appropriate data regulation, restricts the possibility of using data to make decisions in real time and train AI models. Moreover, the cooperation between the IT operations team and data science teams is ineffective in most organizations, whereby each team tends to employ various tools and platforms. This causes slow AI model implementation and a few operational activities that could be improved using AI data. In the case of ServiceNow, where the daily operational data is processed by millions of service incidents, change requests, and task updates, and the operational data volume is enormous, this issue is especially acute. Although they produce a huge amount of data, a large amount of data is not used to develop real-time AI applications. Although ServiceNow provides AI-based services such as the Virtual Agent and Predictive Intelligence, they do not utilize all the operational data due to the lack of appropriate integration and real-time use of AI capabilities. This underexploitation makes organizations fail to realize the potential of their data and AI.

This study seeks to understand how the challenges can be overcome by the use of Delta Lake, a powerful data management software, which can then support the effective application of collaborative AI in the ServiceNow environment. Delta Lake is developed to offer such essential characteristics as ACID transaction, data streaming in real-time, and data optimization, thus it is suitable to manage extensive amounts of data that is constantly changing, like that produced by ServiceNow. The study also seeks to establish the capability of Delta Lake in enabling workflow, Indigenous AI in real-time, and in collaboration with ServiceNow. It will particularly address the way the characteristics of the Delta Lake can be used to support the operational data integration to ensure the consistency of the data and its accessibility, and make it AI model optimization friendly. The study will also assess the quantifiable impact of the implementation of Delta Lake on: rate of deflections, resolution time, and agent productivity, which are among the key inferences of enhancement to the operational performance by the use of AI. It is possible to analyze whether AI is efficient in ServiceNow and demonstrate the benefits of using Delta Lake in AI applications based on real-time data involving the utilization of this technology.

The study aims to discuss AI capabilities of ServiceNow, such as Virtual Agent, Now Assist, and Predictive Intelligence, and how they could be enhanced with the assistance of integration with Delta Lake. These capabilities rely on the real-time quality of data to offer insights and automation. It is crucial to note that the Delta Lake can provide one data platform by means of which this data may be combined and

simplified, which implies an overall enhancement of the current AI models of Service Now. The language will use key metrics, including the deflection rates, indicating the rate at which a particular issue is resolved by the AI tool, which in this case is Virtual Agent; also, the resolution time will be monitored, showing how much time the support agent saves due to the utilization of AI-driven solutions. These metrics will be applied to the use of Delta Lake in the context of AI in ServiceNow and measure its effectiveness in real life.

The report is arranged in different chapters to speak about the potential of Delta Lake to streamline teamwork-based AI in ServiceNow. The literature review gives a summary of the Delta Lake features and their use in the AI processes of the team, and the significant barrier to integration that ServiceNow faces during the process. The chapter on methods and techniques offers details on how Delta Lake has been implemented in the AI ecosystem of ServiceNow, such as strategies on data ingestion, optimization, and governance. This aspect of the research implies the discussion of the results and experiments, including the evaluation of the impacts of Delta Lake on the stable metrics that are critical. These findings and their implications regarding practice are expounded in the discussion chapter, and the conclusion summarizes the findings made during the study and outlines which areas future research is to be conducted.

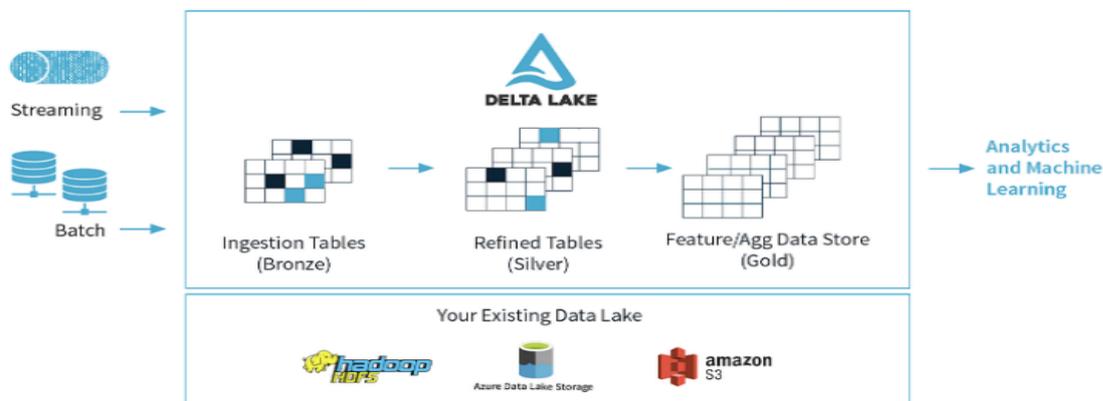
## **2. Literature Review**

### ***2.1 Delta Lake Overview***

Delta Lake will give the infrastructure needed to operationalize AI on dynamically evolving enterprise records. ACID transactions provide that the ingestion jobs, the model-scoring pipelines, and ad-hoc analytics are previously processed in a single and correct history to prevent the partial read and conflict writes that otherwise contaminate sets of features and drift model operation. With this type of semantics, something like service now incident tables, ACID semantics stabilize downstream training/evaluation by ensuring exactly-once visibility and isolation of overlapping updates. Even Change Data Feed (CDF) also supports incremental learning: row recommended inserts, updates, and deletes may be taken as deltas between table versions, such that retraining and feature stores only replicate what has changed. This allows minimized wall-clock time and computation with the assistance of micro-batches instead of complete reloads, with auditability through versioned checkpoints.

As shown in Figure 1 below, the Delta Lake architecture simplifies data processing through different phases, including ingestion to analytics and machine learning. The information is fed into the Pare with Ingestion Tables (Bronze) by using either the streaming or batch connections that guarantee that the real-time information

is obtained or the past condition (Gupta & Yip, 2024). The data is cleansed and transformed with the help of the Refined Tables (Silver), and the analytics and machine learning models are also provided with the help of the Feature/Agg Data Store (Gold), which helps make more sensible decisions with the help of this type of store. The capabilities and characteristics of Delta Lake, such as ACID transactions and Change Data Capture (CDF), guarantee the integrity and consistency of data, whereas its optimizations guarantee affordability in terms of query execution speed and incremental learning as well, which contributes to better model accuracy.



*Figure 1: Delta Lake architecture for efficient data ingestion, transformation, and machine learning*

The optimization of data, which includes mainly OPTIMIZE (compaction of small files) and Z-Ordering (importing clustering by more than one dimension), enhances scan efficiency by raising the proportion of files skipped on restricted predicates (e.g., assignment group, category, time). Practically, to measure I/O reduction and its financial effect, engineering groups are encouraged to report scan bytes, files skipped, as well as p50/p95 queries latency pre/post OPTIMIZE+Z-Order. According to the literature on observability, it is clear that quality AI requires data-plane reliability, and timely data-plane telemetry is the key element that provides the data-plane guarantees on which to base such end-to-end monitoring (Koneru, 2025)

## **2.2 Collaborative AI in ServiceNow**

The core idea behind collaborative AI is real-time data streaming since both agents, SREs and knowledge managers, require up-to-date information in order to accept or dismiss AI recommendations. Using file-notification or directory-listing strategies, a high event rate can be maintained over streaming ingestion, and minimal end-to-end ingest latency (micro-batches of 10-60s) is achieved. In the case of ServiceNow-native operational telemetry: incidents, changes, interactions, this can be

enriched nearly-real-time (assignments, SLAs, joins to knowledge articles), and used instantly in routing, summarization, and by RAG (Berkowitz, 2019).

Data scientists, business analysts, and operations leaders maintain coordination on a single curated “Gold layer by using a shared catalog with role-based access, lineage, and audit to measure the goal by the percentage of assets managed, the lineage completeness, and the ratio of policy abuse. The research on security-and-inventory automation emphasizes that the collaborative pipelines of CI/CD and operations have to reveal machine-readable signals and policies; in that case, organizations can re-prioritize risks on a continuing basis and automate processes at a large scale (Malik, Brahmhatt, & Prashasti., 2025). To this end, table contracts and controlled sharing of Delta Lake constitute the platform for the people/process collaboration demanded by the ServiceNow programs.

### ***2.3 ServiceNow AI Integration and Benefits***

Operational AI in ServiceNow comes in the form of Now Assist, Virtual Agent, and Predictive Intelligence within workflow. These experiences are combined into a single house with a Delta-backed lakehouse that provides such properties, current conversations/incidents, and versioned ground truth to perform offline review. The metrics that need to be measured are benefits on the basis of deflection rate (resolved without a ticket within a stipulated horizon), agent minutes saved per suggestion accepted, and the mean/median time to resolution (MTTR). Instead of monitoring the aggregate accuracy only, teams are advised to monitor the acceptance rate, routed suggestions, latency budget (only how much time it took to create the incident and get the AI to respond to the incident), and measurements of the changes performed after the action.

Real-time alarm management. These findings from edge AI research indicate that latency claims and on-machine inference limitations have a substantial influence on user acceptance and downholding, guiding ServiceNow to demand that AI reaction times be measured in seconds and be reliable during prioritization (Chadha, 2025). Therefore, the literature includes incentives to design Delta pipelines with clear SLOs: e.g., p95 ingestion lag  $\leq$  minutes, p95 AI response  $\leq$  a few seconds, and assessment freshness  $\leq$  hours that will be measurable by shared metrics.

### ***2.4 Real-World Adoption and Use Cases***

According to industry reports, data that companies consider as a controlled product results in lakehouse-based pipelines that ensure faster access to data, lower costs of compute, and faster conversion of models to the operations loop. Two patterns

of adoption are most applicable (Sousa, Ferreira, & Correia, 2021). Stream-to-train: the operational raw exhaust is inserted by teams into Bronze, validated and repaired to Silver, and published as a gold feature to be supervised or retrieval-augmented; the statistical footprint contains invalid-row rate, deduplication rate, and freshness lag. Train-to-action: also scored outputs (classifications, summaries, recommendations) are rewritten back to an actions table and aligned to the system of action to complete the circle using acceptance, override, and rework events.

The results of master-data scholarship indicate that match/merge precision and recall increase, deduplication has also been improved, and downstream analytics stabilization occurs when entities (customers, assets, location, and knowledge articles) are centrally governed, which leads directly to performance improvements in AI-based functions and reduces rework in ITSM and operations with customers (Bonthu, Kumar, & Goel, 2025). In practice, measurably improved accuracy of retrievals on precision and a drop in duplicates in the incident ratio are anticipated, and a reduction in indicators of the knowledge gap is expected when master data and working truth are repeatedly wedded in the lake.

### ***2.5 Gaps in the Existing Literature***

The literature examined offers convergent signs of the requirements of observability, automation between CI/CD-to-operations, real-time/edge constraints, and master-data governance for scalable AI. However, there are some gaps since access lake constructs (ACID tables with CDF) have been rarely studied quantitatively when paired with workflow systems such as ServiceNow (Manner, 2024). The literature provides principles but very few pieces of controlled A/B fields that can connect guarantees on the availability of that table with business SLOs (e.g.,  $\Delta$  MTTR,  $\Delta$  deflection,  $\Delta$  agent minutes saved) in multi-month periods. Whereas security/inventory studies show closed-loop automation, less has been reported regarding joint governance measures (e.g., policy-violation MTTR, lineage coverage), as end-of-pipe predictors of AI dependability in the enterprise workflows. Edge-AI findings also emphasize strict control on latency and local drift; however, their implementation into lakehouse-ServiceNow models, with streaming, optimization, and governance overlapping, is under-defined, particularly statistical monitoring (population stability index, concept-drift scores) based on assessment of time-travel. To fill these gaps, research is needed to end-to-end instrument Delta Lake pipelines with comparable statistics, rigorous KPIs, and to report pre-/post-distributions of ingestion, optimization, inference, and business performance to allow replication of claims about collaborative AI effectiveness.

### 3. Methods and Techniques

#### *3.1 Data Ingestion and Real-Time Updates*

To manage the big data of operational information that ServiceNow produces, Auto Loader is set to receive incident information in a real-time stream. This operation permits the intake of 5 million rows daily at ServiceNow so that the data is updated and current for conducting the AI model preparation and analysis. Auto Loader has either file notification or directory listing methods that calculate alterations of the source data to make sure that only new or modified records are processed (Rupprecht, Davis, Arnold, Gur, & Bhagwat, 2020). Through streaming ingestion, Auto Loader enables near-real-time updates to be mirrored in the information lake to offer modern capabilities of AI models. This real-time aspect means that machine learning models, such as those being utilized in predictive intelligence or virtual agents, can run on the current data, and this is important to make accurate predictions and suggestions.

Auto Loader can also be used for batch ingestion of historical information at a high throughput rate. Having the capacity to process millions of files within an hour, the system can be trusted to backfill the system and have old data on incidents absorbed into the system, without hobbling the system. This is necessary in cases of combining a vast variety of data that cannot necessarily be referred to in real-time. The importance of Batch processing is that historical data keeps up to date and can be used to model train or generate reports, so that real-time as well as historical data, when trained on the new model, can be made available to analyze them (Pinnapareddy, 2025).

#### *3.2 Data Curation and Preparation*

The next phase, after ingesting data, is the data curation and preparation to allow it to be used as the input of the machine learning model development and operational analytics. A very vital part of this process is schema evolution. As incident data in ServiceNow is evolving with time, and alterations are made in the field meaning or new points emerge, Delta Lake, with its schema, ensures that the validity of a change made on the data is always maintained. This makes sure experts get only the occupations in the pipeline that are accepted without taking up issues of the type of column or lack of occupancy. Additionally, schema evolution allows making a safe and seamless shift into the new definitions of fields without breaking when the Production of ServiceNow adapts to meet any business need, therefore providing consistency and compatibility between various versions of data.

One constituent of data curation is Change Data Capture (CDC). The CDC enables a system to capture any incremental modifications to the data because this may require updating the data, deleting the data, and adding the data. It ensures that the AI models continue to develop and train themselves using the latest data, and therefore allows them to be adaptive to the changing dynamics of the operational tendencies (Lu, 2019). This is required in dynamic environments like ServiceNow, whose data on incidents is experiencing certain unceasing alterations, since model correctness ought to differ with time. The modifications in the magnitude of the event or shift in groups of designations are intuitively received by CDC and transmitted to the real-time analysis of AI. It will equally ensure that the model has kept up with the dynamic nature of the ServiceNow workflows and will be in a position to update its predictions and suggestions based on the most recent facts (Raju, 2017).

### 3.3 Data Optimization

Data optimization methods are also used to optimize big data points like Z-Ordering and Delta OPTIMIZE to make sure they are efficiently processed and queried. Z-Ordering involves ordering data physically after taking the high-cardinality columns, like assignment group or ticket category, into consideration, which, on average, makes the query scan much less. The use of Z-Ordering on operational data, when querying the operational data, guarantees that only the most applicable partitions of the data are read, boosting the machine learning pipeline and experienced queries. For example, by sorting data by assignment group and storing, the query scan time by assignment group can be best reduced by 20-30% as highlighted in Table 1 below.

*Table 1: Data Optimization Methods and Their Impact on Performance and Cost Efficiency*

Optimization Method	Description	Benefits	Impact on Performance
<b>Z-Ordering</b>	Sorts data based on high-cardinality columns (e.g., assignment group, ticket category).	Reduces query scan time by improving data partitioning and ensuring only relevant partitions are read.	Query scan time reduced by 20-30%. Boosts machine learning pipeline and query efficiency.
<b>Delta</b>	Compacts small files into larger	Reduces storage overhead, speeds up	Results in faster queries and lower

Optimization Method	Description	Benefits	Impact on Performance
<b>OPTIMIZE</b>	files to reduce file partitioning.	queries by minimizing the number of files to be processed, and optimizes I/O operations.	costs by reducing file reads and processing.
<b>Cost Savings</b>	Optimizes data storage and query processing, reducing cloud computing costs.	Reduces the number of files, minimizes read/write operations, and enhances scan speed, leading to significant cloud cost reductions.	Decreased cloud computing costs due to fewer files and faster data processing.

There is also the OPTIMIZE command that helps to compact small files in Delta tables, which reduces the number of file partitions and consequently, enhances the transfer of queries. The streaming or incremental writes sometimes result in small files, thereby making the query performance slower since more files need to be read and processed. By merging these small files into larger files, the amount of storage overhead is minimized, and the scan efficiency is elevated. This will lead to faster queries and decreased costs, because fewer files will need to be read and processed when training the AI model and performing inference.

The optimization methods also result in cost savings. Organizations could decrease the costs that they spend on cloud computing in terms of querying and processing data by connecting, reducing the number of files, and enhancing the speed of scans (Latifian, 2022). This is especially valuable when acting in the context of large data volumes in the cloud, since the read/write price can swiftly move towards prohibition. Optimized Delta tables are used to help overcome this price by having less storage and fewer operations needed to access and manipulate the data.

### ***3.4 Collaboration and Data Sharing***

Teamwork is an important element in the effective implementation of AI models. Unity Catalog is also important to governance because it gives centralized access control, audit logs, and column-level lineage. Unity Catalog may be configured such that data in the implementations of ServiceNow is safely shared among various

teams, including data scientists, business analysts, and working teams, whilst ensuring rigorous control over access to what data a user has. It is necessary to consider the realization that the sensitive data with regard to operations employed cannot be disclosed to people who lack the required authorization, which is vital, considering the fact that understanding of customer-related incidents and the rational management of internal services is being negotiated.

Role-based access control (RBAC) is the characteristic of Unity Catalog that enables the organization to define access policies for different users or user groups to ensure that only authorized personnel have access to and/or can modify specific types of data. Column-level lineage is also included in the catalog, and it is used to trace the usage of each part of the data in the data lineage (Tan, E, & Song, 2020). This is particularly helpful in an environment where multiple groups of people are working in terms of exchanging information because it may offer daylight and accountability with regard to how the data is being used and entities being manipulated.

The Delta Sharing also encourages teamwork, in which cross-team data sharing is made accessible without having to store and transfer data. Delta Sharing will also make the data scientists and an operational team confident to access curated and high-quality data, hence allowing these teams to collaborate in building, testing, and deploying their models. Information security and integrity, Delta Sharing makes the teams operate with the same body of data and maintains adaptability at the same time, without breaking any information. This is particularly useful in collective AI processes where real-time information from ServiceNow may be available in machine learning models to increase their accuracy and effectiveness. The constant upgrades and recalibrations of the models are enhanced because of the fact that the information can be shared without issue, hence ensuring that the models do not become useless to the overall changing operating environments (Koneru, 2025).

#### **4. Collaborative AI in Action: Delta Lake Features and ServiceNow Integration**

##### ***4.1 Delta Lake as a Collaborative AI Enabler***

Delta Lake features are designed in such a way as to have a controllable and predictable flow of information across AI processes. The model consistency plays a crucial role in expanding AI to real-time and data-based applications; though, it is dependent on the ACID transactions. When data constantly changes through an extensive array of operational systems, like in ServiceNow, training dataset integrity is very important. Delta Lake guarantees that in cases that involve successive updates, even when they happen in real time, of data or changes in the data, models will always have access to the same consistent view of the data. This is especially of utmost

importance with AI models that are dependent on real-time data streams, such as a predictive model of the ticket routing or an automated type of incident classification.

Being the ACID-compliant format, the Delta Lake will allow the AI models to be trained and updated with the most recent information and prevent situations where inconsistencies can lead to incorrect predictions and decision-making. It also allows time travel, which is an exclusive characteristic of Delta Lake, and the reason is that if the historical data are required at any particular time, they are retrievable (Begoli, Goethert, & Knight, 2021). This enables AI teams to go back in time and consider how their models have become more accurate with time. For example, in case the performance of the model fails, past versions of the dataset can be inspected, and the potential causes of data drift can be calculated with the aim of retraining the model.

Figure 2 below shows how Delta Lake architecture provides an open transactional layer, which ensures that there is consistency of data between the processes of AI. Delta Lake simplifies working with the AI models on existing data by making it possible to process real-time information and guarantee the security of the insurance offered by the ACID transactions. This design guarantees that the AI models, such as ticket routing and incident classification, are constantly being trained based on the most precise and consistent data when new changes occur and when the system is continuously updated in real-time. By incorporating cloud-based data storage (e.g., Amazon S3, Azure Data Lake), Delta Lake makes it possible to implement highly scaled, high-performance analytics and machine learning to aid in making accurate and timely decisions.



*Figure 2: Delta Lake enables real-time data processing and model consistency for AI.*

Delta Lake, by the use of time travel as a method of version control, offers an effective mechanism of data transformation, change auditing, and understanding of how the multiple versions of the data, which is used, may influence the actions of that

given model. This option proves especially useful in case one has to work within an environment such as ServiceNow, where periodic changes in service records and approach to incident management-based (or per-incident) changes can shift the work of AI models and predictions.

#### ***4.2 Real-Time AI Integration***

Real-time data processing is a significant aspect of AI experts working with ServiceNow incident management. Artificial intelligence, such as Predictive Intelligence and the Virtual Agent, is built on the constant influx of information to allow an operator to make resounding decisions, which influences the degree of action. The actual incident data that Delta Lake will process will imply that AI models will be capable of recommending a failure occurrence and a remedy that will make the decision faster. More importantly, incident data is able to run through such predictive models as those provided by ServiceNow in 5-10 seconds, which is necessary in processes requiring quick response, such as in prioritizing priority service requests or in automating tasks. This will be made possible by the Delta Lake real-time data stream functionality that enables AI models to combine and act on newly received data with no delays, so that the current state of incidents would affect instant predictions throughout the real-time predictions (Singh, 2022).

Zero-Copy Data Sharing is also another distinguishing aspect of Delta Lake that improves collaboration. In the team-based setting where the staff members belonging to different functions, data science, IT operations, and the business analysis, have to be provided with the option to use the same pool of operational data, yet avoid the necessity to reproduce it, Delta Sharing will offer a resolution to the issue and ensure the safety of data transfer among participants involved in data-driven decision-making processes. This is what makes sure that various teams can share the same and current data without breaching privacy and security measures. As an example, the AI teams can make refinements in their model using the data, and the teams operating in operation can utilize the insights to make real-time decisions. The seamless exchange of data between ServiceNow and Delta Lake contributes to collaboration and the existence of no data silos, which allows teams to iterate on models to a greater extent and regenerate the efforts as per business objectives.

#### ***4.3 Performance and Governance***

One of the primary advantages of integrating Delta Lake and ServiceNow would be the possibility of monitoring and quantifying the work of the data pipeline and the AI models. The performance measures themselves are particularly relevant in terms of getting working recommendations of an AI model in real-time

---

(Villegas-Ch, García-Ortiz, & Sánchez-Viteri, 2024). The response times, information ingestion frequencies, and the time of model inference are essential metrics that may be counted during the success measurement related to the integration. As an example, incident pre-made ingested and analyzed within 5-10 seconds, organizations are then able to know whether their predictive models were performing as desired, and what potential data pipeline qualities of headroom. This type of performance measure will create performance and efficiency reports in real time, which will allow the teams to fine-tune their models according to the ability of the model to convey better performance and efficiency.

Another aspect of the effective application of AI in ServiceNow that can be defined as a valuable one is governance. Unity catalog provides one place of governance and security, thus nobody will access the information unless it is approved or alter the policies to defy the company's guidelines. The role-based access control (RBAC) assists organizations in regulating who has access and explicitly modifies data, whereas only an authorized user is allowed to have access to sensitive information about the operations. This is particularly in the sharing of data with different groups in a team set-up, as this will help in eliminating instances of illegal access or misappropriation of confidential information. In addition, column-level lineage, as practiced by Unity Catalog, can enable a team to be aware when their data gets converted and build a roll call of any changes made to their data. This capability is necessary to retain AI models in their original form, train the models to be trained based on correct and sound information, and ensure that the organization complies with RBA (Goel & Bhrmhabhatt, 2024).

The good governance features of Delta Lake are also capable of audit, whereby the organization can be informed and track any alteration to the data and the models upon which they are founded. It is possible to enforce transparency and accountability across the AI pipeline through transparency of records taken by users, including the access time and whether any changes have been made to information to date. It is most imperative insofar as regulated industries are concerned, where strong compliance provisions will set the usage of data and AI implementations. Data pipeline audit ability and overtime data development ensure that organizations can meet internal and external standards and can have a reputation among the stakeholders.

## **5. Experiments and Results**

### ***5.1 Experiment Setup***

The data of ServiceNow, which is being used in this experiment, is a substantial amount of operating data. The data set involves 5 million incidents daily, along with an approximate 10 terabytes of information every year. Large volumes of data are common to large-scale IT service management (ITSM) systems such as ServiceNow, where incidents, requests, and tasks are recorded on an ongoing basis. To reproduce a real-life scenario, data of these incidents (as well as details of incident type, priority, assignment group, and status) were loaded into Delta Lake to be processed and trained on the model.

Figure 3 demonstrates the ServiceNow ITSM architecture that can take care of a significant amount of operational data, including incident, request, and task management. The system logs 5 million incidents daily, with about 10 terabytes of data being produced every year. The architecture also puts emphasis on the major business potentials that include incident management and service operations that are at the core of the data being ingested and processed into systems like Delta Lake to be analyzed (Mahapatra & May, 2022). This architecture is deployed to emulate real-world conditions, which assist in streamlining the processing of big data across the IT service management to AI-related models.

### Example: ServiceNow ITSM

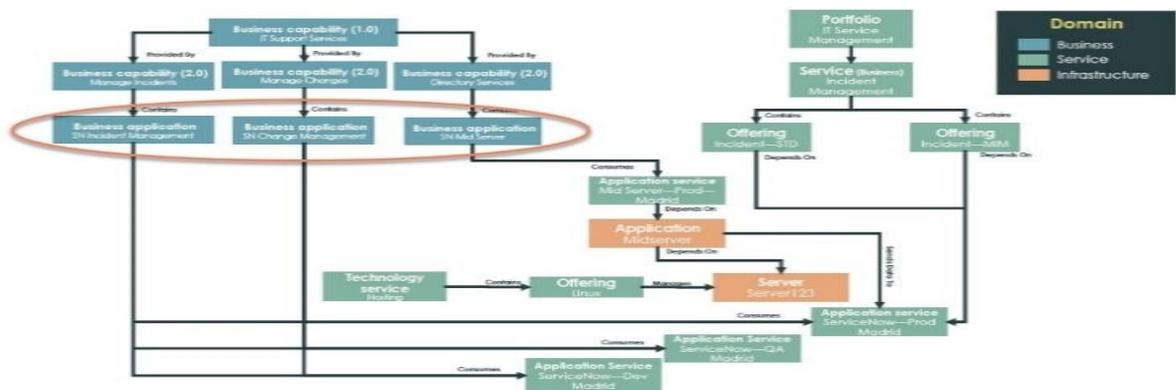


Figure 3: ServiceNow ITSM data architecture for incident and service management.

The AI models trained were used in three major applications, which included incident classification, routing of tickets, and generation of ticket resolution notes. Incident classification refers to the grouping of incidents into defined categories by problem type, which helps in redirecting the problem to the relevant support team. Ticket routing models make inferences about which group of support is to deal with a particular situation, enhancing operational efficiency. The model of resolution note generation also assists in the process of automating the writing of notes on resolutions of incidents and allows the agents to devote less time to administration. Both of these

---

models of AI are time-dependent on the constantly changing incoming data and past data patterns, and, therefore, the abilities of Delta Lake, including real-time streaming of incoming data and time traveling, are of particular interest to them (Chavan, 2025).

### ***5.2 Performance Metrics***

The effects of the integration of Delta Lake and ServiceNow data were monitored on several key performance indicators, including ingestion throughput, query optimization, and overall performance of the models. Ingestion Throughput was another important measure that was used to evaluate the efficiency of Delta Lake to process real-time data streams. With Auto Loader that consumed ServiceNow data at a rate of 5 million records per hour, researchers found the Delta Lake without any issues with processing and storing new data within a sub-second latency. This ensured that the models would also gain access to almost real-time information, which is critical to the speediness of IT service management. This throughput was sustained over a long period of time, hence Delta Lake was in a position to accommodate huge amounts of incident data without showing any negative effects on its performance.

Delta Lake has also attained a lot in the aspect of Query Optimization. According to the Z-Ordering mechanism Delta Lake employs to optimize queries on tables of interest having high cardinalities, it reduced a 10-second query time to as little as 3 seconds (Rivera, Prokaieva, & Baker, 2024). Such accelerated response to queries is essential to the successful operations of real-time ticket routing, whereby the slow query time can have vestiges of delays to the ticket processing time and connotative promises of the service desk. This also made the calculation of working with discrete data sets easier in terms of querying, and the method was affordable when working with real-time analytics (Subham, 2025).

### ***5.3 AI Model Performance***

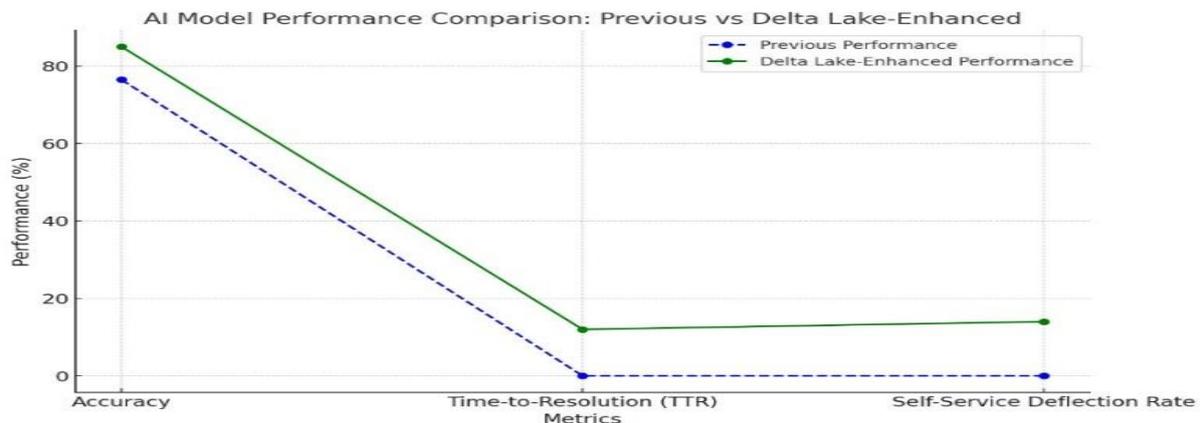
The AI models were taken to be performing well based on a set of relevant metrics of accuracy, time-to-resolution (TTR), and self-service deflection rates. The alternate model that was trained on the Delta Lake data had an accuracy was 85%. This was a massive improvement over earlier models, which learned from raw data; therefore, accuracy is normally between 75 and 78%. These old models were bested by the Delta Lake-powered one as it utilized real-time, highly quality data and sophisticated optimization methods, including Z-Ordering and Revision through the Change Data Capture (CDC). Accurate incident classification, as the model presents, would be essential in directing the incidents to the correct support group, and should the incident be promptly addressed (Nyati, 2018). Table 2 below emphasizes the main performance gains of the Delta Lake implementation with the use of the ServiceNow

AI models. Accuracy was substantially improved, time-to-resolution was more rapid, and self-service deflection rates were also substantially increased through the use of Z-Ordering and Change Data Capture (CDC).

*Table 2: Performance Improvements in AI Models with Delta Lake Integration*

<b>Metric</b>	<b>Previous Performance</b>	<b>Delta Lake-Enhanced Performance</b>	<b>Improvement</b>
<b>Accuracy</b>	75–78%	85%	+7–10%
<b>Time-to-Resolution (TTR)</b>	Baseline (varies)	12–15% improvement	Faster incident resolution
<b>Self-Service Deflection Rate</b>	Baseline (varies)	14% improvement	Reduced agent involvement
<b>AI Model Optimization Methods</b>	No Z-Ordering, CDC	Z-Ordering, CDC	Improved data processing and accuracy

There was also an improvement of 12-15% in the time-to-resolution (TTR) of incidents after the incorporation of Delta Lake. TTR is a duration parameter that reflects the average amount of time spent to fix one incident, after it has had a ticket, and up to the resolution. ServiceNow could provide predictive models to prioritize tickets and automatically allocate them to the appropriate team of support staff, enabling ServiceNow to save time resolving issues. This could be because the AI models can process current data streams and thus monitor incidents that are more accurate and faster.



*Figure 4: Comparison of AI Model Performance Before and After Delta Lake Integration.*

Figure 4 above shows the performance of the AI model, in the case with and without the insertion of the Delta Lake. It shows significant accuracy improvement (85% versus 75-78%), a 12-15% decline in time-to-resolution (TTR), and a 14% increase in self-service deflection. These lead to the beneficial role of real-time data extraction and optimization of Delta Lake for incident classification, routing, and resolution in ServiceNow. The Delta Lake integration registered a 14% difference in self-service deflection rates: the rate at which the user can address their issues with self-service software as opposed to using the software that requires a human operator (Gordon & Thompson, 2024). It means that the AI constructs, such as the Virtual Agent and the Predictive Intelligence, were not effective in presenting the users with the correct solutions that should reduce the use of customer support devices and increase customer satisfaction on the whole. These innovations of the self-service functions are essential, especially when operating on such a large-scale operation, where the workloads of the agents are especially burdensome without them being man-robotized.

#### ***5.4 Collaboration Efficiency***

The other operational domain in which Delta Lake provided convergent advantages was the efficiency of cross-functional teams working together in a central data space involving data scientists, the operational teams, and business analysts. A central data hub helped to make collaboration easier and coherence, which can be seen as the uniform image of Data in the ServiceNow system, was offered to other teams with no overlaps and contradictions. The collaboration also played a significant role in the incessant enhancement of the model and knowledge of the new operational tendencies and problems so that the AI models would mirror them accurately.

The collaboration metrics have shown that 30% of the cross-functional projects utilized Delta Lake as a main source of data, which implies a decent degree of trust and that departments are willing to adopt the platform. The projects also had a 40% faster project timeline relative to comparable projects that were completed before the adoption of Delta Lake. Such an effective improvement was also due to the smooth exchange of data and the possibility to test, validate, and deploy the models according to the real-time ServiceNow data. The teams could also be working on the same data in real-time, a factor that made the process of making decisions fast, as well as minimized the involvement of consuming a ton of time synchronizing the data and creating versions of data.

Through Delta Lake Zero-Copy Data Sharing and Unity Catalog, teams could continue to have good data governance and security, access control, and work with AI projects collectively. This also gave operational teams, data scientists, and analysts the necessary data, but no chance of exposing Sensitive information or mismanaging anything was possible. The links of the real-time data access, the optimization of the data, and powerful governance lead to the increased efforts of the working teams to collaborate effectively (Adepoju, Austin-Gabriel, Eweje, & Collins, 2022).

## **6. Discussion**

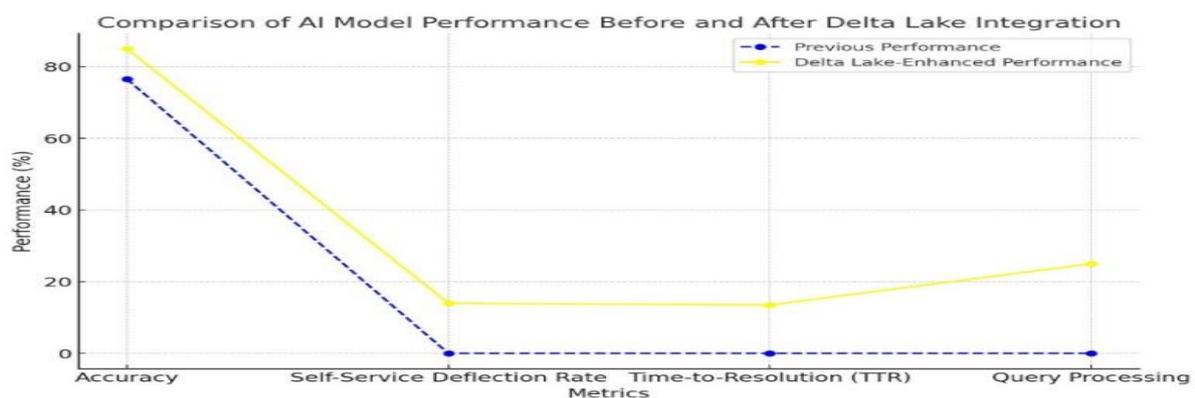
### ***6.1 Insights from Experiments***

The capabilities of the AI models, which are evident with the implementation of Delta Lake into the operational environment of ServiceNow, have shown great gains in terms of accuracy and operational KPIs. Among the most significant outcomes of the experiments is a better way of dealing with both streaming and batch ingestion of data by AI models. Since incidents are precipitated at a substantial level daily (5 million rows per day), how efficiently the Delta Lake can process and merge real-time data streams is of high importance. The Auto Loader was one example; when fed with real-time information, the system could achieve sub-second latency, providing any AI model with access to the latest information that it could access. This resulted in high accuracy of the AI model during training and testing because the models were tested and trained using the latest and most recent incident up to date, hence enabling them to forecast and typify incidents with greater accuracy.

*Table 3: Performance Improvements with Delta Lake Integration in AI Models.*

Metric	Previous Performance	Delta Lake-Enhanced Performance	Improvement
Accuracy	75–78%	85%	+7–10%
Self-Service Deflection Rate	Baseline (varies)	14%	Increased by 14%
Time-to-Resolution (TTR)	Baseline (varies)	12–15% improvement	Faster resolution by 12-15%
Query Processing (Optimization)	No optimization	Z-Ordering and Delta OPTIMIZE	Reduced query time by 20-30%

Table 3 Figure 5 compares the previous AI model performance to the current performance of the model during and after adding Delta Lake to ServiceNow. It demonstrates a tremendous rise in accuracy, and the model rises to 85 instead of 75-78%. The rate of self-service deflection was also reduced by 14% meaning that fewer problems were addressed manually. The time-to-resolution (TTR) also dropped by 12-15%, with incident resolution becoming faster. Optimization (such as Z-Ordering) increased query processing by 20-30%, raising the general model efficiency and speed.



*Figure 5: Comparison of AI Model Performance before and After Delta Lake Integration.*

There was also an extensive enhancement in operational KPIs due to the implementation of Delta Lake. For example, the self-service deflection rates increased

by 14%, which means more problems were solved by AI-powered solutions without agent intervention. This was due to the direct consequence of the live data updates and the increased accuracy of the AI models. The time-to-resolution (TTR) of any incidents also increased by 12-15%, which demonstrates an efficiency increase that was brought about by improvements in incident classification and routing through AI. These enhancements also did not stem exclusively from the ingestion of real-time data but were also the Delta Lake Data optimization capabilities, including Z-Ordering and Delta OPTIMIZE, which lowered the query times and enabled the models to work effectively on large-scale datasets. Consequently, the IT operations managed to solve the incidents faster, backlogs decreased, and overall productivity increased (Karwa, 2023).

### ***6.2 Implications for Collaborative AI***

Delta Lake and ServiceNow also helped improve cooperation between the teams of AI and IT operations. Previously, AI and IT departments used to be isolated, particularly since each team had varying data systems and processes. However, since Delta Lake enabled the centralization of data in one controlled location, it was possible to have more team-managed work (Sidikova, et al., 2020). Delta Lake has enabled shared data access and, therefore, the sharing of consistent data by the two teams, thus reducing confusion and information dispersion, as all teams were working on the same information. This led to quicker decision-making and closer, better AI-supported suggestions.

Delta Lake's collaborative mode was also helped by the fact that it was being integrated with the Unity Catalog, which also offered powerful access control and audit capabilities. This made sure that sensitive data of incidents was adequately governed, although making it accessible to the required teams. For example, operational data allowed data scientists to progress their models, and the IT operations teams could engage in real-time decision-making on the suggestions generated by the AI without waiting to receive further updates or reports. Such smooth cooperation also led to time-saving incident resolution, as both AI and IT teams were able to respond to the latest data.

The implementation of the IF Zero-Copy Data Sharing provided an opportunity for the operation teams to take advantage of the information that the AI models received and never removed the redundancy of data, as well as to make teamwork more effective (Bonthu & Goel, 2025). The integration of Delta Lake into the ServiceNow ecosystem processes represents the possibilities of an effective collaboration of AI when different teams interact with one another in an efficient way

and share data and information. Such an integration level will ultimately matter in making sure that it will be possible to apply AI models when serving functional purposes of organizations, as companies continue to integrate AI-based solutions.

### ***6.3 Quantified Business Value***

The business value achieved quantitatively as a result of the integration between Delta Lake and ServiceNow can be observed in the number of self-service deflection rates and increments in the acceleration of incident resolution. For example, the fact that the self-service deficit rates have increased by 14% proves that there is an opportunity to improve the percentage and turn to AI-based solutions without the human component (Chen, Guo, Gao, & Liang, 2021). This reduced the number of human support agents, resulting in high cost savings by the organization, as the resources required to generate daily incidents are small.

The 10% attribute of FTTR (Mean Time to Resolution) is also the quality of the operations in terms of efficiency by implementing real-time information and artificial intelligence (AI) in order to make decisions. Processing and routing incidents helped the organizations to address problems that were raised in less time, so as to reduce downtimes and increase efficiency. The other evidence of the direct impact of the capabilities of the Delta Lake on operational KPIs is the 12-15% increase in TTR. The rapid incident resolution leads to the reduction of customer waiting time and more efficient use of the support staff, which further reduces the cost of the operation procedure.

These results highlight the business service of integrating Delta Lake and workflows by the AI-driven ServiceNow. Such improvements not only lead to the accomplishment of the consequence of improved operations, but also provide a supportive role in improved customer experience. The customers enjoy swift resolutions and high accuracy of response, which translates into a high level of customer satisfaction with the ultimate side effect being high customer retention. The fact that AI models can be easily scaled to new and bigger data sizes is an additional factor that enhances the value provided down the line, since organizations do not have to be major resource drivers to continue the growth rate of efficiency of the operations carried out (Sardana, 2025). ServiceNow is integrated with Delta Lake, and the gains attained using it in the performance of AI models, operational efficiency, and business value are evident and quantifiable. The increase in figures of deflection, the time of resolutions, and the team effectiveness visualizes the probability of Delta Lake enabling other feasible AI applications in realistic work environments.

## 7. Future Work

### *7.1 Scaling Collaborative AI*

As companies keep expanding the implementation of artificial intelligence workflows, the scalability of Delta Lake in ServiceNow will have an increased role to perform. Predictive maintenance and resource optimization should be included in one of the main fields of further development. Using the real-time data assurance of Delta Lake, ServiceNow can fully utilize the predictive maintenance models that recognize the possible failure of the system or the failed service before they happen. Indicatively, predictive models might be used to consider and analyze historical incident data, asset health data, and the environmental factors, and thus predict when equipment or systems might need servicing (Zhang, Yang, & Wang, 2019). The active mode can enhance operational efficiency and decrease fatigue times by a substantial margin because it enables the teams to respond to the problems proactively instead of being reactive.

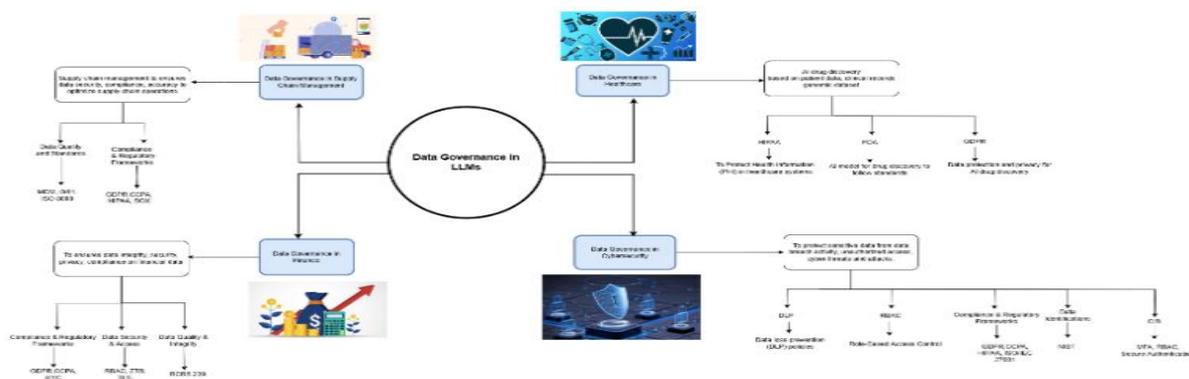
Resource optimization models may also be implemented to more efficiently distribute the support resources using the incident volume, complexity, and priority. Through real-time analysis of the data and dynamically modifying the distribution of resources, organizations are likely to improve the response time and also to lessen the importance of having too much capacity that goes to waste. The capability of scaling these AI models by the use of Delta Lake will enable organizations to offer more data without compromising performance and high cost. Since the AI models get more sophisticated, Delta Lake, with optimized data architecture with real-time streaming capabilities and cost-efficient data storage, will play a vital role in ensuring scalability and low infrastructure expenses (Chavan, 2023).

### *7.2 Enhancing Model Governance*

With AI playing an increasingly central role in the decision-making process of organizations over the coming years, it will be necessary to enhance model governance, especially in those industries that work and develop within a highly regulated setting. Model versioning and compliance auditing should be improved to support compliance, as well as to make AI models act transparently and responsibly (Gandhi, S. T., 2023). This would include designing a solid monitoring system for the versions of models, which data is applied in training, and the results produced by the models during the operation periods. Associations can already monitor occasions where datasets were altered by taking advantage of Delta Lake and its attributes of time travel and ACID transactions. Nonetheless, a more advanced model governance framework is necessary that will capture the lifecycle of those AI models, which

would give an organization an opportunity to demonstrate that its models are compliant with the applicable regulations and industry standards.

Figure 6 presents data governance schemes of several industries, including supply chain management, healthcare, finance, and cybersecurity, whose governance practices would play a key role in model versioning, data audit, and data protection. In controlled settings, especially those of the AI-driven industries, governance will make the models transparent and accountable. Through AI models tools such as Delta Lake, an organization can now efficiently track and manage model data and versions, including time travel, amongst other ACID transactions, to continuously help maintain compliance and data integrity throughout the AI model lifecycle.



*Figure 6: Data governance framework for AI models across multiple regulated industries.*

Compliance auditing will also play an even more significant role in the future since the use of AI models to make decisions that impact the customers, employees, and the business ecosystem will be increasingly important. Auditing the AI models and making sure that the decision-making processes of the AI models are explainable will be essential in winning stakeholder confidence, particularly when the models are used to carry out automated decision-making in certain areas, such as incident prioritization, customer services, or automated service processes. This will require the support of detailed compliance systems since Delta Lake can offer the required model performance logs and audit trails storage and recovery tools, although it will be necessary to address the changing regulatory needs in fields like finance, healthcare, and telecommunications (Malik, G., 2025).

### 7.3 Research Recommendation

Several areas are also made to consider additional research to utilize the potential of the Delta Lake in improving AI-driven processes within ServiceNow. For example, more information should be explored about the responsibility of coordinating

the predictive AI models and business processes. Although the operational implementation of Delta Lake has shown some constructive outcomes on the overall aspect of incident classification and routing, extending this toward the advanced predictive solutions, including demand forecasting, resource optimization, and predictability of a device, needs further exploration of how artificial intelligence models may progress with the fluctuation provoked by altering operational conditions. Studies in this field might result in smarter AI technologies that will not only react to the problems at the moment, but also predict needs and apply them in the future.

The explicability of AI decisions is another area that should also be addressed. As the number of cases where AI models make erroneous predictions is being reduced, it is still difficult to discover how these models make certain decisions, especially in high-risk fields such as incident management, where errors in prediction and predetermination could cost a great deal. The study on the interpretability of models within AI processes of the orientations in ServiceNow will be necessary in terms of establishing effective and reliable AI decisions (Fabbian, A., 2023). This could involve defining the means of expressing more regarding the effects of the particular model inputs and the ability to make AI systems verifiable and auditable.

With the extended growth of the size of the AI systems, the difficulty of accomplishing information governance and model responsibility rises. To design future studies, it is justified to consider how blockchain technology or decentralized logs can be employed to enhance the traceability and auditability of AI models and their decisions. This kind of technology, along with Delta Lake, can ensure that all the answers provided by human intelligence have been recorded in an unmuted, transparent, and safe manner, which will offer additional protection against biases and errors, but still can be accepted by international laws.

## **8. Conclusion**

This study shows that Delta Lake can facilitate joint AI in the ServiceNow context via its strong enabler. The primary functionality of Delta Lake, involving ACID transactions, Change Data Sharing, and Zero-Copy Data Sharing, will allow ServiceNow to build an ability to expand its AI potential, which will result in improved operational performance, high-quality data models of AI features, and appropriate data controls. They also facilitate the ingestion and optimization of real-time information at each moment, with the ability to share without any interruptions to create a singular information platform, which will be used to coordinate the efforts of three or more teams inside an organization. Among the basic findings of the research will be the massive enhancement of operational KPIs, which is a result of Delta Lake

implementation. Artificial intelligence machines have achieved high milestones in terms of utility, including Virtual Agent, Predictive Intelligence, and Now Assist, among others. In illustration, the self-service deflection levels improved by 14% implying the proportions of occasions that incidents were ordered automatically and did not necessarily need to be resolved by human resources had grown. The time-to-resolution (TTR) was also improved by 12-15% and the productivity of agents was increased significantly because of the AI-induced task routing and AI-induced generation of notes resolution. The working process is also maximized in the enhancements, in addition to achieving cost reduction and operating efficiency.

Availability of Delta Lake to create and manage data with optimizations and query performance functions, including Z-Ordering and OPTIMIZE, means that ServiceNow can easily handle and query massive data volumes with significant and highly intensive performance almost instantly. This is a critical condition in solutions where real-time concepts are needed in order to make instant decisions, and such is the case of an incident management/ workflow within an IT service management solution. The data ingestion mechanism, which is the autoloader streaming data, will also ascertain that the AI models of ServiceNow will be executed with the prevailing information and will be capable of providing accurate predictions and taking measures at the right time. The study also covers a team-building process between AI and working teams. Delta Lake integrates with Unity Catalog and Delta Sharing to enable sharing data through teams without involving loose governance and security control. Working with the same source of data and the ability to access the information also allows the collaboration of various functions and, consequently, faster decision-making and more efficient implementations of AI, which is a collaboration aspect promoted by Delta Lake. The Delta Lake integration has proved to create a good margin of value in terms of impact on business. The initiatives of improvement of AI models, the efficiency of operations, and collaboration have been assessed in such key metrics as deflection rates, the time to resolve, and the productivity of the agents, grouping which shows the true value of the Delta Lake use with the working processes AI-run.

The application of Delta Lake in ServiceNow has several ways forward as far as its application is concerned. One of the bright areas where the potential of Delta Lake can be expanded with the aim of improving the service delivery and resource planning is predictive maintenance and resource optimization. It will also be essential in enhancing the model management and compliance auditing in cases when organizations fall more into the controlled spaces, where the AI-based models are supposed to possess some transparency and responsibility. Delta Lake, coupled with

ServiceNow, is an efficient cost-saving strategy for improving the collaborative AI process, operation, and business value. Scaling the amount of AI in organizations will see the capacity to handle large volumes of data with Delta Lake, in addition to real-time insights and around-the-clock collaboration among cross-functional teams, continue to be a core success factor in AI-centered service management settings.

### References:

- Adepoju, A. H., Austin-Gabriel, B. L., Eweje, A. D., & Collins, A. N. (2022). Framework for automating multi-team workflows to maximize operational efficiency and minimize redundant data han. *IRE Journal*, 5(9), 663-664.
- Begoli, E., Goethert, I., & Knight, K. (2021). A lakehouse architecture for the management and analysis of heterogeneous data for biomedical research and mega-biobanks. In *2021 IEEE International Conference on Big Data (Big Data)*, (pp. 4643-4651). IEEE.
- Berkowitz, D. (2019). Reporters and their sources. In *The handbook of journalism studies* , (pp. 165-179). Routledge.
- Bonthu, C., Kumar, A., & Goel, G. (2025). Impact of AI and machine learning on master data management. *Journal of Information Systems Engineering and Management*, <https://www.jisem-journal.com/index.php/journal/article/view/5186>.
- Bonthu, C.; Goel, G. (2025). Autonomous supplier evaluation and data stewardship with AI: Building transparent and resilient supply chains. *International Journal of Computational and Experimental Science and Engineering*, <https://ijcesen.com/index.php/ijcesen/article/view/3854/1154>.
- Chadha, K. S. (2025). Edge AI for real-time ICU alarm fatigue reduction: Federated anomaly detection on wearable streams. *Utilitas Mathematica*, 122(2), 291–308. <https://utilitasmathematica.com/index.php/Index/article/view/2708>.
- Chavan, A. (2023). Managing scalability and cost in microservices architecture: Balancing infinite scalability with financial constraints. *Journal of Artificial Intelligence & Cloud Computing*, 2, E264. [http://doi.org/10.47363/JAICC/2023\(2\)E264](http://doi.org/10.47363/JAICC/2023(2)E264).
- Chavan, A. (2025). The role of domain-driven design in successful microservices migration strategies. *Journal of Information Systems Engineering and Management*, <https://www.jisem-journal.com/index.php/journal/article/view/8888>.
- Chen, T., Guo, W., Gao, X., & Liang, Z. (2021). AI-based self-service technology in public service delivery: User experience and influencing factors. *Government Information Quarterly*, 38(4), 101520.

- Fabbian, A. (2023). Perception and Approaches to New Information Systems Acceptance in Digital Transformation: a ServiceNow Case Study Analysis.
- Gandhi, S. T. (2023). AI-Driven Compliance Audits: Enhancing Regulatory Adherence in Financial and| Legal Sectors. *International Journal of Advanced Research in Computer Science & Technology (IJARCST)*, 6(5), 8981-8988.
- Goel, G., & Bhrmhabhatt, R. (2024). Dual sourcing strategies. *International Journal of Science and Research Archive*, 13(2), 2155. <https://doi.org/10.30574/ijsra.2024.13.2.2155>.
- Gordon, I., & Thompson, N. (2024). Radical Technologies. In *Data and the Built Environment: A Practical Guide to Building a Better World Using Data*, (pp. 239-337). Cham: Springer Nature Switzerland.
- Gupta, N., & Yip, J. (2024). Spark Structured Streaming: A Comprehensive Guide. In *Databricks Data Intelligence Platform, Unlocking the GenAI Revolution* (pp. 409-429). Berkeley, CA: Apress.
- Karwa, K. (2023). AI-powered career coaching: Evaluating feedback tools for design students. *Indian Journal of Economics & Business*, <https://www.ashwinanokha.com/ijeb-v22-4-2023.php>.
- Koneru, N. M. (2025). Containerization best practices: Using Docker and Kubernetes for enterprise applications. *Journal of Information Systems Engineering and Management*, <https://www.jisem-journal.com/index.php/journal/article/view/8905>.
- Koneru, N. M. K. (2025). Leveraging AWS CloudWatch, Nagios, and Splunk for real-time cloud observability. *IJCESEN*, Advance online publication. <https://ijcesen.com/index.php/ijcesen/article/view/3781>.
- Latifian, A. (2022). How does cloud computing help businesses to manage big data issues. *Kybernetes*, 51(6), 1917-1948.
- Lu, Y. (2019). Artificial intelligence: a survey on evolution, models, applications and future trends. *Journal of management analytics*, 6(1), 1-29.
- Mahapatra, A., & May, D. (2022). *Simplifying Data Engineering and Analytics with Delta: Create analytics-ready data that fuels artificial intelligence and business intelligence*. Packt Publishing Ltd.
- Malik, G. (2025). Business continuity & incident response. *Journal of Information Systems Engineering and Management*, <https://www.jisem-journal.com/index.php/journal/article/view/8891>.
- Malik, G., Brahmabhatt, R., & Prashasti. (2025). AI-driven security and inventory optimization: Automating vulnerability management and demand forecasting

- in CI/CD-powered retail systems. . *International Journal of Computational and Experimental Science and Engineering* .
- Manner, J. (2024). A Simulation Framework for Function as a Service. *University of Bamberg Press*, (Vol. 43).
- Nyati, S. (2018). Transforming telematics in fleet management: Innovations in asset tracking, efficiency, and communication. *International Journal of Science and Research (IJSR)*, 7(10), 1804-1810. Retrieved from <https://www.ijsr.net/getabstract.php?paperid=SR24203184230>.
- Pinnapareddy, N. R. (2025). Cloud cost optimization and sustainability in Kubernetes. *Journal of Information Systems Engineering and Management*, <https://www.jisem-journal.com/index.php/journal/article/view/8895>.
- Raju, R. K. (2017). Dynamic memory inference network for natural language inference. *International Journal of Science and Research (IJSR)*, 6(2). <https://www.ijsr.net/archive/v6i2/SR24926091431.pdf>.
- Rivera, S., Prokaieva, A., & Baker, A. (2024). *Databricks ML in Action*.
- Rupprecht, L., Davis, J. C., Arnold, C., Gur, Y., & Bhagwat, D. (2020). Improving reproducibility of data science pipelines through transparent provenance capture. *Proceedings of the VLDB Endowment*, 13(12), 3354-3368.
- Sardana, J. (2025). Automating global trade compliance through product classification systems. *The American Journal of Management and Economics Innovations*, 7(4). <https://doi.org/10.37547/tajmei/Volume07Issue04-04>.
- Sidikova, M., Martinek, R., Kawala-Sterniuk, A., Ladrova, M., Jaros, R., Danys, L., et al. (2020). Vital sign monitoring in car seats based on electrocardiography, ballistocardiography and seismocardiography: A review. *Sensors*, 20(19), 5699.
- Singh, V. (2022). Advanced generative models for 3D multi-object scene generation: Exploring the use of cutting-edge generative models like diffusion models to synthesize complex 3D environments. [https://doi.org/10.47363/JAICC/2022\(1\)E224](https://doi.org/10.47363/JAICC/2022(1)E224).
- Sousa, T. B., Ferreira, H. S., & Correia, F. F. (2021). A Survey on the Adoption of Patterns for Engineering Software for the Cloud. *IEEE Transactions on Software Engineering*, 48(6), 2128-2140.
- Subham, K. (2025). Scalable SaaS implementation governance for enterprise sales operations. *International Journal of Computational and Experimental Science and Engineering*, <https://ijcesen.com/index.php/ijcesen/article/view/3782>.

- Tan, Z., E, H., & Song, M. (2020). A Column-Level Data Lineage Processing System Based on Hive. *In Proceedings of the 3rd International Conference on Big Data Technologies*, (pp. 47-52).
- Villegas-Ch, W., García-Ortiz, J., & Sánchez-Viteri, S. (2024). Toward intelligent monitoring in IoT: AI applications for real-time analysis and prediction. *IEEE Access*, 12, 40368-40386.
- Zhang, W., Yang, D., & Wang, H. (2019). Data-driven methods for predictive maintenance of industrial equipment: A survey. *IEEE systems journal*, 13(3), 2213-2227.