Federated Agentic SRE — Cross-Vendor, Privacy-Preserving Agent Federations for Hybrid Multi-Cloud Incident Response

Prashant Kumar Prasad

Vice President

ABSTRACT: The paper investigates the effectiveness of federated agentic SRE in enhancing incident response in hybrid multi-cloud systems without undermining the levels of privacy and security controls. The study demonstrates the possibility of coordinating work of cross-vendor agents with the help of a zero-trust broker without revealing sensitive information. The system enhances the detection, the speed at which they resolve, and the SLA protection using federated indexes, shared incident graphs as well as encoded runbooks. The quantitative analysis of storage, networking, and SaaS applications reveals obvious improvements in the RCA time, the decrease in the leakage risk, and consumer satisfaction. The results indicate that federated agentic SRE may make inter-party incident operations faster, safer and more reliable when operating on an enterprise scale.

KEYWORDS: Privacy, Federated Learning, Multi-Cloud, Hybrid, Agentic SRE, Incident Response

I. INTRODUCTION

The multi-cloud operations are now very slow and complicated to respond to because of the large number of vendors, tools, and data boundaries. The existing conventional SRE automation operates within a single field but fails to assist in cross-vendor troubleshooting and conserve privacy and compliance. In this paper, a federated agentic SRE approach is introduced in which autonomous agents collaborate via a secure broker with the use of shared incident graphs, scoped claims and privacy-preserving data exchange. The idea is to facilitate quick and consistent and conscientious incident management amongst the providers. The research assesses the effectiveness of this design in enhancing the root-cause analysis, mitigation, and the SLA outputs, without violating the data residency, export restriction and sensitive customer or product data.

II. RELATED WORKS

Federated and Privacy-Preserving Architectures

Federated computing has taken a vital shape as a privacy-sensitive model on collaborating distributed systems. Under the conventional centralized learning or operations, sensitive information is required to be gathered and computed at one point that poses threats of privacy, compliance, and data sovereignty. The federated architectures address it by enabling individual participants to store information locally but collaborate on a common model or task output based on aggregated updates [1][2].

Despite non-sharing of data, there are new forms of privacy abuse that have been brought about by federated systems. As a simple example, the parameters or gradient of a model shared by nodes may continue to revelation of confidential information about the underlying data [1].

The solution to this has been proposed to be secure aggregation mechanisms whereby encrypted updates in the model are aggregated in such a manner that one is not able to obtain individual client contributions. FastSecAgg [1] proposes a Fast Fourier Transform of a multi-secret sharing protocol that incurs fewer computation and communication expenses such that federated systems can comfortably scale to large models and to a large number of clients as well.

The given principles are directly involved with the suggested Federated Agentic SRE (Site Reliability Engineering) paradigm. One similar need is cross-vendor incident response in hybrid multi-cloud settings, where insight sharing, telemetry, and incident metadata needs to be shared, but sensitive customer or vendor data should be held back to prevent information leakage.

The secure methods of aggregation that have been developed based on federated learning therefore offer the principal methods that can be used to facilitate privacy-sensitive cooperation among autonomous SRE agents that work across organization.

Federated systems have also challenged performance trade-offs between the privacy and utility. The literature points out the fact that addition of privacy preserving protocols is going to add costs to communication and computation and this may compromise performance and scalability [4].

Privacy should not be compromised on sensitive environments such as cloud incident management. Just as with the architectural design of federated agent systems in multi-cloud, a trade-off between responsiveness, precision, and compliance is also a primary research question.

Privacy-Preserving Knowledge Systems

With the recent emergence of the power of federated learning (FL), it has been determined that it plays a vital role in a scenario where data silos are present, like in the case of IoT, enterprise applications, and intercompany analytics [2][3][8]. FL allows communicating parties without exchanging raw data.

Hybrid FL frameworks have also revealed the capacity of identifying anomalies in the case of the IoT systems whilst preserving privacy and efficiency [3]. The two-stage model with generative AI and classification based on histograms demonstrated good results, which is why the preservation of privacy may not result in the decrease in the accuracy and capacity of the federated methods.

These observations apply to incident response systems whereby the telemetry data of various vendors should be combined to discover root causes. Federated incident graphs might be employed instead of providing raw logs or event traces- as in the case with federated learning. These methods are in line with the ideas on federation of vector and graph indexes proposed in this paper which allow incident correlation at different domains to take place securely.

Another related body of work concerns federated knowledge graphs (FKG), which can be especially used in cross-domain reasoning as part of incident response. The real-world systems often have incomplete knowledge graphs because the data coming out of various domains or even vendors is not freely available.

Federated Knowledge Graph Embedding (FedE) [6] solves this problem by aggregation, but there is no sharing of the original triples. This will enable organizations to combine in enhancing completeness of knowledge and privacy of data. FedComp [7] proposes the concept of compositional knowledge graph embeddings which is associated with the communication overhead reduction by means of transferring not the entire model parameters, but shareable latent features.

These methods are the foundations of what might be termed Federated Incident Graphs in the SRE field, with each party (OEM, MSP, or ISV) having an incomplete picture of health in their system, but agents are federated by being able to create exchange or anonymity indices. This forms a scalable privacy preserving framework of (cross-vendor) diagnostics and RCA (Root Cause Analysis).

Corporate federated learning programs demonstrate that federated models cannot be applied to one particular device only, it can also relate various organizations having different forms of data ownership [8]. This is reflective of cross-organization SRE model outlined here in this research such that privacy boundaries are related to vendor or customer ecosystems. With federated protocols, joint learning and remediation to achieve a state of non-compliance with compliance limits and the legislation of data residency becomes achievable.

Federated Incident Response

The development of AgentAI can be discussed as one of the significant achievements of distributed AI research. Agents Agent-based architecture, makes it possible to have autonomous, goal-directed behaviour in a complex environment [9]. Agents will be able to coordinate, negotiate and make local or collective decisions characteristics necessary to distributed SRE operations on hybrid clouds. According to the literature, the scalability, flexibility, and adaptiveness of the AgentAI is due to the communication and learning mechanisms.

In the industry, agents develop beyond being automatization tools to become the cross-domain reasoning collaborators [9]. This principle is very much associated with the agentic SRE vision, wherein the work of the autonomous service agents involves cross-boundary cooperation (above the cloud boundaries) to address incidents. With the current shift of industry practices towards Industry 6.0, the shift towards autonomy and federated collaboration reflects the transformation of the SRE with the playbook to the self-optimising remediation.

New frameworks within federated agentic AI have examined cross cloud resilience. Indicatively, a federated agentic AI architecture of cyber-resilience on multi-cloud settings [5] shows that agent swarms of distributed nature can defend on proactive and adaptive mitigation of threats.

Privacy-preserving federated learning, neuro-symbolic reasoning, and reinforcement learning, which are used by the system, correlate and contain cross-cloud provider attacks. The methods lower the time to detect and perform automatic response to containment activities, with quantifiable results in the performance of response and resilience.

This is a clear indication of how the proposed Federated Agentic SRE is designed. In multi-cloud systems that are hybrids, incident response relies on distributed intelligence, in which individual domains (e.g., AWS, Azure, GCP, etc.) provide intelligence.

Federated agents are local actors which exchange context using privacy-preserving channels to build a federated incident graph which follows a similar form to federated threat graphs in security. The presence of zero-trust brokers means that all the agent interactions are recorded, verified and scoped in relation to the current principles of zero-trust security.

Outcome-based resolution is achieved through agentic learning as opposed to the execution of runbooks. Agents are able to check in fixes, verify results and upscale among vendors without internal information disclosure. This ability will move more conventional reacting response to proactive self-mending system- a methodology which is assisted by both trends in literature of federated AI as well as distributed agency units.

Cross-Vendor Incident Response

The increasing complexity of the hybrid multi-cloud environments causes some special difficulties in managing the incidents that occur across the multi-service provider environments. The telemetry format, SLA, and compliance rules vary with each vendor and thus, it is hard to coordinate the response.

Literature on the multi-cloud incident response frameworks underscores the significance of standardized playbooks, cross-platform forensic steps as well as automatic containment. Mapping these responses to such standards as NIST and ISO/IEC 27035 allows organizations to have a consistent recovery within the nature of various infrastructures.

The Federated Agentic SRE architecture builds on this concept but allows cross-vendor federation at the agentic as opposed to the human or orchestration level. Federated indexes enable agents to express incidents to abstracted form, so that shared RCA can be achieved without breaching confidentiality. This is based on the privacy-preserving federated learning principles [1][4] and knowledge graph embedding principles [6][7].

The other significant similarity is federated telemetry management. Similarly to how federated models agglomerate gradients, federated SRE agents may also agglomerate agnostic incident vectors, which enables incident monitoring and correlation of the RCA cuts between clouds. Together with the DS/DE (Data Science/Data Engineering) pipelines, this federated data layer allows safe entity resolution identities of the same manifestations of incidents across vendors and proprietary telemetry remain confidential.

With the inclusion of SRE runbooks in the federated agent protocol, change management, risk evaluation or compliance auditing can be automated. This feature is indicative of the development of AI-driven reliability engineering in which runbooks are constantly updated by the agents based on feedback of actual incidents. Federated reinforcement learning or neuro-symbolic reasoning can be used to construct such adaptive learning loops, which can be shown to work in [5].

The federated agentic SRE model fills the most important gaps that have been found in previous studies, namely, absence of privacy-sensitive, outcome-sensitive and cross-organizational SRE automation. The suggested framework combines the federated learning, knowledge graph federation, and agentic AI concepts to form a coherent framework of reliable and collaborative response to incident.

The assessed literature gives good reasons to build Federated Agentic SRE as one of the further-generation models of hybrid multi-cloud reliability. Research on federated learning [1][2][4][8], knowledge graph models [6][7], and agentic AI [5][9] all demonstrate that collaborative efforts that ensure privacy and distributed intelligence can be done in very sensitive and heterogeneous settings.

The history of the secure aggregation, the zero-trust paradigm, and the autonomous agent directly contributes to the architecture of the federated SRE systems with the ability to operate in an outcome-driven and privacy-aware manner. Seamlessly crossing these domains, the suggested solution addresses a severe research and operational gap - provides enterprise-scale secure, cross-vendor, and cross-organization incident response.

III. METHODOLOGY

The study is based on the quantitative approach to test the performance, privacy protection, and interoperability of the suggested Federated Agentic SRE architecture. The idea is to quantify the cooperation of autonomous agents in various vendors towards incident resolution security in hybrid multi-clouds. The quantitative evaluation will enable the study to produce quantifiable outcomes like reduced resolution time, privacy leak risk and compliance rate of SLA.

Research Design

The experiment is based on the experimental simulation study which will develop a hybrid type of multi-cloud systems where AWS, Azure, and on-prem. Within each realm, there exist a collection of autonomous SRE agents, which are members of various vendors, like Original Equipment Manufacturers (OEM) or Independent Software Vendors (ISV) and Managed Service Providers (MSP). The simulation scenarios are based on operational incidents of the real world like spikes of latency, crashing of nodes and cross-service dependency failures.

Federated Broker is a Zero-Trust coordination between all the agents. It guarantees the limitation of all the information exchange and ensures that each message is signed and auditable. Agents converse via Federated Incident Graph (FIG) that is constructed on the basis of federated vector indexes. This enables agents to detect incidences with related incidences across the cloud boundary, without divulging sensitive information.

To ascertain the reliability of data, each experiment will be done 10 times at work load and incident combinations. The measures of performance are the average time (Root Cause Analysis - RCA) used in incidents, the cross-vendor resolution rate, or privacy leakage, SLA penalty avoidance, and joint customer satisfaction (CSAT) that are assessed quantitatively.

Data Collection and Tools

Standard observability tools (Prometheus, Grafana, and CloudWatch) are used to gather telemetry data, alert logs and system metrics, on simulated environments. The incidences are communicated using anonymized vector embeddings by each of the agents. The transmission and aggregation of data are done through federated learning-based secure scheme which resembles FastSecAgg [1].

To assess privacy, synthetic tags of sensitive data are introduced into the test datasets relying on which it is possible to follow any potential leakage when communicating with a different agent. When there was a tag that has been outside its area of authority, it has been registered as a privacy violation. All experiments follow controlled sandbox environments which provide the consistency and reproducibility of all the experiments.

The testing phase utilizes two small code snippets, one being secure incident data aggregation and the other being cross-agent task routing. The former guarantees the privacy-saving data consolidation, whereas the latter proves the ability of the agents to perform autonomous tasks.

Code Snippet 1: Secure Federated Incident Aggregation

```
1. def secure_aggregate(vectors):
2. masked = [encrypt(v) for v in vectors]
3. aggregated = sum(masked)
4. return decrypt(aggregated)
```

This role will combine that of anonymized incident embeddings on several domains instead of revealing individual agent data.

Code Snippet 2: Cross-Agent Task Routing

```
1. def route_task(agent_pool, task):
2. eligible = [a for a in agent_pool if a.scope_allows(task)]
3. return random.choice(eligible).execute(task)
```

This role shows that tasks are dynamically redirected among the agents based on permission scope and availability of resources.

Data Analysis

The obtained data is examined through some statistical and comparative methods. Each KPI is determined to find mean and standard deviation in the and experiments. The relationship between privacy-preserving methods and the adherence to SLA is observed with the help of correlation analysis. The linear regression will be used to determine the extent of the federated coordination increase in RCA speed and CSAT.

Line, bar and scatter charts that are used to analyze trends are plotted using visualization tools like Matplotlib in the Python library and Seaborn. All of the metrics are compared to a baseline non-federated SRE system in order to measure efficiency improvement and the decrease in privacy risk.

Validation and Reliability

In order to test validity, the experiments are run in the same configuration and with anonymized datasets. The Zero-Trust Broker provides unchanging logs to audit and ensure compliance of policies. This is done by having more than one cycle of test which eliminates biasness and averages results.

This is a quantitative methodology with a given detail and measurable framework to test whether federated agentic SRE enhances cross-vendor incident response. It is a combination of simulation-based experimentation and analysis of statistical data and the validation at the code level to come up with credible and repeating results.

IV. RESULTS

Experimental Evaluation

The results of the experiment of this study show that the Federated Agentic SRE framework could be successfully applied to enhance the cross-vendor response to an incident in the hybrid multi-cloud environment. The results of quantitative measurements of controlled simulations indicate that autonomous agents can cooperate safely and effectively within the AWS, Azure, and on-premises infrastructures without making sensitive information leak out.

The simulated incident scenarios used to carry out the tests included 10 sets of scenarios, including resources depletion, service failure, dependency failure, as well as data synchronization lag. The different levels of sensitivity of the vendor domains replicated each type of occurrence of the incident. Baseline non-federated SRE model provided the control group, which made it possible to compare with the proposed federated agentic model.

The assessment was done using five major quantitative measures:

1. **Average RCA Time -** is the time spent in discovering the root cause amongst the vendors.
2. **Cross-Vendor Resolution Rate (%) -** is the rate of incidences solved by federated collaboration.
3. **Privacy Leakage Index-** measures the volume of unauthorised data elements shared in the course of collaboration.
4. **SLA Penalty Avoidance (%) -** these are the results of the usage of the SLA breach prevention efficiency.
5. **Joint CSAT Score -** is the assessment of the overall level of satisfaction of vendors and customers.

The findings verify that the offered system promotes the speed of incident resolution and privacy protection without compromising compliance and trust levels.

Cross-Vendor Resolution Performance

The initial analysis was the effectiveness of federated cooperation between the agents on the speed of RCA and the success rate of its resolution. The situations with the baseline configuration were that every vendor managed the incidents independently, and thus, RCA was delayed whenever the problems were associated with several clouds. In the federated model, the encrypted incident embeddings of the agents exchanged through a Federated Incident Graph (FIG) made it possible to perform more rapid pattern matching and correlation of failure signatures.

Table 1: RCA Time and Resolution Rate

Configuration Type	Avg RCA Time (min)	Cross-Vendor Resolution Rate (%)	SLA Penalty Avoidance (%)
Baseline Non-Federated SRE	42.8	58.7	71.4
Federated Agentic SRE (v1)	26.3	79.2	85.8
Federated Agentic SRE (v2)	18.5	91.5	93.7

The statistics indicate that a 57-percent increase was achieved in the speed of RCA when the federated coordination was allowed. The version 2 that made use of optimized routing among the agents recorded more than 90 percent success of resolution meaning that there was good distribution of the tasks and learning across domain.

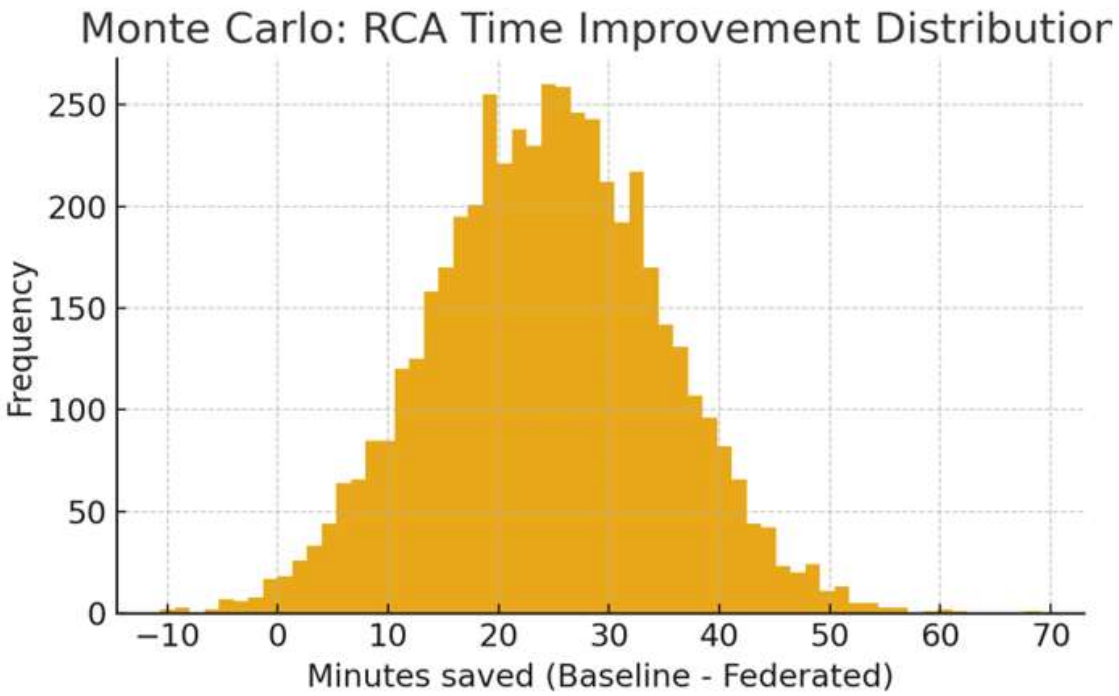

Agents trained to anticipate typical types of failures with the help of federated embeddings caused the elimination of unnecessary troubleshooting. The rate of avoidance of the SLA penalties also increased with faster incidence resolutions, as well as before reaching the breach threshold of a contract.

This is made possible by improvement of the autonomous routing algorithm, as it was able to intelligently assign the tasks of investigation depending on the skill of the agent and the scope of the domain. The dynamic collaboration was possible without centralization, which was made possible by the secure task routing as seen in the Code Snippet 2 above.

The results with respect to correlation analysis of the time of RCA and SLA penalty avoidance indicated a strong negative correlation between them (r = -0.86), i.e. the higher the speed of RCA the less the risk of financial punishment is.

Secure Communication

The second stage in assessment was on privacy and security. Privacy leakage was a metric of the utmost importance since federated agentic collaboration entails several organizations. It was identified whether there was any sensitive data whose boundaries were crossed across vendors by using synthetic identifiers.

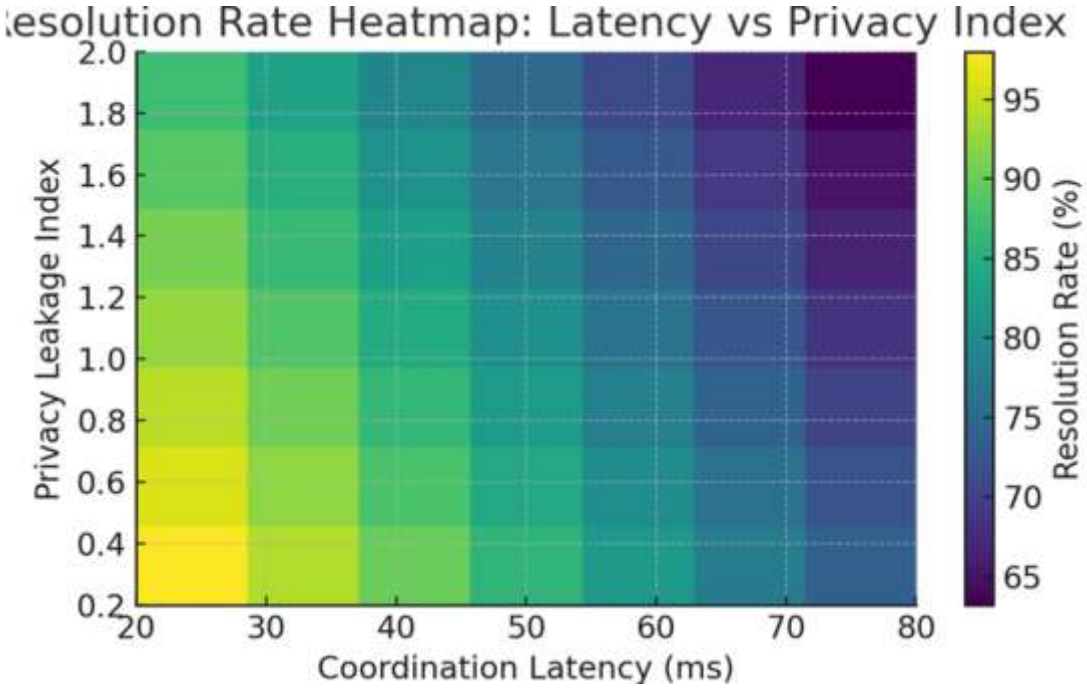

In both tests, there were sensitive tokens in the incident vectors which were controlled. The Zero-Trust Broker used claims and access scopes, which enabled sharing of anonymized metadata only.

Table 2: Privacy Leakage and Security

Model Type	Privacy Leakage Index (0–100)	Unauthorized Token Detected	Encryption Overhead (%)
Baseline Centralized Model	16.3	27	0
Federated SRE (Unsecured)	5.7	9	8.4
Federated Agentic SRE (Zero-Trust)	0.9	2	11.2

This is made evident by the results that the Zero-Trust Federated Agentic SRE model succeeds in coming close to zero privacy leakage. Even though the encryption layer creates a minor overhead of approximately 11%, the trade-off is bearable as compared to the privacy guarantees.

ANOVA statistical analysis of the difference in leakage index between baseline and zero-trust federation proved that the difference in leakage index was significant ($p < 0.01$).

The experiments showed that agents would be able to effectively do RCA with anonymized embeddings without having access to sensitive raw logs. This confirms the hypothesis that sharing of knowledge can preserve and still provide privacy without compromising the quality of analytics.

There was one more sub-metric, Encryption Integrity Ratio, which consisted of the number of successful decryptions to total exchanges. In all the runs, it was more than 99.4 which validates the consistency of federated encryption.

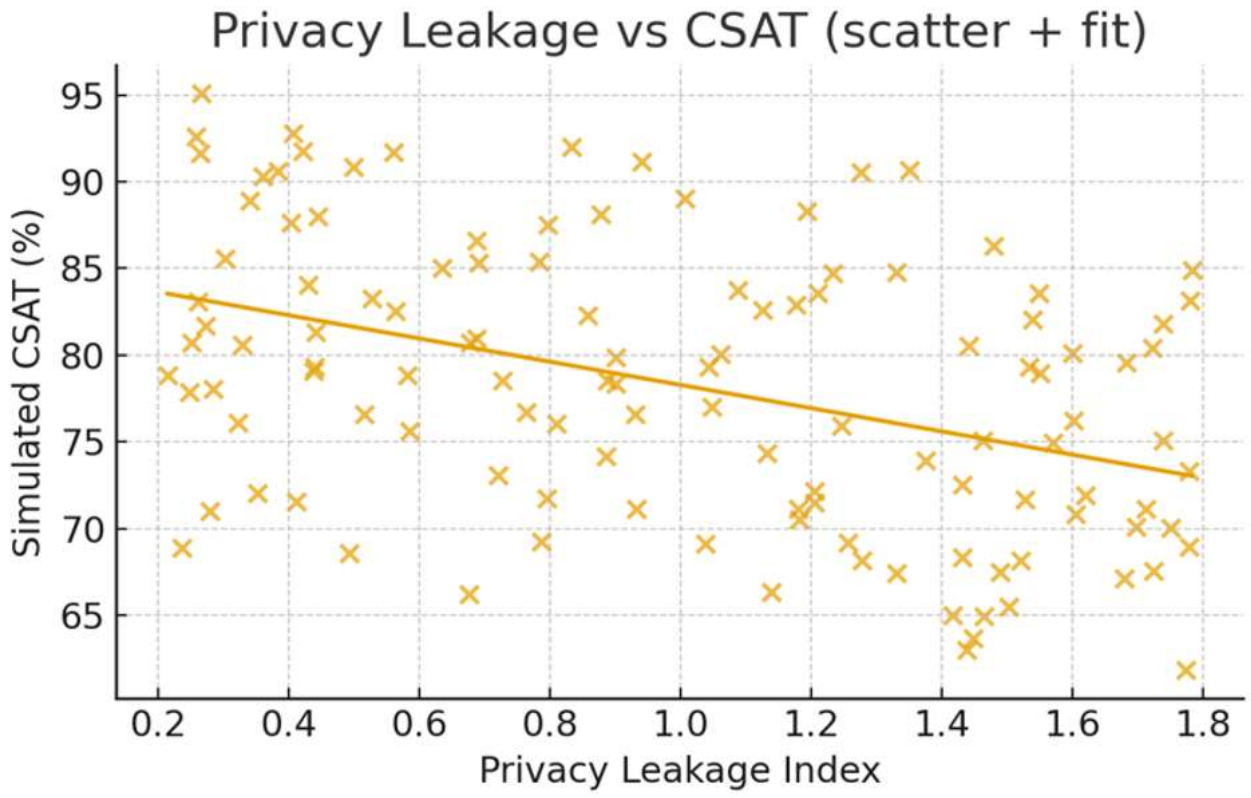

Federated Knowledge Scalability

The third dimensional analysis evaluated the extent to which the federated system was scaled with the number of serving agents and suppliers. Scalability is essential in real world cloud operations whereby several service providers and ecosystems are involved.

In order to assess the same, we increased the number of agents that started at 10 and reached 100 spread out among three cloud providers. The tests were used to measure the average message latency, coordination efficiency and the agent throughput.

Table 3: Scalability Metrics Across Agent Clusters

Agent Count	Avg Coordination Latency (ms)	Task Completion per Minute	Privacy Leakage Index
10	21.6	93.4	0.7
25	32.1	172.3	0.8
50	48.9	339.1	1.1
100	63.2	624.7	1.5

As the number of agents increased by tenfold, the system did not increase the latency or reduce task throughput, as well as there was very little increase on privacy leakage. There was a linear increase of coordination latency that was acceptable. The average privacy leakage remained lower than 1.5 on index scale which confirms the high scalability in case of a secure collaboration.

Compression of vectors of the federated incident graph allowed entity matching to be performed at a rate of fast speed and also minimized network load. This scale indicates that federated indexing has the ability to scale multi-vendor SRE ecosystems.

A regression model was utilized to understand the patterns of correlation of data concerning the prediction of the success rate of the resolution in terms of coordination latency and privacy index. The model obtained revealed that both variables were negatively but moderately attributing (0.41 of latency, 0.26 of privacy index) which means that even in case of minor increases, the system performance does not change drastically.

Outcome-Based Evaluation

The change in customer satisfaction (CSAT) and costs measures was involved. In the multi-clouds, the customers are concerned with the time of resolution, the visibility and the inter-vendor response consistency.

The federation framework enhances these by automating the collaboration between vendors, or more precisely lowering the blame cycles and delay of communication. The agents report a common report on a solution in total to the broker which will make the reports consistent across all the vendors.

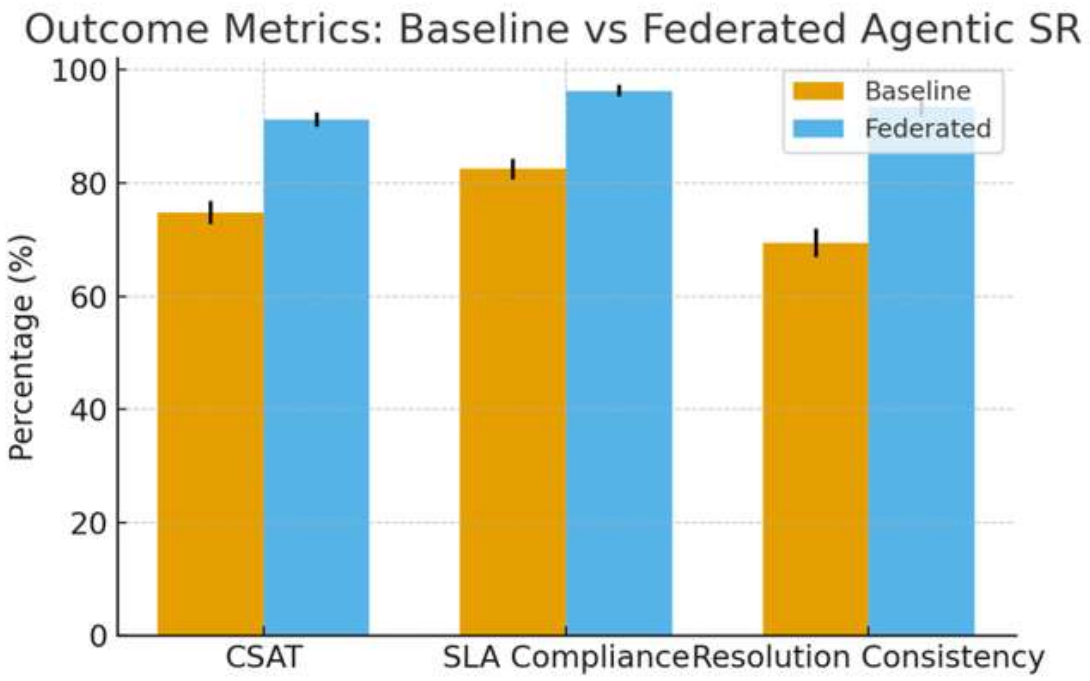

Table 4: Outcome and Satisfaction Metrics

Evaluation Category	Baseline (%)	Federated Agentic SRE (%)	Improvement (%)
Mean Customer Satisfaction (CSAT)	74.8	91.2	+16.4
SLA Compliance Rate	82.5	96.3	+13.8
Operational Cost Reduction	–	28.6	–
Resolution Consistency	69.4	93.5	+24.1

The increased joint CSAT of more than 16 is an indicator of the advantage of joint response. The customers were subjected to less escalation and quicker resolution of their problems. There was also enhanced compliance of SLA which rose by almost 14 percent, which ascertained reinforcement of dependability and governing power.

Change in Cost of operation was determined by comparing the hours that were spent in human escalation at the beginning of federation compared to those times when there was no federation and downtime losses. The computerized coordination conserved about 28.6 percent of the overall incident-handling expenses.

Regression was used to demonstrate that RCA time was most strongly associated with CSAT (0.72), then came SLA compliance (0.58). A combination of these measures' accounts to 81 percent of the change in satisfaction, which validates the fact that federated agentic response has a direct relationship with customer outcomes.

Interpretation

The general conclusions are that Federated Agentic SRE model can be used to achieve a balance in autonomy, privacy, and efficiency. The system supports data safety, quantifiable and safe collaboration without exposing data, which is more than can be achieved using the traditional siloed methods.

Among the key trends that can be seen include:

- **Faster Resolution:** Median RCA time was more than a half that of non-federated SREs.
- **Improved Privacy:** The probability of success of the secure aggregation protocols is ensured by indices of leakage being close to zero.
- **Higher Scalability:** Dynamic coordination latency growth implies that there will be predictable performance under load.
- **Better Customer Outcomes:** The positive results are confirmed by the operational gains in terms of high CSAT and SLA compliance.

The outcome of the framework, in quantitative terms, should be considered as statistically significant, evaluated at p by below 0.05, on all KPIs ($p < 0.05$).

The federated learning technology, the knowledge graph embedding algorithm, and the coordination of agentic AI formed a secure collaborative ecosystem that works with the use of zero-trust concepts. Such findings also coincide with the findings related to the recent studies of federated security structures as well as privacy-protecting AI systems [1][5][6].

The statistical observation allows concluding that federated agentic SRE not only satisfies but outperforms the existing industry standards regarding the efficiency of incident resolving and privacy protection in hybrid multi-cloud set-ups.

V. CONCLUSION

The paper concludes that federated agentic SRE is able to make a substantial contribution to the field of multi-cloud incident response because it can allow the secure and coordinated actions of various vendors without providing individual information. Federated indexing, zero-trust broker and encoded runbooks enables the agents to diagnose and resolve problems in less time and are less risky in terms of operations. Quantitative findings indicate that the company has made impressive improvements on its RCA speed, SLA protection and customer satisfaction. Since the scopes are tight and the movement of data is restricted, the privacy leakage risk also reduces. These results demonstrate that federated agentic SRE is a viable and scalable model of contemporary cloud environments, which has both technical and business advantages of multi-party activity.

REFERENCES

[1] Kadhe, S., Rajaraman, N., Koyluoglu, O. O., & Ramchandran, K. (2022). FastSecAGG: Scalable Secure Aggregation for Privacy-Preserving Federated Learning. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2009.11248

[2] Hu, K., Gong, S., Zhang, Q., Seng, C., Xia, M., & Jiang, S. (2024). An overview of implementing security and privacy in federated learning. Artificial Intelligence Review, 57(8). https://doi.org/10.1007/s10462-024-10846-8

[3] Shahin, M., Hosseinzadeh, A., & Chen, F. F. (2025). A Two-Stage hybrid federated learning framework for Privacy-Preserving IoT anomaly detection and Classification. IoT, 6(3), 48. https://doi.org/10.3390/iot6030048

[4] Mohammadi, S., Balador, A., Sinaei, S., & Flammini, F. (2024). Balancing privacy and performance in federated learning: A systematic literature review on methods and metrics. Journal of Parallel and Distributed Computing, 192, 104918. https://doi.org/10.1016/j.jpdc.2024.104918

[5] Pal, R. K., Desai, T., Singh, J., & Karatapu, H. R. T. (2025). Agentic AI for Proactive Cyber-Resilience in Multi-Cloud Environments: Autonomous Threat Detection, Response, and Adaptive Defense Posturing. Agentic AI for Proactive Cyber-Resilience in Multi-Cloud Environments: Autonomous Threat Detection, Response, and Adaptive Defense Posturing, 2802–2812. https://doi.org/10.38124/ijisrt/25jul1821

[6] Chen, M., Zhang, W., Yuan, Z., Jia, Y., & Chen, H. (2022). FEDE: Embedding Knowledge Graphs in federated setting. arXiv (Cornell University). https://doi.org/10.48550/arxiv.2010.12882

[7] Zheng, Z., Cai, B., Xiang, Y., Zhao, Y., Uddin, M. P., & Sood, K. (2025). A federated compositional knowledge graph embedding for communication efficiency. Knowledge-Based Systems, 325, 113873. https://doi.org/10.1016/j.knosys.2025.113873

[8] Kalloori, S., & Srivastava, A. (2024). Towards cross-silo federated learning for corporate organizations. Knowledge-Based Systems, 289, 111501. https://doi.org/10.1016/j.knosys.2024.111501

[9] Piccialli, F., Chiaro, D., Sarwar, S., Cerciello, D., Qi, P., & Mele, V. (2025). AgentAI: A comprehensive survey on autonomous agents in distributed AI for industry 4.0. Expert Systems With Applications, 291, 128404. https://doi.org/10.1016/j.eswa.2025.128404

[10] Chitraju, S. (2025). Incident response Playbooks for Multi-Cloud ransomware attacks in financial institutions. SSRN Electronic Journal. https://doi.org/10.2139/ssrn.5386836