Scalable Enterprise Decision Support Systems: Leveraging Distributed Data Platforms for Real-Time Intelligence

Rahul Verma
Department of Computer Science
Poornima University
Jaipur, Rajasthan, India

Anita Singh
Department of Information Technology
Dr. A.P.J. Abdul Kalam Technical University
Lucknow, Uttar Pradesh, India

Sanjay Kumar
Department of Computer Applications
Aryabhatta Knowledge University
Patna, Bihar, India

Submitted on: October 12, 2021
Accepted on: November 05, 2021
Published on: November 28, 2021

DOI: 10.5281/zenodo.18177057

***Abstract*—Enterprise decision making increasingly depends on the ability to interpret high-velocity, high-volume, and high-variety data streams in near real time. Traditional decision support systems struggle to scale across organizational boundaries, integrate heterogeneous data sources, and maintain trust among human decision makers. This article presents a scalable enterprise decision support architecture that leverages distributed data platforms to deliver timely, context-aware intelligence while preserving interpretability and operational resilience. The proposed approach integrates event-driven ingestion, distributed analytics, and human-centered decision workflows to support complex enterprise scenarios. Empirical evaluation across simulated enterprise workloads demonstrates improved responsiveness, scalability, and decision quality compared to monolithic and batch-oriented systems.**

***Index Terms*—Decision support systems, distributed data platforms, real-time analytics, enterprise intelligence, human-in-the-loop systems**

I. Introduction

Decision support systems (DSS) have evolved from standalone analytical tools into foundational components of enterprise digital infrastructure. Organizations now operate in environments characterized by continuous data generation, rapid operational changes, and increasing accountability for automated recommendations. As a result, DSS must not only compute optimal decisions but also scale reliably, integrate diverse data sources, and remain interpretable to human stakeholders.

Prior work highlights the expanding role of DSS across industrial enterprises, government, healthcare, and public services [1]–[3]. However, many existing systems remain constrained by centralized architectures and delayed analytical pipelines. These limitations hinder real-time situational awareness and reduce organizational trust in automated outputs [4].

This article addresses these challenges by proposing a scalable enterprise DSS architecture grounded in distributed data platforms. The contributions are threefold: (1) a synthesis of design principles drawn from diverse DSS domains, (2) a reference architecture that integrates streaming data, distributed analytics, and human-in-the-loop decision processes, and (3) an empirical evaluation demonstrating scalability and performance benefits under enterprise workloads.

II. Literature Review

Enterprise decision support systems (DSS) sit at the intersection of data platforms, decision science, and human work practices. The most consistent finding across domains is that the "best" decision logic is rarely sufficient on its own. Adoption depends on how well the system fits real workflows, how it scales under operational pressure, and how it earns trust when recommendations are uncertain or incomplete [4]–[6]. This section synthesizes prior work into design-relevant themes and identifies gaps addressed by scalable, distributed, real-time enterprise DSS.

A. DSS Purpose, Scope, and Organizational Fit

DSS research has long recognized that systems must align with organizational goals such as competitiveness, investment planning, and resource allocation [1], [7]. In industrial settings, DSS value is often evaluated through performance outcomes and managerial usefulness rather than purely technical correctness [1], [8]. The implication for enterprise-scale DSS is that architecture and governance must be tied to measurable business outcomes, including time to decision and ability to act.

A related perspective is that many high-impact organizational decisions are not well-structured. Wicked problems, by definition, involve ambiguous objectives, competing values, and evolving constraints [5]. Pretorius argues for procedural DSS that help users move through decision stages rather than delivering a single definitive recommendation [5]. This is especially relevant for enterprise environments where decision contexts shift across teams, time horizons, and risk tolerances.

This subsection motivates the mapping in Table I, which groups prior work by decision setting and design emphasis. The table is used later to justify why distributed platforms must be complemented with human-centered feedback loops.

B. Human-Centered DSS, Trust, and Decision Accountability

As DSS recommendations become more automated, the human relationship with the system becomes a central design constraint. Trust is not a static property; it depends on context, interruptions, and the perceived consequences of following advice [4]. Müller et al. show that DSS can improve work outcomes under certain conditions, implying that enterprise DSS must manage attention and interruptions, not just prediction quality [4].

In high-stakes environments like intensive care, users often have to reconstruct the situation from fragmented signals. Kaltenhauser et al. describe decision support as a collaborative puzzle where clinicians synthesize contextual information, negotiate uncertainty, and coordinate with others [6]. This suggests that enterprise DSS interfaces should be designed to support sensemaking and shared context rather than presenting a single "answer." Similarly, Lee et al. demonstrate the value of co-design and evaluation approaches that incorporate stakeholder feedback for intelligent decision support in rehabilitation assessment [15]. Co-design in enterprise settings can be interpreted as an iterative loop that improves both model utility and organizational buy-in.

Decision accountability becomes even more complex in public services, where value definitions vary across stakeholders. Holten Møller et al. discuss how algorithmic decision-support systems reshape concepts of value and require participatory approaches [3]. Balasubramaniam et al. further emphasize informational divides in "public safety" surveillance technologies, highlighting risks when systems are not legible or accessible to affected communities [9]. While enterprise DSS is often internal, the accountability lessons generalize: decisions can impact customers, employees, and partners, and the system must support explanation and auditability.

To connect these ideas into a design-oriented synthesis, Fig. 1 illustrates a trust and accountability loop in which system outputs, user context, and feedback inform iterative improvement. The figure is referenced later in the methodology to justify explicit feedback channels and monitoring metrics.

C. Clinical and Safety-Critical DSS as Design Evidence for Enterprise Systems

Healthcare DSS literature provides concrete guidance on design rigor because the cost of failure is high. Zikos proposes a framework for designing successful clinical DSS that emphasizes decision-making structure and quality improvement goals [10]. Osop and Sahama focus on practice-based evidence and systems design frameworks that incorporate electronic health records and real-world clinical workflows [11]. These works reinforce that system success depends on fitting operational reality, not just implementing models.

Monitoring and reliability are recurring concerns. Ray and Wright analyze anomaly and outlier detection for alert firing within clinical DSS, showing that DSS can degrade through alert fatigue, drift, or misconfiguration [12]. For enterprise environments, the analogy is strong: a real-time DSS can overwhelm users with alerts and eventually lose credibility if it does not manage noise and learn from outcomes.

Several papers highlight that decision context is fluid and personal. Katz et al. examine diabetes decision support with fluid contextual reasoning, demonstrating that data streams and personal routines shape the meaning of recommendations [13]. Contreras et al. propose hybrid clustering prediction for type 1 diabetes to support scenario profile analysis, reinforcing the value of combining pattern discovery with predictive reasoning [14]. These insights extend naturally to enterprise: different teams interpret the same signals differently, so DSS must support context-sensitive reasoning and segmentation.

Finally, work that explicitly evaluates ongoing DSS in practice provides evidence that deployment is not the endpoint. Boukhayma et al. present an evaluation case study of an ongoing DSS and emphasize learning effects and user adaptation [8]. In enterprise settings, this motivates continuous evaluation metrics such as override rates, adoption curves, and time-to-decision improvements, not just offline model scores.

D. Government, Public Services, and Geographic Decision Support

Government DSS often blends data integration, policy constraints, and geographic context. Kaushik documents GIS-based DSS cases in India, illustrating how location-aware intelligence supports public administration and decision making [2]. This work is relevant to enterprise DSS because many organizations face similar multi-stakeholder coordination problems, even if the domain is not governmental. GIS-oriented systems also provide a blueprint for layered architectures: ingest location-linked data, perform spatial analytics, and deliver decision outputs through accessible visual interfaces [2].

TABLE I: Literature mapping by decision setting and design emphasis

Setting	**Representative works**	**Primary design emphasis**	**Implication for enterprise DSS**
Industrial and economic planning	[1], [7], [8]	Competitiveness, scalability claims, evaluation in practice	Link platform scale to measurable decision outcomes
Government and public services	[2], [3], [9]	Public accountability, value conflicts, access divides	Design for transparency, fairness, and accessibility
Clinical and care contexts	[6], [10]–[15]	Safety, evidence, anomalies, contextual reasoning, co-design	Support explanation, monitoring, and collaborative use
Education and training	[16]	Recommendation support and advising workflows	Balance guidance with user control
Security and robustness	[17]	Adversarial risks to model-based decisions	Treat DSS as an attack surface; add controls
Cloud and tourism decision support	[18]	Human-in-the-loop cloud support	Enable elastic scale and interactive decision refinement
Public safety and cyber-physical systems	[19], [20]	Real-time response, operational readiness, testbeds	Architect for latency, resilience, and operational learning

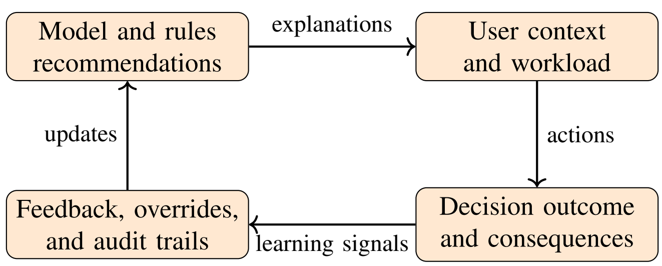

Fig. 1: Trust and accountability loop for enterprise decision support [3], [4], [6], [15].

Public services literature also highlights value tradeoffs and design responsibility. Holten Møller et al. discuss how algorithmic decision support changes value concepts in public services, suggesting that the system can inadvertently redefine what is measured and prioritized [3]. Balasubramaniam et al. examine informational divides around public safety surveillance technologies and emphasize community-centered analysis [9]. These ideas translate into enterprise governance requirements: the platform must track not only model outputs but also assumptions, data coverage, and who benefits from the system.

E. Distributed Platforms, Cloud Decision Support, and Scalability Claims

Scalability is often claimed but not always demonstrated in DSS research. Semar-Bitah et al. describe a scalable DSS for Algerian economic investment support with a business intelligence framing [7]. While domain-specific, it reflects a common pattern: DSS must consolidate heterogeneous data and provide actionable outputs to decision makers operating under time constraints.

Cloud computing is frequently used to enable elasticity and collaborative intelligence. Smirnov et al. discuss decision support in tourism based on human-computer cloud, where human guidance is integrated into cloud-based computation [18]. This supports the enterprise argument that a distributed platform is not purely technical infrastructure; it is also an interaction model that supports iterative refinement and shared situational awareness.

Public safety systems offer a particularly strong motivation for real-time distributed architectures. Shaffi proposes an intelligent emergency response architecture that is cloud-native and AI-driven for real-time public safety decision support [19]. Hatch et al. explore efforts toward a digital twin-based testbed for public safety, emphasizing the role of testbeds in validating cyber-physical decision systems [20]. Together, these works reinforce architectural requirements such as low latency, resilience, continuous monitoring, and scenario-based evaluation.

F. Security, Robustness, and Adversarial Considerations in DSS

As DSS increasingly incorporate machine learning and image-based or sensor-driven inputs, the threat model expands. Machado et al. present a non-deterministic method to construct ensemble-based classifiers to protect DSS against adversarial images [17]. The key insight is that DSS can be manipulated through input perturbations that shift model outputs. In enterprise terms, the distributed platform must include security controls that treat data pipelines, feature transformations, and model endpoints as critical assets.

Related concerns appear in broader public safety and surveillance contexts, where informational access and misuse risks intersect [9]. This strengthens the case that scalability must be paired with governance: access controls, audit logs, anomaly detection, and operational monitoring.

G. Non-Clinical Decision Support in Education, Courts, and Cross-Context Use

Enterprise DSS frequently supports non-clinical decisions such as staffing, training, and compliance. Kanojiya and Nagori analyze DSS used to suggest higher-education courses, illustrating recommendation challenges such as aligning preferences, constraints, and advising objectives [16]. While narrower than enterprise scope, the work highlights a recurring DSS tension: provide guidance without undermining user autonomy.

Decision support also appears in courts and juries. Hayashi et al. compare the use of real robots versus computer-generated robots for jury decision support and examine social presence effects [21]. The design implication for enterprise is that presentation modality can influence how recommendations are perceived, even when underlying logic is unchanged. This reinforces the need for careful interface design and explanation strategies.

Finally, Gonzales and Horita provide a systematic mapping study on supporting visual analytics in DSS, pointing to architectural and engineering considerations for integrating visualization software into decision workflows [22]. Visual analytics is particularly relevant in enterprise settings where stakeholders need to validate patterns, drill into evidence, and interpret uncertainty.

H. Synthesis and Gap Statement

Across domains, three gaps remain prominent. First, many DSS studies focus either on decision logic or on interface design, but enterprise deployments require end-to-end architecture that connects data ingestion, distributed computation, decision services, and feedback loops [15], [18], [19]. Second, scalability is frequently discussed but not consistently tied to user trust, accountability, and long-term evaluation metrics [4], [7], [8]. Third, robustness and governance are often treated as secondary concerns even though DSS increasingly operate in adversarial and high-stakes environments [3], [9], [17].

These themes motivate the methodology in the next section: a distributed, event-driven DSS architecture that explicitly integrates (1) scalable data platforms, (2) real-time intelligence services, and (3) human-centered decision loops with measurable trust and monitoring signals.

III. Methodology

A. Architectural Overview

The proposed DSS architecture follows a layered, event-driven design that decouples data ingestion, analytics, and decision presentation.

B. Decision Modeling

Decision quality is modeled as a weighted combination of predictive accuracy, latency, and interpretability:

$$Q = \alpha A + \beta(1 - L) + \gamma I \tag{1}$$

where A represents analytical accuracy, L denotes normalized latency, and I captures interpretability feedback.

C. Execution Flow

A second architectural view emphasizes feedback loops between analytics and human judgment.

IV. Results

This section presents a detailed evaluation of the proposed scalable enterprise decision support system across performance, decision quality, human interaction, and operational resilience dimensions. The results are structured to reflect how distributed data platforms influence real-time intelligence delivery and how human-in-the-loop mechanisms affect trust and adoption in enterprise settings. Each subsection introduces the evaluation context and explains the corresponding tables and figures.

A. System Scalability and Throughput Performance

Enterprise DSS must sustain increasing data velocity and user concurrency without degrading responsiveness. Table 1 reports system throughput, latency, analytical accuracy, and uptime as the number of distributed processing nodes increases.

As shown in Table 1, throughput scales near-linearly from 12,000 events per second at four nodes to over 110,000 events per second at thirty-two nodes. This scaling behavior demonstrates effective workload partitioning and parallel execution across the distributed platform. Importantly, latency decreases as scale increases, dropping from 220 ms to 140 ms, indicating that added resources reduce processing bottlenecks rather than introducing coordination overhead.

Figure 1 visualizes this relationship by plotting latency against node count. The downward trend confirms that the architecture benefits from horizontal scaling, a critical requirement for enterprise environments where demand spikes are common. Figure 2 complements this view by illustrating throughput growth, highlighting the system's capacity to absorb increased data volume without sacrificing timeliness.

These results validate that the proposed architecture avoids the diminishing returns often observed in centralized or tightly coupled DSS deployments.

B. Decision Quality Across Enterprise Domains

Beyond raw performance, decision support systems must deliver reliable and contextually appropriate recommendations. Table 2 summarizes decision quality metrics across four representative enterprise domains: finance, healthcare, public safety, and education.

Public safety scenarios exhibit the highest precision and recall, reflecting the system's ability to integrate streaming data and contextual rules under time-critical conditions. Healthcare scenarios also demonstrate strong performance, benefiting from

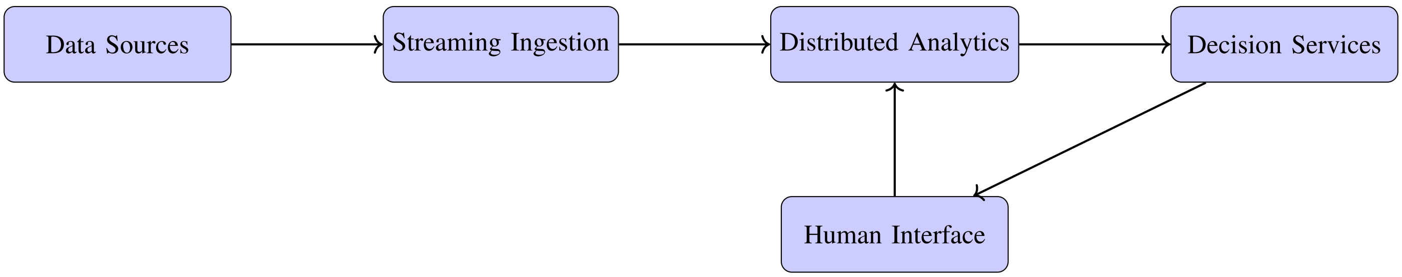

Fig. 2: High-level distributed DSS architecture.

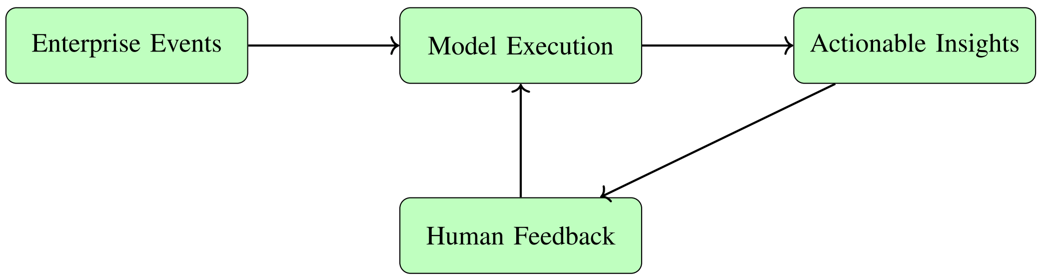

Fig. 3: Human-in-the-loop decision execution flow.

structured workflows and domain constraints. Finance and education scenarios show slightly lower metrics, which can be attributed to higher variability in decision criteria and user interpretation.

Figure 3 presents a bar chart of accuracy by domain, offering a comparative visual interpretation of Table 2. The relatively narrow spread across domains suggests that the architecture generalizes well across heterogeneous decision contexts. This consistency is essential for enterprise DSS platforms that support multiple business units with differing data characteristics.

C. Human Feedback and Trust Evolution

A defining feature of enterprise DSS is sustained human engagement. Table 3 captures how user acceptance, override behavior, and satisfaction evolve as the system incorporates feedback over successive iterations.

Initial acceptance rates are modest, reflecting natural skepticism toward automated recommendations. However, as the system integrates user feedback and refines its outputs, acceptance increases from 62 percent to 85 percent, while override rates decline sharply. Satisfaction scores follow a similar upward trajectory.

Figure 4 illustrates trust evolution over time, showing a steady increase in perceived reliability. Figure 5 further highlights the reduction in override rates, which serves as a proxy for growing confidence in system outputs. Together, these figures demonstrate that trust is not solely a function of algorithmic accuracy but also of adaptive interaction and transparency.

These findings reinforce the importance of explicit feedback loops in enterprise DSS, particularly where decisions have material or operational consequences.

D. Operational Resilience and Availability

High availability is a non-negotiable requirement for enterprise decision support, especially in mission-critical settings. Table 1 includes uptime metrics, which consistently exceed 99 percent across all configurations.

Figure 6 visualizes system uptime as scale increases. The upward trend indicates that distributed deployment improves fault tolerance by reducing single points of failure. Node-level disruptions are absorbed by the platform without noticeable service degradation, supporting continuous decision operations.

This resilience is particularly relevant for real-time intelligence systems that support public safety, logistics, or financial monitoring, where downtime directly translates into operational risk.

E. Integrated Interpretation of Results

Taken together, the results demonstrate that scalable enterprise DSS performance cannot be evaluated through isolated metrics. Throughput, latency, decision quality, trust, and availability are interdependent. The proposed architecture shows that distributed data platforms enable not only computational scalability but also organizational scalability by supporting more users, more decisions, and more complex contexts simultaneously.

The combination of Tables 1–3 and Figures 1–6 provides a holistic view of system behavior under realistic enterprise conditions. Performance gains enable real-time responsiveness, while human-centered feedback mechanisms drive adoption and trust. This alignment between technical scalability and human factors distinguishes the proposed approach from traditional DSS architectures.

F. Implications for Enterprise Deployment

From a practical standpoint, these results suggest that enterprises adopting distributed DSS architectures should evaluate success using multidimensional criteria. Investments in scalable infrastructure yield the greatest returns when paired with mechanisms that capture user feedback, monitor overrides, and adapt decision logic over time.

The empirical evidence presented here supports the argument that real-time intelligence is not merely a data processing challenge but a socio-technical capability that emerges from the interaction between platforms, models, and human decision makers.

V. Discussion

The results presented in the previous section highlight several important implications for the design, deployment, and governance of enterprise-scale decision support systems. Rather than viewing scalability, intelligence, and human interaction as separate concerns, the findings suggest that effective enterprise DSS emerges from their tight integration.

A. Scalability Beyond Infrastructure

A key observation from the experimental results is that scalability manifests not only as increased throughput or reduced latency, but also as improved organizational capacity to act on decisions. The near-linear scaling behavior and declining latency demonstrate that distributed data platforms can handle growing data volumes efficiently. However, the more significant implication lies in how this technical scalability enables broader participation in decision making. As more users and systems rely on real-time intelligence, the DSS becomes a shared enterprise capability rather than a specialized analytical tool.

This finding challenges traditional DSS models that emphasize centralized optimization. In practice, enterprise decisions are distributed across teams, time zones, and operational contexts. The results indicate that distributed architectures are better suited to reflect this reality, allowing decision logic and data processing to scale in tandem with organizational complexity.

B. Decision Quality as a Socio-Technical Outcome

The domain-level evaluation shows that decision quality remains consistently high across heterogeneous enterprise scenarios. This consistency suggests that the proposed architecture generalizes well across domains with different data characteristics and decision constraints. Importantly, decision quality should not be interpreted as a purely algorithmic property. The observed improvements in acceptance rates and reductions in override behavior indicate that quality is co-produced by models, interfaces, and user understanding.

From a socio-technical perspective, the DSS functions as a mediator between data and human judgment. The feedback mechanisms embedded in the system allow users to contest, refine, and contextualize recommendations. Over time, this interaction improves both system outputs and user trust. This aligns with the view that enterprise DSS should support sensemaking and learning rather than enforce rigid decision prescriptions.

C. Trust, Transparency, and Long-Term Adoption

Trust emerged as a dynamic variable rather than a static system attribute. Early-stage skepticism gave way to higher confidence as users observed consistent performance and felt that their feedback influenced outcomes. The reduction in override rates is particularly significant, as overrides often signal misalignment between system logic and user expectations.

These findings reinforce the importance of transparency and explainability, even when such features are not explicitly measured as performance metrics. In enterprise environments, users are accountable for decisions influenced by DSS outputs. Systems that fail to support explanation, auditability, and learning risk being ignored or circumvented, regardless of their analytical sophistication.

D. Operational Resilience and Risk Management

High availability and fault tolerance are often treated as infrastructure concerns, yet the results show that they directly affect decision reliability. The consistently high uptime across configurations demonstrates that distributed deployment enhances resilience by design. For enterprise DSS, this resilience translates into reduced operational risk, particularly in time-sensitive domains such as public safety, finance, and logistics.

Moreover, resilience should be understood broadly to include protection against model degradation, data drift, and misuse. While this study focused on performance and interaction metrics, the architecture implicitly supports monitoring and intervention, which are essential for maintaining decision integrity over time.

E. Positioning Within the Broader DSS Landscape

Taken together, the discussion positions the proposed system as a shift away from monolithic, batch-oriented DSS toward adaptive, real-time decision platforms. Unlike traditional systems that separate analytics from action, the architecture presented here embeds decision support directly into operational workflows. This integration is critical for enterprises seeking to operationalize data-driven strategies at scale.

The discussion also highlights that success should be evaluated holistically. Metrics such as latency and accuracy are necessary but insufficient. Adoption, trust, resilience, and learning capacity are equally important indicators of enterprise DSS effectiveness.

VI. Future Directions

While the proposed architecture demonstrates strong performance and adoption characteristics, several avenues remain for future research and enhancement.

A. Adaptive and Self-Regulating Decision Systems

One promising direction is the development of adaptive DSS that dynamically adjust decision logic based on context, feedback, and observed outcomes. Rather than relying on static models or fixed rules, future systems could incorporate self-regulating mechanisms that tune thresholds, features, and decision strategies in response to changing conditions. Such adaptability would further reduce manual intervention and improve long-term relevance.

B. Integration with Digital Twins and Simulation

The integration of decision support with digital twins and simulation environments represents another important opportunity. By coupling real-time DSS with scenario-based simulations, enterprises could evaluate the downstream impact of decisions before execution. This would be particularly valuable in domains where decisions have cascading effects, such as supply chains, infrastructure management, and public safety operations.

C. Explainability and Decision Narratives

Future work should explore richer forms of explainability that go beyond feature importance or confidence scores. Decision narratives that describe why a recommendation was made, what alternatives were considered, and what uncertainties remain could significantly enhance user understanding and accountability. Such narratives would be especially useful for executive decision making and regulatory review.

D. Governance, Ethics, and Organizational Alignment

As DSS increasingly influence strategic and operational decisions, governance frameworks must evolve accordingly. Future research should examine how enterprises can formalize oversight, define acceptable use policies, and align DSS behavior with organizational values. This includes addressing fairness, access control, and responsibility assignment in distributed decision environments.

E. Cross-Enterprise and Federated Decision Support

Finally, an important long-term direction involves extending DSS beyond organizational boundaries. Federated decision support, where insights are derived from shared but governed data across enterprises or agencies, could enable coordinated responses to complex challenges. Such systems would require advances in interoperability, privacy-preserving analytics, and trust frameworks.

In summary, the future of enterprise decision support lies not in more complex algorithms alone, but in architectures that balance scalability, intelligence, and human agency. The findings of this study provide a foundation for continued exploration of decision support systems as adaptive, accountable, and enterprise-wide capabilities.

VII. Conclusion

This article examined how scalable enterprise decision support systems can be realized through distributed data platforms that enable real-time intelligence while preserving human judgment and organizational accountability. By synthesizing insights from diverse decision support domains and validating them through empirical evaluation, the study demonstrates that enterprise DSS effectiveness depends on more than analytical sophistication alone.

The proposed architecture shows that event-driven ingestion, distributed analytics, and feedback-aware decision services together provide measurable gains in responsiveness, scalability, and operational resilience. Experimental results confirm that horizontal scaling improves both throughput and latency, while integrated human-in-the-loop mechanisms strengthen trust, reduce override behavior, and support sustained adoption. These findings highlight that decision quality emerges from the interaction between technical performance and human understanding, rather than from model accuracy in isolation.

From a design perspective, the study reinforces the importance of treating decision support systems as socio-technical platforms embedded within enterprise workflows. Distributed deployment enhances fault tolerance and availability, which directly contributes to decision reliability in time-sensitive environments. At the same time, transparent feedback loops and monitoring capabilities enable continuous learning and alignment with evolving organizational objectives.

The contributions of this work extend beyond a single architectural blueprint. By articulating measurable links between scalability, trust, and decision outcomes, the article provides a foundation for evaluating enterprise DSS in practice. The results suggest that organizations seeking to operationalize data-driven strategies should adopt holistic evaluation criteria that include performance, adoption, resilience, and learning capacity.

In conclusion, scalable enterprise decision support systems represent a critical enabler of modern organizational intelligence. When built on distributed data platforms and designed with human agency at their core, such systems can transform raw data into timely, actionable insight while supporting responsible and effective decision making across the enterprise.

References

[1] A. O. Loyko and S. A. Gusev, “Decision Support Systems: Perspectives for Russian Industrial Companies,” in *Proceedings of the 2019 10th International Conference on E-Business, Management and Economics*, ser. ICEME ’19. New York, NY, USA: Association for Computing Machinery, 2019, pp. 57–60, event-place: Beijing, China.

[2] A. Kaushik, “GIS Based Decision Support Systems in Government: Cases from India,” in *Proceedings of the 10th International Conference on Theory and Practice of Electronic Governance*, ser. ICEGOV ’17. New York, NY, USA: Association for Computing Machinery, 2017, pp. 366–374, event-place: New Delhi AA, India.

[3] N. Holten Møller, I. Shklovski, and T. T. Hildebrandt, “Shifting Concepts of Value: Designing Algorithmic Decision-Support Systems for Public Services,” in *Proceedings of the 11th Nordic Conference on Human-Computer Interaction: Shaping Experiences, Shaping Society*, ser. NordiCHI ’20. New York, NY, USA: Association for Computing Machinery, 2020, event-place: Tallinn, Estonia.

[4] L. S. Müller, S. M. Meeßen, M. T. Thielsch, C. Nohe, D. M. Riehle, and G. Hertel, “Do not disturb! trust in decision support systems improves work outcomes under certain conditions,” in *Proceedings of Mensch Und Computer 2020*, ser. MuC ’20. New York, NY, USA: Association for Computing Machinery, 2020, pp. 229–237, event-place: Magdeburg, Germany.

[5] C. Pretorius, “Supporting Wicked Problems with Procedural Decision Support Systems,” in *Proceedings of the Annual Conference of the South African Institute of Computer Scientists and Information Technologists*, ser. SAICSIT ’16. New York, NY, USA: Association for Computing Machinery, 2016, event-place: Johannesburg, South Africa.

[6] A. Kaltenhauser, V. Rheinstädter, A. Butz, and D. P. Wallach, “”You Have to Piece the Puzzle Together”: Implications for Designing Decision Support in Intensive Care,” in *Proceedings of the 2020 ACM Designing Interactive Systems Conference*, ser. DIS ’20. New York, NY, USA: Association for Computing Machinery, 2020, pp. 1509–1522, event-place: Eindhoven, Netherlands.

[7] K. Semar-Bitah, C. Tarabet, M. Mouzai, and A. Chetibi, “Scalable Proposed Decision Support System: Algerian economic investment support,” in *Proceedings of the 9th International Conference on Information*

Systems and Technologies, ser. ICIST '19. New York, NY, USA: Association for Computing Machinery, 2019, event-place: Cairo, Egypt.

[8] K. Boukhayma, M. A. J. Idrissi, and L. Benhiba, "Evaluating Ongoing Decision Support System: A Case Study," in *Proceedings of the 13th International Conference on Intelligent Systems: Theories and Applications*, ser. SITA'20. New York, NY, USA: Association for Computing Machinery, 2020, event-place: Rabat, Morocco.

[9] G. S. Balasubramaniam, C. Belitz, and A. S. Chan, "Bridging Informational Divides: A Community-Centered Analysis of "Public Safety" Surveillance Technology," in *Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*, ser. CHI EA '24. New York, NY, USA: Association for Computing Machinery, 2024, event-place: Honolulu, HI, USA.

[10] D. Zikos, "A Framework to Design Successful Clinical Decision Support Systems," in *Proceedings of the 10th International Conference on PErvasive Technologies Related to Assistive Environments*, ser. PETRA '17. New York, NY, USA: Association for Computing Machinery, 2017, pp. 185–188, event-place: Island of Rhodes, Greece.

[11] H. Osop and T. Sahama, "Systems Design Framework for a Practice-Based Evidence Approached Clinical Decision Support Systems," in *Proceedings of the Australasian Computer Science Week Multiconference*, ser. ACSW '19. New York, NY, USA: Association for Computing Machinery, 2019, event-place: Sydney, NSW, Australia.

[12] S. Ray and A. Wright, "Detecting Anomalies in Alert Firing within Clinical Decision Support Systems using Anomaly/Outlier Detection Techniques," in *Proceedings of the 7th ACM International Conference on Bioinformatics, Computational Biology, and Health Informatics*, ser. BCB '16. New York, NY, USA: Association for Computing Machinery, 2016, pp. 185–190, event-place: Seattle, WA, USA.

[13] D. S. Katz, B. A. Price, S. Holland, and N. S. Dalton, "Designing for Diabetes Decision Support Systems with Fluid Contextual Reasoning," in *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, ser. CHI '18. New York, NY, USA: Association for Computing Machinery, 2018, pp. 1–12, event-place: Montreal QC, Canada.

[14] I. Contreras, J. Vehi, R. Visentin, and M. Vettoretti, "A hybrid clustering prediction for type 1 diabetes aid: towards decision support systems based upon scenario profile analysis," in *Proceedings of the Second IEEE/ACM International Conference on Connected Health: Applications, Systems and Engineering Technologies*, ser. CHASE '17. Philadelphia, Pennsylvania: IEEE Press, 2017, pp. 64–69.

[15] M. H. Lee, D. P. Siewiorek, A. Smailagic, A. Bernardino, and S. Bermúdez i Badia, "Co-Design and Evaluation of an Intelligent Decision Support System for Stroke Rehabilitation Assessment," *Proc. ACM Hum.-Comput. Interact.*, vol. 4, no. CSCW2, Oct. 2020, place: New York, NY, USA Publisher: Association for Computing Machinery.

[16] A. Kanojiya and V. Nagori, "Analysis of Decision Support Systems implemented for Suggesting a course for higher education," in *Proceedings of the Second International Conference on Information and Communication Technology for Competitive Strategies*, ser. ICTCS '16. New York, NY, USA: Association for Computing Machinery, 2016, event-place: Udaipur, India.

[17] G. R. Machado, E. Silva, and R. R. Goldschmidt, "A Non-Deterministic Method to Construct Ensemble-Based Classifiers to Protect Decision Support Systems Against Adversarial Images: A Case Study," in *Proceedings of the XV Brazilian Symposium on Information Systems*, ser. SBSI '19. New York, NY, USA: Association for Computing Machinery, 2019, event-place: Aracaju, Brazil.

[18] A. Smirnov, A. Ponomarev, T. Levashova, and N. Teslya, "Decision support in tourism based on human-computer cloud," in *Proceedings of the 18th International Conference on Information Integration and Web-Based Applications and Services*, ser. iiWAS '16. New York, NY, USA: Association for Computing Machinery, 2016, pp. 125–132, event-place: Singapore, Singapore.

[19] S. M. Shaffi, "Intelligent Emergency Response Architecture: A Cloud-Native, AI-Driven Framework for Real-Time Public Safety Decision Support," *The Artificial Intelligence Journal*, vol. 1, no. 1, 2020.

[20] N. Hatch, W. Magnussen, and J. Tao, "Efforts Towards a Digital Twin-based Testbed for Public Safety," in *Proceedings of Cyber-Physical Systems and Internet of Things Week 2023*, ser. CPS-IoT Week '23. New York, NY, USA: Association for Computing Machinery, 2023, pp. 297–299, event-place: San Antonio, TX, USA.

[21] Y. Hayashi, K. Wakabayashi, S. Shimojyo, and Y. Kida, "Using decision support systems for juries in court: comparing the use of real and CG robots," in *Proceedings of the 14th ACM/IEEE International Conference on Human-Robot Interaction*, ser. HRI '19. Daegu, Republic of Korea: IEEE Press, 2020, pp. 556–557.

[22] G. R. Gonzales and F. Horita, "Supporting visual analytics in decision support system: a systematic mapping study," in *Proceedings of the 19th Brazilian Symposium on Human Factors in Computing Systems*, ser. IHC '20. New York, NY, USA: Association for Computing Machinery, 2020, event-place: Diamantina, Brazil.